Websense Manager Help

Websense® Web Security
Websense Web Filter

v7

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

Published 2015
Printed in the United States and Ireland

Trademarks

Contents

1 Getting Started

Websense software gives network administrators in all sectors of industry, from business to education to government and beyond, the ability to control or monitor network traffic to the Internet.

- Minimize employee downtime spent accessing Internet data deemed objectionable, inappropriate, or not work-related.
- Minimize misuse of network resources and the threat of legal action due to inappropriate access.
- Add a solid layer of security to your network, protecting it from potential spyware, malware, hacking, and other intrusions.

From here, you can find information about:

Basic Websense configuration	Implementing Internet filtering
• *Working in Websense Manager*, page 14 • *Your subscription*, page 25 • *The Websense Master Database*, page 28 • *Verifying Network Agent configuration*, page 319	• *Filtering categories and protocols*, page 34 • *Adding a client*, page 61 • *Working with policies*, page 67 • *Assigning a policy to clients*, page 70

You can also learn how to:

Evaluate your configuration	Refine filtering policies
• *Today: Health, Security, and Value Since Midnight*, page 19 • *History: Last 30 Days*, page 22 • *Presentation reports*, page 88 • *Investigative reports*, page 105 • *Using the Toolbox to verify filtering behavior*, page 180	• *Creating a custom category*, page 163 • *Redefining filtering for specific sites*, page 166 • *Restricting users to a defined list of Internet sites*, page 153 • *Filtering based on keyword*, page 164 • *Managing traffic based on file type*, page 176 • *Using Bandwidth Optimizer to manage bandwidth*, page 174

Overview

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

Together, a series of Websense components (described in *Websense product components*, page 246) provide Internet filtering, user identification, alerting, reporting, and troubleshooting capabilities.

An overview of the new features included in this Websense software version can be found in the Release Notes, available from the Websense Support Portal.

After installation, Websense software applies the **Default** policy to monitor Internet usage without blocking requests. This policy governs Internet access for all clients in the network until you define your own policies and assign them to clients. Even after you have created your custom filtering settings, the Default policy applies any time a client is not governed by any other policy. See *The Default policy*, page 66, for more information.

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

- *Internet Usage Filters*, page 33
- *Clients*, page 53
- *Internet Filtering Policies*, page 65

A single, browser-based tool—Websense Manager—provides a central, graphical interface to the general configuration, policy management, and reporting functions of your Websense software. See *Working in Websense Manager*, page 14, for more information.

You can define levels of access to Websense Manager to allow certain administrators to manage only a specific group of clients, or to allow individuals to run reports on their own Internet usage. See *Delegated Administration*, page 215, for more information.

Working in Websense Manager

Related topics:

- *Logging on to Websense Manager*, page 16
- *Navigating in Websense Manager*, page 17
- *Today: Health, Security, and Value Since Midnight*, page 19
- *History: Last 30 Days*, page 22

Websense Manager is the central configuration interface used to customize filtering behavior, monitor Internet usage, generate Internet usage reports, and manage Websense software configuration and settings. This Web-based tool runs on 2 supported browsers:

- Microsoft Internet Explorer 7
- Mozilla Firefox 2

Although it is possible to launch Websense Manager using some other browsers, use the supported browsers to receive full functionality and proper display of the application.

To launch Websense Manager, do one of the following:

- On Windows machines:
  - Go to **Start > All Programs > Websense**, and then select **Websense Manager**.
  - Double-click the Websense Manager desktop icon.
- Open a supported browser on any machine in your network and enter the following:

  ```
  https://<IP address>:9443/mng
  ```

  Replace *<IP address>* with the IP address of the Websense Manager machine.

If you are unable to connect to Websense Manager on the default port, refer to the **tomcat.log** file on the Websense Manager machine (located by default in the **C:\Program Files\Websense\tomcat\logs** or **/opt/Websense/tomcat/logs/** directory) to verify the port.

If you are using the correct port, and are still unable to connect to Websense Manager from a remote machine, make sure that your firewall allows communication on that port.

An SSL connection is used for secure, browser-based communication with Websense Manager. This connection uses a security certificate issued by Websense, Inc. Because the supported browsers do not recognize Websense, Inc., as a known Certificate Authority, a certificate error is displayed the first time you launch Websense Manager from a new browser. To avoid seeing this error, you can install or permanently accept the certificate within the browser. See the Websense Knowledge Base for instructions.

Once the security certificate has been accepted, the Websense Manager logon page is displayed in the browser window (see *Logging on to Websense Manager*).

Logging on to Websense Manager

Related topics:

- *Working in Websense Manager*
- *Navigating in Websense Manager*, page 17
- *Today: Health, Security, and Value Since Midnight*, page 19
- *History: Last 30 Days*, page 22

After installation, the first user to log on to Websense Manager has full administrative access. The user name is **WebsenseAdministrator**, and cannot be changed. The WebsenseAdministrator password is configured during installation.

To log on, first launch Websense Manager (see *Working in Websense Manager*). At the logon page:

1. Select a **Policy Server** to manage.

   If your environment includes only one Policy Server, it is selected by default.
2. Select an **Account Type**:
   - To log on using a Websense user account, such as WebsenseAdministrator, click **Websense account** (default).
   - To log on using your network credentials, click **Network account**.
3. Enter a **User name** and **Password**, and then click **Log On**.

You are logged on to Websense Manager.

- If this is your first time logging on to Websense Manager, you are offered the option of launching a Quick Start tutorial. If you are new to Websense software, or new to this version of Websense software, completing a Quick Start tutorial is highly recommended.
- If you are using Delegated Administration, and have created administrative roles, you may be prompted to select a role to manage. See *Delegated Administration*, page 215, for more information.

A Websense 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 session end.

- If there are uncached changes on the page or cached changes pending, the changes are lost when the session ends. Remember to click **OK** to cache and **Save All** to save and implement any changes.
- If Websense Manager 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 Websense Manager is open in multiple browser windows on the same computer, the instances share the same session **if**:

- You are using Microsoft Internet Explorer and use the Ctrl-N shortcut to open a new instance of Websense Manager.
- You are using Mozilla Firefox.

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

- If you launch multiple Internet Explorer windows independently of one another, and then use them to log on as different Websense Manager administrators, the windows do **not** share a session. If one window times out, the others are not affected.

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

Navigating in Websense Manager

The Websense Manager interface can be divided into 4 main areas:

1. Websense banner
2. Left navigation pane
3. Right shortcut pane
4. Content pane

The **Websense banner** shows:

- Which **Policy Server** you are currently logged on to (see *Working with Policy Server*, page 251)

- Your current administrative **Role** (see *Introducing administrative roles*, page 216)
- A **Log Off** button, for when you're ready to end your administrative session

The content displayed in Websense Manager varies based on the privileges granted to the logged on user. A user with reporting-only privileges, for example, is not shown server configuration settings or policy administration tools. See *Delegated Administration*, page 215, for more information.

This section describes the options available to WebsenseAdministrator, and other users with Super Administrator privileges.

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

The **right shortcut pane** contains links to useful tools and common administrative tasks. This is also where you can review and save any changes that you have made in Websense Manager.

- The top portion of the navigation pane indicates whether there are cached changes waiting to be saved. When you are working in Websense Manager, the Changes bar indicates whether or not there are **Changes Pending**.

  In most cases, when you perform a task in Websense Manager and click **OK**, your changes are cached. (Sometimes you must click OK both on a subordinate page and a main page to cache changes.)

  After caching changes, click **Save All** to save and implement the changes. To view cached changes before saving (see *Reviewing, saving, and discarding changes*, page 18), click the **View Pending Changes** button. This is the smaller button to the left of Save All.
- **Common Tasks** provides shortcuts to frequently-performed administrative tasks. Click an item in the list to jump to the page where the task is performed.
- The **Toolbox** contains quick lookup tools that you can use to verify your filtering setup. See *Using the Toolbox to verify filtering behavior*, page 180, for more information.

Reviewing, saving, and discarding changes

When you perform a task in Websense Manager, and then click **OK**, your changes are cached. Use the **View Pending Changes** page to review cached changes.

Important

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

Changes to a single area of functionality are typically grouped into a single entry in the cache list. For example, if you add 6 clients and delete 2 clients, the cache list

indicates only that changes were made to Clients. Changes to a single Settings page, on the other hand, may result in multiple entries in the cache list. This occurs when a single Settings page is used to configure multiple Websense software functions.

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

After choosing Save All or Cancel All, the Changes bar in the right shortcut pane is updated appropriately, and you are returned to the last page you selected. There is no undo for either the Save All or Cancel All functions.

Use the Audit Log to review the details of changes made in Websense Manager. See *Viewing and exporting the audit log*, page 256, for more information.

Today: Health, Security, and Value Since Midnight

Related topics:

- *Navigating in Websense Manager*, page 17
- *History: Last 30 Days*, page 22
- *Customizing the Today page*, page 21
- *Alerting*, page 259]

The **Status > Today: Health, Security and Value Since Midnight** page appears first when you log on to Websense Manager. It presents the current status of your filtering software and graphically illustrates Internet filtering activity for up to 24-hours, beginning at 12:01 a.m. according to the time on the Log Database machine.

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

- The **Health Alert Summary** shows the status of your Websense software. If an error or warning appears in the summary, click the alert message to open the Alerts page, where more detailed information is available (see *Reviewing current system status*, page 266).

  Information in the Health Alert Summary is updated every 30 seconds.
- Under **Today's Value**, see examples of how Websense filtering has protected your network today, as well as the total number of Internet requests handled, and other important activity totals.

Below the summary information, up to 4 charts provide information about filtering activities. These charts are available to Super Administrators, and to delegated administrators who are granted permission to view reports on the Today page. See *Editing roles*, page 233.

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

Chart Name	Description
Current Filtering Load	See the number of filtered Internet traffic processed into the Log Database, shown in 10-minute intervals.
Top Security Risks by Requests	Find out which Security Risk categories have received the most requests today, and determine whether filtering policies are providing the right protection for your network.
Top Categories by Requests	See the categories that are being accessed most today. Get a high-level overview of potential security, bandwidth, or productivity concerns.
Policy Enforcement by Risk Class	See how many requests to each risk class have been permitted and blocked today (see *Risk classes*, page 37). Evaluate whether the current policies are effective or whether changes are needed.
Top Protocols by Bandwidth	Learn which protocols are using the most bandwidth in your network today. Use this information to evaluate bandwidth needs, and the potential need for policy changes.
Computers Requesting Security Risk Sites	Find out which computers have accessed Security Risk sites today. You may want to check these machines to make sure they are not infected with any viruses or spyware.
Top Blocked Users	See which users have requested the most blocked sites today, gaining insight into compliance with the organization's Internet use standards.
Top Uncategorized Sites	Learn which sites not categorized by the Websense Master Database have been accessed most today. Go to **Common Tasks > Recategorize a URL** to assign a site to a category for filtering.

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

Three buttons appear above the page:

- **Database Download**, available to Super Administrators only, opens the page for viewing the status of Master Database downloads, or initiating a download (see *Review Master Database download status*, page 255).
- **Customize**, available to Super Administrators only, opens a page where you can change which charts appear on the page (see *Customizing the Today page*, page 21).
- **Print**, available to all administrators, opens a secondary window with a printable version of the charts displayed on the Today page. Use browser options to print this page, which omits all the navigation options found in the main Websense Manager window.

Below the Internet activity and filtering charts, the **Filtering Service Summary** shows the status of each Filtering Service associated with the current Policy Server. Click the Filtering Service IP address to see more information about that Filtering Service instance.

For security purposes, a Websense Manager session ends after 30 minutes of inactivity. You can, however, choose to continue to monitor filtering and alerting data: mark **Continue monitoring Today, History, and Alerts status without timing out** at the bottom of the Today page. Information on these 3 pages continues to update normally until you close the browser or navigate to another Websense Manager page.

Important

If you enable the monitoring option and stay within the Today, History, and Alerts pages for more than 30 minutes, an attempt to navigate to another Websense Manager page returns you to the Logon page.

When you enable this option, be sure to save any cached changes before the 30 minute timeout period ends.

Customizing the Today page

Related topics:

- *Today: Health, Security, and Value Since Midnight*, page 19
- *Customize the History page*, page 24

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

The charts that you select appear on the Today page for all Super Administrators, and for delegated administrators who have permission to view charts on the Today page. See *Editing roles*, page 233.

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

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

For a short description of the information displayed in each chart, see *Today: Health, Security, and Value Since Midnight*, page 19.

History: Last 30 Days

Related topics:

- *Today: Health, Security, and Value Since Midnight*, page 19
- *Navigating in Websense Manager*, page 17
- *Customize the History page*, page 24

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

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

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

Mouse over the **Time** or **Bandwidth** item (under Saved) for an explanation of how the estimate was calculated (see *Time and bandwidth saved*, page 23). You can click **Customize** to change the way the values are calculated.

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

Depending on the reporting permissions granted to the role, delegated administrators may not see the charts described below. See *Editing roles*, page 233.

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

Chart Name	Description
Internet Activity by Requests	Review the number of filtered Internet requests processed into the Log Database each day.
Top Security Risks by Requests	See which Security Risk categories have been accessed recently, and determine whether filtering policies are providing the right protection for your network.
Top Categories by Requests	See which categories have been accessed most. Get a high level overview of potential security, bandwidth, or productivity concerns.

Chart Name	Description
Top Uncategorized Sites	Learn which sites, not categorized by the Websense Master Database, have been accessed most. Go to **Common Tasks > Recategorize a URL** to assign a site to a category for filtering.
Top Protocols by Bandwidth	Learn which protocols are using the most bandwidth in your network recently. Use this information to evaluate bandwidth needs and potential policy changes.
Policy Enforcement by Risk Class	See how many requests to each risk class have been permitted and blocked recently (see *Risk classes*, page 37). Evaluate whether the current policies are effective or whether changes are needed.
Top Blocked Users	See which users' Internet requests have been blocked the most. Gain insight into compliance with your organization's Internet use standards.
Policy Enforcement Summary	Get an overview of recently permitted requests, blocked requests to sites in the Security Risk class, and blocked requests to other sites. Consider which aspects of filtering need a more detailed evaluation.

Two buttons appear above the page:

- **Customize**, available to Super Administrators only, opens a page where you can change which charts appear on the page, and to change how estimated savings are calculated (see *Customize the History page*, page 24).
- **Print**, available to all administrators, opens a secondary window with a printable version of the charts displayed on the History page. Use browser options to print this page, which omits all the navigation options found in the main Websense Manager window.

Time and bandwidth saved

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

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

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

See *Customize the History page*, page 24, for information on how to change the values used in these calculations to reflect usage at your organization.

Customize the History page

Related topics:

- *History: Last 30 Days*, page 22
- *Customizing the Today page*, page 21

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

Mark the check box next to each chart name, up to 4, that you want to include on the History page. For a short description of each chart, see *History: Last 30 Days*, page 22. Only Super Administrators with unconditional policy permissions (including WebsenseAdministrator) can customize the charts on the History page.

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

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

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

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

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

Your subscription

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

When you purchase a subscription, a subscription key is provided via email. Each key is valid for one installation of Websense Policy Server. If you install multiple Policy Servers, you need a separate key for each.

Before you can begin filtering, you must enter a valid subscription key (see *Configuring your account information*, page 27). This lets you download the Master Database (see *The Websense Master Database*, page 28), which enables Websense software to filter clients.

After the first successful database download, Websense Manager displays the number of clients your subscription includes.

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

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

To configure filtering behavior when a subscription is exceeded or expires, see *Configuring your account information*, page 27.

To configure Websense software to send email warnings when the subscription approaches or exceeds its limit, see *Configuring system alerts*, page 262.

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

Managing your account through the MyWebsense Portal

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

When you create an account, you are prompted to enter all Websense subscription keys. This helps to ensure your access to information, alerts, and patches relevant to your Websense product and version.

Once you have a MyWebsense account, if you are ever unable to log on to Websense Manager due to a lost WebsenseAdministrator password, just click **Forgot my password** on the Websense Manager logon page. You are prompted to log on to

MyWebsense, and then given instructions for generating and activating a new password.

Important

When you request a new password, the subscription key that you select in the MyWebsense portal must match the key entered on the Account page in Websense Manager.

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

To access the MyWebsense portal from within Websense Manager, go to **Help > MyWebsense**.

Activating Websense Web Protection Services™

Websense Web Security subscriptions include access to Websense Web Protection Services: SiteWatcher™, BrandWatcher™, and ThreatWatcher™. Once you activate these services, they work to protect your organization's Web sites, brands, and Web servers.

Service	Description
SiteWatcher	Alerts you when your organization's Web sites have been infected with malicious code, allowing you to take immediate action to protect customers, prospects, and partners who might visit the site.
BrandWatcher	• Alerts you when your organization's Web sites or brands have been targeted in phishing or malicious keylogging attacks. • Provides Internet security intelligence, attack details, and other security-related information so that you can take action, notify customers, and minimizing any public relations impact.
ThreatWatcher	• Offers a hacker's-eye view of your organization's Web server, scanning for known vulnerabilities and potential threats. • Reports risk levels and provides recommendations through a Web-based portal. • Helps prevent malicious attacks on your Web servers before they happen.

Log on to the MyWebsense portal to activate Websense Protection Services. Once ThreatWatcher is activated, log on to MyWebsense to access threat reports for registered Web servers.

Configuring your account information

Related topics:

- *Your subscription*, page 25
- *Configuring database downloads*, page 29
- *Working with protocols*, page 168

Use the **Settings > Account** page to enter and view subscription information, and change the WebsenseAdministrator password, used to access Websense Manager. WebsenseAdministrator is the default, master administrative account used to manage Websense software.

This is also where you can enable Websense software to send protocol usage data to Websense, Inc., anonymously. This information may be used to update to the Websense Master Database, a collection of more than 36 million Internet sites and more than 100 protocol definitions (see *The Websense Master Database*, page 28, for more information).

1. After installing Websense software, or any time you receive a new subscription key, use the **Subscription key** field to enter the key.

   After you enter a new subscription key and click OK, a Master Database download begins automatically.
2. After the first Master Database download, the following information appears:

Key expires	End date for your current subscription. After this date, you must renew the subscription to continue downloading the Master Database and filtering your network.
Subscribed network users	Number of in-network users that can be filtered.
Subscribed remote users	Number of users that can be filtered outside the network (requires the optional Remote Filtering feature).

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

   If this option is not selected, users have unfiltered Internet access in these situations.
4. To change the WebsenseAdministrator password, first provide the current password, and then enter and confirm the new password.
   - The password must be from 4 to 25 characters long. It is case sensitive, and can include letters, numbers, special characters, and spaces.

- It is a good idea to create a strong password for the WebsenseAdministrator account. The password should be at least 8 characters long and include at least one capital letter, lowercase letter, number, and special character.

5. Mark **Send category and protocol data to Websense, Inc.** to have Websense software collect usage data about Websense-defined categories and protocols, and submit it anonymously to Websense, Inc.

   This usage data helps Websense, Inc., to continually enhance the filtering capabilities of Websense software.

The Websense Master Database

Related topics:

- *Real-time database updates*, page 29
- *Real-Time Security Updates™*, page 29
- *Filtering categories and protocols*, page 34
- *Working with Filtering Service*, page 255
- *Review Master Database download status*, page 255
- *Resumable Master Database downloads*, page 256

The Websense Master Database houses the category and protocol definitions that provide the basis for filtering Internet content (see *Filtering categories and protocols*, page 34).

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

A limited version of the filtering database is installed during Websense software installation, but it is a good idea to download the full Master Database as soon as possible to enable comprehensive Internet filtering capabilities. To download the Master Database for the first time, enter your subscription key on the **Settings > Account** page (see *Configuring your account information*, page 27).

If Websense software must go through a proxy to perform the download, also use the **Settings > Database Download** page to configure proxy settings (see *Configuring database downloads*, page 29).

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

After the initial download, Websense software downloads database changes on a schedule that you establish (see *Configuring database downloads*, page 29). Because

the Master Database is updated frequently, by default, database downloads are scheduled to happen daily.

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

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

Real-time database updates

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

Websense software checks for database updates every hour.

The most recent updates are listed on the **Status > Alerts** page (see *Reviewing current system status*, page 266).

Real-Time Security Updates™

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

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

Filtering Service checks for security updates every 5 minutes, but because updates are sent out only when security threats occur, actual changes are occasional, and tend not to disrupt normal network activity.

Use the **Settings > Database Download** page to enable Real-Time Security Updates (see *Configuring database downloads*, page 29).

Configuring database downloads

Related topics:

- *Configuring your account information*, page 27
- *The Websense Master Database*, page 28
- *Review Master Database download status*, page 255

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

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

   You must download the Master Database at least once every 14 days for Websense software to continue filtering uninterrupted. If you deselect all download days, Websense software automatically attempts a download when the database is 7 days old.

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

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

   Websense software selects a random time during this period to contact the Master Database server. To configure alerts for download failures, see *Configuring system alerts*, page 262.

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

3. (*Websense Web Security*) Select **Enable real-time security updates** to have Websense software check for security updates to the Master Database every 5 minutes. When a security update is detected, it is downloaded immediately.

   Real-time security updates rapidly protect your network from vulnerability to threats like new phishing (identity fraud) scams, rogue applications, and malicious code infecting a mainstream Web site or application.

4. Select **Use proxy server or firewall** if Websense software must access the Internet through a proxy server or a proxying firewall (other than the integration product that Websense software communicates with) to download the Master Database. Then, configure the following.

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

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

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

   By default, the user name and password are encoded to match the character set for the Policy Server machine's locale. This encoding can be configured manually via the **Settings > Directory Services** page (see *Advanced directory settings*, page 59).

Testing your network configuration

In order for Internet request filtering to occur, Websense software must be aware of Internet traffic to and from machines in your network. Use the Network Traffic Detector to ensure this Internet communication is visible to the filtering software. See *Verifying Network Agent configuration*, page 319, for instructions.

If the Traffic Detector is not able to see all segments of your network, see *Network Configuration*, page 311, for configuration instructions.

Websense Technical Support

Websense, Inc., is committed to customer satisfaction. Go to the Websense Technical Support Web site any time for the latest release information, to access the Knowledge Base or product documentation, or to create a support request.

www.websense.com/SupportPortal/

The response time for online requests during business hours is approximately 4 hours. Response to after-hours requests occurs the next business day.

Telephone assistance is also available. For quick and efficient answers to telephone requests, please be ready with the following:

- Websense subscription key
- Access to Websense Manager
- Access to the machines running Filtering Service and Log Server, and the database server (Microsoft SQL Server or MSDE)
- Permission to access the Websense Log Database
- Familiarity with your network's architecture, or access to a person who has this knowledge

- Specifications of the machines running Filtering Service and Websense Manager
- A list of other applications running on the Filtering Service machine

For severe problems, additional information may be needed.

Standard telephone assistance is available during normal business hours Monday through Friday at the following numbers:

- San Diego, California, USA: **+1 858.458.2940**
- London, England: **+44 (0) 1932 796244**

Check the Support Web site listed above for operating hours and other support options.

Customers in Japan should contact their distributor for the most rapid service.

2 Internet Usage Filters

Related topics:

- *Filtering categories and protocols*, page 34
- *Working with filters*, page 43
- *Configuring Websense filtering settings*, page 50
- *Internet Filtering Policies*, page 65
- *Refine Filtering Policies*, page 153

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

- **Category filters**, used to apply actions (permit, block) to Web site categories
- **Protocol filters**, used to apply actions to Internet applications and non-HTTP protocols
- A schedule that determines when each filter is enforced

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

In a first-time installation, Websense software creates a **Default** policy and uses it to begin monitoring Internet requests as soon as a subscription key is entered (see *The Default policy*, page 66). Initially, the Default policy permits all requests.

Note

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

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

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

- A second category filter that permits all Web sites except those that represent a security risk and those containing adult material
- A third category filter that monitors access to Web sites without blocking them (see *Creating a category filter*, page 44)

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 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 *Internet Filtering Policies*, page 65).

Filtering categories and protocols

The Websense Master Database organizes similar Web sites (identified by URLs and IP addresses) into **categories**. Each category has a descriptive name, like Adult Material, Gambling, or Peer-to-Peer File Sharing. You can also create your own, custom categories to group sites of particular interest to your organization (see *Creating a custom category*, page 163). Together, the Master Database categories and user-defined categories form the basis for Internet filtering.

Websense, Inc., does not make value judgments about categories or sites in the Master Database. Categories are designed to create useful groupings of the sites of concern to subscribing customers. They are not intended to characterize any site or group of sites or the persons or interests who publish them, and they should not be construed as such. Likewise, the labels attached to Websense categories are convenient shorthand and are not intended to convey, nor should they be construed as conveying, any opinion or attitude, approving or otherwise, toward the subject matter or the sites so classified.

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

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

To suggest that a site be added to the Master Database, click **Suggest New Category** in the right shortcut pane of Websense Manager, or go to:

www.websense.com/SupportPortal/SiteLookup.aspx

After logging on to the MyWebsense portal, you are taken to the Site Lookup and Category Suggestion tool.

When you create a **category filter** in Websense Manager, 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 AOL Instant Messenger or BitTorrent). The definitions are verified and updated as frequently as nightly.

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

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

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

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

> **Note**
> Network Agent must be installed to enable protocol-based filtering.

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

New categories and protocols

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

- The default action is applied in all active category and protocol filters (see *Working with filters*, page 43). Edit the active filters to change the way the category or protocol is filtered.
- The default action is based on feedback regarding whether or not the sites or protocols in question are generally considered business-appropriate.

You can configure Websense software to generate a system alert and notify you whenever new categories or protocols are added to the Master Database. See *Alerting*, page 259, for more information.

Special categories

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

- The **Special Events** category is used to classify sites considered hot topics to help you manage event-related surges in Internet traffic. For example, the official

World Cup site might generally appear in the Sports category, but be moved to the Special Events category during the World Cup Finals.

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

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

Websense Web Security includes additional security categories:

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

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

When Websense security research detects a potential threat, it is added to the Extended Protection category until Websense is 100% confident of the site's final categorization.

Risk classes

Related topics:

- *Assigning categories to risk classes*, page 278
- *Presentation reports*, page 88
- *Investigative reports*, page 105

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

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

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

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

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

Legal Liability

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

Bandwidth > Peer-to-Peer File Sharing

Gambling

Illegal or Questionable

Information Technology > Hacking and Proxy Avoidance

Militancy and Extremist

Racism and Hate

Tasteless

Violence

Weapons

Network Bandwidth Loss

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

Entertainment > MP3 and Audio Download Services

Productivity > Advertisements and Freeware *and* Software Download

Business Usage

Business and Economy (including Financial Data and Services)

Education > Educational Materials *and* Reference Materials

Government (including Military)

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

Travel

Vehicles

Security Risk

Bandwidth > Peer-to-Peer File Sharing

Extended Protection (including Elevated Exposure, Emerging Exploits, and Potentially Damaging Content) [*Websense Web Security*]

Information Technology > Hacking *and* Proxy Avoidance

Productivity > Freeware and Software Download

Security (including Bot Networks, Keyloggers, Malicious Web Sites, Phishing and Other Frauds, Potentially Unwanted Software, and Spyware)

Productivity Loss

Abortion (including Pro-Choice *and* Pro-Life)

Adult Material > Sex Education

Advocacy Groups

Bandwidth > Internet Radio and TV, Peer-to-Peer File Sharing, and Streaming Media

Drugs (including Abused Drugs, Marijuana, Prescribed Medications, and Supplements and Unregulated Compounds)

Education (including Cultural Institutions *and* Educational Institutions)

Entertainment (including MP3 and Audio Download Services)

Gambling

Games

Government > Political Organizations

Health

Information Technology > Web Hosting

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

Job Search

News and Media (including Alternative Journals)

Productivity Loss

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

Religion (including Non-Traditional Religions and Occult and Folklore *and* Traditional Religions)

Shopping (including Internet Auctions *and* Real Estate)

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

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

Special Events

Sports (including Sport Hunting and Gun Clubs)

Travel

Vehicles

Super Administrators can change the categories assigned to each risk class on the **Settings > Risk Class** page (see *Assigning categories to risk classes*, page 278).

Security protocol groups

In addition to the Security and Extended Protection categories, Websense Web Security includes two protocols 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 - Monitor Only** 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 Malicious Traffic** 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 - Monitor Only protocols can be logged for reporting, but no other filtering action can be applied.

Instant Messaging Attachment Manager

The Instant Messaging (IM) Attachment Manager is an optional feature. If you subscribe to this feature, you can restrict file sharing with IM clients including AOL/

ICQ, Microsoft (MSN), and Yahoo. This lets you permit IM traffic while blocking the transfer of attachments by IM clients.

Instant Messaging File Attachments is a protocol group that includes definitions for multiple IM clients. When the IM Attachment Manager is enabled, these protocols appear in the protocol list in all active protocol filters, and on the Manage Protocols page.

IM attachment filtering can be applied to both internal and external traffic. To enable internal traffic filtering, define the portion of your network to monitor on the **Settings > Network Agent > Global Settings** page (see *Configuring global settings*, page 314).

Filtering actions

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

- **Block** the request. Users receive a block page or block message, and are not able to view the site or use the Internet application.
- **Permit** the request. Users can view the site or use the Internet application.
- Evaluate current **Bandwidth** usage before blocking or permitting the request. When this action is enabled, and bandwidth usage reaches a specified threshold, further Internet requests for a specific category or protocol are blocked. See *Using Bandwidth Optimizer to manage bandwidth*, page 174.

Additional actions can be applied only to categories.

Note

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

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

- **Confirm**—Users receive a block page, asking them to confirm that the site is being accessed for business purposes. If a user clicks **Continue**, she can view the site.

  Clicking Continue starts a timer. During the configured time period (60 seconds by default), the user can visit other sites in Confirm categories without receiving another block page. Once the time period ends, browsing to any other Confirm site results in another block page.

  The default time can be changed on the **Settings > 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 > 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.
- **Block Keywords**—When you define keywords and enable keyword blocking, users requesting a site whose URL contains a blocked keyword are not allowed to access the site. See *Filtering based on keyword*, page 164.
- **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 176.

Using quota time to limit Internet access

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

Note

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

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

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

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

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

If a user is a member of multiple groups, Websense software grants quota time according to the **Use more restrictive blocking** setting on the **Settings > Filtering** page (see *Configuring Websense 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, Websense software detects the communication to the host server, and then blocks the request according to applicable quota restrictions.

Password override

Password override lets users with valid passwords access sites blocked by Websense software. Password override can be granted to individual clients (users, groups, computers, or networks).

When the password override option is enabled, Websense block messages include a password field. Clients who enter a valid password can access blocked sites for a limited amount of time.

Note

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

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

The password override option is enabled via the **Settings > Filtering** page (see *Configuring Websense filtering settings*, page 50).

Grant password override privileges to specific clients via the **Policy Management > Clients** page (see *Adding a client*, page 61, or *Changing client settings*, page 63).

Search filtering

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

Ordinarily, Internet search engine results may include thumbnail images associated with sites matching the search criteria. If those thumbnails are associated with blocked sites, Websense software prevents users from accessing the full site, but does not prevent the search engine from displaying the image.

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

Websense, Inc., maintains a database of search engines with Search Filtering capabilities. When a search engine is added to or removed from the database, an alert is generated (see *Alerting*, page 259).

Search Filtering is activated via the **Settings > Filtering** page. See *Configuring Websense filtering settings*, page 50, for more information.

Working with filters

Related topics:

- *Filtering categories and protocols*, page 34
- *Internet Filtering Policies*, page 65
- *Creating a category filter*, page 44
- *Creating a protocol filter*, page 46
- *Creating a limited access filter*, page 155

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

The Filters page is divided into 3 main sections:

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

  Network Agent must be installed to enable protocol-based filtering.
- **Limited Access Filters** define a restrictive list of permitted Web sites (see *Restricting users to a defined list of Internet sites*, page 153).

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 45
  - *Editing a protocol filter*, page 47
  - *Editing a limited access filter*, page 156

- To create a new category, protocol, or limited access filter, click **Add**. For more information, see:
  - *Creating a category filter*, page 44
  - *Creating a protocol filter*, page 46
  - *Creating a limited access filter*, page 155

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*, page 215), 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 158, for more information.

Creating a category filter

Related topics:

- *Working with filters*, page 43
- *Editing a category filter*, page 45

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

Related topics:

- *Filtering categories and protocols*, page 34
- *Filtering actions*, page 40
- *Using quota time to limit Internet access*, page 41
- *Password override*, page 42
- *Working with filters*, page 43
- *Working with categories*, page 160

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

Important

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

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

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

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

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

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

1. Select an entry in the **Categories** list to view category information or to change the filtering action associated with the selected category.
2. Before making changes to the action applied to a category, use the **Category Details** section to review any special attributes associated with the category.
   - To review recategorized or unfiltered URLs assigned to the category, if any, click **See custom URLs in this category**. See *Redefining filtering for specific sites*, page 166.
   - To review keywords assigned to the category, click **See keywords in this category**. See *Filtering based on keyword*, page 164.

- To review regular expressions used to define custom URLs or keywords for the category, click **See regular expressions in this category**.

3. Use the buttons at the bottom of the category list to change the action applied to the selected category. For more information about the available actions, see *Filtering actions*, page 40.

   Delegated administrators cannot change the action associated with categories that have been locked by a Super Administrator. See *Defining filtering restrictions for all roles*, page 242, for more information.

4. Use the check boxes to the right of the Categories list to apply advanced filtering actions to the selected category:

   - To change the way that keywords are used in filtering the selected category, mark or clear **Block keywords**. *Filtering based on keyword*, page 164
   - 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 176.

     If you have chosen to block file types, select one or more file types to block.
   - To specify whether access to sites in the category is limited based on certain bandwidth thresholds, mark or clear **Block with Bandwidth Optimizer**. See *Using Bandwidth Optimizer to manage bandwidth*, page 174.

     If you have chosen to block based on bandwidth, specify which threshold limits to use.

5. Repeat steps 1 through 3 to make changes to the filtering actions applied to other categories.
6. After editing the filter, click **OK** to cache your changes and return to the Filters page. Changes are not implemented until you click **Save All**.

To activate a new category filter, add it to a policy and assign the policy to clients. See *Internet Filtering Policies*, page 65.

Creating a protocol filter

Related topics:

- *Filtering categories and protocols*, page 34
- *Filtering actions*, page 40
- *Editing a protocol filter*, page 47
- *Working with protocols*, page 168

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

Related topics:

- *Filtering categories and protocols*, page 34
- *Creating a protocol filter*, page 46
- *Filtering actions*, page 40
- *Working with protocols*, page 168
- *Using Bandwidth Optimizer to manage bandwidth*, page 174

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

Important

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

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

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

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

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

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

To change the way that protocols are filtered and logged:

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

   Note

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

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

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

   To apply the same action to the other protocols in the selected protocol group, click **Apply to Group**.
3. If you want information about use of the selected protocol available for alerting or reporting, mark the **Log protocol data** check box.
4. To impose bandwidth limits on the use of this protocol, click **Block with Bandwidth Optimizer**, and then supply the bandwidth thresholds to use. See *Using Bandwidth Optimizer to manage bandwidth*, page 174, for more information.
5. After editing the filter, click **OK** to cache your changes and return to the Filters page. Changes are not implemented until you click **Save All**.

To activate a new protocol filter, add it to a policy and apply the policy to clients (see *Internet Filtering Policies*, page 65).

Note

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

Websense-defined category and protocol filters

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

The predefined category filters are:

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

The Block All and Permit All category filters are not listed on the Filters page, though they can be added to policies. These filters play a special role in filtering, and cannot be deleted or edited. When an Internet request is filtered, Websense software first checks to see if the Block All or Permit All filter applies, before performing any additional filtering checks (see *Filtering a site*, page 72).

The predefined protocol filters are:

- Basic Security
- Default
- Monitor Only
- Permit All

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

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

Category and protocol filter templates

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

Websense software includes 5 category filter templates:

- **Monitor Only** and **Permit All** permits all categories.
- **Block All** blocks all categories.
- **Basic** blocks the most frequently blocked categories and permits the rest.
- **Default** applies the Block, Permit, Continue, and Quota actions to categories.
- **Basic Security** blocks only the default categories in the Security Risk class (see *Risk classes*, page 37).

Websense software also includes 3 protocol filter templates:

- **Monitor Only** and **Permit All** permit all protocols.

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

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

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

For instructions on using a template to create a new filter, see *Creating a category filter*, page 44, or *Creating a protocol filter*, page 46.

Configuring Websense filtering settings

Related Topics:

- *Filtering categories and protocols*, page 34
- *Clients*, page 53
- *Block Pages*, page 77
- *Filtering actions*, page 40
- *Password override*, page 42
- *Filtering order*, page 71
- *Using Bandwidth Optimizer to manage bandwidth*, page 174
- *Filtering based on keyword*, page 164

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

Under **Bandwidth Optimizer**, enter the information needed to filter Internet usage based on available bandwidth. For more information about bandwidth-based filtering, see *Using Bandwidth Optimizer to manage bandwidth*, page 174.

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. Use the **Default bandwidth for network** field to enter a default threshold (percentage of total network traffic) to use when network bandwidth filtering is enabled.
3. Use the **Default bandwidth per protocol** field to enter a default threshold to use when protocol bandwidth filtering is enabled.

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

1. To determine how users are filtered when multiple group policies apply, mark or clear **Use most restrictive group policy** (see *Filtering order*, page 71).
   - When the option is selected, the policy that applies the most restrictive filtering setting is applied. In other words, if one applicable group policy blocks access to a category and another permits access, the user's request for a site in that category is blocked.
   - When the option is not selected, the most permissive setting is used.
2. Select one of the following **Keyword search options** (see *Filtering based on keyword*, page 164).

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

3. In the **Password override timeout** field, enter the maximum number of seconds (up to 3600, default 60) that a user can access sites in all categories after selecting password override (see *Password override*, page 42).
4. In the **Continue timeout** field, enter the maximum time in seconds (up to 3600, default 60) that a user who clicks Continue can access sites in categories governed by the Confirm action (see *Filtering actions*, page 40).
5. 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.

6. Enter the **Default quota time per day** (up to 240 minutes, default 60) for all users.

   To change the quota time for individual users, go to the **Policies > Clients** page.

   As you make changes to the quota session length and the default quota time per day, the **Default quota sessions per day** is calculated and displayed.

Use the **Block Messages** section to enter the URL or path to the alternative HTML block page you created for the top frame of browser-based block messages (see *Creating alternate block messages*, page 83).

- Separate pages can be used for the different protocols: **FTP**, **HTTP** (including **HTTPS**), and **Gopher**.
- Leave these fields blank to use the default block message provided with the Websense software, or a customized version of that message (see *Customizing the block message*, page 79).

Under **Search Filtering**, select **Enable search filtering** to have Websense software activate a setting built into certain search engines so thumbnail images and other explicit content associated with blocked sites are not displayed in search results (see *Search filtering*, page 42).

The search engines for which this feature is supported are displayed at the bottom of the section.

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

3 Clients

You can customize how Websense software filters requests from specific users or machines by adding them as **clients** in Websense Manager. Clients can be:

- **Computers**: Individual machines in your network, defined by IP address.
- **Networks**: Groups of computers, defined collectively as an IP address range.
- **Users**: User, group, or domain accounts in a supported directory service.

Initially, Websense software filters all clients in the same manner, using the **Default** policy (see *The Default policy*, page 66). Once you add a client to the Clients page in Websense Manager, you can assign that client a specific filtering policy.

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

1. Apply the policy assigned to the **user** making the request. If that policy has no filters scheduled at the time of the request, use the next applicable policy.
2. If there is no user-specific policy, or the policy has no active filters at the time of the request, look for a policy assigned to the **computer** (first) or **network** (second) from which the request was made.
3. If there is no computer or network-specific policy, or the policy has no active filters at the time of the request, look for a policy assigned to any **group** to which the user belongs. If the user belongs to multiple groups, Websense software considers all group policies that apply (see *Filtering order*, page 71).
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 information about how Websense software applies filtering policies to clients, see *Filtering a site*, page 72.

Working with clients

Related topics:

- *Clients*, page 53
- *Working with computers and networks*, page 55
- *Working with users and groups*, page 56
- *Adding a client*, page 61
- *Changing client settings*, page 63

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, you must add clients in your managed clients list to see them on the Clients page. See *Adding a client*, page 61, for instructions.

Clients are divided into 3 groups:

- **Directory**, which includes users, groups, and domains from your directory service (see *Working with users and groups*, page 56).
- **Networks**, IP address ranges within the filtered network that can be governed by a single policy (see *Working with computers and networks*, page 55).
- **Computers**, individual machines in the filtered network, identified by IP address (see *Working with computers and networks*, page 55).

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

- The client name, IP address, or IP address range.
- The **policy** currently assigned to the client. The **Default** policy is used until you assign another policy (see *Internet Filtering Policies*, page 65).
- Whether or not the client can use a **password override** option to view blocked sites (see *Password override*, page 42).
- 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 63, 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 61, for more information.

If you have created delegated administration roles (see *Delegated Administration*, page 215), 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 63, for more information.

If you have configured Websense software to communicate with an LDAP-based directory service, the **Manage Custom LDAP Groups** button appears in the toolbar at the top of the page. Click this button to add or edit groups based on an LDAP attribute (see *Working with custom LDAP groups*, page 60).

To remove a client from Websense Manager, select the client and click **Delete**.

Working with computers and networks

Related topics:

- *Working with clients*, page 54
- *Working with users and groups*, page 56
- *Adding a client*, page 61
- *Assigning a policy to clients*, page 70

In Websense Manager, a **computer** is the IP address (for example, 10.201.3.1) associated with a filtered machine. A **network** is the IP address range (for example, 10.201.3.2 - 10.201.3.44) associated with a group of filtered machines.

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

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

When you assign a policy to a computer or network, that policy is enforced regardless of who is logged on to the filtered machine, **unless** you have assigned a policy to the logged-on user. This computer or network policy takes precedence over any **group** policies that may apply to the user.

Working with users and groups

Related topics:

- *Working with clients*, page 54
- *Directory services*, page 56
- *Working with custom LDAP groups*, page 60
- *Working with computers and networks*, page 55
- *Adding a client*, page 61
- *Assigning a policy to clients*, page 70

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

Websense software can communicate with Windows NT Directory / Active Directory (Mixed Mode), and with Windows Active Directory, Novell eDirectory, and Sun Java System Directory accessed via Lightweight Directory Access Protocol (LDAP).

Note

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.

Also, if you are using Windows Active Directory or Sun Java System Directory, user names with blank passwords are not supported. Make sure that all users have passwords assigned.

The Websense User Service conveys information from the directory service to Policy Server and Filtering Service for use in applying filtering policies.

Websense, Inc., recommends installing User Service on a Windows machine (though it can reside on a Linux machine). Typically, this is the machine where Policy Server is installed.

To configure Websense software to communicate with your directory service, see *Directory services*.

Directory services

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

Use the **Settings > 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.

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

See the appropriate section for configuration instructions:

- *Windows NT Directory / Active Directory (Mixed Mode)*, page 57
- *Windows Active Directory (Native Mode)*, page 57
- *Novell eDirectory and Sun Java System Directory*, page 58

Windows NT Directory / Active Directory (Mixed Mode)

If your directory service is Windows NT Directory or Active Directory in Mixed Mode, no further configuration is necessary.

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

- DC Agent is being used for transparent identification (see *DC Agent*, page 193)

  and
- User Service runs on a Linux machine

If this matches your configuration, provide the administrative credentials listed under Windows NT Directory / Active Directory (Mixed Mode). If your installation does not use this configuration, the administrative credential fields are disabled.

Windows Active Directory (Native Mode)

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

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

1. Click **Add**, next to the Global catalog servers list. The Add Global Catalog Server page appears.
2. Use the **Server IP or name** field to identify the global catalog server:
   - If you have multiple global catalog servers configured for failover, enter the DNS domain name.
   - If your global catalog servers are not configured for failover, enter the IP address or host name (if name resolution is enabled in your network) of the server to add.
3. Enter the **Port** that Websense software should use to communicate with the global catalog (by default, **3268**).

4. Optionally, enter the **Root context** that Websense software should use to search for user information. If you supply a value, it must be a valid context in your domain.
   - If you have specified a communications port of 3268 or 3269, you do not need to supply a root context.
   - If the specified port is 389 or 636, you must provide a root context.
   - If the Root context field is left blank, Websense software begins searching at the top level of the directory service.

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

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

   Select **Distinguished name by components** or **Full distinguished name** to specify how you prefer to enter the account information.
   - If you selected Distinguished name by components, enter the **Display name**, account **Password**, **Account folder**, and **DNS domain name** for the administrative account. Use the common name (cn) form of the administrative user name, and not the user ID (uid) form.

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

   - If you selected Full distinguished name, enter the distinguished name as a single string in the **User distinguished name** field (for example, *cn=Admin, cn=Users, ou=InfoSystems, dc=company, dc=net*), and then supply the **Password** for that account.
6. Click **OK**.
7. Repeat the process above for each global catalog server.
8. Click **Advanced Directory Settings**, and then go to *Advanced directory settings*, page 59.

Novell eDirectory and Sun Java System Directory

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

1. Enter the IP address of the directory server machine in the **Server IP** field.
2. Enter the **Port** number that Websense software will use to communicate with the directory. The default is 389.
3. If your directory requires administrator privileges for read-only access, enter the **Administrator distinguished name** and **Password**.
4. Optionally, enter the **Root Context** that Websense software should use when searching for user information. For example, *o=domain.com*.

   Narrowing the context increases speed and efficiency in retrieving user information.

   Note

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

5. Click **Advanced Directory Settings**, and then go to *Advanced directory settings*, page 59.

Advanced directory settings

Related topics:

- *Windows Active Directory (Native Mode)*, page 57
- *Novell eDirectory and Sun Java System Directory*, page 58

These settings can be used to define:

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

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

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

   - **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 Groups Search Filter** determines how User Service associates users with groups.

3. To secure communications between Websense software and your directory service, check **Use SSL**.
4. To determine which character set Websense software uses to encode LDAP information, select **UTF-8** or **MBCS**.

   MBCS, or multibyte character set, is commonly used for encoding East Asian languages such as Chinese, Japanese, and Korean.
5. Click **OK** to cache your changes. Changes are not implemented until you click **Save All**.

Working with custom LDAP groups

Related topics:

- *Working with users and groups*, page 56
- *Directory services*, page 56
- *Adding or editing a custom LDAP group*, page 61

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

Important

When you add custom LDAP groups to Websense Manager, the group definitions are stored by the active Policy Server, and do not affect other Policy Server instances. To add custom LDAP groups to multiple Policy Servers, use Websense Manager to log on to each Policy Server and enter the information.

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

- To add a group, click **Add** (see *Adding or editing a custom LDAP group*, page 61).
- To change an entry in the list, click on its group name (see *Adding or editing a custom LDAP group*).
- To remove an entry, first select it, and then click **Delete**.

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

Adding or editing a custom LDAP group

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

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

1. Enter or change the **Group name**. Use a descriptive name that clearly indicates the purpose of the LDAP group.

   Group names are case-insensitive, and must be unique.
2. Enter or change the description that defines this group in your directory service. For example:

   ```
   (WorkStatus=parttime)
   ```

   In this example, **WorkStatus** is a user attribute that indicates employment status, and **parttime** is a value indicating that the user is a part-time employee.
3. Click **OK** to return to the Manage Custom LDAP Groups page. The new or revised entry appears in the list.
4. Add or edit another entry, or click **OK** to cache changes and return to the previous page. Changes are not implemented until you click **Save All**.

Adding a client

Related topics:

- *Working with clients*, page 54
- *Working with computers and networks*, page 55
- *Working with users and groups*, page 56
- *Searching the directory service*, page 62
- *Changing client settings*, page 63

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

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

1. Identify one or more clients:
   - To add a user, group, or domain client, browse the **Directory** tree to find entries in your directory service. If you are using an LDAP-based directory service, you can also click **Search** to enable a directory search tool (see *Searching the directory service*, page 62).
   - To add a computer or network client, enter an **IP address** or **IP address range**. No two network definitions can overlap, but a network client can include an IP address identified separately as a computer client. In the case of such an overlap, the policy assigned to the computer takes precedence over the policy assigned to the network.
2. Click an arrow button (>) to add each client to the **Selected Clients** list.

   To remove an entry from the Selected Clients list, select the client, and then click **Remove**.
3. Select a **Policy** to assign to all clients in the Selected Clients list.
4. When you are finished, click **OK** to cache your changes. Changes are not implemented until you click **Save All**.

The clients are added to the appropriate list on the **Policy Management > Clients** page. To change the policy assigned to one or more clients, or to configure additional client settings, select each client entry, and then click **Edit**. See *Changing client settings*, page 63, for more information.

Searching the directory service

If you have configured Websense software to communicate with an LDAP-based directory service, you can use a search function to identify users to be added as clients in Websense Manager. Search is also available for adding managed clients and administrators to delegated administration roles.

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

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

   In a large directory service, selecting **All** may cause the search to take a very long time.
4. Browse the **Search context** tree to specify which portion of the directory to search. A more precise context helps to speed the search.
5. Click **Go**.

   A list of search results is displayed.

6. Select one or more entry in search results, and then click the right arrow (>) to add each selection as a client or administrator.
7. Click **New Search** to enter another set of search criteria.
8. Click **Browse** to return to browsing the directory.
9. When you are finished making changes, click **OK** to cache your changes. Changes are not implemented until you click **Save All**.

Changing client settings

Use the **Policy Management > Clients > Edit Clients** 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 Clients 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. To allow users to override a Websense block page by entering a password, click **On** under Password Override, and then enter and confirm the password.

   To remove a client's password override privileges, click **Off**.
3. To allocate a custom amount of **Quota Time** to the selected clients, click **Custom**, and then enter the number of minutes of quota time to assign.

   To revert to the default quota settings, click **Default**.
4. Click **OK** to cache your changes and return to the Clients page. Changes are not implemented until you click **Save All**.

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

Moving clients to roles

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

- The policy applied to the client in the Super Administrators role and the filters that it enforces are copied to the delegated administration role.
- Delegated administrators can change the policies applied to their managed clients.
- Filter Lock restrictions do not affect clients managed by Super Administrators, but do affect managed clients in delegated administration roles.
- If a group, domain, or organizational unit is added to a role as a managed client, delegated administrators in that role can assign policies to individual users in the group, domain, or organizational unit.

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

To move the selected clients to a delegated administration role:

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

   A pop-up dialog box indicates that the selected clients are being moved. The move process may take a while.
3. Changes are not implemented until you click **Save All**.

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

4 Internet Filtering Policies

Related topics:

- *Internet Usage Filters*, page 33
- *Clients*, page 53
- *The Default policy*, page 66
- *Working with policies*, page 67
- *Filtering order*, page 71

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

- Category filters, used to apply actions (permit, block) to Web site categories (see *Filtering categories and protocols*, page 34)
- Limited access filters, used to permit access to only a restricted list of Web sites (see *Restricting users to a defined list of Internet sites*, page 153)
- Protocol filters, used to apply actions to Internet protocols (see *Filtering categories and protocols*, page 34)
- A schedule that determines when each category or limited access filter and protocol filter is enforced

A new Websense software installation includes 3 predefined policies:

- **Default** filters Internet access for all clients not governed by another policy. Websense software begins enforcing this policy as soon as a subscription key is entered (see *The Default policy*, page 66).
- **Unrestricted** provides unlimited access to the Internet. This policy is not applied to any clients by default.
- **Example - Standard User** shows how multiple category and protocol filters can be applied in a policy to provide different degrees of filtering restriction at different times. This policy is used in the New User Quick Start tutorial to demonstrate the process of editing a policy and applying it to clients.

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

The Default policy

Related topics:

- *Internet Filtering Policies*, page 65
- *Working with policies*, page 67
- *Filtering order*, page 71

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

Note

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

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

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

Important

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

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

Working with policies

Related topics:

- *Internet Filtering Policies*, page 65
- *Creating a policy*
- *Editing a policy*
- *Internet Usage Filters*
- *Refine Filtering Policies*

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 68, for more information.
- To edit a policy, click the policy name in the list, and then see *Editing a policy*, page 68, for more information.
- To see which clients are filtered by the policy, click a number in the Users, Networks, or Computers column. The client information appears in a pop-up dialog box.

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

If you have created delegated administration roles (see *Delegated Administration*, page 215), 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, the Permit All filters and policies that enforce them cannot be copied to roles.

To copy policies to another role, first mark the check box next to the policy name, and then click **Copy to Role**. See *Copying filters and policies to roles*, page 158, for more information.

Creating a policy

Related topics:

- *Internet Filtering Policies*, page 65
- *Working with policies*, page 67
- *Editing a policy*, page 68
- *Working with filters*, page 43
- *Restricting users to a defined list of Internet sites*, page 153

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 68.

Editing a policy

Related topics:

- *Internet Filtering Policies*, page 65
- *Working with policies*, page 67
- *Creating a policy*, page 68
- *Working with filters*, page 43
- *Restricting users to a defined list of Internet sites*, page 153

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 (user, 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 70.

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

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

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

   The **Example - Standard User** policy, included with your Websense software, demonstrates how to define a filtering period that spans midnight.
3. Use the **Days** column to define which days of the week are included in this time block. To select days from a list, click the down arrow in the right portion of the column. When you are finished selecting days, click the up arrow.
4. Use the **Category / Limited Access Filter** column to select a filter to enforce during this time block.

   To add a new filter to enforce in this policy, select **Create category filter** or **Create limited access filter**. See *Creating a category filter*, page 44, or *Creating a limited access filter*, page 155, 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	• Use the **Permit**, **Block**, **Confirm**, or **Quota** button to change the action applied to the selected categories (see *Filtering 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	• Use the **Add Sites** and **Add Expressions** button to add permitted URLs, IP addresses, or regular expressions to the filter (see *Restricting users to a defined list of Internet sites*, page 153). • 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	• Use the **Permit** or **Block** button to change the action applied to the selected protocols (see *Filtering 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 All**.

Assigning a policy to clients

Related topics:

- *Internet Filtering Policies*, page 65
- *Creating a policy*, page 68
- *Editing a policy*, page 68
- *Clients*, page 53
- *Adding a client*, page 61

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 user, computer, and network clients, as well as the policy currently assigned to each client.

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

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

Filtering order

Websense software uses multiple filters, applied in a specific order, to determine whether to permit, block, or limit requested Internet data.

For each request it receives, Websense software:

1. Verifies subscription compliance, making sure that the subscription is current and the number of subscribed clients has not been exceeded.
2. Determines which policy applies, searching in this order:
   a. Policy assigned to the **user**.
   b. Policy assigned to the **IP address** (computer or network) of the machine being used.
   c. Policies assigned to **groups** the user belongs to.
   d. Policies assigned to the user's **domain**.
   e. The **Default** policy.

   The first applicable policy found is used.
3. Filters the request according to the policy's restrictions.

In some cases, a user belongs to more than one group or domain, and no user, computer, or network policy applies. In these cases, Websense software checks the policies assigned to each of the user's groups.

- If all the groups have the same policy, Websense software filters the request according to that policy.
- If one of the groups has a different policy, Websense software filters the request according to the **Use more restrictive blocking** selection on the **Settings > Filtering** page.

  If **Use more restrictive blocking** is checked, and any of the applicable policies blocks access to the requested category, Websense software blocks the site.

  If the option is not checked, and any of the applicable policies permits access to the requested category, Websense software permits the site.

If one of the applicable policies enforces a limited access filter, the **Use more restrictive blocking** option can have different effects than expected. See *Limited access filters and filtering precedence*, page 154.

Filtering a site

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

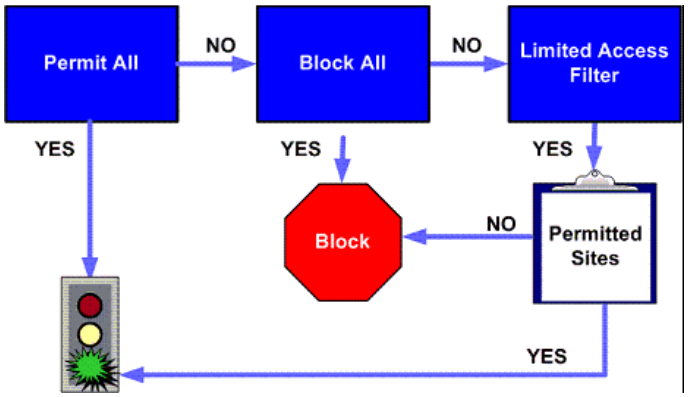

1. Determines which **category filter** or **limited access filter** the policy enforces for the current day and time.
   - If the active category filter is **Permit All**, permit the site.
   - If the active category filter is **Block All**, block the site.
   - If the filter is a **limited access filter**,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

- If any other category filter applies, continue to Step 2.

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

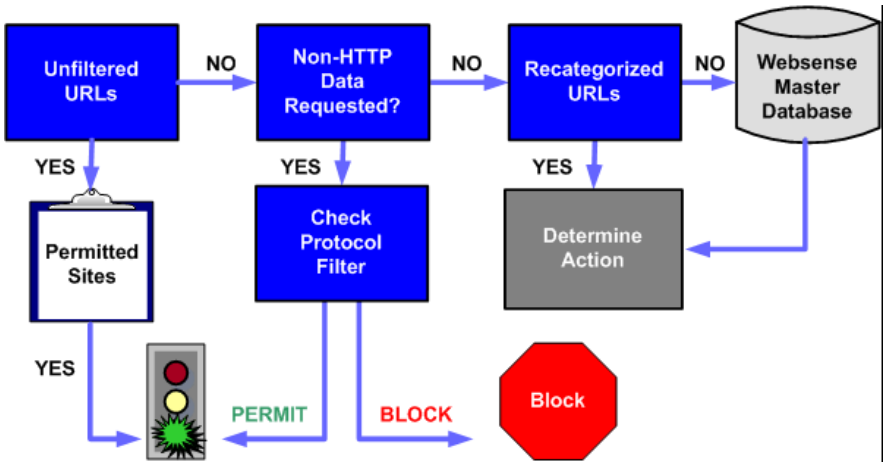

2. Tries to match the site to an entry in the **Unfiltered URLs** list.
   - If the URL appears on the list, permit the site.
   - If the URL does not appear on the list, continue to Step 3.
3. Checks the active **protocol filter** and determines whether any non-HTTP protocols are associated with the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

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

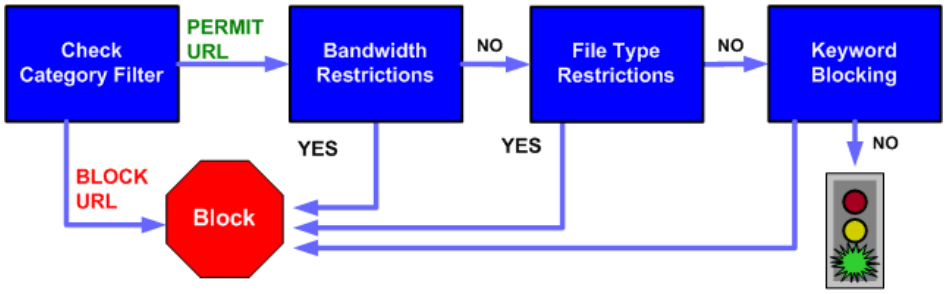

6. Checks the active category filter and identifies the action applied to the category containing the requested site.
   - If the action is **Block**, block the site.
   - If any other action is applied, continue to Step 7.
7. Checks for **Bandwidth Optimizer** settings in the active category filter (see *Using Bandwidth Optimizer to manage bandwidth*, page 174).
   - 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 176).
   - If the site contains files whose extensions are blocked, block access to those files. If the site itself is comprised of a blocked file type, block access to the site.
   - If the site does not contain files whose extensions are blocked, go to Step 9.
9. Checks for blocked **keywords** in the URL and CGI path, if keyword blocking is enabled (see *Filtering based on keyword*, page 164).
   - If a blocked keyword is found, block the site.

- If a blocked keyword is not found, continue to Step 10.

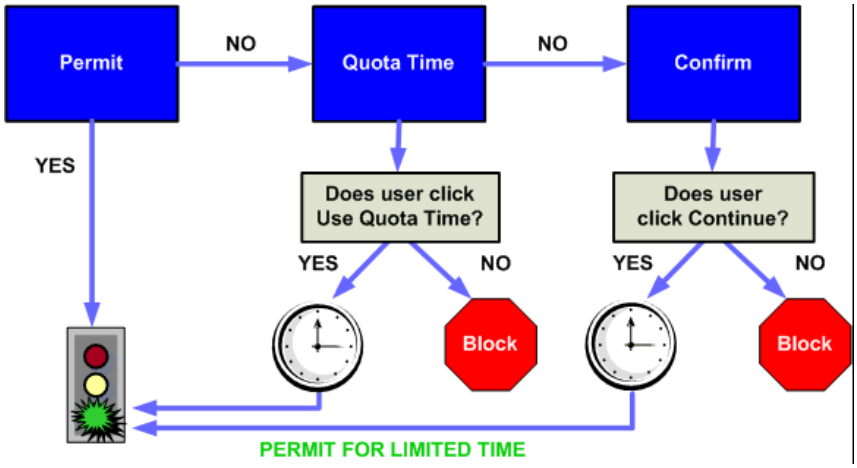

10. Handles the site according to the action applied to the category.
    - **Permit**: Permit the site.
    - **Limit by Quota**: Display the block message with an option to view the site using quota time or go back to the previous page.
    - **Confirm**: Display the block message with the option to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

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

Note

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

5 Block Pages

Related topics:

- *Protocol block messages*, page 78
- *Working with block pages*, page 79
- *Creating alternate block messages*, page 83
- *Using an alternate block page on another machine*, page 84

When Websense software blocks a Web site, it displays a block page in the client's browser. If the site is blocked because it belongs to a category in the Security Risk class (see *Risk classes*, page 37), a special version of the block page is displayed.

By default, a block page is made up of 3 main sections.

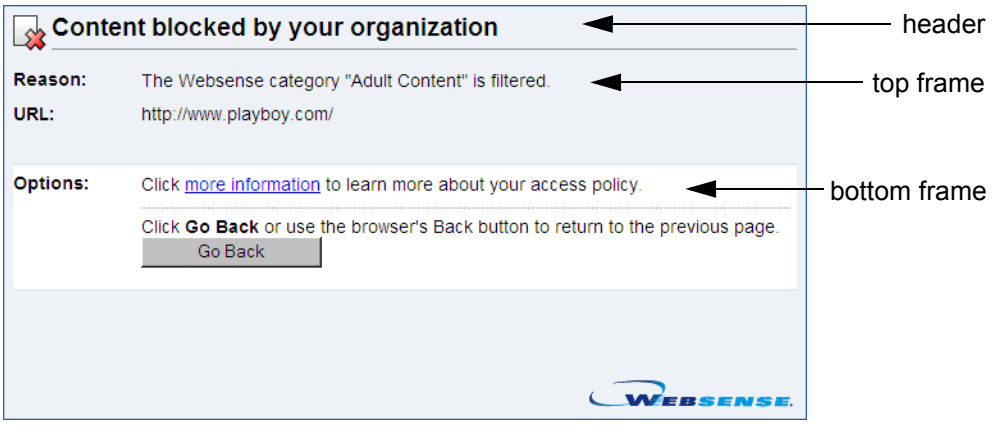

- The **header** explains that the site is blocked.
- The **top frame** contains a block message showing the requested URL and the reason the URL was blocked.
- The **bottom frame** presents any options available to the user, such as the option to go back to the previous page, or to click a Continue or Use Quota Time button to view the site.

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

- Customize the default files to change the block message (see *Working with block pages*, page 79).
- 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 84).

Protocol block messages

Related topics:

- *Working with block pages*, page 79
- *Creating alternate block messages*, page 83
- *Using an alternate block page on another machine*, page 84

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

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

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

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

- To display a protocol block message on client machines running Windows NT, XP, or 200x, the Windows Messenger service must be enabled. This service is disabled by default. You can use the Windows Services dialog box to find out if the service is running on a given machine (see *The Windows Services dialog box*, page 361).
- To display protocol block messages on a Windows 98 machine, you must start **winpopup.exe**, located in the Windows directory. Run the application from the command prompt, or configure it to launch automatically by copying it to the Startup folder.

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

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

Working with block pages

Related topics:

- *Protocol block messages*, page 78
- *Customizing the block message*, page 79
- *Creating alternate block messages*, page 83
- *Using an alternate block page on another machine*, page 84

The files used to create Websense block pages are stored in the **Websense\BlockPages\en\Default** directory:

- **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.

Customizing the block message

Related topics:

- *Changing the size of the message frame*, page 80
- *Changing the logo that displays on the block page*, page 81
- *Using block page content variables*, page 81
- *Reverting to the default block pages*, page 83

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.

- Add information about your organization's Internet use policies.
- Provide a method for contacting Human Resources or a Websense administrator about Internet use policies.

1. Navigate to the Websense block page directory:

   `<installation path>\BlockPages\en\Default`

2. Copy the block page files to the custom block page directory:

   `<installation path>\BlockPages\en\Custom`

   Note

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

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

   Warning

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

4. Modify the text. The files contain comments that guide you in making changes.

   Do not modify the tokens (enclosed by $* and *$ symbols), or the structure of the HTML code. These enable Websense software to display specific information in the block message.

5. Save the file.
6. Restart Filtering Service (see *Stopping and starting Websense services*, page 258, for instructions).

Changing the size of the message frame

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

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

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

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

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

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

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

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

5. Save and close the file.
6. Restart Filtering Service to implement the change (see *Stopping and starting Websense services*, page 258).

Changing the logo that displays on the block page

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

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

   ```
   <img title="Websense" src="/en/Custom/wslogo_block_page.png"
   ...>
   ```

   - 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 258).

Using block page content variables

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

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

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

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

Here, *WS_USERNAME* is the variable.

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

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

Variable Name	Purpose
WS_TOPFRAMESIZE	Indicates the size (as a percentage) of the top portion of a block page sent by a custom block server, if one is configured
WS_BLOCKMESSAGE_PAGE	Indicates the source to be used for a block page's top frame
WS_CATEGORY	Displays the category of the blocked URL
WS_CATEGORYID	The unique identifier for the requested URL category

Reverting to the default block pages

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

1. Delete all the files from the **Websense\BlockPages\en\Custom** directory. By default, Websense software will return to using the files in the Default directory.
2. Restart Filtering Service (see *Stopping and starting Websense services*, page 258).

Creating alternate block messages

Related topics:

- *Working with block pages*, page 79
- *Customizing the block message*, page 79

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 84).

After creating alternate block message files, you must configure Websense software to display the new messages (see *Configuring Websense 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

Related topics:

- *Working with block pages*, page 79
- *Customizing the block message*, page 79
- *Creating alternate block messages*, page 83

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 258).
2. On the Filtering Service machine, navigate to the Websense **bin** directory (by default, \Program Files\Websense\bin or /opt/websense/bin).
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 host name or the IP address of the server hosting the block page in the following format:

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

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

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

6 Using Reports to Evaluate Filtering Policies

Related topics:

- *Reporting overview*, page 86
- *Presentation reports*, page 88
- *Investigative reports*, page 105
- *Accessing self-reporting*, page 130

Websense Manager can provide several reporting tools for use in evaluating the effectiveness of your filtering policies. (Websense Manager and Websense reporting components must be installed on Windows servers.)

- The **Today** page appears first when you open Websense Manager. It shows the operating status of Websense software, and can display charts of filtering activities in the network since midnight. (See *Today: Health, Security, and Value Since Midnight*, page 19.)
- The **History** page shows charts of filtering activities in the network for up to 30 days, depending on the amount of information in the Log Database. These charts do not include today's activities. (See *History: Last 30 Days*, page 22.)
- **Presentation reports** and **Investigative reports** offer many options for generating, customizing, and scheduling reports. See *Reporting overview*, page 86, for more information.

If your organization has installed Websense Manager on a Linux server, or chooses the Websense Explorer for Linux reporting program instead of the reporting components that run on Windows, reporting options do not appear in Websense Manager. No Internet filtering charts are shown on the Today and History pages. See the *Explorer for Linux Administrator's Guide* for information on installing that program and running reports.

Reporting overview

Related topics:

- *Using Reports to Evaluate Filtering Policies*, page 85
- *Presentation reports*, page 88
- *Investigative reports*, page 105
- *Accessing self-reporting*, page 130

In addition to the charts that appear on the Today and History pages, Websense software offers 2 reporting options: presentation reports and investigative reports.

> **Note**
> In organizations that use delegated administration, some administrators may not be able to access all reporting features. See *Delegated Administration*, page 215.

Presentation reports offer a list of report definitions. Some are tabular reports, some combine a bar chart and a table. To generate a presentation report:

1. Select a report from the list.
2. Click **Run**.
3. Select the date range.
4. Click **Run Now**.

In addition to generating predefined charts, you can copy them and apply a customized report filter that identifies specific clients, categories, protocols, or actions to include. Mark report definitions that you use frequently as Favorites to make them easier to find.

You can schedule any presentation report to run at a particular time or on a repeating cycle. See *Presentation reports*, page 88, 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.

- Click the risk class name and then select a finer level of detail related to that risk class. For example, you might choose to show activity by user for the Legal Liability risk class.
- Click a user name on the resulting chart to view more detail about that user.
- Choose a different option from the **Internet use by** list to change the summary bar chart.

- Fill in the fields just above the bar chart to display two levels of information at one time. For example, starting with a summary chart of categories, you might choose **10**, **User**, and **5** to display activity for the top 5 users in the top 10 categories.
- Click a bar or number to open a detail report for that item (risk class, category, user, or other).
- Click **Favorite Reports** to save a particularly useful report format for future use, or to generate a previously saved Favorite.

The possibilities are almost endless. See *Investigative reports*, page 105, for details on the many ways you can view Internet use data.

What is Internet browse time?

Related topics:

- *Database jobs*, page 293
- *Configuring Internet browse time options*, page 297

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

Websense software includes a Log Database job to calculate the Internet browse time (IBT), using a formula based on certain configurable values. This job runs once a day, so browse time information can lag the actual log data.

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

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

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

See *Database jobs*, page 293, and *Configuring Internet browse time options*, page 297, for information about the Internet browse time job and the associated configuration options.

Presentation reports

Related topics:

- *Copying a presentation report*, page 90
- *Copying a presentation report*, page 90
- *Working with Favorites*, page 97
- *Generating presentation reports*, page 98
- *Scheduling presentation reports*, page 99
- *Viewing the scheduled jobs list*, page 103

The **Reporting > Presentation Reports** page presents a list of predefined charts and tabular reports, each showing specific information from the Log Database (see *Introducing the Log Database*, page 292). Select a report from this Report Catalog to display a brief description.

You can copy a predefined report and customize the report filter, specifying which clients, categories, protocols, and actions to include. Reports that are used frequently can be marked as Favorites to help you find them more quickly.

Run any report now, or schedule selected reports to run on a delayed or periodic basis. Choose the output format, and distribute the scheduled reports to a selected group of recipients.

If you generate a report directly from the Presentation Reports page in HTML format, the report is not saved when you move to a different page. If you generate and immediately view a report in PDF or XLS format, the report is not saved when you close the viewing program (Adobe Reader or Microsoft Excel).

Alternatively, you can choose to save the PDF or XLS file instead of displaying it immediately, or use the Save option within the viewing program. In these cases, be sure to delete or move report files periodically to avoid disk space problems.

Scheduled reports are automatically saved to the following directory:

```
<install_path>\ReportingOutput
```

The default installation path is C:\Program Files\Websense.

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. Note that the name of the attachment does not match the name of the file stored in the ReportingOutput directory. To find a specific report in this directory, search for files created on the date that the scheduled job ran.

Reports are automatically deleted from the ReportingOutput directory after 15 days. If you want to retain the reports for a longer time, include them in your backup routine or schedule them and save the emailed files in a location that permits long term storage.

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 Websense Manager machine. If the ReportingOutput directory grows too large before the files are automatically deleted, you can delete the files manually.

Websense software generates the report in the format you choose: PDF (Adobe Reader), XLS (Microsoft Excel), or HTML. If you choose HTML format, the report displays in the Websense Manager content pane. These reports cannot be printed or saved to a file. To print or save a report to file, choose the PDF or XLS output format.

If you choose PDF or XLS format, you have the option to save the report file to disk or display it in a separate window.

Important

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

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

On the Presentation Reports page, navigate through the Report Catalog and select a report of interest. Then, use the controls on the page to run the report, create a copy for which you can customize the report filter, and more.

Button	Action
Show Favorites only	Select this option to limit the Report Catalog to displaying only those reports marked as Favorites. Deselect this option to restore the full list of reports.
Edit Report Filter	Available only when a copy of a predefined report is selected, this option lets you select specific categories, protocols, users, and actions to include in the report. See *Copying a presentation report*, page 90.
Copy	Makes a copy of the selected report and adds it to the Report Catalog as a custom report. See *Copying a presentation report*, page 90. Select the custom report, and then set specific parameters for it by clicking **Edit Report Filter**.
Favorite	Marks the selected report as a Favorite, or removes the Favorite designation. See *Working with Favorites*, page 97. The Report Catalog shows a star symbol beside the report name for any report marked as a Favorite. Use the **Show Favorites only** check box to control which reports appear in the Report Catalog.

Button	Action
Delete	Deletes the selected report copy from the Report Catalog. You cannot delete predefined reports installed with the software. If the deleted report appears in any scheduled jobs, it will continue to be generated with that job.
Run	Generates the selected report after you set the date range and output format. See *Generating presentation reports*, page 98. To control other aspects of a custom report (copy of a predefined report), see *Copying a presentation report*, page 90. To schedule the report to run at a different time, or on a repeating schedule, click Scheduler.

The buttons above the page provide additional options for presentation reports.

Button	Action
Job Queue	Displays a page listing scheduled jobs that have been created, along with the status of each job. See *Viewing the scheduled jobs list*, page 103
Scheduler	Lets you 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 99.

Copying a presentation report

Related topics:

- *Copying a presentation report*, page 90
- *Presentation reports*, page 88

Initially, the **Presentation Reports** page shows a Report Catalog listing all the predefined reports installed with the software. You can generate any of these reports for a specific time period by selecting the report, and then clicking Run.

These predefined reports also work as templates that can be copied to create a custom report filter. Create a report filter to control elements such as which users, categories, protocols, and actions are to be included when you generate a report from the copy.

After copying a report and editing the report filter, you can copy the new report to create variations based on that copy.

1. Select any report in the Report Catalog.
2. Click **Copy**.

   A duplicate of the report name appears in the Report Catalog, with a code appended to indicate that it is a copy.

3. Select the copy in the Report Catalog, and then click **Edit Report Filter** to modify the elements of the report. See *Copying a presentation report*, page 90.

Defining the report filter

Related topics:

- *Copying a presentation report*, page 90
- *Generating presentation reports*, page 98

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

Note

Using a custom logo requires some preparation before you define the report filter. You must create the desired graphic in a supported graphic format and place the file in the appropriate location. See *Customizing the report logo*, page 96.

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

The filter for predefined reports cannot be changed. You can edit the filter for a copy of a predefined report:

1. Select a report in the Report Catalog.

   If the Edit Report Filter button is disabled, continue with step 2.

   If the Edit Report Filter button is enabled, skip to step 3.

2. Click **Copy** to make a copy that you can customize.

   A duplicate of the report name appears in the Report Catalog, with a code appended to show that it is a copy.

3. Click the **Edit Report Filter** button.

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

   - *Selecting clients for a report*, page 92
   - *Selecting categories for a report*, page 93
   - *Selecting protocols for a report*, page 94

- *Selecting actions for a report*, page 94
- *Setting report options*, page 95

4. On the **Confirm** tab, choose whether to run or schedule the report, in addition to saving the report filter. See *Confirming report filter definition*, page 96.

Selecting clients for a report

Related topics:

- *Selecting categories for a report*, page 93
- *Selecting protocols for a report*, page 94
- *Selecting actions for a report*, page 94
- *Setting report options*, page 95
- *Confirming report filter definition*, page 96

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

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

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

1. Select a client type from the drop-down list.
2. Set the maximum number of search results from the **Limit search** list.

   Depending on the traffic in your organization, there may be large numbers of users, groups, or domains in the Log Database. This option manages the length of the results list, and the time required to display the search results.
3. Enter one or more characters for searching, and then click **Search**.

   Use asterisk (*) as a wildcard to signify missing characters. For example, J*n might return Jackson, Jan, Jason, Jon, John, and so forth.

   Define your search string carefully, to assure that all desired results are included within the number selected for limiting the search.
4. Highlight one or more entries in the results list, and click the right arrow button (>) to move them to the **Selected** list.
5. Repeat steps 2-4 as needed to conduct additional searches and add more clients to the Selected list.
6. After you are finished making selections, click **Next** to open the Categories tab. See *Selecting categories for a report*, page 93.

Selecting categories for a report

Related topics:

- *Selecting clients for a report*, page 92
- *Selecting protocols for a report*, page 94
- *Selecting actions for a report*, page 94
- *Setting report options*, page 95
- *Confirming report filter definition*, page 96

The **Categories** tab of the Presentation Reports > Report Filter page lets you control the information included in the report on the basis of categories or risk classes. See *Risk classes*, page 37.

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

1. Select a classification: **Category** or **Risk Class**.

   Expand a parent category to display its subcategories. Expand a risk class to see a list of the categories currently assigned to that risk class.

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

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

2. Mark the check box for each category or risk class to be reported.

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

3. Click the right arrow button (>) to move your selections to the **Selected** list.

   When you mark a risk class, clicking the right arrow places all the associated categories into the Selected list.

4. After all selections are complete, click **Next** to open the Protocols tab. See *Selecting protocols for a report*, page 94.

Selecting protocols for a report

Related topics:

- *Selecting clients for a report*, page 92
- *Selecting categories for a report*, page 93
- *Selecting actions for a report*, page 94
- *Setting report options*, page 95
- *Confirming report filter definition*, page 96

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 94.

Selecting actions for a report

Related topics:

- *Selecting clients for a report*, page 92
- *Selecting categories for a report*, page 93
- *Selecting protocols for a report*, page 94
- *Setting report options*, page 95
- *Confirming report filter definition*, page 96

The **Actions** tab of the Presentation Reports > Report Filter page lets you control which precise filtering actions, such as permitted by limited access filter, blocked by quota, are included in the report. If the report specifies a particular type of action, such as Blocked, you are limited to selecting actions of that type for the report.

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

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

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

3. Click the right arrow button (>) to move your selections to the **Selected** list.
4. After all selections are complete, click **Next** to open the Options tab. See *Setting report options*, page 95.

Setting report options

Related topics:

- *Customizing the report logo*, page 96
- *Selecting clients for a report*, page 92
- *Selecting categories for a report*, page 93
- *Selecting protocols for a report*, page 94
- *Selecting actions for a report*, page 94
- *Setting report options*, page 95
- *Confirming report filter definition*, page 96

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

1. Modify the **Report catalog name** to appear in the Report Catalog. The name can have up to 85 characters.

   This name does not appear on the report itself; it is used only for identifying the unique combination of report format and filter in the Report Catalog.
2. Modify the **Report title** that appears on the report. The title can have up to 85 characters.
3. Modify the **Description** to appear in the Report Catalog. The description can have up to 336 characters.

   The description should help you identify this unique combination of report format and filter in the Report Catalog.
4. Select a logo to appear on the report.

   All supported image files in the appropriate directory are listed. See *Customizing the report logo*, page 96.
5. Mark the **Save as Favorite** check box to have the report listed as a Favorite.

   The Report Catalog shows a star symbol beside Favorite reports. You can select **Show only Favorites** on the Report Catalog page to reduce the number of reports listed, which enables you to move more quickly to a particular report.
6. Mark the **Show only top** check box and then enter a number from 1 to 20 to limit the number of items reported.

This option appears only if the selected report is formatted as a Top N report, designed to show a limited number of items. The item that is limited depends on the report. For example, for a Top Categories Visited report, this entry determines how many categories are reported.

7. After all entries and selections are complete, click **Next** to open the Confirm tab. See *Confirming report filter definition*, page 96.

Customizing the report logo

Predefined presentation reports display the Websense logo in the upper left corner. When you copy a predefined report and define 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
   - .ttf

2. Use a maximum of 25 characters for the image file name, including extension.
3. Place the image file into the following directory:

   `<install_path>\Manager\ReportingTemplates\images`

   The default installation path is C:\Program Files\Websense.

All supported image files in this directory automatically appear in the drop-down list on the Options tab of the Report Filter page. The image is automatically scaled to fit within the space allocated for the logo. (See *Setting report options*, page 95.)

Note

Do not remove images that are active in report filters from this directory. If the specified logo file is missing, the report cannot be generated.

Confirming report filter definition

Related topics:

- *Selecting clients for a report*, page 92
- *Selecting categories for a report*, page 93
- *Selecting protocols for a report*, page 94
- *Selecting actions for a report*, page 94
- *Setting report options*, page 95

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

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

   If any changes are needed, click **Back** to return to the Options tab, where you can make those changes. (See *Setting report options*, page 95.)

2. Indicate how you want to proceed:

Option	Description
Save	Saves the report filter and returns to the Report Catalog. See *Presentation reports*, page 88.
Save and Run	Saves the report filter and opens the Run Report page. See *Generating presentation reports*, page 98.
Save and Schedule	Saves the report filter and opens the Schedule Report page. See *Scheduling presentation reports*, page 99.

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

Working with Favorites

Related topics:

- *Presentation reports*, page 88
- *Generating presentation reports*, page 98
- *Scheduling presentation reports*, page 99

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

1. On the **Presentation Reports** page, highlight a report that you generate frequently, or want to be able to locate quickly.
2. Click **Favorite**.

   A star symbol appears beside Favorite report names in the list, letting you quickly identify them when all reports are shown.

3. Mark the **Show only Favorites** check box above the Report Catalog to limit the list to those marked as Favorites. Clear this check box to restore the full list of reports.

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

1. Highlight a report that shows the star symbol of a Favorite.
2. Click **Favorite**.

The star symbol is removed from that report name in the Report Catalog. The report is now omitted from the list if you choose **Show only Favorites**.

Generating presentation reports

Related topics:

- *Presentation reports*, page 88
- *Scheduling presentation reports*, page 99

Generating a single report immediately involves the few steps shown below.

> **Note**
>
> Before generating a report in PDF format, make sure that Adobe Reader v7.0 or later is installed on the machine from which you are accessing Websense Manager.
>
> Before generating a report in XLS format, make sure that Microsoft Excel 2003 or later is installed on the machine from which you are accessing Websense Manager.
>
> If the appropriate software is not installed, you have the option to save the file.

To create jobs with one or more reports to run once or on a repeating cycle with the presentation reports scheduling feature. See *Scheduling presentation reports*, page 99.

1. On the **Presentation Reports** page, highlight a report in the Report Catalog tree, and then click **Run**.
2. Select the **Start date** and **End date** for the report data.
3. Select an **Output format** for the report.

Format	Description
PDF	Portable Document Format. PDF files are viewed in Adobe Reader.
HTML	HyperText Markup Language. HTML files can be viewed directly in your Internet Explorer or Firefox browser.
XLS	Excel spreadsheet. XLS files are viewed in Microsoft Excel.

4. If you selected a **Top N** report, choose the number of items to be reported.
5. Click **Run**.

   HTML reports appear in the content pane. If you selected PDF or XLS output, you have a choice of whether to open the report in a separate window or save the report to disk.

6. To print a report, use the print option that is part of the program displaying the report.

   For best results, generate PDF or XLS output for printing. Then, use the print options in Adobe Reader or Microsoft Excel, respectively.

You can save a report that is output in PDF or XLS format by using the Save feature in Adobe Reader or Microsoft Excel.

Scheduling presentation reports

Related topics:

- *Presentation reports*, page 88
- *Generating presentation reports*, page 98
- *Viewing the scheduled jobs list*, page 103
- *Copying a presentation report*, page 90

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.

The 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.

To access the Scheduler:

- Click the **Scheduler** button at the top of the Presentation Reports page (above the Report Catalog).
- When adding or editing a report filter for a report, choose **Save and schedule** in the Confirm tab, and then click **Finish**. (See *Copying a presentation report*, page 90.)
- Click the job name link on the Job Queue page to edit a job.
- Click **Add** on the Job Queue page to create a new job.

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

- *Setting the schedule*, page 100
- *Selecting reports to schedule*, page 101
- *Selecting output options*, page 102
- *Setting the date range*, page 101
- *Selecting output options*, page 102

After creating jobs, you can view a list of jobs that displays their status and other helpful information. See *Viewing the scheduled jobs list*, page 103.

Setting the schedule

Related topics:

- *Scheduling presentation reports*, page 99
- *Selecting reports to schedule*, page 101
- *Selecting output options*, page 102
- *Setting the date range*, page 101

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 machine running Websense Manager.

> **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 scheduled jobs list*, page 103.
End after	Select the number of times to run the job. After that number of occurrences, the job does not run again, but it stays in the Job Queue until you delete it. See *Viewing the scheduled jobs list*, page 103.
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 reports to schedule*, page 101.

Selecting reports to schedule

Related topics:

- *Scheduling presentation reports*, page 99
- *Setting the schedule*, page 100
- *Selecting output options*, page 102
- *Setting the date range*, page 101

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*, page 101.

Setting the date range

Related topics:

- *Scheduling presentation reports*, page 99
- *Setting the schedule*, page 100
- *Selecting reports to schedule*, page 101
- *Selecting output options*, page 102

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 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*, page 102.

Selecting output options

Related topics:

- *Scheduling presentation reports*, page 99
- *Setting the schedule*, page 100
- *Selecting reports to schedule*, page 101
- *Setting the date range*, page 101

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

1. Select the file format for the finished report.

Format	Description
PDF	Portable Document Format. Recipients must have Adobe Reader v7.0 or later to view the PDF reports.
XLS	Excel Spreadsheet. Recipients must have Microsoft Excel 2003 or later to view the XLS reports.

2. Enter email addresses for distributing the report.

   Enter each address on a separate line.

3. Mark the **Customize subject and body of email** check box, if desired. Then, enter the custom **Subject** and **Body** text for this job's distribution email.
4. Click **Save Job** to save and implement the job definition, and display the Job Queue page.
5. Review this job and any other scheduled jobs. See *Viewing the scheduled jobs list*, page 103.

Viewing the scheduled jobs list

Related topics:

- *Presentation reports*, page 88
- *Scheduling presentation reports*, page 99
- *Selecting output options*, page 102
- *Scheduling investigative reports*, page 125

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 127.)

The list provides the following information for each job.

Column	Description
Job Name	The name assigned when the job was created.
State	One of the following: • ENABLED indicates a job that runs according to the established recurrence pattern. • DISABLED indicates a job that is inactive, and does not run.

Column	Description
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 job history*, page 104.
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 99.
Add Job	Opens the Scheduler page where you can define a new job. See *Scheduling presentation reports*, page 99.
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 job history

Related topics:

- *Scheduling presentation reports*, page 99
- *Viewing the scheduled jobs list*, page 103

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.

Investigative reports

Related topics:

- *Summary reports*, page 107
- *Multi-level summary reports*, page 111
- *Flexible detail reports*, page 112
- *User Activity Detail reports*, page 116
- *Standard reports*, page 121
- *Favorite investigative reports*, page 122
- *Scheduling investigative reports*, page 125
- *Outliers reports*, page 128
- *Output to file*, page 129
- *Database connection and report defaults*, page 303

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

Initially, the main Investigative Reports page shows a summary report of activity by risk class. Work in the summary report view by clicking the available links and elements to explore areas of interest and gain general insight into your organization's Internet usage. See *Summary reports*, page 107.

Multi-level summary reports (see *Multi-level summary reports*, page 111) and flexible detail reports (see *Flexible detail reports*, page 112) let you analyze the information from different perspectives.

Other report views and investigative reports features can be accessed from links at the top of the page. See the table below for a list of links and the features they access. (Not all links are available on all pages.)

Option	Action
User by Day/Month	Displays a dialog box that lets you define a report of a specific user's activity, covering either a day or a month. For more information, see *User Activity Detail reports*, page 116.
Standard Reports	Displays a list of predefined reports so you can quickly see a specific combination of data. See *Standard reports*, page 121.
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 122.
Job Queue	Displays the list of scheduled investigative reports jobs. See *Scheduling investigative reports*, page 125.
View Outliers	Displays reports showing Internet usage that is significantly different from average. See *Outliers reports*, page 128.
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 303.
	Click this button, at the right of the Search fields, to export the current report to a spreadsheet file compatible with Microsoft Excel. You are prompted to either open or save the file. To open the file, Microsoft Excel 2003 or later must be installed. See *Output to file*, page 129.
	Click this button, at the right of the Search fields, to export the current report to a PDF file compatible with Adobe Reader. You are prompted to either open or save the file. To open the file, Adobe Reader version 7.0 or later must be installed. See *Output to file*, page 129.

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 Filtering Service for logging*, page 280), 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 full URL logging*, page 296).

Websense investigative reports are limited by the processor and available memory of the machine running Websense Manager, as well as some network resources. Some large reports may take a very long time to generate. The progress message includes an

option for saving the report as a Favorite so you can schedule it to run at a separate time. See *Scheduling investigative reports*, page 125.

Summary reports

Related topics:

- *Multi-level summary reports*, page 111
- *Flexible detail reports*, page 112
- *User Activity Detail reports*, page 116
- *Standard reports*, page 121
- *Favorite investigative reports*, page 122
- *Scheduling investigative reports*, page 125
- *Outliers reports*, page 128
- *Output to file*, page 129

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 303.

Quickly change the information reported, or drill down into the report details, by clicking the various links and options available on the page.

1. Select 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. See *Configuring log cache files*, page 286.
Bandwidth [KB]	The amount of data, in kilobytes, contained in both the initial request from the user and the response from the Web site. This is the combined total of the Sent and Received values. Keep in mind that some integration products do not send this information to Websense software. Two examples are Check Point FireWall-1 and Cisco PIX Firewall. If your integration does not send this information, and Websense Network Agent is installed, activate the **Log HTTP requests (enhanced logging)** option for the appropriate NIC to enable reporting on bandwidth information. See *Configuring NIC settings*, page 316.

Option	Description
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 a more significant submission if the user is registering for a Web site, for example.
Received [KB]	The number of kilobytes received in response to the request. This includes all text, graphics, and scripts that make up the site. For sites that are blocked, the number of kilobytes 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 Websense block page. If the log record is created by Websense Security Gateway, as a result of real-time scanning, the kilobytes received represents the size of the page scanned. See *Analyze Content with the Real-Time Options*, page 131, for more information on real-time scanning. If another integration product creates the log records, the kilobytes received for a blocked site may be zero (0), may represent the size of the block page, or may be a value obtained from the requested site.
Browse Time	An estimate of the amount of time spent viewing the site. See *What is Internet browse time?*, page 87.

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 305.)

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	Select a period for the report: One Day, One Week, One Month, or All. The report updates to show data for the selected period. Use the adjacent arrow buttons to move through the available data, one period (day, week, month) at a time. As you change this selection, the **View from** fields update to reflect the time period being viewed. 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.
View from... to...	The dates in these fields update automatically to reflect the time period being viewed when you make changes in the **View** field. Alternatively, enter exact start and end dates for the reports, or click the calendar icon to select the desired dates. Click the adjacent right arrow button to update the report after selecting dates.
Pie Chart / Bar Chart	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. When the pie chart is active, click **Bar Chart** to display the current summary report as a bar chart.
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	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 display user names in these locations. Under some circumstances, user names cannot be displayed. For more information, see *Configuring Filtering Service for logging*, page 280. If 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. If individual administrators should never have access to user names in reports, assign them to a role in which the reporting permissions prevent viewing of user names in investigative reports and access to presentation reports.
Search for	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. 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.

8. Add a subset of information for all or selected entries in the left column by creating a multi-level summary report. See *Multi-level summary reports*, page 111.
9. Create a tabular report for a specific item in the left column by clicking the adjacent number or measurement bar. This detailed report can be modified to meet your specific needs. See *Flexible detail reports*, page 112.

Multi-level summary reports

Related topics:

- *Investigative reports*, page 105
- *Summary reports*, page 107
- *Flexible detail reports*, page 112
- *User Activity Detail reports*, page 116
- *Standard reports*, page 121
- *Favorite investigative reports*, page 122
- *Scheduling investigative reports*, page 125
- *Outliers reports*, page 128
- *Output to file*, page 129

Multi-level summary reports show a second level of information to supplement the primary information displayed. For example, if the primary display shows risk classes, you can define a second level to learn which categories have been requested most within each risk class. As another example, if the primary report shows requests for each category, you might show the top 5 categories and the 10 users who made the most requests to each.

Use the settings immediately above the summary report to create a multi-level summary report.

1. In the **Select top** list, choose a number to designate how many primary entries (left column) to report. The resulting report includes the primary entries with the largest values. (This shows the earliest dates if Day is the primary entry.)

   Alternatively, mark the check box beside the desired individual entries in the left column to report only those entries. The **Select top** field displays **Custom**.
2. From the **by** list, choose the secondary information to report.
3. In the **Display** field, choose the number of secondary results to report for each primary entry
4. Click **Display Results** to generate the multi-level summary report.

   The summary report updates to show only the selected number of primary entries. Below the bar for each primary entry, a list of secondary entries appears.
5. Use the arrows beside a column heading to change the report's sort order.

To return to a single-level summary report, select a different option under **Internet Use by**. Alternatively, click one of the primary or secondary entries, and select an option to generate a new investigative report of that information.

Flexible detail reports

Related topics:

- *Investigative reports*, page 105
- *Summary reports*, page 107
- *Multi-level summary reports*, page 111
- *Favorite investigative reports*, page 122
- *Scheduling investigative reports*, page 125
- *Outliers reports*, page 128
- *Output to file*, page 129
- *Database connection and report defaults*, page 303
- *Columns for flexible detail reports*, page 114

Detail reports give you a tabular view of the information in the Log Database. Access the detail report view from the main page after viewing a summary report for which you want more detail.

You can request a detail view from any row. However, when requesting a detail report based on hits, it is best to start from a row that shows fewer than 100,000 hits. If there are more than 100,000 hits for a particular row, the hits value displays in red to alert you that a detail report may be slow to generate.

Detail report view is considered *flexible* because it lets you design your own report. You can add or delete columns of information, and change the order of the columns displayed. The information is sorted according to order of the columns. You can even reverse the sort order within any column from ascending to descending, or vice versa.

Websense investigative reports are limited by the processor and available memory of the machine running Websense Manager, as well as some network resources. Requests for large reports may time out. When you request a large report, you are given options for generating the report without timeouts.

Important

In any drop-down or values list, some options may appear in red. The red lettering indicates that selecting this option may result in a very large report. It is generally more effective to drill down further into the details before selecting that option.

1. Generate a summary report or multi-level report on the investigative reports main page. (See *Summary reports*, page 107, or *Multi-level summary reports*, page 111.)
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 pop-up message shows progress while the detail report loads.

   > **Note**
   > If the report takes a long time to generate, consider saving it as a Favorite report by clicking the link in the Loading message, and scheduling it to run later. See *Favorite investigative reports*, page 122.

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 303.)

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 reports*, page 114, 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 305.
    - Click the URL to open the requested Web site in a new window.
11. Click **Favorite Report** if you want to save the report so that you can generate it again quickly or on a recurring basis (see *Saving a report as a Favorite*, page 123).

Columns for flexible detail reports

Related topics:

- *Flexible detail reports*, page 112
- *Favorite investigative reports*, page 122
- *Scheduling investigative reports*, page 125

The table below describes the columns available for detail reports (see *Flexible detail reports*, page 112).

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 on user-based reports.
Day	Date the 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 278.
Directory Object	Directory path for the user who made the request, excluding the user name. Typically, this results in multiple rows for the same traffic, because each user belongs in multiple paths. If you are using a non-LDAP directory service, this column is not available.
Disposition	Action Websense software took as a result of the request, such as category permitted or category blocked.
Source Server	IP address of the machine sending requests to Filtering Service. This is the machine running either the integration product or Websense Network Agent.
Protocol	Protocol of the request.
Protocol Group	Master Database group in which the requested protocol falls.
Source IP	IP address of the machine from which the request was made.
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 full URL logging*, page 296.
Month	Calendar month the request was made.
Port	TCP/IP port over which the user communicated with the site.
Bandwidth	The amount of data, in kilobytes, contained in both the initial request from the user and the response from the Web site. This is the combined total of the Sent and Received values. Keep in mind that some integration products do not send this information to Websense software. Two examples are Check Point FireWall-1 and Cisco PIX Firewall. If your integration does not send this information, and Websense Network Agent is installed, activate the **Log HTTP requests (enhanced logging)** option for the appropriate NIC to enable reporting on bandwidth information. See *Configuring NIC settings*, page 316.
Bytes Sent	Number of bytes sent as the Internet request. This represents the amount of data transmitted, which may be a simple request for a URL, or may be a more significant submission if the user is registering for a Web site, for example.

Column Name	Description
Bytes Received	Number of bytes received from the Internet in response to the request. This includes all text, graphics, and scripts that make up the site. 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. If the log record is created by Websense Security Gateway, as a result of real-time scanning, the bytes received represents the size of the page scanned. See *Analyze Content with the Real-Time Options*, page 131, for more information on real-time scanning. If another integration product creates the log records, the bytes received for a blocked site may be zero (0), may represent the size of the block page, or may be a value obtained from the requested site.
Time	Time of day when the site was requested, shown in the HH:MM:SS format, using a 24-hour clock.
Category	Category under which the request was filtered. This may be a category from the Websense Master Database or a custom category.

User Activity Detail reports

Related topics:

- *Investigative reports*, page 105

Click the **User by Day/Month** link to generate a User Activity Detail report for one user. This report gives a graphical interpretation of the user's Internet activity for a single day or a full month.

First, generate a report for a specific user for a selected day. From that report, you can generate a report of the same user's activity for a full month. For detailed instructions, see:

- *User activity detail by day*, page 117
- *User activity detail by month*, page 118

User activity detail by day

Related topics:

- *Investigative reports*, page 105
- *User Activity Detail reports*, page 116
- *User activity detail by month*, page 118

The User Activity Detail by Day report gives a more in-depth view of a specific user's activity on one day.

1. Select **User by Day/Month** at the top of the main page. The User Detail by Day dialog box appears.
2. Enter a user's name, or a portion of the name, in the **Search for a user** field, and then click **Search**.

   The search displays a scrolling list of up to 100 matching user names from the Log Database.
3. Make a selection from the **Select user** list.
4. In the **Select day** field, either accept the last activity date that appears by default, or choose a different date.

   You can type the new date or click the calendar icon to select a date. The calendar selection box indicates the date range covered by the active Log Database.
5. Click **Go to User by Day** to see a detailed report of activity for that user on the requested date.

   The initial report shows a timeline of the user's activity in 5-minute increments. Each request appears as an icon, which corresponds to a Websense Master Database category. A single icon represents all custom categories. (The color of the icons corresponds to the risk grouping shown on the User Activity by Month reports. See *User activity detail by month*, page 118.)

   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.

Option	Description
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 305. After you click the link, it becomes View All Hits, which restores the original list of each request.
Category View Control	Displays a list of each category in the current report, showing both the category name and the icon representing that category. Control which categories appear in the report by marking the check boxes for the categories to be included. Then, click **Accept** to update the report according to your selections.

6. Click **User Activity Detail by Month**, above the report, to view the same user's activity for the full month. See *User activity detail by month*, page 118,. for more information.

User activity detail by month

Related topics:

- *Investigative reports*, page 105
- *User Activity Detail reports*, page 116
- *User activity detail by day*, page 117
- *Category mapping*, page 119

While the User Activity Detail by Day report is open, you can switch to see the monthly activity for that user.

1. Open a User Activity Detail by Day report. See *User activity detail by day*, page 117.
2. Click **User Activity Detail by Month** at the top.

   The new report displays a calendar image, with each day's area showing small colored blocks representing the user's Internet activity for that day. Requests to sites in custom categories are shown as gray blocks.
3. Click **Database Category Legend** at the top left to see how the colors represent low to high potential risk for the requested site.

   The category assignments are fixed, and cannot be changed. See *Category mapping*, page 119.
4. Click **Prev** or **Next** to display this user's Internet activity for the previous or the next month.

Category mapping

Related topics:

- *Investigative reports*, page 105
- *User Activity Detail reports*, page 116
- *User activity detail by month*, page 118

The following list identifies which categories are represented by each of the colors on the User Activity by Day and User Activity Detail by Month reports.

Keep in mind that category names in the Master Database are subject to change. Additionally, categories may be added or deleted at any time.

Color	Categories
Gray	Custom Categories Non-HTTP traffic
Dark Blue	**Business and Economy** and all its subcategories **Education**, and all its subcategories **Health** **Information Technology**, including the Search Engines and Portals, and Web Hosting subcategories **Miscellaneous** subcategories Content Delivery Networks, Dynamic Content, Images (Media), Image Servers, and Private IP Addresses **Productivity**/Advertisements
Light Blue	**Drugs**/Prescribed Medications **Government** and its Military subcategory **Information Technology**/URL Translation Sites **Miscellaneous**, parent category only **News and Media**, parent category only **Special Events**

Color	Categories
Yellow Green	**Abortion** and all its subcategories **Adult Material**/Sex Education **Bandwidth**, including the subcategories Internet Radio and TV, Personal Network Storage and Backup, and Streaming Media **Entertainment**, including its subcategory MP3 **Games** **Government**/Political Organizations **Information Technology**/Computer Security **Internet Communication**/Web-based Email **Miscellaneous**/File Download Servers **Miscellaneous**/Network Errors **News and Media**/Alternative Journals **Productivity**, including its subcategories Instant Messaging, Message Boards and Clubs, and Online Brokerage and Trading **Religion** and its subcategories Non-Traditional Religions and Occult and Folklore, and Traditional Religions **Security**, parent category only **Shopping** and all its subcategories **Social Organizations** and all its subcategories **Society and Lifestyles**, including its subcategories Gay or Lesbian or Bisexual Interest, Hobbies, Personal Web Sites, and Restaurants and Dining **Sports** and all its subcategories **Travel** **User-Defined** **Vehicles**

Color	Categories
Orange	**Adult Material**/Nudity **Advocacy Groups** **Bandwidth**/Internet Telephony **Drugs** and its subcategories Abused Drugs, Marijuana, and Supplements and Unregulated Compounds **Information Technology**/Proxy Avoidance **Internet Communication** and its subcategory Web Chat **Job Search** **Miscellaneous**/Uncategorized **Productivity** subcategories Freeware and Software Download, and Pay-to-Surf **Religion** **Society and Lifestyles** subcategories Alcohol and Tobacco, and Personals and Dating **Tasteless** **Weapons**
Red	**Adult Material** and these subcategories: Adult Content, Lingerie and Swimsuit, and Sex **Bandwidth**/Peer-to-Peer File Sharing **Gambling** **Illegal or Questionable** **Information Technology**/Hacking **Militancy and Extremist** **Racism and Hate** **Security** subcategories Keyloggers, Malicious Web Sites, Phishing, and Spyware **Violence**

Standard reports

Related topics:

- *Investigative reports*, page 105
- *Favorite investigative reports*, page 122
- *Scheduling investigative reports*, page 125

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 122.

Favorite investigative reports

Related topics:

- *Investigative reports*, page 105
- *Scheduling investigative reports*, page 125

You can save most investigative reports as **Favorites**. This includes reports you generate by drilling down to specific information, standard reports, and detail reports that you have modified to meet your specific needs. Then, run the Favorite report at any time, or schedule it to run on specific days and times.

In organizations that use delegated administration, permission to save and schedule Favorites is set by the Super Administrator. Administrators who are granted this permission can run and schedule only the Favorites they saved; they do not have access to Favorites saved by other administrators.

For detailed instructions on working with Favorite reports, see:

- *Saving a report as a Favorite*, page 123
- *Generating or deleting a Favorite report*, page 123
- *Modifying a Favorite report*, page 124

Saving a report as a Favorite

Related topics:

- *Favorite investigative reports*, page 122
- *Modifying a Favorite report*, page 124

Use the following procedure to save a report as a Favorite.

1. Generate an investigative report with the desired format and information.
2. Click **Favorite Reports**.
3. Accept or modify the name displayed by Websense Manager.

   The name may contain letters, numbers and underscore characters (_). No blanks or other special characters can be used.
4. Click **Add**.

   The report name is added to the list of Favorites.
5. Select a report on this list, then select an option for managing the report. Depending on the option you choose, see:
   - *Generating or deleting a Favorite report*, page 123
   - *Scheduling investigative reports*, page 125

Generating or deleting a Favorite report

Related topics:

- *Favorite investigative reports*, page 122
- *Modifying a Favorite report*, page 124

You can generate a Favorite report at any time, or delete one that has become obsolete.

1. Click **Favorite Reports** to display a list of reports saved as favorites.

   Note
   If your organization uses delegated administration, this list does not include favorite reports saved by other administrators.

2. Select the desired report from the list.

   If the desired report has not been saved as a Favorite, see *Saving a report as a Favorite*, page 123.

3. Depending on your need:
   - 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 125, for more information.
   - Click **Delete** to remove the report from the Favorites list.

Modifying a Favorite report

Related topics:

- *Investigative reports*, page 105
- *Favorite investigative reports*, page 122

You can easily create a new Favorite report that is similar to an existing Favorite report, as follows.

1. Click **Favorite Reports** to display a list of reports saved as favorites.

   Note
   If your organization uses delegated administration, this list does not include favorite reports saved by other administrators.

2. Select and run the existing Favorite report that most closely resembles the new report you want to create. (See *Generating or deleting a Favorite report*, page 123.)
3. Modify the displayed report as desired.
4. Click **Favorite Reports** to save the revised display as a Favorite report with a new name. (See *Saving a report as a Favorite*, page 123.)

Scheduling investigative reports

Related topics:

- *Favorite investigative reports*, page 122
- *Saving a report as a Favorite*, page 123
- *Managing scheduled investigative reports jobs*, page 127

You must save an investigative report as a Favorite before it can be scheduled to run at a later time or on a repeating cycle. When the scheduled report job runs, the resulting reports are sent via email to the recipients you designate. As you create scheduled jobs, consider whether your email server will be able to handle the size and quantity of the attached report files.

Scheduled report files are stored in the following directory:

```
<install_path>\webroot\Explorer\<name>\
```

The default installation path is C:\Program Files\Websense. If the scheduled job has only one recipient, <name> is the first portion of the email address (before the @). In the case of multiple recipients, the reports are saved in a directory called Other.

Note

The reports saved from a repeating job use the same file name each time. If you want to save files for longer than a single cycle, be sure to change the file name or copy the file to another location.

Depending on the size and number of reports scheduled, this directory could become very large. Be sure to clear the directory periodically, eliminating unneeded report files.

1. Save one or more reports as Favorites. (See *Saving a report as a Favorite*, page 123).
2. Click **Favorite Reports** to display a list of reports saved as favorites.

Note

If your organization uses delegated administration roles, this list does not include favorite reports saved by other administrators.

3. Highlight up to 5 reports to run as part of the job.
4. Click **Schedule** to create a scheduled report job, and then provide the information requested on the Schedule Report page.

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

Field	Description
Recurrence	Select the frequency (Once, Daily, Weekly, Monthly) for running the report job.
Start Date	Choose the day of the week or calendar date for running the job the first (or only) time.
Run Time	Set the time of day for running the job.
Email to	Use the **Additional Email Addresses** field to add the appropriate addresses to this list. Highlight one or more email addresses to receive the reports in the job. (Be sure to deselect any that should not receive the reports.)
Additional Email Addresses	Enter an email address, and then click **Add** to put it on the **Email to** list. The new email address is automatically highlighted with the other selected email addresses.
Customize email subject and body text	Mark this check box to customize your email notification subject line and body text. If this box is not checked, the default subject and body text are used.
Email Subject	Enter the text to appear as the email subject line when scheduled reports are distributed. The default email subject reads: Investigative Reports scheduled job
Email Text	Enter text to be added to the email message for distributing scheduled reports. The email reads as shown below, with your text in place of <CUSTOM TEXT>. Report scheduler generated the attached file or files on <date time>. <CUSTOM TEXT> To view the generated report(s), click on the following link(s). Note: The link will not work if the recipient does not have access to the web server from which the job was sent.
Schedule Job Name	Assign a unique name for the scheduled job. The name identifies this job in the Job Queue. See *Managing scheduled investigative reports jobs*, page 127.

Field	Description
Output Format	Choose the file format for the scheduled reports: **PDF**: Portable Document Format files are viewed in Adobe Reader. **Excel**: Excel spreadsheet files are viewed in Microsoft Excel.
Date Range	Set the date range to be covered by reports in this job. **All Dates**: all available dates in the Log Database. **Relative**: Choose a time period (Days, Weeks, or Months) and the specific period to include (This, Last, Last 2, and so on). **Specific**: set specific dates or a date range for the reports in this job.

5. Click **Next** to display the Schedule Confirmation page.
6. Click **Save** to save your selections and go to the Job Queue page (see *Managing scheduled investigative reports jobs*, page 127).

Managing scheduled investigative reports jobs

Related topics:

- *Investigative reports*, page 105
- *Scheduling presentation reports*, page 99

When you create a scheduled job for investigative reports, the **Job Queue** page appears, showing the new job and a list of existing scheduled jobs. You can also access the page by clicking the **Job Queue** link on the main investigative reports page.

Note

If your organization uses delegated administration, this page does not show jobs scheduled by other administrators.

The **Schedule Report Detail** section lists each scheduled job in the order it was created showing an overview of the defined schedule and the job status. In addition, the following options are available.

Option	Description
Edit	Displays the schedule defined for this job, and allows you to modify it, as needed.
Delete	Deletes the job and adds an entry to the Status Log section showing the job as Deleted.

The **Status Log** section lists each job that has changed in some way, showing the scheduled start time for the job, the actual end time, and the status.

Click **Clear Status Log** to remove all entries in the Status Log section.

Outliers reports

Related topics:

- *Investigative reports*, page 105
- *Summary reports*, page 107

An Outliers report shows which users have the most unusual Internet activity in the database. Websense software calculates the average activity for all users per category, per day, per action (sometimes called 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 **View Outliers**.

   The rows are sorted in descending order with the highest variance shown first. Each row shows:

   - Total (hits or bandwidth) for the user, category, protocol, day, and action.
   - Average (hits or bandwidth) for all users, for that category, protocol, day, and action.
   - Variance from the average for the user.

3. To see an individual user's activity in this category over time, click the user name.

   For example, if one user's activity is noticeably high for a certain day, click that user's name to see a report that gives a more in-depth understanding of the user's overall activity.

Output to file

Related topics:

- *Investigative reports*, page 105
- *Printing investigative reports*, page 129

After you generating an investigative report, you can use the buttons above the report to save it to a file. The button you click determines the format of the file.

Option	Description
	Saves the report in XLS format. If Microsoft Excel 2003 or later is installed on the machine from which you are accessing Websense Manager, you are prompted to view or save the report. Otherwise, you are prompted to select a directory and file name for the saved report. Use the options in Microsoft Excel to print, save, or email the report.
	Generates a report in PDF format. If Adobe Reader v7.0 or later is installed on the machine from which you are accessing Websense Manager, you are prompted to view or save the report. Otherwise, you are prompted to select a directory and file name for the saved report. Use the options in Adobe Reader to print, save, or email the report.

Printing investigative reports

Related topics:

- *Investigative reports*, page 105
- *Output to file*, page 129

You can print investigative reports by:

- Using the Web browser print function while the report is displayed.
- Creating a PDF or XLS file, and then using the print function in Adobe Reader or Microsoft Excel (see *Output to file*, page 129).

Although reports have been set up to print successfully from the browser, you may want to test printing to check the result.

User Activity Detail by Month reports are configured to print in landscape mode. All other reports are configured for portrait mode.

When you design your own report (see *Flexible detail reports*, page 112), 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 305.)

Accessing self-reporting

Related topics:

- *Investigative reports*, page 105
- *Configuring reporting preferences*, page 279
- *Self-reporting*, page 307

Websense self-reporting allows you to evaluate your own Internet browsing activities and adjust them, as needed, to meet organizational guidelines. It also accommodates government regulations that require organizations to let users see the type of information being collected.

If self-reporting is enabled in your organization, access it from your browser:

1. Enter the URL supplied by your Websense Administrator, or click the Self-Reporting link on the main Websense Manager logon page to access the self-reporting logon page.
2. If **Policy Server** shows a drop-down list, choose the IP address for the Policy Server that logs information on your Internet activity.

   Contact your Websense Administrator for assistance.
3. Enter the **User name** and **Password** you use to log on to the network.
4. Click **Log On**.

Websense Manager opens to an investigative report showing your Internet activity by risk class. Click the various links and elements on the page to access other options for alternative views of the information stored on your activity. Use the **Help** system for assistance when working with the reports.

7 Analyze Content with the Real-Time Options

Related topics:

- *Scanning options*, page 133
- *Categorizing content and scanning for threats*, page 134
- *File scanning*, page 135
- *Stripping content*, page 136
- *Reporting on real-time scanning activity*, page 139

Websense filtering software filters Internet activity based on your active policy and information stored in the Master Database. If you subscribe to Websense Content Gateway or Websense Web Security Gateway, you can also analyze Web site and file content in real time.

Depending on your subscription, 2 real-time analysis options are available: content categorization and Security real-time scanning.

- Use **content categorization** to review the content of URLs that are not already blocked (based on your active policy and on the URL's Websense Master Database categorization), and return a category for use in filtering.
- If you subscribe to Websense Web Security Gateway, 3 **Security real-time scanning** options are available.
  - **Content scanning** looks at Web content to find security threats such as phishing, URL redirection, Web exploits, and proxy avoidance.
  - **File scanning** inspects file content to determine a threat category, such as viruses, Trojan horses, or worms.
  - **Content stripping** removes active content from requested Web pages.

When any of these options are activated, only sites **not** already blocked based on your active policy and their Websense Master Database categorization are analyzed. For more information, see *Scanning options*, page 133.

Important

Limited access filters and unfiltered URLs supersede real-time categorization.

If a user requests a site in an active limited access filter (see *Restricting users to a defined list of Internet sites*, page 153) or the Unfiltered URLs list (see *Redefining filtering for specific sites*, page 166), the request is permitted, even when real-time scanning is performed and threats are found.

To take advantage of these real-time security features, enter a subscription key that includes support for Websense Content Gateway or Websense Web Security Gateway in 2 places:

- In Websense Manager (go to **Settings > Account**).
- In the Websense Content Gateway management interface (go to the **Configure > My Proxy > Subscription > Subscription Management** tab).

It takes several minutes for the 2 products to download the necessary databases, synchronize, and display all real-time features in both management tools.

Websense real-time options

Websense real-time options help ensure network security. Use these options to scan Internet content and assign it to a filtering category. The real-time result is sent to Filtering Service, which filters the site based on the action assigned to its real-time categorization in the active policy.

Database download

The real-time options rely on small databases installed with Websense Web Security Gateway, which checks for database updates at a regular interval. Updates to these databases occur independently of all Master Database updates (including real-time database updates and Real-Time Security Updates).

Every time you use the **./WCGAdmin start** command to start Websense Security Gateway, a database download is initiated. If the download fails, a new download is attempted every 15 minutes until a successful download occurs.

The default interval for database update checks is 15 minutes. You can change this interval by editing the **PollInterval** value in the **/opt/bin/downloadservice.ini** file on the Websense Content Gateway machine.

After editing the **downloadservice.ini** file, you must stop and then restart Websense Content Gateway from the command line.

- To stop, enter: **/opt/WCG/WCGAdmin stop**
- To restart, enter: **/opt/WCG/WCGAdmin start**

Scanning options

Use the **Settings > Real-Time Scanning** page to enable and configure real-time options. The individual scanning options are detailed in the following sections.

- *Categorizing content and scanning for threats*, page 134
- *File scanning*, page 135
- *Stripping content*, page 136

For each option, you have at least 2 choices:

- **Off.** No real-time scanning or blocking occurs. This option provides no additional security.
- **Recommended** or **On.** If your site is configured for real-time scanning, this setting provides your best performance. Scans are performed based on 2 factors:
  - The Always Scan and Never Scan lists on the **Settings > Real-Time Scanning > Exceptions** tab (see *Refining scanning*, page 137).
  - Whether Websense software has identified the site as including dynamic content. Sites flagged as including dynamic content are scanned. The marker that identifies a site as including dynamic content is not user-configurable.

    Sites with dynamic content that appear on the Never Scan list are not scanned.
- **All.** All requested Web pages are scanned. The only exceptions are those listed on the Never Scan list.

  This option provides the highest security, but can slow system performance significantly.

  Warning

  Sites on the Never Scan list are not analyzed under any circumstances. If a site on the Never Scan list is compromised, real-time options do not analyze and detect the malicious code.

Categorizing content and scanning for threats

Related topics:

- *Scanning options*, page 133
- *File scanning*, page 135
- *Stripping content*, page 136
- *Refining scanning*, page 137
- *Reporting on real-time scanning activity*, page 139

Web content changes rapidly. Statistics have shown that a significant majority of Web content is dynamic. In addition, the Internet is hosting more user-generated content, such as that found on social-networking sites. This material is not subject to the content and style guidelines that govern corporate Web sites.

When content categorization is enabled, selected sites are categorized in real-time, and the resulting category is forwarded to the Websense filtering software to be blocked or permitted based on the active policy.

Important

Enable full URL logging (see *Configuring full URL logging*, page 296) if you plan to generate reports of real-time scanning activity. Otherwise, log records include only the domain (www.domain.com) of the site categorized, and individual pages within a site may fit into different categories.

If your site uses WebCatcher to report uncategorized URLs to Websense, Inc. (see *Configuring WebCatcher*, page 289), URLs categorized through content categorization are forwarded for inclusion in the Master Database.

If your subscription includes the Websense Security Gateway, you can also specify that sites be scanned for security threats.

Use the **Settings > Real-Time Scanning > Common Options** page to specify when to use content categorization and content scanning.

1. In the Content Categorization area, select **Off** or **On** (default) to determine whether scanning is performed. See *Scanning options*, page 133.

   After the category is determined, any other real-time options that you have configured are applied to provide additional security.
2. (*Websense Security Gateway*) In the Content Scanning area, select **Off** (default), **Recommended**, or **All** to determine the level of scanning.

3. Do one of the following:
   - To add sites to the Never Scan or Always Scan lists, select the **Exceptions** tab. See *Refining scanning*, page 137.
   - To change the settings for other real-time options, continue on the **Common Options** page. See *File scanning*, page 135, and *Stripping content*, page 136.
4. When you are finished, click **OK** to cache your changes. Changes are not implemented until you click **Save All**.

Presentation reports can provide details about attempts to access sites containing threats. See *Presentation reports*, page 88, for details on running Websense reports.

File scanning

Related topics:

- *Scanning options*, page 133
- *Categorizing content and scanning for threats*, page 134
- *Stripping content*, page 136
- *Refining scanning*, page 137
- *Reporting on real-time scanning activity*, page 139

File scanning looks at content in incoming application files that users attempt to download or open remotely. This real-time option returns a category to Websense filtering software so that the file is permitted or blocked appropriately.

As a best practice, scan all **executable** files (for example, **.exe** and **.dll** files). You can also identify additional types of files to scan, and set a maximum size for scanning.

Note

Only Windows 32-bit portable application files are scanned.

Use the **Settings > Real-Time Scanning > Common Options** tab to specify when to use file scanning.

1. In the File Scanning area, select **Off**, **Recommended** (default), or **All** to determine the level of scanning. See *Scanning options*, page 133.
2. Click **Advanced Settings**.
3. **Scan all types of files with executable content** is selected by default. Clear this check box if you prefer to list individual file extensions to scan.

4. To specify additional file types to scan, enter the file extension (such as **ppt** or **wmv**), and then click **Add**. The file extension can contain only alphanumeric characters, an underscore (_), or a dash (-). Do not include the dot that precedes the extension.

   To remove a file extension from the Selected file extensions list, select the extension, and click **Remove**.
5. Under Options, enter the maximum size for files to be scanned (by default, 10 MB). Select **Custom** to enter a size up to 4096 MB (4 GB). Files larger than the specified size are not scanned.
6. Do one of the following:
   - If you want to add sites to the Never Scan or Always Scan lists, select the **Exceptions** tab. See *Refining scanning*, page 137.
   - If you want to change the settings for other real-time options, continue on the **Common Options** tab. See *Categorizing content and scanning for threats*, page 134, and *Stripping content*, page 136.
7. When you are finished, click **OK** to cache your changes. Changes are not implemented until you click **Save All**.

Several presentation reports provide details about attempts to download files containing security risks. See *Presentation reports*, page 88, for instructions on running Websense reports.

See *Managing traffic based on file type*, page 176, for information about blocking files based on type and URL category.

Stripping content

Related topics:

- *Scanning options*, page 133
- *Categorizing content and scanning for threats*, page 134
- *File scanning*, page 135
- *Refining scanning*, page 137
- *Reporting on real-time scanning activity*, page 139

Threats to your system can be hiding in active content sent via Web pages. One way to preserve the integrity of your system is to ensure that such content never arrives.

The Websense real-time options make it possible to specify that content in particular scripting languages (ActiveX, JavaScript, or VB Script) be stripped from incoming Web pages. If content stripping is enabled, all content in the specified scripting languages is removed from sites flagged as containing dynamic content or appearing on the Always Scan list (see *Scanning options*, page 133).

Content is removed only after the real-time options have categorized the site and Websense filtering software has determined which policy applies.

Important

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 > Real-Time Scanning > Common Options** tab to specify when to strip content from sites with dynamic content.

1. In the Content Stripping area, select the types of active content that should be removed from incoming Web pages.
2. To change the settings for other real-time options, see:
   - *Categorizing content and scanning for threats*, page 134
   - *File scanning*, page 135.
3. When you are finished, click **OK** to cache your changes. Changes are not implemented until you click **Save All**.

To disable content stripping any selected language, clear the associated check box.

Refining scanning

Related topics:

- *Scanning options*, page 133
- *Categorizing content and scanning for threats*, page 134
- *File scanning*, page 135
- *Stripping content*, page 136

Use the Always Scan and Never Scan lists to customize the behavior of the Recommended and All scanning options.

- When a real-time option is set to Recommended or On, sites with dynamic content and sites on the Always Scan list are scanned (see *Scanning options*, page 133). Sites on the Never Scan list are ignored.
- When a real-time option is set to All, sites on the Never Scan list are ignored. This can improve performance.

Use the Never Scan list with caution. If a site on this list is compromised, Websense Security Gateway does not scan that site to catch the security problem.

Use the **Settings > Real-Time Scanning > Exceptions** page to populate and edit the Always Scan and Never Scan lists.

To add sites to the Always Scan or Never Scan list:

1. Enter site names in the **URLs** box.

   Enter only the host name (for example, **thissite.com**). It is not necessary to enter the full URL. Be sure to enter both the domain and the extension; **thissite.com** and **thissite.net** are distinct entries.

   You can enter more than one host name at a time.
2. In the **Options** column, select which real-time options apply to all of the sites that you have entered. You can select one or more options. Note that **Security threats** refers only to content scanning, not to file scanning. File scanning is not affected by the Always Scan and Never Scan lists.

   To apply different options to different sites, enter the sites separately.
3. Select **Add to Always Scan** or **Add to Never Scan**.

   A site can appear in only one of the 2 lists. You cannot, for example, specify that the same site should always be scanned for threats and never for content stripping.

   - To change which list a site appears in, first select the site, and then use the right arrow (>) and left arrow (>) buttons to move the site to a new list.
   - To delete a site from either list, select the site, and then click **Remove**.
4. When you are finished, click **OK** to cache your changes. Changes are not implemented until you click **Save All**.

To change the scanning options associated with a site:

1. Select the site in the Always Scan or Never Scan list, and then click **Edit**.
2. In the Edit Rules box, select the new options for that host name:
   - **No change** maintains the current setting.
   - **On** indicates that content is scanned for the specified option, such as content categorization.
   - **Off** indicates that no scanning occurs for the specified option. If an option is turned off, performance can improve, but security can be compromised.
3. When you are finished making changes, click **OK** in the Edit Rules box to return to the Exceptions tab.
4. Click **OK** again to cache your changes. Changes are not implemented until you click **Save All**.

Reporting on real-time scanning activity

Related topics:

- *Scanning options*, page 133
- *Categorizing content and scanning for threats*, page 134
- *File scanning*, page 135
- *Stripping content*, page 136

If your subscription includes real-time scanning features, you can analyze the effects of these features with presentation reports and investigative reports.

On the Presentation Reports page, a group of reports called Real Time Security Threats is available. These reports focus specifically on threat-related activity. As with all presentation reports, you can copy a security threat report and edit its report filter to refine the information included when you generate a report from that copy.

Some security threat reports include a Threat ID column. You can click the individual threat ID to open a Websense Security Labs page that describes the type of threat identified.

Additionally, other presentation reports contain information on real-time scanning activities, as well as standard filtering activities. Copy a predefined report and edit its filter to create a report specific to real-time scanning activities.

Important

Enable full URL logging (see *Configuring full URL logging*, page 296) to ensure that reports of real-time scanning 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 contain different threats.

As an 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 real-time scanning, copy the Detail of Full URLs by Category report and edit its report filter. On the Actions tab, select only permitted and blocked actions that relate to real-time scanning. On the Options tab, change the report catalog title and report name to identify this as a real-time scanning report. For example, you might change the name and title to Real-Time: Detail of Full URLs by Category.

Investigative reports can also be used to gain insight into real-time scanning activities.

1. In the **Internet use by** drop-down list, select Action.

2. In the resulting report, click a real-time 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 105, for details on using all the investigative reports features.

How real-time scanning is logged

When you use real-time scanning options, note that there are differences in the way that standard Web filtering activity and real-time scanning activity are logged.

For standard Web filtering, you have several options to reduce the size of the Log Database.

- Enable **visits** to log only one record for each Web site requested. See *Configuring log cache files*, page 286.
- Enable **consolidation** to combine into a single log record multiple requests with certain common elements. See *Configuring consolidation options*, page 287.
- 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 full URL logging*, page 296.
- Enable **selective category logging** to limit logging to selected categories that are crucial for your organization. See *Configuring Filtering Service for logging*, page 280.

Real-time scanning features, however, are bound only partially by these settings. When real-time scanning analyzes a site, it creates 2 separate log records.

- **Web filter records** take advantage of any size reduction settings that have been implemented, and are available for all Web filter reports.
- **Real-time records** ignore most size reduction settings. Every separate hit is logged, requests to all categories are logged, and no records are consolidated. A real-time record is generated regardless of whether the site is blocked or permitted as a result of real-time scanning. Only the setting for full URL logging is honored for real-time records.

If you have enabled any Log Database size reduction options, the numbers that appear in real-time reports may **not** match those that appear in standard filtering reports, even when the reports are configured for the same users, time periods, and categories. For example, if you have chosen to log visits, and a user requests a site analyzed by real-time scanning features, that user request appears as one visit in standard filtering reports, but may show as multiple hits in real-time reports.

To see comparable data for standard and real-time filtering, **disable** the Log Database size reduction settings. Because this may result in a very large and fast-growing

database, make sure that the Log Database machine has adequate hard disk, processing, and memory capacity.

See *Reporting Administration*, page 275, for more information on configuring size reduction settings. See *Presentation reports*, page 88, and *Investigative reports*, page 105, for information on generating reports.

8 Filter Remote Clients

Related topics:

- *How Remote Filtering works*, page 144
- *Configuring Remote Filtering settings*, page 150

Many organizations have users who sometimes take their laptop computers outside the network. For remote users who run a Microsoft Windows operating system, you can filter Internet requests by implementing Websense Remote Filtering, an optional feature available for both Websense Web Security and Websense Web Filter.

Remote Filtering monitors HTTP, SSL, and FTP traffic, applying either the policy assigned to the individual user or group, or the Default policy, depending on how the user logs on to the remote computer. Remote Filtering does not filter on the basis of policies assigned to computers or network ranges. See *Identifying remote users*, page 147, for more information.

Bandwidth-based filtering is not supported for remote clients (see *Using Bandwidth Optimizer to manage bandwidth*, page 174). Bandwidth generated by remote traffic is not included in bandwidth measurements and reports.

Remote Filtering of FTP and SSL requests, such as HTTPS, can be blocked or permitted only. If a remote user requests an FTP site or HTTPS site, for example, from a category assigned the quota or confirm action, the site is blocked for Remote Filtering clients. When these computers browse from inside the network, the quota and confirm filtering actions are applied normally.

To implement Remote Filtering, you must install the following components:

- Remote Filtering Server must be inside the outermost firewall, and the remote computers must be permitted to communicate with it. Typically, it is installed in the network's *demilitarized zone*, or DMZ, outside the firewall that protects the rest of the network. You can install up to 3 Remote Filtering Servers to provide failover capabilities.

- Remote Filtering Client must be on each computer that runs a Windows operating system and is used outside the network.

  Notes

  Follow the recommendations in the *Deployment Guide* carefully to deploy these components. See the *Installation Guide* for instructions on installing them.

  If you are using Websense software in stand-alone mode (without an integration product), configure Network Agent **not** to monitor the Remote Filtering Server machine (see *Configuring global settings*, page 314).

All communication between Remote Filtering Client and Remote Filtering Server is authenticated and encrypted.

How Remote Filtering works

Related topics:

- *Inside the network*, page 145
- *Outside the network*, page 146
- *Identifying remote users*, page 147
- *When server communication fails*, page 147
- *Virtual Private Network (VPN)*, page 149
- *Configuring Remote Filtering settings*, page 150

Whenever a remote computer makes an HTTP, SSL, or FTP request, its Remote Filtering Client communicates with Remote Filtering Server. Remote Filtering Server communicates with Websense Filtering Service to determine what action applies. Remote Filtering Server then responds to Remote Filtering Client, either permitting the site, or sending the appropriate block message.

When the browser on a computer running Remote Filtering Client makes a request via HTTP, SSL or FTP, Remote Filtering Client must decide whether to query Remote Filtering Server about the request. This determination is controlled by the location of the computer relative to the network.

Inside the network

Related topics:

- *How Remote Filtering works*, page 144
- *Outside the network*, page 146
- *Identifying remote users*, page 147
- *When server communication fails*, page 147
- *Virtual Private Network (VPN)*, page 149
- *Configuring Remote Filtering settings*, page 150

When a computer is started *inside* the network, the Remote Filtering Client attempts to send a **heartbeat** to the Remote Filtering Server in the DMZ. The heartbeat is successful because the heartbeat port is open on the internal firewall.

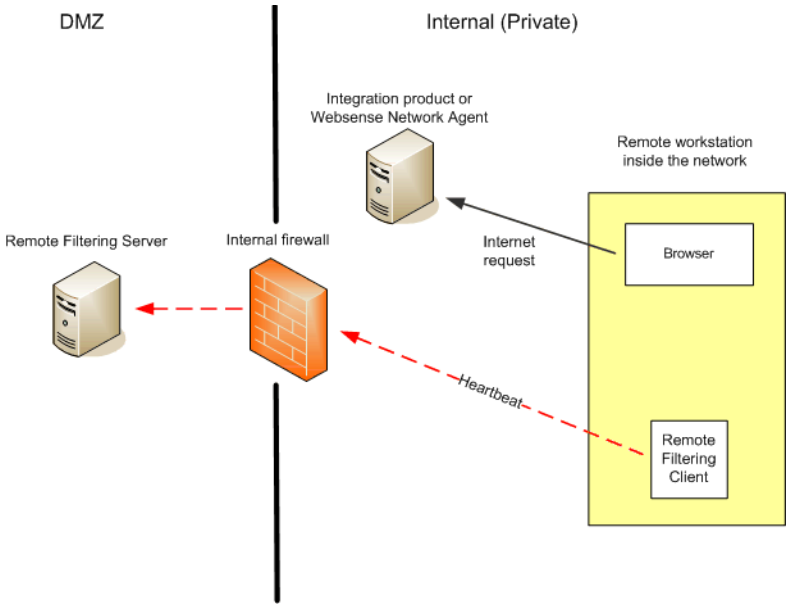

In this case, Remote Filtering Client becomes passive and does not query Remote Filtering Server about Internet requests. Instead, these requests are passed directly to the integration product (such as Cisco Pix, Microsoft ISA Server), or to Websense Network Agent. The request is filtered like any other internal request.

Outside the network

Related topics:

- *How Remote Filtering works*, page 144
- *Inside the network*, page 145
- *Identifying remote users*, page 147
- *When server communication fails*, page 147
- *Virtual Private Network (VPN)*, page 149
- *Configuring Remote Filtering settings*, page 150

When a computer is started *outside* the network, Remote Filtering Client attempts to send a heartbeat to Remote Filtering Server. The heartbeat is unsuccessful because the heartbeat port is blocked at the external firewall.

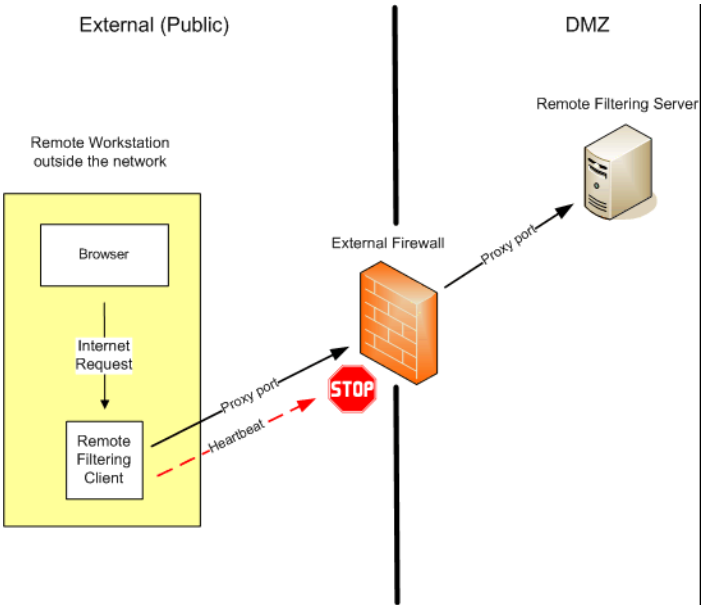

This heartbeat failure prompts Remote Filtering Client to send a query about each HTTP, SSL, or FTP request over the configured port (default 80) to Remote Filtering Server in the DMZ. Remote Filtering Server then forwards the filtering request to Websense Filtering Service inside the network. Filtering Service evaluates the request and sends a response to Remote Filtering Server. The response is then sent to the remote computer. If the site is blocked, Remote Filtering Client requests and receives the appropriate block page, which is then displayed to the user.

Remote Filtering Client delays each filtered request until it receives a response from Remote Filtering Server. Depending on the response received, Remote Filtering Client either permits the site or displays the block page.

A log file tracks Remote Filtering activities, such as entering and leaving the network, failing open or closed, and restarting the client. Remote Filtering Client creates the log

file when it starts for the first time. You control the presence and size of this log file. See *Configuring Remote Filtering settings*, page 150.

Identifying remote users

Related topics:

- *How Remote Filtering works*, page 144
- *Inside the network*, page 145
- *Outside the network*, page 146
- *When server communication fails*, page 147
- *Virtual Private Network (VPN)*, page 149
- *Configuring Remote Filtering settings*, page 150

How a user logs on to a remote computer determines which policy is enforced.

If a user logs on using cached domain credentials (network directory logon information), Websense Filtering Service is able to resolve the user name, and applies appropriate user and group-based policies to the remote computer. Additionally, Internet activity is logged under the network user name.

If the user logs on with a user account that is local to the computer, Filtering Service cannot resolve the user name and applies the Default policy instead. Internet activity is logged under the local user name. Remote Filtering does not filter on the basis of policies assigned to computers or network ranges.

Note

Remote users are always filtered according to their logon credentials, as described here. Selective authentication settings do not apply to these users.

When server communication fails

Related topics:

- *How Remote Filtering works*, page 144
- *Inside the network*, page 145
- *Outside the network*, page 146
- *Identifying remote users*, page 147
- *Virtual Private Network (VPN)*, page 149
- *Configuring Remote Filtering settings*, page 150

Filtering occurs when Remote Filtering Client, outside the network, successfully communicates with Remote Filtering Server in the network DMZ. However, there may be times when that communication is unsuccessful.

The action Remote Filtering Client takes if it cannot contact Remote Filtering Server is configurable. By default, Remote Filtering Client uses the **fail open** setting, which permits all HTTP, SSL, and FTP requests when communication between these components cannot be established. Remote Filtering Client continues attempting to contact Remote Filtering Server. When the communication is successful, the appropriate filtering policy is enforced.

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 Web access at startup, no filtering occurs (all requests are permitted) until Remote Filtering Client connects to the Remote Filtering Server. When this occurs, the appropriate filtering policy is enforced.

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 Websense Filtering Service for any reason, an error is returned to the Remote Filtering Client, and filtering always fails open.

This timeout period allows users who pay for Internet access when travelling to start the computer and arrange for connection without being locked out. If the user does not establish Web access before the 15 minute timeout period expires, Web access cannot be established during that session. When this occurs, the user must restart the computer to begin the timeout interval again.

To change the fail open/fail closed setting, and change the timeout value, see *Configuring Remote Filtering settings*, page 150.

Virtual Private Network (VPN)

Related topics:

- *How Remote Filtering works*, page 144
- *Inside the network*, page 145
- *Outside the network*, page 146
- *Identifying remote users*, page 147
- *When server communication fails*, page 147
- *Configuring Remote Filtering settings*, page 150

Websense Remote Filtering supports VPN connections, including split-tunneled VPN. When a remote computer connects to the internal network via VPN (non split-tunneled), Remote Filtering Client is able to send a heartbeat to Remote Filtering Server. As a result, Remote Filtering Client becomes passive and all HTTP, SSL, and FTP requests from the remote computer are filtered by the internal integration product or Network Agent, like other in-network computers.

If the remote computer connects to the internal network via a split-tunneled VPN client, Remote Filtering Client detects this and does not send a heartbeat to Remote Filtering Server. Remote Filtering Client assumes that it is operating externally and submits requests to Remote Filtering Server for filtering.

Websense software supports split-tunneling for the following VPN clients:

- Checkpoint SecureClient
- Cisco
- Juniper/Netscreen
- Microsoft PPTP
- Nokia
- Nortel
- SonicWALL

Configuring Remote Filtering settings

Related topics:

- *How Remote Filtering works*, page 144
- *Inside the network*, page 145
- *Outside the network*, page 146
- *Identifying remote users*, page 147
- *When server communication fails*, page 147
- *Virtual Private Network (VPN)*, page 149

Unconditional Super Administrators can use the **Settings > General > Remote Filtering** page to configure options that affect all Remote Filtering Clients associated with this installation.

For details on how Remote Filtering operates, see *How Remote Filtering works*, page 144.

1. Select the **Fail closed** check box to block Remote Filtering Clients from all Internet access unless their computer is communicating with Remote Filtering Server.

   Be default, this is not selected, which means remote users have unfiltered access to the Internet when their computers cannot communicate with the Remote Filtering Server.

2. If you marked the Fail closed option, use the **Fail closed timeout** field to select a number of minutes up to 60 (default 15), or choose **No timeout**.

   During the timeout period, all HTTP, SSL, and FTP requests are permitted.

   If the Remote Filtering Client cannot communicate with Remote Filtering Server during the timeout interval, all Internet access will be blocked (fail closed).

   Choosing **No timeout** may lock out a remote computer before the user can establish Internet connection from a hotel or other pay-for-use-provider. Additionally, Remote Filtering Client attempts to communicate with Remote Filtering Server continuously.

   

   Warning

   Websense, Inc., does not recommend choosing **No timeout** or setting the timeout period to a very low number.

3. Select a **Maximum size for the local log cache** size (in megabytes), up to 10. Choose **No Log** to disable logging.

   This controls the size and existence of the log file the remote computer creates when it is initially disconnected from the Remote Filtering Server. This log file tracks the following events:

- The computer leaves the network
- The computer rejoins the network
- The Remote Filtering Client is restarted
- Fail open condition occurs
- Fail closed condition occurs
- Remote Filtering Client receives a policy update

The computer retains the 2 most recent logs. These logs can be used to troubleshoot connection issues or other problems with Remote Filtering.

9 Refine Filtering Policies

At it's simplest, Internet usage filtering requires a single policy that applies one category filter and one protocol filter 24 hours a day, 7 days a week. Websense software offers tools, however, for going far beyond this basic filtering, to achieve precisely the level of granularity you need to manage Internet usage. You can:

- Create **limited access filters** to block access to all but a specified list of sites for certain users (see *Restricting users to a defined list of Internet sites*, page 153).
- Create **custom categories** to redefine how selected sites are filtered (see *Working with categories*, page 160).
- **Recategorize URLs** to move specific sites from their default, Master Database category to another Websense-defined or custom category (see *Recategorizing URLs*, page 168).
- Define **unfiltered URLs** to allow users to access specific sites, even when the sites are assigned to a blocked category in the active category filter (see *Defining unfiltered URLs*, page 167).
- Implement **bandwidth** restrictions, blocking users from accessing otherwise permitted categories and protocols when bandwidth usage reaches a specified threshold.
- Define **keywords** used to block sites in otherwise permitted categories when keyword blocking is enabled and activated (see *Filtering based on keyword*, page 164).
- 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 176).

Restricting users to a defined list of Internet sites

Related topics:

- *Limited access filters and filtering precedence*, page 154
- *Creating a limited access filter*, page 155
- *Editing a limited access filter*, page 156

Limited access filters provide a very precise method of filtering Internet access. Each limited access filter is a list of individual Web sites. Like category filters, limited access filters are added to policies and enforced during a specified time period. When a limited access filter is active in a policy, users assigned that policy can visit only sites in the list. All other sites are blocked.

For example, if the First Grade policy enforces a limited access filter that includes only certain educational and reference sites, students governed by the First Grade policy can visit only those sites, and no others.

Important

When a limited access filter is in effect, Websense software checks to see only if a requested site appears in the filter. No other checking is performed.

This means that if a site permitted by the filter becomes infected with malicious code, user requests for that site are still permitted, regardless of the site's Master Database or Real-Time Scanning categorization.

When a limited access filter is active, a block page is returned for any requested URL not included in that filter.

Websense software can support up to 2,500 limited access filters containing 25,000 URLs in total.

Limited access filters and filtering precedence

In some cases, more than one filtering policy could apply to a single user. This happens when a user belongs to more than one group, and the groups are governed by different policies. In addition, a URL might both appear in a limited access filter and be defined as an unfiltered URL.

When multiple group policies apply to a user, the **Use more restrictive blocking** setting (see *Filtering order*, page 71) determines how the user is filtered. By default, this setting is off.

Websense software determines which filtering setting is less restrictive at the filter level. In cases where a user might be filtered by multiple policies, one of which is enforcing a limited access filter, "less restrictive" may sometimes seem counterintuitive.

When **Use more restrictive blocking** is **OFF**:

- If the **Block All** category filter and a limited access filter could apply, the limited access filter is always considered less restrictive.
- If any other category filter and a limited access filter could apply, the category filter is considered less restrictive.

  This means that even when the limited access filter permits the site and the category filter blocks the site, the site is blocked.

When **Use more restrictive blocking** is **ON**, a limited access filter is considered more restrictive than any category filter except Block All.

The table below summarizes how the **Use more restrictive blocking** setting affects filtering when multiple policies could apply:

	***Use more restrictive blocking* OFF**	***Use more restrictive blocking* ON**
limited access filter + **Block All** category filter	limited access filter (request permitted)	**Block All** (request blocked)
limited access filter + permitted category	category filter (request permitted)	limited access filter (request permitted)
limited access filter + blocked category	category filter (request blocked)	limited access filter (request permitted)
limited access filter + Quota/Confirm category	category filter (request limited by quota/ confirm)	limited access filter (request permitted)
limited access filter + unfiltered URL	unfiltered URL (request permitted)	limited access filter (request permitted)

Creating a limited access filter

Related topics:

- *Working with filters*, page 43
- *Restricting users to a defined list of Internet sites*, page 153
- *Editing a limited access filter*, page 156

Use the **Add Limited Access Filter** page (accessed via the **Filters** or **Edit Policy** page) to give your new filter a unique name and a description. After creating the filter, enter a list of permitted URLs, assign the filter to a policy, and apply the policy to clients.

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

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

   Filter names can include spaces, dashes, and apostrophes.
2. Enter a short **Description** of the filter. This description appears next to the filter name in the Limited Access Filters section of the Filters page, and should explain the filter's purpose to help administrators manage policies over time.

   The character restrictions that apply to filter names also apply to descriptions, with 2 exceptions: descriptions can include periods (.) and commas (,).
3. To see and edit the new filter, click **OK**. To abandon your changes and return to the Filters page, click **Cancel**.

When you create a new limited access filter, it is added to the **Policy Management > Filters > Limited Access Filters** list. Click a filter name to edit the filter.

To finish customizing your new filter, continue with *Editing a limited access filter*.

Editing a limited access filter

Related topics:

- *Restricting users to a defined list of Internet sites*, page 153
- *Limited access filters and filtering precedence*, page 154
- *Creating a limited access filter*, page 155
- *Editing a policy*, page 68

A limited access filter is a list of Web sites (URLs or IP addresses) and regular expressions, used to identify specific sites that users can access. When the filter is applied to clients, those clients cannot visit any site that is not in the list.

Important

When a limited access filter is in effect, Websense software checks to see only if a requested site appears in the filter. No other checking is performed.

This means that if a site permitted by the filter becomes infected with malicious code, user requests for that site are still permitted, regardless of the site's Master Database or Real-Time Scanning categorization.

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 polices that enforce the filter, and manage which sites are included in the filter.

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

1. Verify the filter name and description. To change the filter name, click **Rename**, and then enter the new name. The name is updated in all policies that enforce the selected limited access filter.
2. Use the **Policies using this filter** field to see how many policies currently enforce this filter. If 1 or more policies enforce the filter, click **View policies** to list them.
3. Under Add or Remove Sites, enter the URLs and IP addresses that you want to add to the limited access filter. Enter one URL or IP address per line.

   It is not necessary to include the HTTP:// prefix.

When a site is filtered according to its Master Database category, Websense software matches the URL with its equivalent IP address. This is not the case for limited access filters. To permit a site's URL and IP address, add both to the filter.

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 179, for detailed information about using regular expressions for filtering.
6. Review the URLs, IP addresses, and regular expressions in the **Permitted sites** list.
   - To make changes to a site or expression, select it and click **Edit**.
   - To remove a site or expression from the list, select it and click **Delete**.
7. After editing the filter, click **OK** to cache your changes and return to the Filters page. Changes are not implemented until you click **Save All**.

Adding sites from the Edit Policy page

Related topics:

- *Restricting users to a defined list of Internet sites*, page 153
- *Limited access filters and filtering precedence*, page 154
- *Creating a limited access filter*, page 155
- *Editing a policy*, page 68

Use the **Policies > Edit Policy > Add Sites** page to add sites to a limited access filter.

Enter one URL or IP address per line. If you do not specify a protocol, Websense software automatically adds the **HTTP://** prefix.

When you are finished making changes, click **OK** to return to the Edit Policy page. You must also click **OK** on the Edit Policy page to cache the changes. Changes are not implemented until you click **Save All**.

Changes made to a limited access filter affect all policies that enforce the filter.

Copying filters and policies to roles

Related topics:

- *Creating a category filter*, page 44
- *Creating a protocol filter*, page 46
- *Creating a limited access filter*, page 155
- *Creating a policy*, page 68

Super Administrators can use the **Filters > Copy Filters To Role** and **Policies > Copy Policies To Role** pages to copy one or more filters or policies to a delegated administration role. Once the filter or policy has been copied, delegated administrators can use the filters or policies to filter their managed clients.

- In the target role, the tag "(Copied)" is added to the end of the filter or policy name. A number is added if the same filter or policy is copied multiple times.
- Delegated administrators can rename or edit filters or policies that have been copied to their role.
- Category filters copied to a delegated administration role set the filtering action to Permit for custom categories created in the role. Delegated administrators should update the copied category filters to set the desired action for their role-specific custom categories.
- Changes made by a delegated administrator to a filter or policy copied to their role by a Super Administrator do not affect the Super Administrator's original filter or policy, or any other role that received a copy of the filter or policy.
- Filter Lock restrictions do not affect the Super Administrator's original filter or policy, but they do affect the delegated administrator's copy of the filter or policy.
- Because delegated administrators are affected by Filter Lock restrictions, the Permit All category and protocol filters cannot be copied to a delegated administration role.

To copy a filter or policy:

1. On the Copy Filters to Role or Copy Policies to Role page, verify that the correct policies or filters appear in the list at the top of the page.
2. Use the **Select a role** drop-down list to select a destination role.
3. Click **OK**.

   A pop-up dialog box indicates that the selected filters or policies are being copied. The copy process may take a while.

   The changes are not implemented until you click **Save All**.

After the copy process is complete, the copied filters or policies will be available to delegated administrators in the selected role the next time they log on to Websense 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

Use the **Policy Management > Filter Components** page to access tools used to refine and customize the way that Websense software enforces your organization's Internet access policies. The 4 buttons on the screen are associated with the following tasks:

Edit Categories	• Recategorize a URL (see *Redefining filtering for specific sites*, page 166). For example, if the Shopping category is blocked by your Internet filtering policies, but you want to permit access to specific supplier or partner sites, you could move those sites to a permitted category, like Business and Economy. • Define or edit custom categories (see *Creating a custom category*, page 163). Create additional subcategories within Websense-defined parent categories, or within the User-Defined parent category, and then assign URLs to the new categories. • Assign keywords to a category (see *Filtering based on keyword*, page 164). 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 179), patterns or templates that can be used to match multiple URLs and assign them to a category.
Edit Protocols	Define or edit custom protocol definitions (see *Creating a custom protocol*, page 172, and *Editing custom protocols*, page 170). For example, if members of your organization use a custom messaging tool, you could create a custom protocol definition to permit use of that tool while blocking other Instant Messaging / Chat protocols.
File Types	Create or edit file type definitions, used to block specific types of files within otherwise permitted categories (see *Managing traffic based on file type*, page 176).
Unfiltered URLs	Define specific sites to permit for all clients, even when they belong to a blocked category (see *Defining unfiltered URLs*, page 167). Note that adding a URL to this list does not override the Block All category filter or limited access filters.

Working with categories

Related topics:

- *Editing categories and their attributes*, page 160
- *Creating a custom category*, page 163
- *Filtering based on keyword*, page 164
- *Redefining filtering for specific sites*, page 166

Websense software provides multiple methods for filtering sites that are not in the Master Database, and for changing the way that individual sites in the Master Database are filtered.

- Create **custom categories** for more precise filtering and reporting.
- Use **recategorized URLs** to define categories for uncategorized sites, or to change the category for sites that appear in the Master Database.
- Define **keywords** to recategorize all sites whose URL contains a certain string.

Editing categories and their attributes

Related topics:

- *Creating a custom category*, page 163
- *Reviewing all customized category attributes*, page 161
- *Making global category filtering changes*, page 162
- *Filtering based on keyword*, page 164
- *Redefining filtering for specific sites*, page 166

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 161, for more information.

- To create a new category, click **Add**, and then go to *Creating a custom category*, page 163, 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 162).
- To change the filtering action associated with a category in all category filters, click **Override Action** (see *Making global category filtering changes*, page 162).
- The **Recategorized URLs** list shows which recategorized sites (URLs and IP addresses) have been assigned to this category.
  - To add a site to the list, click **Add URLs**. See *Recategorizing URLs*, page 168, for further instructions.
  - To change an existing recategorized site, select the URL or IP address, and then click **Edit**.
- The **Keywords** list shows which keywords have been associated with this category.
  - To define a keyword associated with the selected category, click **Add Keywords**. See *Filtering based on keyword*, page 164, 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 179).
  - To change an existing regular expression, select the expression, and then click **Edit**.
- To delete a recategorized URL, keyword, or regular expression, select the item to remove, and then click **Delete**.

When you are finished making changes on the Edit Categories page, click **OK** to cache the changes and return to the Filter Components page. Changes are not implemented until you click **Save All**.

Reviewing all customized category attributes

Use the **Filter Components > Edit Categories > View All Custom URLs and Keywords** page to review custom URL, keyword, and regular expression definitions. You can also delete definitions that are no longer needed.

The page contains 3 similar tables, one for each category attribute: custom URLs, keywords, or regular expressions. In each table, the attribute is listed next to the name of the category with which it is associated.

To delete a category attribute, mark the appropriate check box, and then click **Delete**.

To return to the Edit Categories page, click **Close**. If you deleted any items on the View All Custom URLs and Keywords page, click **OK** on the Edit Categories page to cache the changes. Changes are not implemented until you click **Save All**.

Making global category filtering changes

Use the **Filter Components > Edit Categories > Override Action** page to change the action applied to a category in all existing category filters. This also determines the default action applied to the category in new filters.

Although this change overrides the action applied to the category in all existing filters, administrators can later edit those filters to apply a different action.

Before changing the filtering settings applied to a category, first verify that the correct category name appears next to **Selected Category**. Next, you can:

1. Chose a new **Action** (Permit, Block, Confirm, or Quota). See *Filtering actions*, page 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 *Filtering based on keyword*, page 164, for more information.
3. Specify whether or not to **Block File Types**, and customize blocking settings. See *Managing traffic based on file type*, page 176, for more information.
4. Under **Advanced Filtering**, specify whether or not to use Bandwidth Optimizer to manage access to HTTP sites, and customize blocking settings. See *Using Bandwidth Optimizer to manage bandwidth*, page 174, 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 160). The changes are not cached until you click **OK** on the Edit Categories page.

Renaming a custom category

Use the **Filter Components > Edit Categories > Rename Category** page to change the name or description associated with a custom category.

- Use the **Filter name** field to edit the category name. The new name must be unique, and cannot exceed 50 characters.

  The name cannot include any of the following characters:

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

- Use the **Description** field to edit the category description. The description cannot exceed 255 characters.

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

When you are finished making changes, click **OK** to return to the Edit Categories page. The changes are not cached until you click **OK** on the Edit Categories page.

Creating a custom category

Related topics:

- *Editing categories and their attributes*, page 160
- *Filtering based on keyword*, page 164
- *Redefining filtering for specific sites*, page 166

In addition to using the more than 90 Websense-defined categories in the Master Database, you can define your own **custom categories** to provide more precise filtering and reporting. For example, create custom categories like:

- **Business Travel**, to group sites from approved vendors that employees can use to buy airplane tickets and make rental car and hotel reservations
- **Reference Materials**, to group online dictionary and encyclopedia sites deemed appropriate for elementary school students
- **Professional Development**, to group training sites and other resources that employees are encouraged to use to build their skills

Use the **Policy Management > Filter Components > Edit Categories > Add Category** page to add custom categories to any parent category. You can create up to 100 custom categories.

1. Enter a unique, descriptive **Category name**. The name cannot include any of the following characters:

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

2. Enter a **Description** for the new category.

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

3. Select a parent category from the **Add to** list. By default, **All Categories** is selected.
4. Enter the sites (URLs or IP addresses) that you want to add to this category. See *Recategorizing URLs*, page 168, for more information.

   You can also edit this list after creating the category.

5. Enter the keywords that you want to associate with this category. See *Filtering based on keyword*, page 164, for more information.

   You can also edit this list after creating the category.

6. Define a default filtering **Action** to apply to this category in all existing category filters. You can edit this action in individual filters later.

   Note
   Category filters copied to a delegated administration role set the filtering action to Permit for custom categories created in the role. Delegated administrators should update the copied category filters to set the desired action for their role-specific custom categories.

7. Enable any **Advanced Filtering** actions (keyword blocking, file type blocking, or bandwidth blocking) that should be applied to this category in all existing category filters.
8. When you are finished defining the new category, click **OK** to cache changes and return to the Edit Categories page. Changes are not implemented until you click **Save All**.

The new category is added to the Categories list and custom URL and keyword information for the category is displayed.

Filtering based on keyword

Related topics:

- *Recategorizing URLs*, page 168
- *Configuring Websense filtering settings*, page 50
- *Creating a category filter*, page 44
- *Editing a category filter*, page 45
- *Working with categories*, page 160

Keywords are associated with categories, and then used to offer protection against sites that have not explicitly been added to the Master Database or defined as a custom URL. Three steps are necessary to enable keyword blocking:

1. Enable keyword blocking at a global level (see *Configuring Websense filtering settings*, page 50).
2. Define keywords associated with a category (see *Defining keywords*, page 165).
3. Enable keyword blocking for the category in an active category filter (see *Editing a category filter*, page 45).

When keywords have been defined and keyword blocking is enabled for a specific category, Websense software blocks any site whose URL contains a keyword, and logs the site as belonging to the specified category. The site is blocked even if other URLs in the category are permitted.

For example, if the Sports category is permitted in an active category filter, but you want to block access to basketball sites, you might associate the keyword “nba” with Sports, and enable keyword blocking. This means that the following URLs are blocked, and logged as belonging to the Sports category:

- sports.espn.go.com/**nba**/
- moder**nba**kery.com
- moder**nba**biesandchildren.com
- fashio**nba**r.com

Be cautious when defining keywords to avoid unintended overblocking.

Important

If you are using Websense Web Security, avoid associating keywords with any of the Extended Protection subcategories. Keyword blocking is not enforced for these categories.

When a request is blocked based on a keyword, this is indicated on the Websense block page that the user receives.

Defining keywords

Related topics:

- *Editing a category filter*, page 45
- *Working with categories*, page 160
- *Filtering based on keyword*, page 164
- *Using regular expressions*, page 179

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 m**sex**change.org (Information Technology), vega**sex**perience.com (Travel), and sci.esa.int/mar**sex**press (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.
- If you are using Websense Web Security, avoid associating keywords with any of the Extended Protection subcategories. Keyword blocking is not enforced for these categories.

When you are finished adding or editing keywords, click **OK** to cache your changes and return to the Edit Categories page. Changes are not implemented until you click **Save All**.

In order for keyword blocking to be enforced, you must also:

1. Enable keyword blocking via the **Settings > Filtering** page (see *Configuring Websense filtering settings*, page 50).
2. Enable keyword blocking in one or more active category filters (see *Editing a category filter*, page 45).

Redefining filtering for specific sites

Related topics:

- *Creating a custom category*, page 163
- *Filtering based on keyword*, page 164
- *Defining unfiltered URLs*, page 167
- *Recategorizing URLs*, page 168

With custom URLs, you can:

- Apply more precise filtering to sites that are not in the Websense Master Database. By default, the action applied to the **Miscellaneous\Uncategorized** category is used to filter these sites.
- Filter sites differently than their Master Database category.

Websense software looks for custom URL definitions for a site before consulting the Master Database, and therefore filters the site according to the category assigned to the custom URL.

There are 2 types of custom URLs: unfiltered and recategorized.

- Unfiltered URLs are permitted for all users not governed by the Block All category filter or a limited access filter (see *Defining unfiltered URLs*, page 167).
- Recategorized URLs have been moved from their Master Database category to another Websense-defined or custom category (see *Recategorizing URLs*, page 168).

A recategorized URL is not blocked by default. It is filtered according to the action applied to its new category in each active category filter.

When a site is filtered according to its Master Database category, Websense software matches the URL with its equivalent IP address. This is not the case for custom URLs. To change the way a site is filtered, define both its URL and its IP address as a custom URL.

If a site can be accessed via multiple URLs, define each URL that can be used to access the site as a custom URL to ensure that the site is permitted or blocked as intended.

If a site is moved to a new domain, and an HTTP redirect is used to send users to the new URL, the new URL is not automatically filtered the same way as the redirecting site. To make sure that the site is filtered appropriately at its new address, create a new custom URL.

Defining unfiltered URLs

Related topics:

- *Working with categories*, page 160
- *Redefining filtering for specific sites*, page 166
- *Recategorizing URLs*, page 168

Use the **Policy Management > Filter Components > Unfiltered URLs** page to define a list of sites that any user can access, except when governed by the Block All category filter or a limited access filter.

The **Permitted sites** list in the right portion of the content pane lists the unfiltered sites (URLs and IP addresses) and regular expressions that you have defined (see *Using regular expressions*, page 179). Each site is associated with a category.

- The URL may be associated with its Master Database category, or recategorized.
- When a user requests access to the unfiltered URL, the request is logged as a permitted custom URL in the category to which it has been assigned.

To add an unfiltered URL:

1. Under **Define Unfiltered URLs**, enter one URL or IP address per line, and then click the right arrow (>).

   Websense software does not match a custom URL with its equivalent IP address. To permit both the URL and the IP address for a site, add both to the Unfiltered URLs list.
2. To add regular expressions that match multiple sites, click **Advanced**. Enter one regular expression per line, and then click the right arrow to move the expressions to the Unfiltered URLs list. To verify that a pattern matches the intended sites, click **Test**.

   See *Using regular expressions*, page 179, for detailed information.

3. When you are finished, click **OK** to cache your changes and return to the Edit Categories page. Changes are not implemented until you click **Save All**.

To remove a site from the Unfiltered URLs list, select the URL, IP address, or regular expression, and then click **Delete**.

Recategorizing URLs

Related topics:

- *Working with categories*, page 160
- *Redefining filtering for specific sites*, page 166
- *Defining unfiltered URLs*, page 167

Use the **Policy Management > Filter Components > Edit Categories > Recategorize URLs** page to add individual sites to any category. Make changes to existing recategorized sites on the **Edit Recategorized URL** page.

Recategorize URLs to change the way that individual sites are filtered and logged. When you add recategorized sites:

- Enter each URL or IP address on a separate line.
- Include the protocol for any non-HTTP site. If the protocol is omitted, Websense software filters the site as an HTTP site.

  For HTTPS sites, also include the port number (https://63.212.171.196:443/, https://www.onlinebanking.com:443/).
- Websense software recognizes custom URLs exactly as they are entered. If the Search Engines and Portals category is blocked, but you recategorize **www.yahoo.com** in a permitted category, the site is permitted only if users type the full address. If a user types images.search.yahoo.com, or just yahoo.com, the site is still blocked. If you recategorize **yahoo.com**, however, all sites with yahoo.com in the address are permitted.

When you are finished adding or editing recategorized sites, click **OK** to cache your changes and return to the Edit Categories page. Changes are not implemented until you click **Save All**.

After saving recategorized URLs, use the **URL Category** tool in the right shortcut pane to verify that the site is assigned to the correct category. See *Using the Toolbox to verify filtering behavior*, page 180.

Working with protocols

The Websense Master database includes protocol definitions used to filter Internet protocols other than HTTP, HTTPS and FTP. These definitions include Internet applications and data transfer methods such as those used for instant messaging,

streaming media, file sharing, file transfer, Internet mail, and other network and database operations.

These protocol definitions can even be used to filter protocols or applications that bypass a firewall by tunneling through ports normally used by HTTP traffic. Instant messaging data, for example, can enter a network whose firewall blocks instant messaging protocols by tunneling through HTTP ports. Websense software accurately identifies these protocols, and filters them according to policies you configure.

> **Note**
> Network Agent is required to enable protocol-based filtering.

In addition to using Websense-defined protocol definitions, you can define custom protocols for filtering. Custom protocol definitions can be based on IP addresses or port numbers, and can be edited.

To block traffic over a specific port, associate that port number with a custom protocol, and then assign that protocol a default action of **Block**.

To work with custom protocol definitions, go to **Policy Management > Filter Components**, and then click **Protocols**. See *Editing custom protocols*, page 170, and *Creating a custom protocol*, page 172, for details.

Filtering protocols

Related topics:

- *Working with protocols*, page 168
- *Editing custom protocols*, page 170
- *Creating a custom protocol*, page 172
- *Adding or editing protocol identifiers*, page 171
- *Adding to a Websense-defined protocol*, page 174

When Network Agent is installed, Websense software can block Internet content transmitted over particular ports, or using specific IP addresses, or marked by certain signatures, regardless of the nature of the data. By default, blocking a port intercepts all Internet content entering your network over that port, regardless of source.

> **Note**
> Occasionally, internal network traffic sent over a particular port may not be blocked, even though the protocol using that port is blocked. The protocol may send data via an internal server more quickly than Network Agent can capture and process the data. This does not occur with data originating outside the network.

When a protocol request is made, Websense software uses the following steps to determine whether to block or permit the request:

1. Determine the protocol (or Internet application) name.
2. Identify the protocol based on the request destination address.
3. Search for related port numbers or IP addresses in custom protocol definitions.
4. Search for related port numbers, IP addresses, or signatures in Websense-defined protocol definitions.

If Websense software is unable to determine any of this information, all content associated with the protocol is permitted.

Editing custom protocols

Related topics:

- *Working with protocols*, page 168
- *Creating a custom protocol*, page 172
- *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 172.

To edit a protocol definition:

1. Select the protocol in the Protocols list. The protocol definition appears to the right of the list.
2. Click **Override Action** to change the filtering action applied to this protocol in all protocol filters (see *Making global protocol filtering changes*, page 172).
3. Click **Add Identifier** to define additional protocol identifiers for this protocol (see *Adding or editing protocol identifiers*, page 171).
4. Select an identifier in the list, and then click **Edit** to make changes to the **Port**, **IP Address Range**, or **Transport Method** defined by that identifier.
5. When you are finished, click **OK** to cache your changes. Changes are not implemented until you click **Save All**.

To delete a protocol definition, select an item in the Protocols list, and then click **Delete**.

Adding or editing protocol identifiers

Use the **Filter Components > Edit Protocols > Add Protocol Identifier** page to define additional protocol identifiers for an existing custom protocol. Use the **Edit Protocol Identifier** page to make changes to a previously-defined identifier.

Before creating or changing an identifier, verify that the correct protocol name appears next to **Selected Protocol**.

When working with protocol identifiers, remember that at least one criterion (port, IP address or transport type) must be unique for each protocol.

1. Specify which **Ports** are included in this identifier.
   - If you select **All Ports**, that criterion overlaps with other ports or IP addresses entered in other protocol definitions.
   - Port ranges are not considered unique if they overlap. For example, the port range 80-6000 overlaps with the range 4000-9000.
   - Use caution when defining a protocol on port 80 or 8080. Network Agent listens for Internet requests over these ports.

     Since custom protocols take precedence over Websense protocols, if you define a custom protocol using port 80, all other protocols that use port 80 are filtered and logged like the custom protocol.
2. Specify which **IP Addresses** are included in this identifier.
   - If you select **All external IP addresses**, that criterion overlaps with any other IP addresses entered in other protocol definitions.
   - IP address ranges are not considered unique if they overlap.
3. Specify which **Protocol Transport** method is included in this identifier.
4. Click **OK** to cache your changes and return to the Edit Protocols page. Changes are not implemented until you click **Save All**.

Renaming a custom protocol

Use the **Filter Components > Edit Protocols > Rename Protocol** page to change the name of a custom protocol, or move it to a different protocol group.

- Use the **Name** field to edit the protocol name. The new name cannot exceed 50 characters.

  The name cannot include any of the following characters:

  ```
  * < > { } ~ ! $ % & @ # . " | \ & + = ? / ; : ,
  ```
- To move the protocol to a different protocol group, select the new group from the **In group** field.

When you are finished making changes, click **OK** to return to the Edit Protocols page. You must also click **OK** on the Edit Protocols to cache the changes.

Making global protocol filtering changes

Use the **Filter Components > Edit Protocols > Override Action** page to change the way a protocol is filtered in all existing protocol filters. This also determines the default action applied to the protocol in new filters.

Although this change overrides the filtering action applied in all existing protocol filters, administrators can later edit those filters to apply a different action.

1. Verify that the correct protocol name appears next to **Selected Protocol**.
2. Select a new **Action** (Permit or Block) to apply to this protocol. By default, **No change** is selected. See *Filtering actions*, page 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 174, 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 170). You must also click **OK** on the Edit Protocols page to cache the changes.

Creating a custom protocol

Related topics:

- *Working with protocols*, page 168
- *Filtering protocols*, page 169
- *Editing custom protocols*, page 170
- *Adding to a Websense-defined protocol*, page 174

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 174, for more information.

2. Expand the **Add protocol to this group** drop-down list, and then select a protocol group. The new protocol appears in this group in all protocol lists and filters.
3. Define a unique **Protocol Identifier** (set of **ports**, **IP addresses**, and **transport methods**) for this group. You can add additional identifiers later, from the Edit Protocols page.

   Follow these guidelines for creating protocol identifiers:

   - At least one criterion (port, IP address or transport type) must be unique for each protocol definition.
   - If you select **All Ports** or **All external IP addresses**, that criterion overlaps with any other ports or IP addresses entered in other protocol definitions.
   - Port ranges or IP address ranges are not considered unique if they overlap. For example, the port range 80-6000 overlaps with the range 4000-9000.

   Note

   Use caution when defining a protocol on port 80 or 8080. Network Agent listens for Internet requests over these ports.

   Since custom protocols take precedence over Websense protocols, if you define a custom protocol using port 80, all other protocols that use port 80 are filtered and logged like the custom protocol.

   The following tables provide examples of valid and invalid protocol definitions:

Port	IP Address	Transport Method	Accepted combination?
70	ANY	TCP	Yes - the port number makes each protocol identifier unique.
90	ANY	TCP	

Port	IP Address	Transport Method	Accepted combination?
70	ANY	TCP	No - the IP addresses are not unique. 10.2.1.201 is included in the “ANY” set.
70	10.2.1.201	TCP	

Port	IP Address	Transport Method	Accepted combination?
70	10.2.3.212	TCP	Yes - the IP addresses are unique.
70	10.2.1.201	TCP	

4. Under Default Filtering Action, specify the default action (**Permit** or **Block**) that should be applied to this protocol in all active protocol filters:
   - Indicate whether traffic using this protocol should be **Logged**. Protocol traffic must be logged to appear in reports and enable protocol usage alerts.
   - Indicate whether access to this protocol should be regulated by **Bandwidth Optimizer** (see *Using Bandwidth Optimizer to manage bandwidth*, page 174).
5. When you are finished, click **OK** to return to the Edit Protocols page. The new protocol definition appears in the Protocols list.
6. Click **OK** again to cache your changes. Changes are not implemented until you click **Save All**.

Adding to a Websense-defined protocol

You cannot add a port number or IP address directly to a Websense-defined protocol. You can, however, create a custom protocol with the same name as the Websense-defined protocol, and then add ports or IP addresses to its definition.

When a custom protocol and a Websense-defined protocol have the same name, Websense software looks for protocol traffic at the ports and IP addresses specified in both definitions.

In reports, custom protocol names have a "C_" prefix. For example, if you created a custom protocol for SQL_NET and specified additional port numbers, reports display C_SQL_NET when the protocol uses the port numbers in the custom protocol.

Using Bandwidth Optimizer to manage bandwidth

Related topics:

- *Working with categories*, page 160
- *Working with protocols*, page 168
- *Configuring the default Bandwidth Optimizer limits*, page 175

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.

For example:

- Block the AOL Instant Messaging protocol if total network bandwidth usage exceeds 50% of available bandwidth, or if current bandwidth usage for AIM exceeds 10% of the total network bandwidth.
- Block the Sports category when total network bandwidth usage reaches 75%, or when bandwidth usage by all HTTP traffic reaches 60% of available network bandwidth.

Protocol bandwidth usage includes traffic over all ports, IP addresses, or signatures defined for the protocol. This means that if a protocol or Internet application uses multiple ports for data transfer, traffic across all of the ports included in the protocol definition are counted toward that protocol's bandwidth usage total. If an Internet application uses a port not included in the protocol definition, however, traffic over that port is not included in bandwidth usage measurements.

Websense software records bandwidth used by filtered TCP- and UDP-based protocols.

Websense, Inc., updates Websense protocol definitions regularly to ensure bandwidth measurement accuracy.

Network Agent sends network bandwidth data to Filtering Service at a predetermined interval. This ensures that Websense software accurately monitors bandwidth usage, and receives measurements that are closest to an average.

When bandwidth-based filtering options are active, Websense software begins bandwidth-based filtering 10 minutes after initial configuration, and 10 minutes after each Websense Policy Server restart. This delay ensures accurate measurement of bandwidth data and use of this data in filtering.

When a request is blocked based on bandwidth limitations, the Websense block page displays this information in the **Reason** field. For more information, see *Block Pages*, page 77.

Configuring the default Bandwidth Optimizer limits

Related topics:

- *Editing a category filter*, page 45
- *Editing a protocol filter*, page 47
- *Using Bandwidth Optimizer to manage bandwidth*, page 174

Before specifying bandwidth settings in policies, verify the default bandwidth thresholds that trigger bandwidth-based filtering settings. The Websense-defined values are:

- Default bandwidth for network: **50%**
- Default bandwidth per protocol: **20%**

Default bandwidth values are stored by Policy Server, and enforced by all associated instances of Network Agent.

To change the default bandwidth values:

1. In Websense Manager, go to **Settings > Filtering**.
2. Enter the bandwidth usage thresholds that will trigger bandwidth-based filtering, when bandwidth filtering is enabled.
   - When a category or protocol is blocked based on traffic for the entire network, **Default bandwidth for network** defines the default filtering threshold.
   - When a category or protocol is blocked based on traffic for the protocol, the **Default bandwidth per protocol** defines the default filtering threshold.

   You can override the default threshold values for each category or protocol in any category or protocol filter.
3. When you are finished, click **OK** to cache your changes. Changes are not implemented until you click **Save All**.

Any changes to the defaults have the potential to affect any category and protocol filters that enforce Bandwidth Optimizer restrictions.

- To manage bandwidth usage associated with a particular protocol, edit the active protocol filter or filters.
- To manage bandwidth usage associated with a particular URL category, edit the appropriate category filter or filters.

  When you filter categories based on HTTP bandwidth usage, Websense software measures total HTTP bandwidth usage over all ports specified as HTTP ports for Websense software.

Managing traffic based on file type

When you create a category filter, you can define filtering based on file extensions, restricting access to particular file types from sites in certain categories. For example, permit the category Sports, but block video files from sites in the Sports category.

Websense software provides several predefined file types, or groupings of file extensions used for similar purposes. These file type definitions are maintained in the Master Database, and may be changed as part of the Master Database update process.

You can implement filtering using predefined file types, modify the existing file type definitions, or create new file types. Note, however, that you cannot delete Websense-defined file types, or delete the file extensions associated with them.

When a user requests a site, Websense software first determines the site category, and then checks for filtered file extensions.

> **Note**
> To implement full filtering for video and audio Internet media, combine protocol-based filtering with file type filtering. In this case, protocol filtering handles streaming media, while file type filtering handles files that can be downloaded and then played.

When a user tries to access a file whose extension is blocked, the **Reason** field on the Websense block page indicates that the file type was blocked. For more information, see *Block Pages*, page 77.

> **Note**
> The standard block page is not displayed if a blocked GIF or JPEG image comprises just a portion of a permitted page. Instead, the image region appears blank. This avoids the possibility of displaying a small portion of a block page in multiple locations on an otherwise permitted page.

File type definitions may contain as many or as few file extensions as are useful for filtering purposes. Websense-defined file types, for example, include the following file extensions:

Audio	Compressed Files		Executables	Video	
.aif	.ace	.mim	.bat	.asf	.mpg
.aifc	.arc	.rar	.exe	.asx	.mpv2
.aiff	.arj	.tar		.avi	.qt
.m3u	.b64	.taz		.ivf	.ra
.mid	.bhx	.tgz		.m1v	.ram
.midi	.cab	.tz		.mov	.wm
.mp3	.gz	.uu		.mp2	.wmp
.ogg	.gzip	.uue		.mp2v	.wmv
.rmi	.hqx	.xxe		.mpa	.wmx
.snd	.iso	.z		.mpe	.wxv
.wav	.jar	.zip			
.wax	.lzh				
.wma					

Any of the file extensions associated with a Websense-defined file type can be added to a custom file type. The file extension is then filtered and logged according to the settings associated with the custom file type.

To view existing file type definitions, edit file types, or create custom file types, go to **Policy Management > Filter Components**, and then click **File Types**. See *Working with file types*, page 178, for more information.

Working with file types

Related topics:

- *Managing traffic based on file type*, page 176
- *Editing a category filter*, page 45
- *Filtering a site*, page 72

Use the **Policy Management > Filter Components > Edit File Types** page to create and manage up to 32 **file types**. File types are groups of file extensions that can be explicitly blocked in category filters (see *Managing traffic based on file type*, page 176).

- 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 179, for further instructions.
- To create a new file type, click **Add File Type**, and then see *Adding custom file types*, page 178, for further instructions.
- To delete a custom file type or extension, select an item, and then click **Delete**.

  You cannot delete Websense-defined file types, or delete the file extensions associated with them.

  You can, however, add file extensions associated with a Websense-defined file type to a custom file type. The file extension is then filtered and logged according to the settings associated with the custom file type. You cannot add the same extension to multiple custom file types.

When you are finished making changes to file type definitions, click **OK**. Changes are not implemented until you click **Save All**.

Adding custom file types

Use the **Filter Components > Edit File Types > Add File Type** page to define custom file types.

1. Enter a unique **File type name**.

   You can create a custom file type with the same name as a Websense-defined file type in order to add additional file extensions to the existing file type.
2. Enter file extensions, one per line, in the **User-defined file extensions** list. You do not need to include the dot (“.”) before each extension.
3. Click **OK** to return to the Edit File Types screen. The new file type appears in the File Types list.

4. When you are finished working with file type definitions, click **OK** on the Edit File Types page. Changes are not implemented until you click **Save All**.

Adding file extensions to a file type

Use the **Filter Components > Edit File Types > Add File Extensions** page to add file extensions to the selected file type.

1. Verify that the expected file type name appears next to **Selected file type**.
2. Enter file extensions, one per line, in the **File extensions** list. You do not need to include the dot (".") before each extension.
3. Click **OK** to return to the Edit File Types screen. The new file extensions appear in the Custom file extensions list.
4. When you are finished working with file type definitions, click **OK** on the Edit File Types page. Changes are not implemented until you click **Save All**.

Using regular expressions

A **regular expression** is a template or pattern used to match multiple strings, or groups of characters. You can use regular expressions in limited access filters, or to define custom URLs or keywords. Websense filtering then tries to match the general pattern, rather than a specific, single URL or keyword.

Consider this simple regular expression:

```
domain.(com|org|net)
```

This expression pattern matches the URLs:

- domain.com
- domain.org
- domain.net

Use regular expressions with caution. They provide a powerful filtering tool, but can easily result in the blocking or permitting of unexpected sites. Also, poorly constructed regular expressions can result in excessive filtering overhead.

Important

Using regular expressions as filtering criteria may increase CPU usage. Tests have shown that with 100 regular expressions, the average CPU usage on the Filtering Service machine increased by 20%.

Websense software supports most Perl regular expression syntax, with a few exceptions. Some of the unsupported syntax is not useful for matching strings that could be found in a URL.

Unsupported regular expression syntax includes:

```
(?<=pattern)string                  (?<!pattern)string
\N{name}                            (?imsx-imsx)
(?(condition)pat1)                  \pP
(?(condition)pat1|pat2)             \PP
(?{code})                           ??{code})
```

For further help with regular expressions, see:

en.wikipedia.org/wiki/Regular_expression

www.regular-expressions.info/

Using the Toolbox to verify filtering behavior

The right shortcut pane in Websense Manager includes a **Toolbox** that allows you to perform quick checks of your filtering setup.

Click a tool name to access the tool. Click the name again to see the list of tools. For more information about using a tool, see:

- *URL Category*, page 180
- *Check Policy*, page 181
- *Test Filtering*, page 181
- *URL Access*, page 181
- *Investigate User*, page 182

You can also click **Support Portal** to access the Websense Technical Support Web site in a new browser tab or window. From the Support Portal, you can use the Knowledge Base to access tutorials, tips, articles, and documentation.

URL Category

To find out how a site is currently categorized:

1. Click **URL Category** in the Toolbox.
2. Enter a URL or IP address.
3. Click **Go**.

The site's current category is displayed in a popup window. If your organization has recategorized the URL, the new category is shown.

The site's categorization may depend on which version of the Master Database (including real-time updates) you are using.

Check Policy

Use this tool to determine which policies apply to a specific client. The results are specific to the current day and time.

1. Click **Check Policy** in the Toolbox.
2. To identify a directory or computer client, enter either:
   - A fully qualified user name

     To browse or search the directory to identify the user, click **Find User** (see *Identifying a user to check policy or test filtering*, page 182).
   - An IP address
3. Click **Go**.

The name of one or more policies is displayed in a popup window. Multiple policies are displayed only when no policy has been assigned to the user, but policies have been assigned to multiple groups, domains, or organizational units to which the user belongs.

Even if multiple policies are shown, only one policy is enforced for a user at any given time (see *Filtering order*, page 71).

Test Filtering

To find out what happens when a specific client requests a particular site:

1. Click **Test Filtering** in the Toolbox.
2. To identify a directory or computer client, enter either:
   - A fully qualified user name

     To browse or search the directory to identify the user, click **Find User** (see *Identifying a user to check policy or test filtering*, page 182).
   - An IP address
3. Enter the URL or IP address of the site you want to check.
4. Click **Go**.

The site category, the action applied to the category, and the reason for the action are displayed in a popup window.

URL Access

To see whether users have attempted to access a site in the past 2 weeks, including today:

1. Click **URL Access** in the Toolbox.
2. Enter all or part of the URL or IP address of the site you want to check.
3. Click **Go**.

An investigative report shows whether the site has been accessed, and if so, when.

You might use this tool after receiving a security alert to find out if your organization has been exposed to phishing or virus-infected sites.

Investigate User

To review a client's Internet usage history for the last 2 weeks, excluding today:

1. Click **Investigate User** in the Toolbox.
2. Enter all or part of a user name or computer IP address.
3. Click **Go**.

An investigative report shows the client's usage history.

Identifying a user to check policy or test filtering

Use the **Find User** page to identify a user (directory) client for the Check Policy or Test Filtering tool.

The page opens with the **User** option selected. Expand the **Directory Entries** folder to browse the directory, or click **Search**. The search feature is available only if you are using an LDAP-based directory service.

To search the directory to find a user:

1. Enter all or part of the user **Name**.
2. Expand the **Directory Entries** tree and browse to identify a search context.

   You must click a folder (DC, OU, or CN) in the tree to specify the context. This populates the field below the tree.
3. Click **Search**. Entries matching your search term are listed under **Search Results**.
4. Click a user name to select a user, or click **Search Again** to enter a new search term or context.

   To return to browsing the directory, click **Cancel Search**.
5. When the correct fully qualified user name appears in the **User** field, click **Go**.

If you are using the Test Filtering tool, make sure that a URL or IP address appears in the **URL** field before you click **Go**.

To identify a computer client instead of a user, click **IP address**.

10 User Identification

To apply policies to users and groups, Websense software must be able to identify the user making a request, given the originating IP address. Various identification methods are available:

- An integration device or application identifies and authenticates users, and then passes user information to Websense software. For more information, see the *Installation Guide*.
- A Websense transparent identification agent works in the background to communicate with a directory service and identify users (see *Transparent identification*).
- Websense software prompts users for their network credentials, requiring them to log on when they open a Web browser (see *Manual authentication*, page 185).

Transparent identification

Related topics:

- *Manual authentication*, page 185
- *Configuring user identification methods*, page 185

In general, **transparent identification** describes any method that Websense software uses to identify users in your directory service without prompting them for logon information. This includes integrating Websense software with a device or application that provides user information for use in filtering, or using optional Websense transparent identification agents.

- Websense *DC Agent*, page 193, is used with a Windows-based directory service. The agent periodically queries domain controllers for user logon sessions and polls client machines to verify logon status. It runs on a Windows server and can be installed in any domain in the network.
- Websense *Logon Agent*, page 196, identifies users as they log on to Windows domains. The agent runs on a Linux or Windows server, but its associated Logon Application runs only on Windows machines.

- Websense *RADIUS Agent*, page 199, can be used in conjunction with either Windows- or LDAP-based directory services. The agent works with a RADIUS server and client to identify users logging on from remote locations.
- Websense *eDirectory Agent*, page 203, is used with Novell eDirectory. The agent uses Novell eDirectory authentication to map users to IP addresses.

For instructions on installing each agent, see the *Installation Guide*. Agent can be used alone, or in certain combinations (see *Configuring multiple agents*, page 209).

Notes

If you are using an integrated NetCache appliance, NetCache must send user names to Websense software in WinNT, LDAP, or RADIUS format for transparent identification to work.

If you have a proxy server and are using a transparent identification agent, it is best to use anonymous authentication in your proxy server.

Both general user identification settings and specific transparent identification agents are configured in Websense Manager. Click the **Settings** tab in the left navigation pane, and then click **User Identification**.

See *Configuring user identification methods*, page 185, for detailed configuration instructions.

In some instances, Websense software may not be able to obtain user information from a transparent identification agent. This can occur if more than one user is assigned to the same machine, or if a user is an anonymous user or guest, or for other reasons. In these cases, you can prompt the user to log on via the browser (see *Manual authentication*, page 185).

Transparent identification of remote users

In certain configurations, Websense software can transparently identify users logging on to your network from remote locations:

- If you have deployed the Websense Remote Filtering Server and Remote Filtering Client, Websense software can identify any remote user logging on to a cached domain using a domain account. For more information, see *Filter Remote Clients*, page 143.
- 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 193).
- If you are using a RADIUS server to authenticate users logging on from remote locations, RADIUS Agent can transparently identify these users so you can apply filtering policies based on users or groups (see *RADIUS Agent*, page 199).

Manual authentication

Related topics:

- *Transparent identification*, page 183
- *Setting authentication rules for specific machines*, page 187
- *Secure manual authentication*, page 190
- *Configuring user identification methods*, page 185

Transparent identification is not always available or desirable in all environments. For organizations that do not use transparent identification, or in situations when transparent identification is not available, you can still filter based on user and group-based policies using **manual authentication**.

Manual authentication prompts users for a user name and password the first time they access the Internet through a browser. Websense software confirms the password with a supported directory service, and then retrieves policy information for that user.

You can configure Websense software to enable manual authentication any time transparent identification is not available (see *Configuring user identification methods*, page 185), or 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 187).

When manual authentication is enabled, users may receive HTTP errors and be unable to access the Internet if:

- They make 3 failed attempts to enter a password. This occurs when the user name or password is invalid.
- They click **Cancel** to bypass the authentication prompt.

When manual authentication is enabled, users who cannot be identified are prevented from browsing the Internet.

Configuring user identification methods

Related topics:

- *Transparent identification*, page 183
- *Manual authentication*, page 185
- *Working with users and groups*, page 56

Use the **Settings > User Identification** page to manage when and how Websense software attempts to identify users in the network in order to apply user- and group-based policies.

- Configure Policy Server to communicate with transparent identification agents.
- Review and update transparent identification agent settings.
- Set a global rule to determine how Websense software responds when users cannot be identified by a transparent identification agent or integration device.
- Identify machines in your network to which global user identification rules do not apply, and specify whether and how users of those machines should be authenticated.

If you are using Websense transparent identification agents, the agents are listed under **Transparent Identification Agents**:

- **Server** shows the IP address or name of the machine hosting the transparent identification agent.
- **Port** lists the port that Websense software uses to communicate with the agent.
- **Type** indicates whether the specified instance is a DC Agent, Logon Agent, RADIUS Agent, or eDirectory Agent. (See *Transparent identification*, page 183, 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 194
- *Configuring Logon Agent*, page 197
- *Configuring RADIUS Agent*, page 201
- *Configuring eDirectory Agent*, page 205

To remove an agent instance from the list, mark the checkbox next to the agent information in the list, and then click **Delete**.

Under **Additional Authentication Options**, specify the default response of Websense software when users are not identified transparently (by an agent or integration):

- Click **Apply computer or network policy** to ignore user and group-based policies in favor of computer and network-based policies, or the Default policy.
- Click **Prompt user for logon information** to require users to provide logon credentials when they open a browser. User and group-based policies can then be applied (see *Manual authentication*, page 185).
- Specify the default domain **Context** that Websense software should use any time a user is prompted for log on credentials. This is the domain in which users' credentials are valid.

  If you use the Exceptions list to specify any machines on which users are prompted for logon information, you must provide a default domain context, even if the global rule is to apply a computer or network-based policy.

After establishing the general rule that determines when and how users are identified by Websense software, you can create exceptions to the rule.

For example, if you use a transparent identification agent or integration product to identify users, and have enabled manual authentication to prompt users for their credentials when they cannot be identified transparently, you can identify specific machines on which:

- Users who cannot be identified are never be prompted for their credentials. In other words, when transparent identification fails, manual authentication is not attempted, and the computer or network policy, or the Default policy, is applied.
- User information is always ignored, even when it is available, and users are always prompted for their credentials.
- User information is always ignored, even when it is available, and users are never prompted for their credentials (the computer or network policy, or the Default policy, is always applied).

To create an exception, click **Exceptions**, and then see *Setting authentication rules for specific machines*, page 187.

When you are finished making changes on this page, click **OK** to save your changes. To avoid saving changes, click **Cancel**.

Setting authentication rules for specific machines

Related topics:

- *Configuring user identification methods*, page 185
- *Manual authentication*, page 185
- *Secure manual authentication*, page 190

Selective authentication lets you determine whether users requesting Internet access from a specific client machine (identified by IP address) are prompted to provide their logon credentials via the browser. This can be used to:

- Establish different authentication rules for a machine in a public kiosk than for employees of the organization supplying the kiosk.
- Ensure that users of an exam-room computer in a medical office are always identified before getting Internet access.

Machines with special user identification settings applied are listed on the **Settings > User Identification** page. Click **Exceptions** to establish specific user identification settings for some machines in your network, or see if special settings have been defined for a specific machine.

To add a machine to the list, click **Add**, and then see *Defining exceptions to user identification settings*, page 188, for further instructions.

When you are finished adding machines or network ranges to the list, click **OK**. Changes are not implemented until you click **Save All**.

Defining exceptions to user identification settings

Related topics:

- *Transparent identification*, page 183
- *Manual authentication*, page 185
- *Configuring user identification methods*, page 185

Use the **Settings > User Identification > Add IP Addresses** page to identify machines to which specific user identification rules should be applied.

1. Enter an **IP address** or **IP address range** to identify machines to which to apply a specific authentication method, and then click the right-arrow button to add them to the **Selected** list.

   If the same rules should be applied to multiple machines, add them all to the list.
2. Select an entry in the **User identification** drop-down list to indicate whether Websense software should attempt to identify users of these machines transparently.
   - Select **Try to identify user transparently** to request user information from a transparent identification agent or integration device.
   - Select **Ignore user information** to avoid using any transparent method to identify users.
3. Indicate whether users should be prompted to provide logon credentials via the browser. This setting applies when user information is not available, either because other identification failed, or because user information was ignored.
   - Select **Prompt user for logon information** to require users to provide logon credentials.

     If **Try to identify user transparently** is also selected, users receive a browser prompt only if they are not identified transparently.
   - Select **Apply computer or network policy** to ensure that users are never required to provide logon credentials.

     If **Try to identify user transparently** is also selected, users whose credentials can be verified transparently are filtered by the appropriate user-based policy.
4. Click **OK** to return to the User Identification page.
5. When you are finished updating the Exceptions list, click **OK** to cache your changes. Changes are not implemented until you click **Save All**.

Revising exceptions to user identification settings

Related topics:

- *Transparent identification*, page 183
- *Manual authentication*, page 185
- *Configuring user identification methods*, page 185

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** to request user information from a transparent identification agent or integration device.
   - Select **Ignore user information** to avoid using any transparent method to identify users.
2. Indicate whether users should be prompted to provide logon credentials via the browser. This setting applies when user information is not available, either because transparent identification failed, or because transparent identification was ignored.
   - Select **Prompt user for logon information** to require users to provide logon credentials.

     If **Try to identify user** is also selected, users receive a browser prompt only if they are not identified transparently.
   - Select **Apply computer or network policy** to ensure that users are never prompted to provide logon credentials.

     If **Try to identify user** is also selected, users whose credentials can be verified transparently are filtered by the appropriate user-based policy.
3. Click **OK** to return to the User Identification page.
4. When you are finished updating the Exceptions list, click **OK** to cache your changes. Changes are not implemented until you click **Save All**.

Secure manual authentication

Related topics:

- *Configuring user identification methods*, page 185
- *Manual authentication*, page 185
- *Setting authentication rules for specific machines*, page 187
- *Activating secure manual authentication*, page 191

Websense secure manual authentication uses Secure Sockets Layer (SSL) encryption to protect authentication data being transmitted between client machines and Websense software. An SSL server built into Filtering Service provides encryption of user names and passwords transmitted between client machines and Filtering Service. By default, secure manual authentication is disabled.

Note

Secure manual authentication cannot be used with Remote Filtering. 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 190).
2. Enable secure manual authentication (see *Activating secure manual authentication*, page 191) and secure communication with the directory service.
3. Import certificates into the browser (see *Accepting the certificate within the client browser*, page 192).

Generating keys and certificates

Related topics:

- *Manual authentication*, page 185
- *Setting authentication rules for specific machines*, page 187
- *Secure manual authentication*, page 190
- *Activating secure manual authentication*, page 191
- *Accepting the certificate within the client browser*, page 192

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 httpd.apache.org/docs/2.2/ssl/ssl_faq.html#aboutcerts.
- To learn more about generating your own private key, CSR, and certificate, see www.akadia.com/services/ssh_test_certificate.html.

There are many tools that you can use to generate a self-signed certificate, including the OpenSSL toolkit (available from www.openssl.org).

Regardless of the method you choose for generating the certificate, use the following general steps.

1. Generate a private key (**server.key**).
2. Generate a Certificate Signing Request (CSR) with the private key.

   

   Important

   When prompted for the CommonName, enter the IP address of the Filtering Server machine. If you skip this step, client browsers will display a security certificate error.

3. Use the CSR to create a self-signed certificate (**server.crt**).
4. Save the **server.crt** and **server.key** files in a location that Websense software can access, and where they can be read by Filtering Service.

Activating secure manual authentication

Related topics:

- *Manual authentication*, page 185
- *Setting authentication rules for specific machines*, page 187
- *Secure manual authentication*, page 190
- *Generating keys and certificates*, page 190
- *Accepting the certificate within the client browser*, page 192

1. Stop Websense Filtering Service (see *Stopping and starting Websense services*, page 258).
2. Navigate to the Websense installation directory on the Filtering Service machine (by default, **C:\Program Files\Websense\bin** or **/opt/Websense/bin/**).
3. Locate **eimserver.ini** and make a backup copy of the file in another directory.
4. Open the original INI file in a text editor.

5. Find the **[WebsenseServer]** section, and then add the line:

   ```
   SSLManualAuth=on
   ```

6. Below the previous line, add the following:

   ```
   SSLCertFileLoc=[path]
   ```

   Replace **[path]** with the full path to the SSL certificate, including the certificate file name (for example, C:\secmanauth\server.crt).

7. Also add:

   ```
   SSLKeyFileLoc=[path]
   ```

   Replace **[path]** with the full path to the SSL key, including the key file name (for example, C:\secmanauth\server.key).

8. Save and close **eimserver.ini**.
9. Start Websense Filtering Service.

After starting, Filtering Service listens for requests on the default secure HTTP port (**15872**).

The preceding steps ensure secure communication between the client machine and Websense software. To also secure communication between Websense software and the directory service, make sure that **Use SSL** is selected on the **Settings > Directory Services** page. See *Advanced directory settings*, page 59, for details.

Accepting the certificate within the client browser

Related topics:

- *Manual authentication*, page 185
- *Setting authentication rules for specific machines*, page 187
- *Secure manual authentication*, page 190
- *Generating keys and certificates*, page 190
- *Activating secure manual authentication*, page 191

The first time you try to browse to a Web site, the browser will display a warning about the security certificate. To avoid seeing this message in the future, install the certificate in the certificate store.

Microsoft Internet Explorer (Version 7)

1. Open the browser and go to a Web site.

   A warning appears, stating that there is a problem with the site's security certificate.

2. Click **Continue to this website (not recommended)**.

   If you receive an authentication prompt, click **Cancel**.

3. Click the **Certificate Error** box to the right of the address bar (at the top of the browser window), and then click **View certificates**.
4. On the General tab of the Certificate dialog box, click **Install Certificate**.
5. Select **Automatically select the certificate store based on the type of certificate**, and then click **Next**.
6. Click **Finish**.
7. When asked whether to install the certificate, click **Yes**.

Users will no longer receive certificate security warnings related to Filtering Service on this machine.

Mozilla Firefox (Version 2.x)

1. Open the browser and go to a Web site.

   A warning message appears.
2. Click **Accept the certificate permanently**.
3. Enter your credentials, if prompted.
4. Go to **Tools > Options**, and then click **Advanced**.
5. Select the **Encryption** tab, and then click **View Certificates**.
6. Select the **Websites** tab and verify that the certificate is listed.

Users will no longer receive certificate security warnings related to Filtering Service on this machine.

Mozilla Firefox (Version 3.x)

1. Open the browser and go to a Web site.

   A warning message appears.
2. Click **Or you can add an exception**.
3. Click **Add Exception**.
4. Make sure that **Permanently store this exception is selected**, and then click **Confirm Security Exception**.

Users will no longer receive certificate security warnings related to Filtering Service on this machine.

DC Agent

Related topics:

- *Transparent identification*, page 183
- *Configuring DC Agent*, page 194
- *Configuring different settings for an agent instance*, page 210

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 *Installing Websense Components Separately* in the *Installation Guide*.

   Note

   Run DC Agent using domain administrator privileges. The domain administrator account must also be a member of the Administrators group on the DC Agent machine.

   This is required for DC Agent to retrieve user logon information from the domain controller. If you cannot install DC Agent with such privileges, configure administrator privileges for these services after installation. For more information, see *Websense software is not applying user or group policies*, page 332.

2. Configure DC Agent to communicate with other Websense components and with domain controllers in your network (see *Configuring DC Agent*).
3. Use Websense Manager to add users and groups to filter (see *Adding a client*, page 61).

Websense software can prompt users for identification if DC Agent is unable to identify users transparently. For more information, see *Manual authentication*, page 185.

Configuring DC Agent

Related topics:

- *Transparent identification*
- *Manual authentication*
- *Configuring user identification methods*
- *DC Agent*
- *Configuring multiple agents*

Use the **Settings > User Identification > DC Agent** page to configure a new instance of DC Agent, as well as to configure the global settings that apply to all instances of DC Agent.

To add a new instance of DC Agent, first provide basic information about where the agent is installed, and how Filtering Service should communicate with it. These settings may be unique to each agent instance.

1. Under Basic Agent Configuration, enter the IP address or name of the **Server** on which the agent is installed.

   Note

   Machine names must start with an alphabetical character (a-z), not a numeric or special character.

   Machine names containing certain extended ASCII characters may not resolve properly. If you are using a non-English version of Websense software, enter an IP address instead of a machine name.

2. Enter the **Port** that DC Agent should use to communicate with other Websense components. The default is 30600.
3. To establish an authenticated connection between Filtering Service and DC Agent, check **Enable Authentication**, and then enter a **Password** for the connection.

Next, customize global DC Agent communication and troubleshooting, domain controller polling, and computer polling settings. By default, changes that you make here affect all DC Agent instances. Settings marked with an asterisk (*), however, can be overriden in an agent's configuration file to customize the behavior of that agent instance (see *Configuring different settings for an agent instance*, page 210).

1. Under DC Agent Communication, enter the **Communications port** to be used for communication between DC Agent and other Websense components. The default is 30600.

   Unless instructed to do so by Websense Technical Support, do not make changes to the **Diagnostic port** setting. The default is 30601.
2. Under Domain Controller Polling, mark **Enable domain controller polling** to enable DC Agent to query domain controllers for user logon sessions.

   You can specify which domain controllers each instance of DC Agent polls in the agent's configuration file. See *Configuring multiple agents*, page 209, for details.
3. Use the **Query interval** field to specify how often (in seconds) DC Agent queries domain controllers.

   Decreasing the query interval may provide greater accuracy in capturing logon sessions, but also increases overall network traffic. Increasing the query interval decreases network traffic, but may also delay or prevent the capture of some logon sessions. The default is 10 seconds.
4. Use the **User entry timeout** field to specify how frequently (in hours) DC Agent refreshes the user entries in its map. The default is 24 hours.
5. Under Computer Polling, check **Enable computer polling** to enable DC Agent to query computers for user logon sessions. This may include computers that are outside the domains that the agent already queries.

DC Agent uses WMI (Windows Management Instruction) for computer polling. If you enable computer polling, configure the Windows Firewall on client machines to allow communication on port **135**.

6. Enter a **User map verification interval** to specify how often DC Agent contacts client machines to verify which users are logged on. The default is 15 minutes.

   DC Agent compares the query results with the user name/IP address pairs in the user map it sends to Filtering Service. Decreasing this interval may provide greater user map accuracy, but increases network traffic. Increasing the interval decreases network traffic, but also may decrease accuracy.

7. Enter a **User entry timeout** period to specify how often DC Agent refreshes entries obtained through computer polling in its user map. The default is 1 hour.

   DC Agent removes any user name/IP address entries that are older than this timeout period, and that DC Agent cannot verify as currently logged on. Increasing this interval may lessen user map accuracy, because the map potentially retains old user names for a longer time.

   Note

   Do not make the user entry timeout interval shorter than the user map verification interval. This could cause user names to be removed from the user map before they can be verified.

8. Click **OK** to immediately save and implement your changes.

Logon Agent

Related topics:

- *Transparent identification*, page 183
- *Configuring Logon Agent*, page 197
- *Configuring different settings for an agent instance*, page 210

Websense Logon Agent identifies users in real time, as they log on to domains. This eliminates the possibility of missing a user logon due to a query timing issue.

Logon Agent (also called Authentication Server) can reside on a Windows or Linux machine. The agent works with the Websense Logon Application (LogonApp.exe) on Windows client machines to identify users as they log on to Windows domains.

In most cases, using either DC Agent or Logon Agent is sufficient, but you can use both agents together. In this case, Logon Agent takes precedence over DC Agent. DC Agent only communicates a logon session to Filtering Service in the unlikely event that Logon Agent has missed one.

Install Logon Agent, and then deploy the Logon Application to client machines from a central location. For more information, see the *Installation Guide*.

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 running on Linux*, page 338, for additional configuration steps.

Configuring Logon Agent

Related topics:

- *Transparent identification*, page 183
- *Manual authentication*, page 185
- *Configuring user identification methods*, page 185
- *Logon Agent*, page 196
- *Configuring multiple agents*, page 209

Use the **Settings > User Identification > Logon Agent** page to configure a new instance of Logon Agent, as well as to configure the global settings that apply to all instances of Logon Agent.

To add a new instance of Logon Agent:

1. Under Basic Agent Configuration, enter the IP address or name of the **Server** on which the agent is installed.

   > **Note**
   > Machine names must start with an alphabetical character (a-z), not a numeric or special character.
   >
   > Machine names containing certain extended ASCII characters may not resolve properly. If you are using a non-English version of Websense software, enter an IP address instead of a machine name.

2. Enter the **Port** that Logon Agent should use to communicate with other Websense components. The default is 30602.
3. To establish an authenticated connection between Filtering Service and Logon Agent, check **Enable Authentication**, and then enter a **Password** for the connection.

4. Either click **OK** to save your changes, or continue to the next section of the screen to enter additional configuration information.

Next, customize global Logon Agent communications settings. By default, changes that you make here affect all Logon Agent instances.

1. Under Logon Agent Communication, enter the **Communications port** that should be used for communication between Logon Agent and other Websense components. The default is 30602.
2. Unless instructed to do so by Websense Technical Support, do not make changes to the **Diagnostic port** setting. The default is 30603.
3. Under Logon Application Communication, specify the **Connection port** that the logon application uses to communicate with Logon Agent. The default is 15880.
4. Enter the **Maximum number of connections** that each Logon Agent instance allows. The default is 200.

   If your network is large, you may need to increase this number. Increasing the number does increase network traffic.
5. Either click **OK** to save your changes, or continue to the next section of the screen to enter additional configuration information.

To configure the default settings that determine how user entry validity is determined, you must first determine whether Logon Agent and the client logon application will operate in **persistent mode** or **nonpersistent mode** (default).

Nonpersistent mode is activated by including the /NOPERSIST parameter when launching **LogonApp.exe**. (More information is available in the **LogonApp_ReadMe.txt** file, which is included with your Logon Agent installation.)

- In persistent mode, the logon application contacts Logon Agent periodically to communicate user logon information.

  If you are using persistent mode, specify a **Query interval** to determine how frequently the logon application communicates logon information.

  > **Note**
  > If you change this value, the change does not take effect until the previous interval period has elapsed. For example, if you change the interval from 15 minutes to 5 minutes, the current 15-minute interval must end before the query starts occurring every 5 minutes.

- In nonpersistent mode, the logon application sends user logon information to Logon Agent only once for each logon.

  If you are using nonpersistent mode, specify a **User entry expiration** time period. When this timeout period is reached, the user entry is removed from the user map.

When you are finished making configuration changes, click **OK** to save your settings.

RADIUS Agent

Related topics:

- *Transparent identification*, page 183
- *Processing RADIUS traffic*, page 200
- *Configuring the RADIUS environment*, page 200
- *Configuring RADIUS Agent*, page 201
- *Configuring the RADIUS client*, page 202
- *Configuring the RADIUS server*, page 203
- *Configuring different settings for an agent instance*, page 210

Websense RADIUS Agent lets you apply user and group-based policies using authentication provided by a RADIUS server. RADIUS Agent enables transparent identification of users who access your network using a dial-up, Virtual Private Network (VPN), Digital Subscriber Line (DSL), or other remote connection (depending on your configuration).

RADIUS Agent works together with the RADIUS server and RADIUS client in your network to process and track Remote Access Dial-In User Service (RADIUS) protocol traffic. This enables you to assign particular filtering policies to users or groups that access your network remotely, as well as to local users.

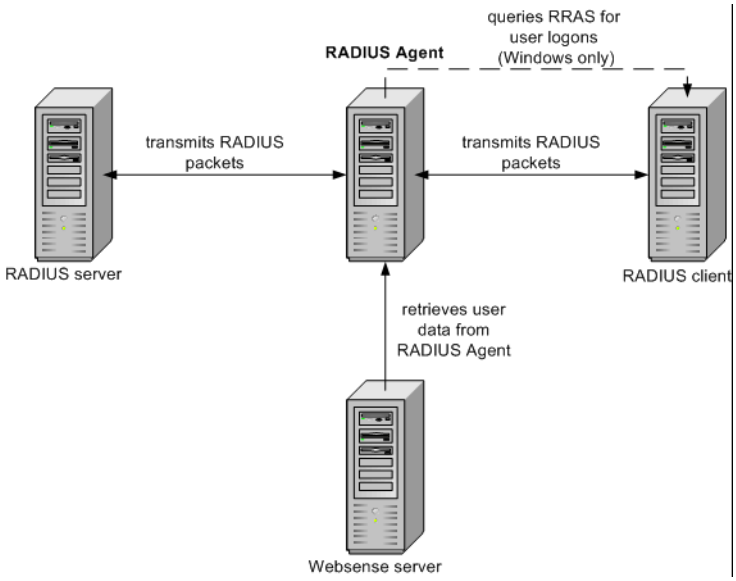

When you install RADIUS Agent, the Agent integrates with existing Websense components. However, RADIUS Agent, your RADIUS server, and your RADIUS client must be configured appropriately (see *Configuring RADIUS Agent*, page 201).

Processing RADIUS traffic

The Websense RADIUS Agent acts as a proxy that forwards RADIUS messages between a RADIUS client and a RADIUS server (or multiple clients and servers).

RADIUS Agent does not authenticate users directly. Instead, the agent identifies remote users and associates them with IP addresses so a RADIUS server can authenticate those users. Ideally, the RADIUS server passes authentication requests to an LDAP-based directory service.

RADIUS Agent stores user name-to-IP-address pairings in a user map. If your RADIUS client supports accounting (or user logon tracking), and accounting is enabled, RADIUS Agent gleans more detail about user logon sessions from the RADIUS messages it receives.

When properly configured, Websense RADIUS Agent captures and processes all RADIUS protocol packets of these types:

- **Access-Request**: Sent by a RADIUS client to request authorization for a network access connection attempt.
- **Access-Accept**: Sent by a RADIUS server in response to an Access-Request message; tells the RADIUS client that the attempted connection is authorized and authenticated.
- **Access-Reject**: Sent by a RADIUS server in response to an Access-Request message; tells the RADIUS client that the attempted connection is rejected.
- **Accounting-Stop-Request**: Sent by a RADIUS client to tell the RADIUS server to stop tracking user activity.

Configuring the RADIUS environment

Websense RADIUS Agent serves as a proxy between a RADIUS client and a RADIUS server. This diagram shows a simplified view of how using RADIUS Agent differs from a standard RADIUS setup.

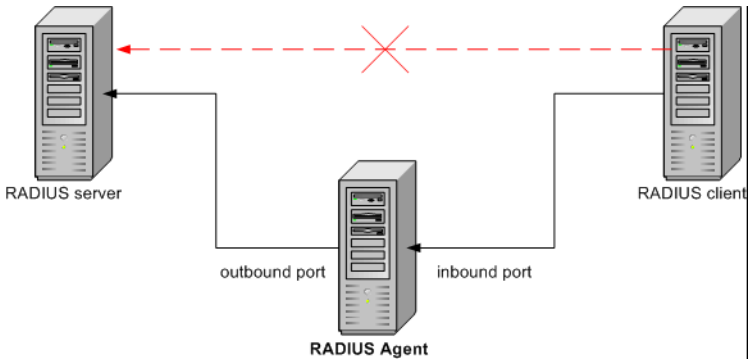

RADIUS Agent and the RADIUS server should be installed on separate machines. The agent and server cannot have the same IP address, and must use different ports.

After installing RADIUS Agent, configure the RADIUS Agent in Websense Manager (see *Configuring RADIUS Agent*, page 201). You must also:

- Configure the RADIUS client (typically a Network Access Server [NAS]) to communicate with RADIUS Agent instead of directly with your RADIUS server.
- Configure the RADIUS server to use RADIUS Agent as a proxy (see the RADIUS server documentation). If you have multiple RADIUS servers, configure each one separately.

> **Note**
> If you use Lucent RADIUS Server and RRAS, you must configure the RADIUS server to use Password Authentication Protocol (PAP), and the RRAS server to accept only PAP requests. For more information, see the related product documentation.

Configuring RADIUS Agent

Related topics:

- *Transparent identification*, page 183
- *Manual authentication*, page 185
- *Configuring user identification methods*, page 185
- *RADIUS Agent*, page 199
- *Configuring multiple agents*, page 209

Use the **Settings > User Identification > RADIUS Agent** page to configure a new instance of RADIUS Agent, as well as to configure the global settings that apply to all instances of RADIUS Agent.

To add a new instance of RADIUS Agent:

1. Under Basic Agent Configuration, enter the IP address or name of the **Server** on which the agent is installed.

   > **Note**
   > Machine names must start with an alphabetical character (a-z), not a numeric or special character.
   >
   > Machine names containing certain extended ASCII characters may not resolve properly. 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 RADIUS Agent should use to communicate with other Websense components. The default is 30800.
3. To establish an authenticated connection between Filtering Service and RADIUS Agent, check **Enable Authentication**, and then enter a **Password** for the connection.

4. Either click **OK** to save your changes, or continue to the next section of the screen to enter additional configuration information.

Next, customize global RADIUS Agent settings. By default, changes that you make here affect all RADIUS Agent instances. Settings marked with an asterisk (*), however, can be overriden in an agent's configuration file to customize the behavior of that agent instance (see *Configuring different settings for an agent instance*, page 210).

1. Enter the **Communications port** used for communication between RADIUS Agent and other Websense components. The default is 30800.
2. Unless instructed to do so by Websense Technical Support, do not make changes to the **Diagnostic port** setting. The default is 30801.
3. Under RADIUS Server, enter the **RADIUS server IP or name**. RADIUS Agent forwards authentication requests to the RADIUS server, and must know the identity of this machine.
4. If Microsoft RRAS is in use, enter the IP address of the **RRAS machine**. Websense software queries this machine for user logon sessions.
5. Enter the **User entry timeout** interval, used to determine how often RADIUS Agent refreshes its user map. Typically, the default query value (24 hours) is best.
6. Use the **Authentication Ports** and **Accounting Ports** settings to specify which ports RADIUS Agent uses to send and receive authentication and accounting requests. For each type of communication, you can specify which port is used for communication between:
   - RADIUS Agent and the RADIUS server
   - RADIUS Agent and the RADIUS client
7. When you are finished, click **OK** to immediately save your settings.

Configuring the RADIUS client

Your RADIUS client must be configured to transmit authentication and accounting requests to the RADIUS server via RADIUS Agent.

Modify your RADIUS client configuration so that:

- The RADIUS client sends authentication requests to machine and port on which RADIUS Agent listens for authentication requests. This is the **Authentication Port** specified during RADIUS Agent configuration.
- The RADIUS client sends accounting requests to the machine and port on which RADIUS Agent listens for accounting requests. This is the **Accounting Port** specified during RADIUS Agent configuration.

The exact procedure for configuring a RADIUS client differs by client type. For details, see your RADIUS client documentation.

> **Note**
> The RADIUS client should include the attributes **User-Name** and **Framed-IP-Address** in authentication and accounting messages it sends. RADIUS Agent uses the values of these attributes to interpret and store user name/IP address pairs. If your RADIUS client does not generate this information by default, configure it to do so (see the RADIUS client documentation).

Configuring the RADIUS server

To enable proper communication between Websense RADIUS Agent and your RADIUS server:

- Add the IP address of the RADIUS Agent machine to your RADIUS server's client list. For instructions, see your RADIUS server documentation.
- Define shared secrets between the RADIUS server and all RADIUS clients that use the agent to communicate with the RADIUS server. Shared secrets are usually specified as authentication security options.

  Configuring a shared secret for RADIUS clients and the RADIUS server provides secure transmission of RADIUS messages. Typically, the shared secret is a common text string. For instructions, see your RADIUS server documentation.

> **Note**
> The RADIUS server should include the attributes **User-Name** and **Framed-IP-Address** in authentication and accounting messages. RADIUS Agent uses the values of these attributes to interpret and store user name/IP address pairs. If your RADIUS server does not generate this information by default, configure it to do so (see the RADIUS server documentation).

eDirectory Agent

Related topics:

- *Transparent identification*, page 183
- *Configuring eDirectory Agent*, page 205
- *Configuring different settings for an agent instance*, page 210

Websense eDirectory Agent works together with Novell eDirectory to transparently identify users so Websense software can filter them according to policies assigned to users, groups, domains, or organizational units.

eDirectory Agent gathers user logon session information from Novell eDirectory, which authenticates users logging on to the network. The agent then associates each authenticated user with an IP address, and records user name-to-IP-address pairings to a local user map. eDirectory Agent then communicates this information to Filtering Service.

Note

From a Novell client running Windows, multiple users can log on to a single Novell eDirectory server. This associates one IP address with multiple users. In this scenario, eDirectory Agent's user map only retains the user name/IP address pairing for the last user logged on from a given IP address.

One instance of Websense eDirectory Agent can support one Novell eDirectory master, plus any number of Novell eDirectory replicas.

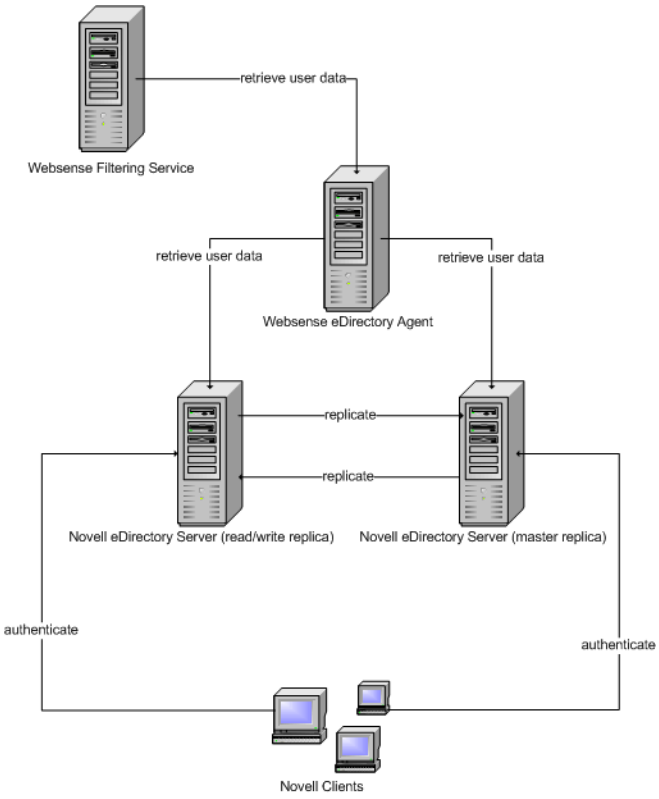

Special configuration considerations

- If you have integrated Cisco Content Engine v5.3.1.5 or higher with Websense software:

- Run the following Websense services on the same machine as Cisco Content Engine:

  Websense eDirectory Agent
  Websense User Service
  Websense Filtering Service
  Websense Policy Server

- Ensure that all Novell eDirectory replicas are added to the **wsedir.ini** file on the same machine.
- Delete the **eDirAgent.bak** file.

Run Websense Reporting Tools services on a machine **separate** from Cisco Content Engine and Websense software.

- Websense software supports using NMAS with eDirectory Agent. To use eDirectory Agent with NMAS enabled, eDirectory Agent must be installed on a machine that is also running the Novell Client.

Configuring eDirectory Agent

Related topics:

- *Transparent identification*, page 183
- *Manual authentication*, page 185
- *Configuring user identification methods*, page 185
- *eDirectory Agent*, page 203
- *Configuring eDirectory Agent to use LDAP*, page 207
- *Configuring multiple agents*, page 209

Use the **Settings > User Identification > eDirectory Agent** page to configure a new instance of eDirectory Agent, as well as to configure the global settings that apply to all instances of eDirectory Agent.

To add a new instance of eDirectory Agent:

1. Under Basic Agent Configuration, enter the IP address or name of the **Server** on which the agent is installed.

   Note

   Machine names must start with an alphabetical character (a-z), not a numeric or special character.

   Machine names containing certain extended ASCII characters may not resolve properly. 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 eDirectory Agent should use to communicate with other Websense components. The default is 30700.
3. To establish an authenticated connection between Filtering Service and eDirectory Agent, check **Enable Authentication**, and then enter a **Password** for the connection.
4. Either click **OK** to save your changes, or continue to the next section of the screen to enter additional configuration information.

Next, customize global eDirectory Agent communication settings. By default, changes that you make here affect all eDirectory Agent instances. Settings marked with an asterisk (*), however, can be overriden in an agent's configuration file to customize the behavior of that agent instance (see *Configuring different settings for an agent instance*, page 210).

1. Enter the default **Communications port** used for communication between eDirectory Agent and other Websense components. The default is 30700.
2. Unless instructed to do so by Websense Technical Support, do not make changes to the **Diagnostic port** setting. The default is 30701.
3. Under eDirectory Server, specify a **Search base** (root context) for eDirectory Agent to use as a starting point when searching for user information in the directory.
4. Provide the administrative user account information that eDirectory Agent should use to communicate with the directory:
   a. Enter the **Administrator distinguished name** for a Novell eDirectory administrative user account.
   b. Enter the **Password** used by that account.
   c. Specify a **User entry timeout** interval to indicate how long entries remain in the agent's user map.

      This interval should be approximately 30% longer than a typical user logon session. This helps prevent user entries from being removed from the map before the users are done browsing.

      Typically, the default value (24 hours) is recommended.

      

      Note

      In some environments, instead of using the User entry timeout interval to determine how frequently eDirectory Agent updates its user map, it may be appropriate to query the eDirectory Server at regular intervals for user logon updates. See *Enabling full eDirectory Server queries*, page 208.

5. Add the eDirectory Server master, as well as any replicas, to the **eDirectory Replicas** list. To add an eDirectory Server master or replica to the list, click **Add**, and the follow the instructions in *Adding an eDirectory server replica*, page 207.

When you are finished making configuration changes, click **OK** to save your settings.

Adding an eDirectory server replica

One instance of the Websense eDirectory Agent can support one Novell eDirectory master, plus any number of Novell eDirectory replicas running on separate machines.

eDirectory Agent must be able to communicate with each machine running a replica of the directory service. This ensures that the agent gets the latest logon information as quickly as possible, and does not wait for eDirectory replication to occur.

Novell eDirectory replicates the attribute that uniquely identifies logged-on users only every 5 minutes. Despite this replication time lag, eDirectory Agent picks up new logon sessions as soon as a user logs on to any eDirectory replica.

To configure eDirectory Agent installation to communicate with eDirectory:

1. In the Add eDirectory replica screen, enter the IP address or name for eDirectory **Server** (master or replica).
2. Enter the **Port** that eDirectory Agent uses to communicate with the eDirectory machine.
3. Click **OK** to return to the eDirectory page. The new entry appears in the eDirectory Replicas list.
4. Repeat the process for any additional eDirectory server machines.
5. Click **OK** to cache changes, and then click **Save All**.
6. Stop and start eDirectory Agent so that the agent can begin communicating with the new replica. See *Stopping and starting Websense services*, page 258, for instructions.

Configuring eDirectory Agent to use LDAP

Websense eDirectory Agent can use Netware Core Protocol (NCP) or Lightweight Directory Access Protocol (LDAP) to get user logon information from Novell eDirectory. By default, eDirectory Agent on Windows uses NCP. On Linux, eDirectory Agent must use LDAP.

If you are running eDirectory Agent on Windows, but want the agent to use LDAP to query Novell eDirectory, set the agent to use LDAP instead of NCP. Generally, NCP provides a more efficient query mechanism.

To set eDirectory Agent on Windows to use LDAP:

1. Ensure that you have at least one Novell eDirectory replica containing all directory objects to monitor and filter in your network.
2. Stop the Websense eDirectory Agent service (see *Stopping and starting Websense services*, page 258).
3. Navigate to the eDirectory Agent installation directory (by default, **\Program Files\Websense\bin**), and then open the **wsedir.ini** file in a text editor.
4. Modify the **QueryMethod** entry as follows:

   ```
   QueryMethod=0
   ```

This sets the Agent to use LDAP to query Novell eDirectory. (The default value is 1, for NCP.)

5. Save and close the file.
6. Restart the Websense eDirectory Agent service.

Enabling full eDirectory Server queries

In small networks, you can configure Websense eDirectory Agent to query the eDirectory Server for all logged-on users at regular intervals. This allows the agent to detect both newly logged-on users and users who have logged off since the last query, and to update its local user map accordingly.

Important

Configuring eDirectory Agent to use full queries is not recommended for larger networks, because the length of time required to return query results depends on the number of logged on users. The more logged-on users there are, the higher the performance impact.

When you enable full queries for eDirectory Agent, the **User entry timeout** interval is not used, because users who have logged off are identified by the query. By default, the query is performed every 30 seconds.

Enabling this feature increases eDirectory Agent processing time in 2 ways:

- Time needed to retrieve the names of logged-on users each time a query is performed
- Time required to process user name information, remove obsolete entries from the local user map, and add new entries based on the most recent query

eDirectory Agent examines the entire local user map after each query, rather than identifying only new logons. The time required for this process depends on the number of users returned by each query. The query process can therefore affect both eDirectory Agent and Novell eDirectory Server response times.

To enable full queries:

1. On the eDirectory Agent machine, navigate to the Websense **bin** directory (C:\Program Files\Websense\bin or /opt/Websense/bin, by default).
2. Locate the file **wsedir.ini** and make a backup copy in another directory.
3. Open **wsedir.ini** in a text editor (like Notepad or vi).
4. Go to the **[eDirAgent]** section of the file and find the following entry:

   ```
   QueryMethod=<N>
   ```

   Make a note of the QueryMethod value, in case you want to revert to the default setting later.
5. Update the **QueryMethod** value as follows:

- If the current value is 0 (communicate with the directory via LDAP), change the value to **2**.
- If the current value is 1 (communicate with the directory via NCP), change the value to **3**.

> **Note**
> If changing this query value slows system performance, return the QueryMethod entry to its previous value.

6. If the default query interval (30 seconds) is not appropriate for your environment, edit the **PollInterval** value appropriately.

   Note that the interval time is set in **milliseconds**.
7. Save and close the file.
8. Restart the Websense eDirectory Agent service (see *Stopping and starting Websense services*, page 258).

Configuring multiple agents

Related topics:

- *DC Agent*, page 193
- *Logon Agent*, page 196
- *RADIUS Agent*, page 199
- *eDirectory Agent*, page 203

It is possible to combine multiple transparent identification agents within the same network. If your network configuration requires multiple agents, it is best to install each agent on a separate machine. In some cases, however, you can configure Websense software to work with multiple agents on a single machine.

The following transparent identification agent combinations are supported:

Combination	Same machine?	Same network?	Configuration required
Multiple DC Agents	No	Yes	Ensure that all instances of DC Agent can communicate with Filtering Service.
Multiple RADIUS Agents	No	Yes	Configure each instance to communicate with Filtering Service.
Multiple eDirectory Agents	No	Yes	Configure each instance to communicate with Filtering Service.

Combination	Same machine?	Same network?	Configuration required
Multiple Logon Agents	No	Yes	Configure each instance to communicate with Filtering Service.
DC Agent + RADIUS Agent	Yes	Yes	Install these agents in separate directories. Configure each agent to communicate with Filtering Service using a different communication port.
DC Agent + eDirectory Agent	No	No	Websense software does not support communication with both Windows and Novell directory services in the same deployment. However, you can have both agents installed, with only 1 active agent.
DC Agent + Logon Agent	Yes	Yes	Configure both agents to communicate with Filtering Service. By default, each agent uses a unique port, so port conflicts are not an issue unless these ports are changed.
eDirectory Agent + Logon Agent	No	No	Websense software does not support communication with both Windows and Novell directory services in the same deployment. However, you can have both agents installed, with only 1 active agent.
RADIUS Agent + eDirectory Agent	Yes	Yes	Configure each agent to communicate with Filtering Service using a different communication port.
DC Agent + Logon Agent + RADIUS Agent	Yes	Yes	Though this combination is rarely required, it is supported. Install each agent in a separate directory. Configure all agents to communicate with Filtering Service using different communication ports.

Configuring different settings for an agent instance

The Websense Manager transparent identification agent configuration settings are global, and apply to all instances of the agent you have installed. If you have multiple instances of any agent, however, you can configure one instance independently of the others.

Unique settings you specify for a particular agent instance override the global settings in the Settings dialog box. Settings that can be overridden are marked with an asterisk (*).

1. Stop the transparent identification agent service (see *Stopping and starting Websense services*, page 258).
2. On the machine running the agent instance, navigate to the agent installation directory and open the appropriate file in a text editor:
   - for DC Agent: **transid.ini**
   - for Logon Agent: **authserver.ini**
   - for eDirectory Agent: **wsedir.ini**
   - for RADIUS Agent: **wsradius.ini**
3. Locate the parameter to change for this agent instance (see *INI file parameters*, page 212).

   For example, you can enable an authenticated connection between this agent instance and other Websense services. To do this, enter a value for the **password** parameter in the INI file:

   ```
   password=[xxxxxx]
   ```
4. Modify any other values as desired.
5. Save and close the INI file.
6. If you made a change to **DC Agent** settings, you must remove 2 files from the Websense **bin** directory (C:\Program Files\Websense\bin, by default):
   a. Stop all Websense services on the DC Agent machine (see *Stopping and starting Websense services*, page 258).
   b. Delete the following files:

      ```
      Journal.dat
      XidDcAgent.bak
      ```

      These files are recreated when you start the Websense DC Agent service.
   c. Restart the Websense services (including DC Agent), and then skip to **step 8**.
7. Restart the transparent identification agent service.
8. Update the agent settings in Websense Manager:
   a. Go to **Settings > User Identification**.
   b. Under **Transparent Identification Agents**, select the agent and then click **Edit**.

      Note

      If you modified the **port** value for this agent instance, remove and then re-add the agent. First select the existing agent entry and click **Delete**, and then click **Add Agent**.

   c. Verify the **Server IP or name** and **Port** this agent instance uses. If you specified a unique port number in the INI file, ensure that your entry matches that value.

d. If you specified a unique authentication password in the INI file, ensure that the **Password** entry shown here is correct.

e. Click **OK** to cache your changes. Changes are not implemented until you click **Save All**.

INI file parameters

Websense Manager field label	.ini parameter name	Description
Communications port (*all agents*)	port	The port over which the agent communicates with other Websense services.
Diagnostic Port (*all agents*)	DiagServerPort	The port over which the agent troubleshooting tool listens for data from the agent.
Password (*all agents*)	password	The password the agent uses to authenticate connections to other Websense services. Specify a password to enable authentication.
Query interval (*DC Agent*)	QueryInterval	The interval at which DC Agent queries domain controllers.
Server IP or name Port (*eDirectory Agent*)	Server=IP:port	The IP address and port number of the machine running eDirectory Agent.
Search base (*eDirectory Agent*)	SearchBase	The root context of the Novell eDirectory server.
Administrator distinguished name (*eDirectory Agent*)	DN	The name of the administrative user for Novell eDirectory server.
Password (*eDirectory Agent*)	PW	The password for the Novell eDirectory server administrative user.
RADIUS server IP or name	RADIUSHost	The IP address or name of your RADIUS server machine.
RRAS machine IP (Windows Only) (*RADIUS Agent*)	RRASHost	The IP address of the machine running RRAS. Websense queries this machine for user logon sessions.
Authentication Ports: Between RADIUS Agent and RADIUS server	AuthOutPort	The port on which the RADIUS server listens for authentication requests.
Authentication Ports: Between RADIUS clients and RADIUS Agent	AuthInPort	The port over which RADIUS Agent accepts authentication requests.

Accounting Ports: Between RADIUS Agent and RADIUS server	AccOutPort	The port over which the RADIUS server listens for RADIUS accounting messages.
Accounting Ports: Between RADIUS clients and RADIUS Agent	AccInPort	The port over which RADIUS Agent accepts accounting requests.

Configuring an agent to ignore certain user names

You can configure a transparent identification agent to ignore logon names that are not associated with actual users. This feature is often used to deal with the way that some Windows 200x and XP services contact domain controllers in the network.

For example, **user1** logs on to the network, and is identified by the domain controller as **computerA/user1**. That user is filtered by a Websense policy assigned to **user1**. If a service starts up on the user's machine that assumes the identity **computerA/ ServiceName** to contact the domain controller, this can cause filtering problems. Websense software treats **computerA/ServiceName** as a new user with no policy assigned, and filters this user by the computer policy, or by the **Default** policy.

To address this issue:

1. Stop the agent service (see *Stopping and starting Websense services*, page 258).
2. Navigate to the **\Websense\bin** directory, and open the **ignore.txt** file in a text editor.
3. Enter each user name on a separate line. Do not include wildcard characters, such as "*":

   ```
   maran01
   WindowsServiceName
   ```

   Websense software ignores these user names, regardless of which machine they are associated with.

   To prompt Websense software to ignore a user name within a specific domain, use the format **username, domain**.

   ```
   aperez, engineering1
   ```
4. When you are finished, save and close the file.
5. Restart the agent service.

The agent ignores the specified user names, and Websense software does not consider these names in filtering.

11 Delegated Administration

Related topics:

- *Introducing administrative roles*, page 216
- *Introducing administrators*, page 216
- *Getting started with administrative roles*, page 220
- *Enabling access to Websense Manager*, page 228
- *Using delegated administration*, page 231
- *Multiple administrators accessing Websense Manager*, page 241
- *Defining filtering restrictions for all roles*, page 242

Delegated administration provides powerful, flexible methods for managing Internet filtering and reporting for particular groups of clients. It is an effective way to distribute responsibility for Internet access management and reporting to individual managers when all users are centrally located. It is especially effective in larger organizations that include multiple locations and geographical regions, allowing the local administrators to manage Internet access and report on filtering activity for the users in their locale.

Implementing delegated administration involves creating an administrative role for each group of clients to be managed by the same administrators. Individual administrators in each role can be granted permissions to manage policy or generate reports for their clients, or both. See *Getting started with administrative roles*, page 220.

The Super Administrator role comes preinstalled, and includes the default administrative user: WebsenseAdministrator. Super Administrators have access to a wider range of policy and configuration settings than administrators in other roles. See *Super Administrators*, page 217.

Introducing administrative roles

Related topics:

- *Introducing administrators*, page 216
- *Getting started with administrative roles*, page 220

An administrative role is a collection of managed clients — users, groups, domains, organizational units, computers and network ranges — managed by one or more administrators. You grant the individual administrators permissions to apply policies to the role's clients, to generate reports, or both.

Websense software comes with a Super Administrator role predefined. There is also a default user, WebsenseAdministrator, which is automatically a member of the Super Administrator role. You can add administrators to this role, but you cannot delete the default administrator.

Important

You cannot delete the predefined Super Administrator role. The default user, WebsenseAdministrator, is an administrator in the Super Administrator role, but is not listed in the role. You cannot delete or change the permissions for WebsenseAdministrator.

Create as many roles as are appropriate for your organization. For example, you might create a role for each department, with the department manager as administrator and the department members as managed clients. In a geographically distributed organization, you might 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.

See *Introducing administrators*, page 216, for information on the options available for defining administrators.

See *Using delegated administration*, page 231, for instructions on creating roles and configuring permissions.

Introducing administrators

Administrators are the individuals who can access Websense Manager to manage policies or generate reports for a group of clients. The specific permissions available depend on the type of role.

- Super Administrator is a special role predefined in Websense Manager. This role offers the most flexibility for defining access permissions. See *Super Administrators*, page 217.
- Delegated administration roles must be created by a Super Administrator. Administrators of these roles have more limited access permissions. See *Delegated administrators*, page 219.

Additionally, you might create some delegated administration roles for reporting only, allowing various individuals to generate reports without giving them responsibility for policy management.

You can assign administrators to roles using their network logon credentials, or you can create special accounts used only to access Websense Manager. See *Enabling access to Websense Manager*, page 228.

Super Administrators

Related topics:

- *Introducing administrators*, page 216
- *Delegated administrators*, page 219
- *Administrators in multiple roles*, page 220

The Super Administrator role is created during installation. The default user, WebsenseAdministrator, is automatically assigned to this role. So, when you first log on with that user name and the password set during installation, you have full administrative access to all policy, reporting, and configuration settings in Websense Manager.

To preserve full access for this account, WebsenseAdministrator does not appear on the list of administrators for the Super Administrator role. It cannot be deleted, and the permissions cannot be modified.

You can add administrators to the Super Administrator role, as needed. Each administrator can be granted permissions as follows:

- **Policy** permissions allow Super Administrators to create and edit delegated administration roles, and copy filters and policies to these roles, as appropriate. They also can create and edit filtering components, filters, and policies, and apply policies to clients that are not managed by any other role.

  Additionally, Super Administrators with policy permissions can view the audit log, and are granted access to Websense configuration and other options, as follows:

  - **Unconditional** permissions gives the Super Administrator access to all system configuration settings for the Websense installation, such as account, Policy Server, and Remote Filtering Server settings, risk class assignments, and logging options.

Unconditional Super Administrators have the option of creating a Filter Lock that blocks certain categories and protocols for all users managed by delegated administration roles. See *Defining filtering restrictions for all roles*, page 242, for more information.

Unconditional Super Administrators can modify the Super Administrator role, adding and deleting administrators, as needed. They also can delete delegated administration roles or delete administrators or clients from these roles.

- **Conditional** permissions gives the Super Administrator access to database download, directory services, user identification, and Network Agent configuration settings. Conditional Super Administrators who also have reporting permissions can access configuration settings for the reporting tools.

  Conditional Super Administrators can add Websense user accounts, but cannot delete them. They can create and edit delegated administration roles, but cannot delete roles or the administrators or managed clients assigned to them. They also cannot delete administrators from the Super Administrator role.

- **Reporting** permissions enables Super Administrators to access all reporting features and report on all users. Unconditional Super Administrators are automatically given reporting permission.

  If an administrator is granted reporting permissions only, the Create Policy, Recategorize URL, and Unblock URL options in the Common Tasks list are unavailable. Additionally, the Check Policy option in the Toolbox is unavailable.

Creating multiple unconditional Super Administrators ensures that if the primary Super Administrator is not available, another administrator has access to all Websense policy and configuration settings.

Keep in mind that 2 administrators cannot log on at the same time to manage policy for the same role. See *Multiple administrators accessing Websense Manager*, page 241, for information on preventing conflicts.

The unique privileges of the Super Administrator role allow an administrator in the role access to all roles. To switch to another role after logon, go to the **Role** drop-down list in the banner, and select a role.

After changing roles, your policy permissions are limited to those available for the delegated administration role. Filters and policies you create are available only to administrators in that role. They can be applied only to managed clients in that role. See *Delegated administrators*, page 219.

Reporting permissions are cumulative, meaning that you get the combined permissions of all roles in which you are an administrator. Unconditional Super Administrators have full reporting permissions, regardless of which role is accessed.

Delegated administrators

Related topics:

- *Introducing administrators*, page 216
- *Super Administrators*, page 217
- *Administrators in multiple roles*, page 220

Delegated administrators manage clients assigned to a specific role. Assign each administrator policy permissions, reporting permissions, or both.

Delegated administrators who have **policy** permissions apply policies to the clients assigned to their role, thereby determining the Internet access available to each client. As part of this responsibility, delegated administrators can create, edit, and delete policies and filters, which are subject to the limitations of the Filter Lock established by the Super Administrator. See *Defining filtering restrictions for all roles*, page 242.

Note

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 256.

Delegated administrators cannot delete the Default policy.

Delegated administrators can edit filter components, with some limitations. See *Create policies and filters*, page 226, for more information.

Administrators with policy permissions who log on to Websense Manager with a Websense user account can also change their own Websense password. (See *Websense user accounts*, page 229.)

The options available to delegated administrators with **reporting** permissions vary according to the way the role is configured. They may be able to report on only those clients managed by their role, or they may be allowed to report on all clients. They may have access to all reporting features, or may have more limited reporting access. See *Editing roles*, page 233, for more information.

An administrator who has only reporting permissions has limited options available in the right shortcut pane (Common Tasks and Toolbox).

Administrators in multiple roles

Related topics:

- *Introducing administrators*, page 216
- *Super Administrators*, page 217
- *Delegated administrators*, page 219

Depending on the needs of your organization, the same administrator may be assigned to multiple roles. Administrators assigned to multiple roles must choose a single role to manage at logon.

After logon, your permissions are as follows:

- **Policy**: you can add and edit filters and policies for the role selected during logon, and apply policies to that role's managed clients. The Delegated Administration page lists all the roles to which you are assigned, allowing you to view each role's managed clients and reporting permissions.
- **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: report on managed clients only, investigative reports only
  - Role 3: report on all clients, full access to all reporting features

  In this situation, regardless of which role you choose during logon, you are permitted to view reports on the Today and History pages, and report on all clients, using all reporting features.

  If you are logged on for reporting only, the Role field in the banner bar indicates whether you have Full Reporting (report on all clients) or Limited Reporting (report on managed clients only) permissions.

Getting started with administrative roles

Related topics:

- *Introducing administrative roles*, page 216
- *Notifying administrators*, page 223
- *Delegated administrator tasks*, page 224

Getting started with delegated administration requires that the Super Administrator complete the following tasks:

- Decide how administrators will log on to Websense Manager. See *Enabling access to Websense Manager*, page 228.
- Add roles and configure them. See *Using delegated administration*, page 231.
- Inform administrators of their responsibilities and options. See *Notifying administrators*, page 223.

In addition to these required tasks, there are some optional tasks associated with delegated administration.

Creating the Filter Lock

Unconditional Super Administrators can create a Filter Lock, which designates specific categories and protocols as blocked for managed clients in all delegated administration roles. These restrictions are automatically enforced for all filters created in or copied to a delegated administration role, and cannot be modified by the delegated administrator.

> **Note**
> The Filter Lock does not apply to clients managed by the Super Administrator role.

The Filter Lock can also block and lock file types and keywords associated with selected categories, and enforce logging of selected protocols. See *Creating a Filter Lock*, page 242.

Moving clients

Adding a client to the Clients page while you are logged on as Super Administrator assigns that client to the Super Administrator role. That client cannot be added to a delegated administration role on the Edit Role page. Ideally, you should add the clients directly to the role, rather than assigning a policy within the Super Administrator role. However, this is not always possible.

To transfer clients from the Super Administrator role to another role, use the **Move to Role** option on the Clients page. See *Moving clients to roles*, page 63.

As part of the move, the policy applied in the Super Administrator role is copied to the delegated administration role. The filters that policy enforces are 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.

> **Note**
>
> Each time a filter or policy is copied to the same role, the (Copied) tag receives a number that is incremented with each new copy: (Copied 1), (Copied 2), and so on. Each becomes a separate filter or policy within the role.
>
> 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.

The Permit All filters in the Super Administrator role permit access to all categories or protocols, and cannot be edited. To preserve the Super Administrator's ability to implement a Filter Lock, these filters cannot be copied to a delegated administration role.

If the policy assigned to the client being moved enforces a Permit All filter, the client cannot be moved until you apply a policy that does not use a Permit All filter.

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.

Copying filters and policies

Initially, filters and policies created by a Super Administrator are available only to administrators in the Super Administrator role. You can use the **Copy to Role** option to copy filters and policies to a delegated administration role without moving a client to the role. See *Copying filters and policies to roles*, page 158.

When copying filters and policies 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 cannot be copied.
- 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.

Applying policies to remaining clients

Clients who are not specifically assigned to a delegated administration role are managed by Super Administrators. There is no Managed Clients list for the Super Administrator role.

To apply policies to these clients, add them to the Policy Management > Clients page. See *Adding a client*, page 61. Clients who have not been assigned a specific policy are governed by the Default policy for their role.

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, domain, or organizational unit 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.

Notifying administrators

Related topics:

- *Introducing administrative roles*, page 216
- *Getting started with administrative roles*, page 220

After assigning individuals as administrators in any administrative role, make sure to give them the following information.

- The URL for logging on to Websense Manager. By default:

  ```
  https://<ServerIP>:9443/mng/
  ```

  In place of <ServerIP>, use the IP address of the machine running Websense Manager.
- What Policy Server to choose during logon, if applicable. In an environment with multiple Policy Servers, administrators must choose a Policy Server during logon. They must choose 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 Websense user account when logging on to Websense Manager. If administrators log on with Websense user accounts, provide the user name and password.
- Their permissions, either to create and apply policies to clients in the role, or generate reports, or both.

  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 Websense Manager and manage policy for that role.
- How to find the list of clients managed by their role. Administrators can go to Policy Management > Delegated Administration, and then click their role name to display the Edit Role page, which includes a list of managed clients.
- Limitations imposed by the Filter Lock, if any categories or protocols have been blocked and locked.

- The tasks that are generally performed by administrators. See *Delegated administrator tasks*, page 224.

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.

Delegated administrator tasks

Related topics:

- *Introducing administrative roles*, page 216
- *Getting started with administrative roles*, page 220
- *Notifying administrators*, page 223

Delegated administrators who have **policy** permissions can perform the following tasks.

- *View your user account*, page 224
- *View your role definition*, page 225
- *Add clients to the Clients page*, page 225
- *Create policies and filters*, page 226
- *Apply policies to clients*, page 227

Reporting permissions can be granted at a granular level. The specific reporting permissions granted to your role determine which of the following tasks are available to administrators with reporting permissions. See *Generate reports*, page 227.

View your user account

Related topics:

- *Delegated administrator tasks*, page 224
- *View your role definition*, page 225
- *Add clients to the Clients page*, page 225
- *Create policies and filters*, page 226
- *Apply policies to clients*, page 227

If you log on to Websense Manager with network credentials, password changes are handled through your network directory service. Contact your system administrator for assistance.

If you have been assigned a Websense user name and password, view information about your account and change your password within Websense Manager.

1. Go to **Policy Management > Delegated Administration**.
2. Click **Manage Websense User Accounts** at the top of the page.
3. Click **Change Password** if you want to change your password. See *Changing a Websense user's password*, page 231.
4. Click **View** to display a list of roles in which you are an administrator.

View your role definition

Related topics:

- *Delegated administrator tasks*, page 224
- *View your user account*, page 224
- *Add clients to the Clients page*, page 225
- *Create policies and filters*, page 226
- *Apply policies to clients*, page 227

Open the Delegated Administration page and click your role name to display the Edit Role page, which lists the role's managed clients. This page also shows the reporting features available to administrators who have reporting permissions in this role.

Administrators who have only reporting permissions are unable to view this page. Only the specified reporting features are available to these administrators.

Add clients to the Clients page

Related topics:

- *Delegated administrator tasks*, page 224
- *View your user account*, page 224
- *View your role definition*, page 225
- *Create policies and filters*, page 226
- *Apply policies to clients*, page 227

Super Administrators assign managed clients to a role, but delegated administrators must add them to the Clients page before applying policies. See *Adding a client*, page 61, for instructions.

As soon as clients are added to the role's managed clients list, they are filtered by that role's Default policy. Clients who were moved to the role from the Super Administrator's Clients page are governed by the policy the Super Administrator applied, which was copied to the role when the client was moved.

Any client listed on the Delegated Administration > Edit Role page for your role can be added to the Clients page and assigned a policy. You can also can add individual users or computers who are members of a group, domain, organizational unit, or network range assigned as a managed client in your role.

Because a user may be part of multiple groups, domains, or organizational units, adding individuals from a larger client grouping has the potential to create conflicts when different roles manage groups, domains, or organizational units with common members. If administrators in different roles access Websense Manager at the same time, they might add the same client (individual member of a group, for instance) to their Clients page. In that situation, Internet filtering for that client is governed by the priority established for each role. See *Managing role conflicts*, page 239.

Create policies and filters

Related topics:

- *Delegated administrator tasks*, page 224
- *View your user account*, page 224
- *View your role definition*, page 225
- *Add clients to the Clients page*, page 225
- *Apply policies to clients*, page 227

When your role was created, it automatically inherited the preinstalled Default policy, category filter, and protocol filter, as they were defined at that time.There also may be policies and filters that the Super Administrator has chosen to copy to your role.

In addition to policies and filters, you also inherit any custom file types, and protocols created by the Super Administrator.

You are free to edit the policies and filters you inherit from the Super Administrator. Changes you make affect your role only. Any changes the Super Administrator makes to the policies and filters you inherited previously do not affect your role.

Note

Changes the Super Administrator makes to custom file types and protocols do automatically affect the filters and policies in your role.

When your Super Administrator informs you of changes in 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 67. For instructions on creating filters, see *Working with filters*, page 43.

You can edit filter components for your role, with some limitations.

- **Categories**: add custom categories, and edit both Master Database and custom categories, defining recategorized URLs and keywords for use within their role; change the action and advanced filtering option applied by default in category filters they create. (Changes to a category's default action are implemented only if the category is not locked by the Filter Lock.)
- **Protocols**: change the action and advance filtering options applied by default in protocol filters they create. (Changes to a protocol's default action are implemented only if the protocol is not locked by the Filter Lock.) Delegated administrators cannot add or delete protocol definitions.
- **File types**: view the file extensions assigned to each file type. Delegated administrators cannot add file types or change the extensions assigned to a file type.
- **Unfiltered URLS**: add URLs and add regular expressions that represent sites to be permitted for all managed clients in their role only.

For more information, see *Building filter components*, page 159.

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. See *Defining filtering restrictions for all roles*, page 242.

Apply policies to clients

Related topics:

- *Delegated administrator tasks*, page 224
- *View your user account*, page 224
- *View your role definition*, page 225
- *Add clients to the Clients page*, page 225
- *Create policies and filters*, page 226

After creating a policy, you can apply that policy directly to clients who have already been added to the Clients page by clicking the **Apply to Clients** button. See *Assigning a policy to clients*, page 70.

Alternatively, you can go to the Clients page and add the clients who should be governed by this policy. See *Working with clients*, page 54.

Generate reports

If you have reporting permissions, the specific reporting options available are set by the Super Administrator. To learn which features you can use, go to the Delegated Administration page, and click the role name. The Edit Role page shows the reporting features for which you have permissions. See *Editing roles*, page 233, for more information.

Enabling access to Websense Manager

When you configure delegated administration roles, you determine which Websense Manager features the administrators can access. To assure that the right features are available to individuals who log on to Websense Manager, each person must log on with a user name and password. Two types of accounts can be used:

- **Network accounts** use the credentials already established in your network directory service (see *Directory accounts*, page 228).
- **Websense user accounts** let you create a user name and password specifically for use within Websense Manager (see *Websense user accounts*, page 229).

Directory accounts

Related topics:

- *Enabling access to Websense Manager*, page 228
- *Websense user accounts*, page 229

Unconditional Super Administrators can use the **Settings > General > Logon Directory** page to enter the directory service information needed to allow administrators to log on to Websense Manager with their network credentials.

Note

This information is used to authenticate Websense Manager users only. It is not applied to filtering clients. Client directory service information is configured on the Settings > Directory Services page (see *Directory services*, page 56).

Websense Manager users' network credentials must be authenticated against a single directory service. If your network includes multiple directory services, a trusted relationship must exist between the Logon Directory service you configure in Websense Manager and the others.

If it is not possible to define a single directory service for use with Websense Manager, consider creating Websense user accounts for administrators (see *Websense user accounts*, page 229).

To define the directory service that Websense Manager should use to authenticate administrators, first verify that the check box for using a directory service to authenticate administrators is selected, and then select a **Directory service** type from the list.

If you select the default, **Windows NT Directory / Active Directory (Mixed Mode)**, no further configuration is needed. Click **OK** to cache your changes. Changes are not implemented until you click **Save All**.

If you select **Active Directory (Native Mode)** or **Other LDAP Directory**, provide the following additional information:

1. Enter the IP address or name of the machine on which the directory service is installed.

   If you are using Active Directory (Native Mode), and you have configured your global catalog servers for failover, you can instead enter the DNS domain name.
2. Enter the **Port** used for directory service communication.
3. To encrypt communication with the directory service, mark **Use SSL**.
4. Enter the **User distinguished name** and **Password** that Websense software should use to connect to the directory service.
5. Enter the **Default domain context** that Websense software should use when authenticating administrators.
   - If you are using Active Directory (Native Mode), configuration is complete. Click **OK** to cache your changes. Changes are not implemented until you click **Save All**.
   - If you are using another LDAP-based directory service, continue.
6. Supply the **User logon ID attributes** and the **User search filter**, if any, that Websense software should use to speed user authentication.

   This information also appears on the **Settings > Directory Services** page, under **Advanced Directory Settings**. You can copy and paste the values, if needed.
7. Under Group Options, specify whether or not your LDAP schema includes the **memberOf** attribute:
   - If memberOf is not used, specify the **User group search filter** that Websense software should apply to authenticate administrators.
   - If memberOf is used, specify the **Group attribute** that should be applied.
8. If your LDAP schema includes nested groups, mark **Perform additional nested group search**.
9. If your directory service uses LDAP referrals, indicate whether Websense software should use or ignore the referrals.
10. Click **OK** to cache your changes. Changes are not implemented until you click **Save All**.

Websense user accounts

Related topics:

- *Enabling access to Websense Manager*, page 228
- *Adding Websense user accounts*, page 230

Super Administrators use the **Delegated Administration > Manage Websense User Accounts** page to create accounts for administrators to access Websense Manager without entering network directory credentials. This page also lets Super Administrators change the password for Websense user accounts, and view the roles to which a Websense user is assigned as administrator.

Unconditional Super Administrators can also delete Websense user accounts from this page.

Delegated administrators use this page to change their Websense password and view the roles to which they are assigned as administrators.

Option	Description
Add	Opens the page for creating a new Websense user account. See *Adding Websense user accounts*, page 230.
Change Password	Opens the page for changing the password for the associated account. See *Changing a Websense user's password*, page 231.
View	Displays a list of roles to which this user is assigned as administrator.
Delete	Mark the check box for one or more obsolete user accounts, then click this button to delete them.
Close	Returns to the Delegated Administration page.

Adding Websense user accounts

Related topics:

- *Enabling access to Websense Manager*, page 228
- *Websense user accounts*, page 229
- *Changing a Websense user's password*, page 231

Use the **Delegated Administration > Manage Websense User Accounts > Add Websense User** page to add Websense user accounts.

1. Enter a unique **User name**, up to 50 characters.

   The name must be between 1 and 50 characters long, and cannot include any of the following characters:

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

   User names can include spaces and dashes.

2. Enter and confirm a **Password** (4-255 characters) for this user.

   Strong passwords are recommended: 8 characters or longer, including at least one each of the following:

   - upper case letter
   - lower case letter

- number
- special character (such as hyphen, underscore, or blank)

3. When you are finished making changes, click **OK** to cache the changes and return to the Manage Websense User Accounts page. Changes are not implemented until you click **Save All**.

Changing a Websense user's password

Related topics:

- *Enabling access to Websense Manager*, page 228
- *Websense user accounts*, page 229
- *Adding Websense user accounts*, page 230

The **Delegated Administration > Manage Websense User Accounts > Change Password** page lets delegated administrators change the password for their own Websense user account. Super Administrators can use this page to change the password for any Websense user account.

1. Verify that the correct **User Name** appears at the top of the page.
2. Enter and confirm the new **Password** (4-255 characters) for this user.

   Strong passwords are recommended: 8 characters or longer, including at least one each of the following:

   - upper case letter
   - lower case letter
   - number
   - special character (such as hyphen, underscore, or blank)

3. When you are finished making changes, click **OK** to cache the changes and return to the Manage Websense User Accounts page. Changes are not implemented until you click **Save All**.

Using delegated administration

Related topics:

- *Introducing administrative roles*, page 216
- *Managing role conflicts*, page 239

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 232.
Role	Click to view or configure the role. See *Editing roles*, page 233.
Delete	Click to delete any roles that are marked in the list. This option is available to unconditional Super Administrators only. See *Special considerations*, page 239, 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 239.
Manage Websense User Accounts	Click to add, edit, and delete user names and passwords for accounts used only to access Websense Manager. See *Websense user accounts*, page 229.
Manage Custom LDAP Groups	Click to add, edit, and delete custom LDAP groups, which can be assigned as managed clients in delegated administration roles. See *Working with custom LDAP groups*, page 60. This option is not available if the configured directory service is Windows NT/Active Directory (Mixed Mode).

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 233.
Manage Websense User Accounts	Click to access options for changing your Websense Manager password and viewing your assigned roles. See *Websense user accounts*, page 229.

Adding roles

Related topics:

- *Editing roles*, page 233
- *Special considerations*, page 239

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. Click **OK** to display the **Edit Role** page and define the characteristics of this role. See *Editing roles*, page 233.

   The new role is added to the Role drop-down list in the banner the next time you log on to Websense Manager.

Editing roles

Related topics:

- *Using delegated administration*, page 231
- *Adding roles*, page 232
- *Managing role conflicts*, page 239

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.

   Note
   The name of the Super Administrator role cannot be changed.

2. Add and delete administrators for this role. (Available to Super Administrators only, this section does not appear if you logged on as a delegated administrator.)

Item	Description
User Name	Administrator's user name.
Account Type	Indicates whether the user is defined in the network directory service (Directory) or as a Websense user account (Websense).

Item	Description
Reporting	Mark this check box to give the administrator permission to use reporting tools.
Policy	Mark this check box to 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 Administrators*, page 217.
Unconditional	Available only for the Super Administrator role, mark this check box to give the administrator permissions to manage all Websense configuration settings, and the Filter Lock. Only unconditional Super Administrators can grant unconditional permissions to a new administrator.
Add	Opens the **Add Administrator** page. See *Adding Administrators*, page 236.
Delete	Removes from the role any administrators marked in the Administrators list. (Available to unconditional Super Administrators only.)

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 *Delegated administrator tasks*, page 224.
Add	Opens the **Add Managed Clients** page. See *Adding managed clients*, page 237.
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 *Special considerations*, page 239, for more information.

4. Use the **Reporting Permissions** area to select the features available to administrators in this role who have reporting access.

a. Choose the general level of reporting permissions:

Option	Description
Report on all clients	Select this option to give administrators permission to generate reports on all network users. Use the remaining options in the Reporting Permissions area to set the specific permissions for administrators in this role.
Report on managed clients only	Select this option to limit administrators to reporting on the managed clients assigned to this role. Then, select the investigative reports features these administrators can access. Administrators limited to reporting on managed clients only cannot access presentation reports or user-based reports on the Today and History pages. They are also prevented from managing Log Database settings.

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 88.
View reports on Today and History pages	Enables display of charts showing Internet activity on these pages. See *Today: Health, Security, and Value Since Midnight*, page 19 and *History: Last 30 Days*, page 22. If this option is deselected, administrators can view only the Health Alert and Value areas of the Today page, and the Value Estimates on the History page.
Access investigative reports	Enables access to basic investigative reports features. When this option is selected, additional investigative reports features can be selected, also. See *Investigative reports*, page 105.
View user names in investigative reports	Enables administrators in this role to view user names, if they are logged. See *Configuring Filtering Service for logging*, page 280. Deselect this option to show only system-generated identification codes, instead of names. This option is available only when administrators are granted access to investigative reports.
Save investigative reports as favorites	Enables administrators in this role to create favorite investigative reports. See *Favorite investigative reports*, page 122. This option is available only when administrators are granted access to investigative reports.

Option	Description
Schedule investigative reports	Enables administrators in this role to schedule investigative reports to run at a future time or on a repeating cycle. See *Scheduling investigative reports*, page 125. This option is available only when administrators are granted permissions to save investigative reports as favorites.
Manage the Log Database	Enables administrators to access the Settings > Log Database page. See *Log Database administration settings*, page 294. This option is available only when administrators can report on all clients.

5. When you are finished making changes, click **OK** to cache the changes and return to the Delegated Administration page. Changes are not implemented until you click **Save All**.

Adding Administrators

Related topics:

- *Editing roles*, page 233
- *Enabling access to Websense Manager*, page 228

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 256.

1. If you plan to add directory accounts as delegated administrators, make sure you are logged on to the Policy Server whose Directory Service configuration (see *Directory services*, page 56) matches the Logon Directory configuration (see *Directory accounts*, page 228).

   If you are adding only Websense user accounts as administrators, you can be logged on to any Policy Server.

2. Under **Directory 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.

   If your environment uses Active Directory (Native Mode) or another LDAP-based directory service, you can search the directory to find specific user, group, domain, or organizational unit names. See *Searching the directory service*, page 62.

3. Under **Websense 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.
4. Set the **Permissions** for administrators in this role.

Option	Description
Policy	Check this option to let administrators in this role apply policies to their managed clients. This also grants access to certain Websense configuration settings.
Unconditional	Check this option to grant access to all Websense configuration settings. This option is available only when an unconditional Super Administrator add administrators to the Super Administrator role with policy permissions.
Reporting	Check this option to grant access to reporting tools. Use the Edit Role page to set the specific reporting features permitted.

5. When you are finished making changes, click **OK** to return to the Edit Role page.
6. Click **OK** on the Edit Role page to cache your changes. Changes are not implemented until you click **Save All**.

Adding managed clients

Related topics:

- *Using delegated administration*, page 231
- *Editing roles*, page 233

Managed clients are the users and computers assigned to a role, whose policies are set by the role's administrators. Directory clients (users, groups, domains, and organizational units), computers, and networks 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 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, domain, or organizational unit to 2 different roles. However, you can assign a user to one role, and then assign to a different role a group, domain, or organizational unit 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. 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 307).

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 managed clients in the same role be from the same directory service.

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, domain, or organizational unit names. See *Searching the directory service*, page 62.
   - Under **Computer**, enter the IP address of a computer to be added to this role.
   - Under **Network**, enter the first and last IP addresses in a range of computers to be added as a unit.
2. Click the right arrow (>) button adjacent to the client type to move the clients to the **Selected** list.
3. When you are finished making changes, click **OK** to return to the Edit Role page.
4. Click **OK** on the Edit Role page to cache your changes. Changes are not implemented until you click **Save All**.

Managing role conflicts

Related topics:

- *Using delegated administration*, page 231
- *Adding managed clients*, page 237

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 and organizational units.

Additionally, it is possible for a user to be managed by one role, and belong to a group, domain, or organizational unit 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 > Set Role Priority** page to tell Websense software what to do if different policies apply to the same user because of an overlap. When a conflict occurs, Websense software applies the filtering policy from the role that appears highest on this list.

1. Select any role on the list, except Super Administrator.

   Note
   The Super Administrator role is always first on this list. It cannot be moved.

2. Click **Move Up** or **Move Down** to change its position in the list.
3. Repeat steps 1 and 2 until all roles have the desired priority.
4. When you are finished making changes, click **OK** to cache the changes and return to the Delegated Administration page. Changes are not implemented until you click **Save All**.

Special considerations

Related topics:

- *Using delegated administration*, page 231
- *Editing roles*, page 233

Review the following information before deleting delegated administration roles or deleting managed clients from a role.

Deleting roles

On the **Delegated Administration** page, unconditional Super Administrators can delete any roles that have become obsolete.

Deleting a role also removes all clients that the role's administrators have added to the Clients page. After the role is deleted, if those clients belong to any networks, groups, or domains managed by other roles, they are governed by the appropriate policy applied in those roles (see *Filtering order*, page 71). Otherwise, they are governed by the Super Administrator's Default policy.

1. On the **Delegated Administration** page, mark the check box beside each role to be deleted.

   Note
   You cannot delete the Super Administrator role.

2. Click **Delete**.
3. Confirm the delete request to remove the selected roles from the Delegated Administration page. Changes are not permanent until you click **Save All**.

   The deleted role is cleared from Role drop-down list in the banner the next time you log on to Websense Manager.

Deleting managed clients

Clients cannot be deleted directly from the managed clients list (Delegated Administration > Edit Role) if:

- the administrator has applied a policy to the client
- the administrator has applied a policy to one or more members of a network, group, domain, or organizational unit

There may also be problems if, during Websense Manager logon, the Super Administrator chooses a different Policy Server than the one that communicates with the directory service containing the clients to be deleted. In this situation, the current Policy Server and directory service do not recognize the clients.

An unconditional Super Administrator can assure that the appropriate clients can be deleted, as follows.

1. Log on to Websense Manager selecting the Policy Server whose directory service contains the managed clients to be deleted. You must log on with unconditional Super Administrator permissions.
2. Open the **Role** list in the banner, and select the role from which managed clients are to be deleted.
3. Go to **Policy Management > Clients** to see a list of all the clients to which the delegated administrator has explicitly assigned a policy.

   This may include both clients that are specifically identified on the role's managed clients list, and clients who are members of networks, groups, domains, or organizational units on the managed clients list.

4. Delete the appropriate clients.
5. Click **OK** to cache the changes.
6. Open the **Role** list in the banner, and select the **Super Administrator** role.
7. Go to **Policy Management > Delegated Administration > Edit Role**.
8. Delete the appropriate clients from the managed clients list, and then click **OK** to confirm the delete request.
9. Click **OK** on the Edit Role page to cache the changes. Changes are not implemented until you click **Save All**.

Multiple administrators accessing Websense Manager

Related topics:

- *Introducing administrators*, page 216
- *Enabling access to Websense Manager*, page 228

Administrators in different roles can access Websense Manager simultaneously to perform whatever activities their role permissions allow. For example, administrators in Role A and Role B who both have policy permissions may log on to Websense Manager at the same time. Since they manage different clients, they can create and apply policies without conflict.

The situation is different if administrators who have policy permissions in the same role log on at the same time. To preserve the integrity of the policy structure and assignments, only one administrator from a role can access Websense Manager with policy permissions at any one time. If a second administrator with policy permissions for the same role tries to log on while the first administrator is still logged on, the second administrator is given a choice.

- 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.
- Try again later, after the first administrator logs off.

When administrators with both policy and reporting permissions log on to generate reports, they should immediately release their policy permissions so that other administrators in the role can perform policy management activities.

- Go to the **Role** drop-down list in the banner, and choose **Release Policy Permissions**.

An alternative approach is to create a special Websense user account (see *Websense user accounts*, page 229) for each role, and give that user only reporting permissions. Provide those logon credentials (user name and password) to administrators in the role who have both policy and reporting permissions. When administrators need to run

reports, they can log on as this reporting administrator, leaving policy access open for a different administrator.

Defining filtering restrictions for all roles

Related topics:

- *Introducing administrators*, page 216
- *Creating a Filter Lock*, page 242

Websense software allows unconditional Super Administrators to establish a Filter Lock that blocks categories and protocols for all clients managed by delegated administration roles. See *Creating a Filter Lock*, page 242, for more information.

Administrators of those roles have freedom to apply any filtering action to other categories and protocols in their policies, but those categories and protocols blocked in the Filter Lock cannot be permitted.

Changes to the Filter Lock are implemented for all managed clients as soon as the changes are saved. Delegated administrators who are working in Websense Manager when the changes take effect will not see the changes in their filters until the next time they log on.

Note

When a filter is copied from Super Administrator role to another role, the copy takes on the constraints of the Filter Lock.

Super Administrators are not limited by the Filter Lock. They can define policies that permit access to categories and protocols blocked and locked for delegated administration roles. Therefore, individuals who require special access rights should be managed by the Super Administrator role.

Creating a Filter Lock

Related topics:

- *Defining filtering restrictions for all roles*, page 242
- *Locking categories*, page 243
- *Locking protocols*, page 244

The **Policy Management > Filter Lock** page gives you the choice of whether to edit the categories or protocols to be blocked for all managed clients in delegated

administration roles. Any category or protocol feature 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 243.
- Click the **Protocols** button to block and lock protocols, or logging for protocols. See *Locking protocols*, page 244.

Locking categories

Related topics:

- *Defining filtering restrictions for all roles*, page 242
- *Creating a Filter Lock*, page 242
- *Locking protocols*, page 244

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 All**.

Locking protocols

Related topics:

- *Defining filtering restrictions for all roles*, page 242
- *Creating a Filter Lock*, page 242
- *Locking categories*, page 243

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 264.

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 Web sites using this protocol.
Lock protocol logging	Logs information about access to this protocol, and prevents delegated administrators from disabling logging.
Apply to Group	Applies the same settings to all protocols in the group.

3. When you are finished making changes, click **OK** to cache the changes and return to the Filter Lock page. Changes are not implemented until you click **Save All**.

12 Websense Server Administration

Related topics:

- *Websense product components*, page 246
- *Working with Policy Server*, page 251
- *Viewing and exporting the audit log*, page 256
- *Stopping and starting Websense services*, page 258
- *Alerting*, page 259
- *Backing up and restoring your Websense data*, page 267

Internet usage filtering requires interaction between several Websense software components:

- User requests for Internet access are received by Network Agent or a third-party integration product.
- The requests are sent to Websense Filtering Service for processing.
- Filtering Service communicates with Policy Server and Policy Broker to apply the appropriate policy in response to the request.

In most environments, a single Policy Database holds client, filter, policy, and general configuration information, whether there is a single Policy Server or multiple Policy Servers.

Each instance of Websense Manager is associated with a single Policy Database, and can be used to configure each Policy Server associated with that database.

Because the policy configuration performed in Websense Manager is stored in the central database, policy information is automatically available to all Policy Servers associated with that Policy Database.

Websense product components

Related topics:

- *Filtering components*, page 247
- *Reporting components*, page 249
- *User identification components*, page 249
- *Working with Policy Server*, page 251
- *Stopping and starting Websense services*, page 258
- *Reviewing current system status*, page 266

Websense software is made up of several components that work together to provide user identification, Internet filtering, and reporting capabilities. This section provides an overview of each component to help you understand and manage your filtering environment.

The primary Websense components include:

- Policy Database
- Policy Broker
- Policy Server
- Filtering Service
- Network Agent
- Master Database
- Websense Manager
- Usage Monitor
- User Service
- Log Server
- Log Database

Websense software also includes optional transparent identification agents:

- DC Agent
- RADIUS Agent
- eDirectory Agent
- Logon Agent

Additional optional components include:

- Remote Filtering Server
- Remote Filtering Client
- Websense Content Gateway

Filtering components

Component	Description
Policy Database	Stores Websense software settings and policy information.
Policy Broker	Manages requests from Websense components for policy and general configuration information.
Policy Server	• Identifies and tracks the location and status of other Websense components. • Stores configuration information specific to a single Policy Server instance. • Communicates configuration data to Filtering Service, for use in filtering Internet requests. Configure Policy Server settings in Websense Manager (see *Working with Policy Server*, page 251). Policy and most configuration settings are shared between Policy Servers that share a Policy Database (see *Working in a multiple Policy Server environment*, page 252).
Filtering Service	Provides Internet filtering in conjunction Network Agent or a third-party integration product. When a user requests a site, Filtering Service receives the request and determines which policy applies. • Filtering Service must be running for Internet requests to be filtered and logged. • Each Filtering Service instance downloads its own copy of the Websense Master Database. Configure filtering and Filtering Service behavior in Websense Manager (see *Internet Usage Filters*, page 33, and *Configuring Websense filtering settings*, page 50).
Network Agent	• Enhances filtering and logging functions • Enables protocol management • Enables filtering in a stand-alone environment For more information, see *Network Configuration*, page 311.
Master Database	• Includes more than 36 million Web sites, sorted into more than 90 categories and subcategories • Contains more than 100 protocol definitions for use in filtering protocols Download the Websense Master Database to activate Internet filtering, and make sure that the database is kept up to date. If the Master Database is more than 2 weeks old, no filtering can occur. See *The Websense Master Database*, page 28, for more information.
Websense Manager	Serves as the configuration and management interface to Websense software. Use Websense Manager to define and customize Internet access policies, add or remove filtering clients, configure Websense software components, and more. See *Working in Websense Manager*, page 14, for more information.

Component	Description
Usage Monitor	Enables alerting based on Internet usage. Usage Monitor tracks URL category and protocol access, and generates alert messages according to the alerting behavior you have configured. See *Alerting*, page 259, for more information.
Remote Filtering Client	• Resides on client machines outside the network firewall. • Identifies the machines as clients to be filtered, and communicates with Remote Filtering Server. See *Filter Remote Clients*, page 143, for more information.
Remote Filtering Server	• Allows filtering of clients outside a network firewall. • Communicates with Filtering Service to provide Internet access management of remote machines. See *Filter Remote Clients*, page 143, for more information.
Websense Content Gateway	• Provides a robust proxy and cache platform. • Can analyze the content of Web sites and files in real time to categorize previously uncategorized sites. See *Analyze Content with the Real-Time Options*, page 131.
Websense Security Gateway	In addition to standard Websense Content Gateway functionality: • 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. See *Analyze Content with the Real-Time Options*, page 131.

Reporting components

Component	Description
Log Server	Logs Internet request data, including: • 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 With Network Agent and some integration products, Log Server also stores information about the amount of bandwidth used. Log Server must be installed on a Windows machine to enable investigative and presentation reports, and Today and History page charts, in Websense Manager. After installing Log Server, configure Filtering Service to pass logging data to the correct location (see *Configuring Filtering Service for logging*, page 280).
Log Database	Stores Internet request data collected by Log Server for use by Websense reporting tools.

User identification components

Component	Description
User Service	• Communicates with your directory service. • Conveys user-related information, including user-to-group and user-to-domain relationships, to Policy Server and Filtering Service, for use in applying filtering policies. If you have installed and configured a Websense transparent identification agent (see *Transparent identification*, page 183), User Service helps to interpret user logon session information, and uses this information to provide user name-to-IP-address associations to Filtering Service. When you add users and groups as Websense clients (see *Adding a client*, page 61), User Service provides name and path information from the directory service to Websense Manager. For information about configuring directory service access, see *Directory services*, page 56.
DC Agent	• Offers transparent user identification for users in a Windows-based directory service. • Communicates with User Service to provide up-to-date user logon session information to Websense software for use in filtering. For more information, see *DC Agent*, page 193.

Component	Description
Logon Agent	• 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. Logon Agent communicates with the logon application on client machines to ensure that individual user logon sessions are captured and processed directly by Websense software. For more information, see *Logon Agent*, page 196.
eDirectory Agent	• 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. For more information, see *eDirectory Agent*, page 203.
RADIUS Agent	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. For more information, see *RADIUS Agent*, page 199.

Understanding the Policy Database

The Websense Policy Database stores both the policy data (including clients, filters, filter components, and delegated administration settings) and global configuration settings specified in Websense Manager. Settings specific to a single Policy Server instance are stored separately.

In most multiple Policy Server environments, a single Policy Database holds policy and general configuration data for multiple Policy Servers.

1. At startup, each Websense component requests applicable configuration information from the Policy Database via the Policy Broker.
2. Running components frequently check for changes to the Policy Database.
3. The Policy Database is updated each time administrators make changes in Websense Manager and click Save All.
4. After a change to the Policy Database, each component requests and receives the changes that affect its functioning.

Back up the Policy Database on a regular basis to safeguard important configuration and policy information. See *Backing up and restoring your Websense data*, page 267, for more information.

Working with Policy Server

Policy Server is the Websense software component that manages policy information and communicates with Filtering Service to aid in policy enforcement. Policy Server is also responsible for identifying other components and tracking their location and status.

When you log on to Websense Manager, you are logging onto a graphical interface to Policy Server.

- You cannot log on to Websense Manager until it is configured to communicate with Policy Server.
- If your Websense software installation includes multiple Policy Servers, you can choose between Policy Server instances at logon time.
- You can add and remove Policy Server instances within Websense Manager.

By default, communication between Websense Manager and a central Policy Server instance is established during Websense Manager installation.

Most environments require only one Policy Server. A single Policy Server can communicate with multiple Filtering Service and Network Agent instances for load balancing. In very large organizations (10,000+ users), however, it may help to install multiple instances of Policy Server. If you install additional Policy Servers, add each instance to Websense Manager (see *Adding and editing Policy Server instances*, page 251).

Adding and editing Policy Server instances

Use the **Settings > Policy Server** page to add Policy Server instances to Websense Manager, or to configure or remove existing Policy Servers.

To add a Policy Server instance:

1. Click **Add**. The Add Policy Server page opens.
2. Enter the IP address or host name of the Policy Server machine in the **Server name or IP** field.
3. Enter the **Port** that Websense Manager should use to communicate with that Policy Server instance. The default is **55806**.
4. Click **OK** to return to the Policy Server page. The new Policy Server instance appears in the list.
5. Click **OK** to cache all changes to the Policy Servers page. Changes are not implemented until you click **Save All**.

To edit a Policy Server instance (for example, if the Policy Server machine IP address or name changes), select an IP address or host name in the Policy Server list, and then click **Edit**.

To delete a Policy Server instance, select an IP address or host name in the Policy Server list, and then click **Delete**. Clicking Delete removes the Policy Server instance

from Websense Manager, but does not uninstall or stop the Websense Policy Server service. If there is only one instance of Policy Server listed, you cannot delete that instance.

Working in a multiple Policy Server environment

In some distributed environments with a large number of users, it may be appropriate to install multiple Policy Servers. This entails some special considerations.

- If you implement a configuration that allows the same client to be managed by different Policy Servers, depending on current load, do **not** implement time-based policy actions:
  - Password Override
  - Confirm
  - Quota

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

  Remember that the Default policy is enforced whenever no other policy applies to a client. If clients can be governed by more than one Policy Server, you may want to make sure that the Default policy does not enforce category filters that apply time-based actions.
- Because policy information is stored in the Policy Database, policy changes are automatically shared between all Policy Servers when you click **Save All**.
- Many global configuration settings (like risk class definitions and alerting options) are also shared between Policy Servers.
- Configuration settings that are specific to a single Policy Server (like its Filtering Service and Network Agent connections) are stored locally by each Policy Server and not distributed.

To switch between Policy Servers in Websense Manager to review or configure settings that apply to a single Policy Server instance:

1. In the Websense banner, expand the **Policy Server** list and select an IP address.
2. If there are unsaved changes to the current Policy Server instance, a list of changes is displayed. Do one of the following:
   - Click **Save All and Log Out** to save the changes and log out of the current Policy Server.
   - Click **Cancel Changes and Log Out** to abandon the changes and log out of the current Policy Server.
   - Click **Go Back** to continue configuring the current Policy Server.

   If there are no unsaved changes, you are taken directly to the logon screen.
3. At the logon screen, enter a user name and password to log on to the selected Policy Server, and then click **Log On**.

Changing the Policy Server IP address

Before changing the IP address of the Policy Server machine, **stop all Websense services** on the machine. If Websense Manager is also installed on the machine, this includes the Apache2Websense and ApacheTomcatWebsense services.

After changing the IP address, you must manually update Websense configuration files used by Websense Manager, Policy Server, and other Websense services before filtering resumes.

Step 1: Update Websense Manager configuration

Update Websense Manager to use the new IP address to connect to Policy Server.

1. On the Websense Manager machine, stop the **Apache2Websense** and **ApacheTomcatWebsense** services (if necessary).

   If Websense Manager and Policy Server are installed on this same machine, the Apache services should already be stopped.
2. Navigate to the following directory:
   - Windows:

     `C:\Program Files\Websense\tomcat\conf\Catalina\localhost\`
   - Linux:

     `/opt/Websense/tomcat/conf/Catalina/localhost/`
3. Locate the **mng.xml** file, and then make a backup copy of the file in another directory.
4. Open **mng.xml** in a text editor (like Notepad or vi) and replace each instance of the old Policy Server IP address with the new one.

   The Policy Server IP address appears twice: as the **ps/default/host** value and the **psHosts** value.
5. When you are finished, save and close the file.

Do not restart the Apache 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 258).
2. Navigate to the Websense **bin** directory.
   - Windows:

     `C:\Program Files\Websense\bin`
   - Linux

     `/opt/Websense/bin`

3. Locate the **config.xml** file, and then make a backup copy of the file in another directory.
4. Open **config.xml** in a text editor and replace each instance of the old Policy Server IP address with the new one.
5. When you are finished, save and close the file.
6. In the **bin** directory, locate the **websense.ini** file, and then make a backup copy in another directory.
7. Open **websense.ini** in a text editor and replace each instance of the old Policy Server IP address with the new one.
8. When you are finished, save and close the file.

Step 3: Verify the Log Database connection

Use the Windows ODBC Data Source Administrator on the Policy Server machine to verify the ODBC connection to the Log Database.

1. Go to **Start > Settings > Control Panel > Administrative Tools > Data Sources (ODBC)**.
2. On the **System DSN** tab, select the appropriate data source name (by default, **wslogdb70**), and then click **Configure**.
3. Verify that the correct database server machine is selected, and then click **Next**.
4. Enter the credentials used to connect to the database, and then click **Next**.
5. Accept the defaults on the next 2 screens, and then click **Test Data Source**.

> **Note**
>
> If the test fails, check the database server machine name and try again.
>
> If the machine name is correct, but the test continues to fail, verify that the correct connection port is being used, and that the firewall allows communication on the selected port.

Step 4: Restart Websense services

1. Reboot the Policy Server machine. Make sure that all Websense services on the machine restart normally.
2. If the Websense Manager used to configure this Policy Server is installed on another machine, restart the **Apache2Websense** and **ApacheTomcatWebsense** services on that machine.

> **Note**
>
> If Websense Manager is installed on the same machine as Policy Server, administrators must use the new IP address to log on.

Working with Filtering Service

Filtering Service is the Websense software component that works with Network Agent or a third-party integration product to filter Internet activity. When a user requests a site, Filtering Service receives the request, determines which policy applies, and uses the applicable policy to determine how the site is filtered.

Each Filtering Service instance downloads its own copy of the Websense Master Database to use in determining how to filter Internet requests.

Filtering Service also sends information about Internet activity to Log Server, so that it can be recorded and used for reporting.

When you log on to Websense Manager, a **Filtering Service Summary** on the Status > Today page lists the IP address and current status of each Filtering Service instance associated with the current Policy Server. Click a Filtering Service IP address for more detailed information about the selected Filtering Service.

Review Filtering Service details

Use the **Status > Today > Filtering Service Details** page to review the status of an individual Filtering Service instance.

The page lists:

- The Filtering Service IP address
- Whether or not the selected instance is running
- The Filtering Service version

  This should match your Websense software version, including any hotfixes that have been applied.
- The operating system running on the Filtering Service machine
- The Websense software platform

  This indicates whether Websense software is running in stand-alone mode or integrated with a third-party product.
- The IP address and status of any Network Agent instances with which the selected Filtering Service communicates.

Click **Close** to return to the Today page.

Review Master Database download status

Each Filtering Service instance in your network downloads its own copy of the Master Database. When you are working in Websense Manager, the Health Alert Summary on the Status > Today page displays a status message when a Master Database download is in progress, or if a download attempt fails.

For detailed information about recent or ongoing database downloads, click **Database Download** on the Today page toolbar. The Database Download page includes an entry for each Filtering Service instance associated with the current Policy Server.

Initially, the Database Download page displays a quick download summary, showing where the database was downloaded, which database version was downloaded, and whether the download was successful. From this summary view, you can:

- Initiate a database download for a single Filtering Service (click **Update**).
- Initiate database downloads for all listed Filtering Service instances (click **Update All**).
- Cancel one or all ongoing updates.

Click an IP address in the list on the right to review more detailed database download status for the selected Filtering Service.

- If the selected Filtering Service has encountered download problems, a recommendation for addressing the problem may be displayed.
- To manually initiate a database download for the selected Filtering Service, click **Update**.

During database download, the status screen shows detailed progress information for each stage of the download process. Click **Close** to hide progress information and continue working in Websense Manager.

Resumable Master Database downloads

If a Master Database download is interrupted, Websense software attempts to resume the download automatically. If Filtering Service is able to reconnect to the download server, the download resumes from where it was interrupted.

You can manually restart a failed or interrupted download. This does not resume the download from the point of interruption, but instead restarts the process from the beginning.

1. In Websense Manager, go to **Status > Today** and click **Database Downloads**.
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.

Viewing and exporting the audit log

Websense software provides an audit trail showing which administrators have accessed Websense 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 Administrators*, page 217).

Delegated administrators have significant control over the Internet activities of their managed clients. Monitoring their changes through the audit log enables you to ensure that this control is handled responsibly and in accordance with your organization's acceptable use policies.

Use the **Status > Audit Log** page to view the audit log, and to export selected portions of it to an Excel spreadsheet (XLS) file, if desired.

Audit records are saved for 60 days. To preserve audit records longer than 60 days, use the export option to export the log on a regular basis. Exporting does not remove records from the audit log.

When the Audit Log page opens, the most recent records are shown. Use the scroll bar and the paging buttons above the log to view older records.

The log displays the following information. If an item is truncated, click the partial entry to display the full record in pop-up dialog box.

Column	Description
Date	Date and time of the change, adjusted for time zones. To assure consistent data in the audit log, be sure all machines running Websense components have their date and time settings synchronized.
User	User name of the administrator who made the change.
Server	IP address or name of machine running the Policy Server affected by the change. This appears only for changes that affect the Policy Server, such as changes made on the Settings tab.
Role	Delegated administration role affected by the change. When a change affects a client explicitly assigned as a managed client in the delegated administrator's role, that change shows as affecting the Super Administrator role. If the change affects a client that is a member of a network range, group, domain or organizational unit assigned to the role, the change shows as affecting the delegated administrator's role.
Type	Configuration element that was changed, such as policy, category filter, or logon/logoff.
Element	Identifier for the specific object changed, such as the category filter name or role name.
Action	Type of change made, such as add, delete, change, log on, and so on.
Previous	Value before the change.
Current	New value after the change.

Not all items are shown for all records. For example, the role is not displayed for logon and logoff records.

To export audit log records:

1. Select a time period from the **Export range** list.

   Choose **Last 60 days** to export the entire audit log file.

2. Click **Go**.

   If Microsoft Excel is installed on the machine running Websense Manager, the exported file opens. Use options in Excel to save or print the file.

   If Microsoft Excel is not installed on the machine running Websense Manager, follow the on-screen instructions to either locate the software or save the file.

Stopping and starting Websense services

Websense services are configured to start each time the machine restarts. However, in some cases you need to stop or start one or more product components separately from a machine restart.

> **Note**
> If Filtering Service is in the process of downloading the Master Database, it does not stop running until the download is complete.

When you stop all Websense services, always end with the following 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 following services, in the order shown:

1. Websense Policy Database
2. Websense Policy Broker
3. Websense Policy Server

Windows

1. Open the Windows Services dialog box (**Start > Settings > Control Panel > Administrative Tools > Services**).
2. Right-click the Websense service name, and then select **Stop** or **Start**.

Linux

On Linux machines, all services stop and start together when you use this procedure.

1. Go to the **/opt/Websense** directory.
2. Check the status of the Websense services with the command:
   - `./WebsenseAdmin status`
3. Stop, start, or restart all Websense services with the commands:
   - `./WebsenseAdmin stop`
   - `./WebsenseAdmin start`
   - `./WebsenseAdmin restart`

Warning

Do not use the **kill** command to stop a Websense service, as it may corrupt the service.

Alerting

Related topics:

- *Flood control*, page 260
- *Configuring general alert options*, page 260
- *Configuring system alerts*, page 262
- *Configuring category usage alerts*, page 263
- *Configuring protocol usage alerts*, page 264

To facilitate tracking and management of both Websense software and client Internet activity, Super Administrators can configure alerts to be sent when selected events occur.

- **System alerts**: Notification regarding subscription status and Master Database activity.
- **Usage alerts**: Notification when Internet activity for particular categories or protocols reaches configured thresholds.

Alerts can be sent to selected recipients via email, on-screen pop-up messages (Windows **net send** messaging), or SNMP messages.

Note

On-screen pop-up alerts cannot be sent to Linux machines. However, they can be sent from a Linux machine running Policy Server to Windows machines, provided that the Samba client is installed on the Linux machine. See the *Deployment Guide*.

Usage alerts can be generated for both Websense-defined and custom categories or protocols.

Flood control

Related topics:

- *Alerting*, page 259
- *Configuring general alert options*, page 260
- *Configuring category usage alerts*, page 263
- *Configuring protocol usage alerts*, page 264

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 260, for more information.

You can also set threshold limits for each category and protocol usage alert. For example, if you set a threshold limit of 10 for a certain category, an alert is generated after 10 requests for that category (by any combination of clients). See *Configuring category usage alerts*, page 263, and *Configuring protocol usage alerts*, page 264, for more information.

Suppose that the maximum daily alerts setting is 20, and the category alert threshold is 10. Administrators are only alerted the first 20 times category requests exceed the threshold. That means that only the first 200 occurrences result in alert messages (threshold of 10 multiplied by alert limit of 20).

Configuring general alert options

Related topics:

- *Alerting*, page 259
- *Configuring system alerts*, page 262
- *Configuring category usage alerts*, page 263
- *Configuring protocol usage alerts*, page 264

Websense software can notify administrators of various kinds of system events, such as updates to Master Database categories and subscription issues, as well as Internet usage that exceeds defined thresholds.

Use the **Settings > Alerts and Notifications > Alerts** page to select and configure the desired notification methods, as described below. Then, use the other pages in the Settings > Alerts and Notifications section to enable the alerts you want to receive.

1. Enter a number in the **Maximum daily alerts per usage type** field to limit the total number of alerts generated daily for each category and protocol usage alert.

For example, you might configure usage alerts to be sent every 5 times (threshold) someone requests a site in the Sports category. Depending on the number of users and their Internet use patterns, that could generate hundreds of alerts each day.

If you enter 10 as the maximum daily alerts per usage type, only 10 alert messages are generated each day for the Sports category. In this example, these messages alert you to the first 50 requests for Sports sites (5 requests per alert multiplied by 10 alerts).

2. Mark the **Enable email alerts** check box to deliver alerts and notifications by email. Then, configure these email settings.

SMTP server IP or name	IP address or name for the SMTP server through which email alerts should be routed.
From email address	Email address to use as the sender for email alerts.
Administrator email address (To)	Email address of the primary recipient of email alerts.
Recipient email addresses (Cc)	Email address for up to 50 additional recipients. Each address must be on a separate line.

3. Mark the **Enable pop-up alerts** check box to display pop-up messages on specific computers. Then, enter the IP address or machine name for up to 50 **Recipients**, each on a separate line.

> **Note**
>
> Pop-up alerts cannot be sent to Linux machines. However, they can be sent from a Linux machine running Policy Server to Windows machines, provided that the Samba client is installed on the Linux machine. See the *Deployment Guide*.

4. Mark the **Enable SNMP alerts** check box to deliver alert messages through an SNMP Trap system installed in your network. Then, provide information about your SNMP Trap system.

Community name	Name of the trap community on your SNMP Trap server.
Server IP or name	IP address or name of the SNMP Trap server.
Port	Port number SNMP messages use.

5. When you are finished, click **OK** to cache your changes. Changes are not implemented until you click **Save All**.

Configuring system alerts

Related topics:

- *Alerting*, page 259
- *Configuring general alert options*, page 260
- *Reviewing current system status*, page 266

Websense Manager displays detailed system health and status information via the **Status > Alerts** (detailed information) page, described in *Reviewing current system status*, page 266.

To assure that administrators are notified of significant system events, like a database download failure or a subscription that is about to expire, when they are not logged on to Websense Manager, configure Websense system alerts to be distributed by email, pop-up message, or through your SNMP Trap system.

On the Settings tab, use the **Alerts and Notifications > System** page to select the method used to send these alerts to Websense administrators, as well as which alerts to send.

1. For each alert, mark the delivery methods to be used. Depending on what methods are enabled on the Alerts page, you may be able to choose **Email**, **Pop-up**, and **SNMP** methods.

   Note
   In addition to generating an alert, information about Master Database download failures and exceeded subscription levels is logged in the Windows Event Viewer (Windows only) and in the Websense.log file (Windows and Linux).

   Alerts are available for events such as:

   - Your subscription expires in one week.
   - Search engines supported for Search Filtering have changed.
   - A Websense Master Database download failed.
   - A category or protocol was added to or removed from the Master Database.
   - The number of current users exceeds your subscription level.
   - The number of current users has reached 90% of your subscription level.
   - Your subscription expires in one month.
   - Websense Master Database has been updated.
2. When you are finished, click **OK** to cache your changes. Changes are not implemented until you click **Save All**.

Configuring category usage alerts

Related topics:

- *Alerting*, page 259
- *Flood control*, page 260
- *Configuring general alert options*, page 260
- *Adding category usage alerts*, page 263

Websense software can notify you when Internet activity for particular URL categories reaches a defined threshold. You can define alerts for permitted requests or for blocked requests to the category.

For example, you might want to be alerted each time 50 requests for sites in the Shopping category have been permitted to help decide whether to place restrictions on that category. Or, you might want to receive an alert each time 100 requests for sites in the Entertainment category have been blocked, to see whether users are adapting to a new Internet use policy.

On the Settings tab, use the **Alerts and Notifications > Category Usage** page to view the alerts that have already been established, and to add or delete usage alert categories.

1. View the **Permitted Category Usage Alerts** and **Blocked Category Usage Alerts** lists to learn which categories are configured for alerts, the threshold for each, and the selected alert methods.
2. Click **Add** below the appropriate list to open the Add Category Usage Alerts page (see *Adding category usage alerts*, page 263) 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 Alerts page. Changes are not implemented until you click **Save All**.

Adding category usage alerts

Related topics:

- *Alerting*, page 259
- *Configuring general alert options*, page 260
- *Configuring category usage alerts*, page 263

The **Add Category Usage Alerts** page appears when you click Add on the Category Usage Alerts page. Here, you can select new categories for usage alerts, establish the threshold for these alerts, and select the alert methods.

1. Mark the check box beside each category to be added with the same threshold and alert methods.

   Note
   You cannot add usage alerts for any category that is excluded from logging. See *Configuring Filtering Service for logging*, page 280.

2. Set the **Threshold** by selecting the number of requests that cause an alert to be generated.
3. Mark the check box for each desired alert method (**Email**, **Pop-up**, **SNMP**) for these categories.

   Only the alert methods that have been enabled on the Alerts page (see *Configuring general alert options*, page 260) are available for selection.
4. Click **OK** to cache your changes and return to the Category Usage Alerts page (see *Configuring category usage alerts*, page 263). Changes are not implemented until you click **Save All**.

Configuring protocol usage alerts

Related topics:

- *Alerting*, page 259
- *Flood control*, page 260
- *Configuring general alert options*, page 260
- *Adding protocol usage alerts*, page 265

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 and Notifications > Protocol Usage Alerts** page to view the alerts that have already been established, and to add or delete protocols for usage alerts.

1. View the **Permitted Protocol Usage Alerts** and **Blocked Protocol Usage Alerts** lists to learn which protocols are configured for alerts, the threshold for each, and the selected alert methods.
2. Click **Add** below the appropriate list to open the Add Protocol Usage Alerts page (see *Adding protocol usage alerts*, page 265) 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 Alerts page. Changes are not implemented until you click **Save All**.

Adding protocol usage alerts

Related topics:

- *Alerting*, page 259
- *Configuring general alert options*, page 260
- *Configuring protocol usage alerts*, page 264

Use the **Protocol Usage Alerts > Add Protocol Usage Alerts** page to select new protocols for usage alerts, establish the threshold for these alerts, and select the alert methods.

1. Mark the check box beside each protocol to be added with the same threshold and alert methods.

   Notes

   You cannot select a protocol for alerting unless it is configured for logging in one or more protocol filters.

   Protocol alerts only reflect usage by clients governed by a protocol filter that logs the protocol.

2. Set the **Threshold** by selecting the number of requests that cause an alert to be generated.
3. Select each desired alert method (**Email**, **Pop-up**, **SNMP**) for these protocols.

   Only the alert methods that have been enabled on the Alerts page (see *Configuring general alert options*, page 260) are available for selection.
4. Click **OK** to cache changes and return to the Protocol Usage Alerts page (see *Configuring protocol usage alerts*, page 264). Changes are not implemented until you click **Save All**.

Reviewing current system status

Use the **Status > Alerts** page to find information about problems affecting the health of your Websense software, get troubleshooting help, and review the details of recent real-time updates to the Websense Master Database.

The **Active Alerts** list shows the status of monitored Websense software components.

- For detailed information about which components are monitored, click **What is monitored?** above the list of alert messages.
- To troubleshoot a problem, click the **Solutions** button next to the error or warning message.
- To hide an alert message, click **Advanced**. If your organization does not use Log Server, Network Agent, or User Service, or if you do not plan to enable WebCatcher, mark a check box to hide the associated alert. When you are finished, click **OK** to enact the change.

  Click **Advanced** again to hide the advanced options.

The **Real-Time Database Updates** list provides information about emergency updates to the Websense Master Database, showing:

- When the update occurred
- The update type
- The new database version number
- The reason for the update
- The IP address of the Filtering Service instance that received the update

These supplemental updates occur in addition to regular, scheduled Master Database updates, and can be used, for example, to recategorize a site that has been temporarily miscategorized. Websense software checks for database updates every hour.

For Websense Web Security users, the Alerts page includes a third list: **Real-Time Security Updates**. This list has the same format as the Real-Time Database Updates list, but specifically shows security-related database updates.

Installing security updates as soon as they are created eliminates vulnerability to threats such as new phishing (identity fraud) scams, rogue applications, or malicious code infecting a mainstream Web site or application.

For more information about Real-Time Security Updates, see *Real-Time Security Updates™*, page 29.

Use the **Print** button, above the page, to open a secondary window with a printable version of the Alerts area. Use browser options to print this page, which omits all the navigation options found in the main Websense Manager window.

Backing up and restoring your Websense data

Related topics:

- *Scheduling backups*, page 269
- *Running immediate backups*, page 270
- *Maintaining the backup files*, page 271
- *Restoring your Websense data*, page 271
- *Discontinuing scheduled backups*, page 272
- *Command reference*, page 273

The Websense Backup Utility makes it easy to back up your Websense software settings and policy data, and to revert to a previous configuration. Data saved by the utility can also be used to import Websense configuration information after an upgrade.

The Backup Utility saves:

- Global configuration information, including client and policy data, stored in the Policy Database.
- Local configuration information, such as Filtering Service and Log Server settings, stored by each Policy Server.
- Websense component initialization and configuration files.

The backup process works as follows:

1. You initiate an immediate backup (see *Running immediate backups*, page 270) or define a backup schedule (see *Scheduling backups*, page 269).
   - Manually launch a backup at any time.
   - Backup files are stored in a directory you specify when you run or schedule the backup.
2. The Backup Utility checks all Websense components on the machine, collects the data eligible for backup, and creates an archive file. The file name is given the format:

   `wsbackup_yyyy-mm-dd_hhmmss.tar.gz`

   Here, *yyyy-mm-dd_hhmmss* represents the date and time of the backup. **tar.gz** is a portable compressed file format.

   Only root (Linux) and members of the Administrators group (Windows) can access the backup files.

Run the Websense Backup Utility on each machine that includes Websense components. The tool identifies and saves any of the following files that it finds on the current machine:

Path	File name
\Program Files\Websense\bin *or* **/opt/Websense/bin**	authserver.ini BrokerService.cfg config.xml eimserver.ini LogServer.ini netcache.conf securewispproxy.ini transid.ini upf.conf websense.ini WebUI.ini wsauthserver.ini wscitrix.ini WSE.ini wsedir.ini wsradius.ini wsufpserver.ini
bin/i18n	i18n.ini
bin/postgres/data	postgresql.conf pg_hba.conf
BlockPages/*/Custom	All custom block page settings
tomcat/conf/Catalina/ Localhost	mng.xml
Windows\system32	isa_ignore.txt
Windows\system32\bin	ignore.txt
/etc/wsLib	wsSquid.ini

Store Websense backup files in a safe and secure location. These files should be part of your organization's regular backup procedures.

To revert to an earlier configuration:

1. Retrieve the backup files from their storage site.
2. Copy each backup file to the Websense machine on which it was created.

3. Run the Backup Utility in restore mode.

Important

Always use the Backup Utility to restore a Websense software configuration. Do not extract the files from the archive using other extraction utilities.

If the backup file is corrupted, you will not be able to restore your settings.

During the restore process, any error messages or warnings are displayed on the machine where the restore is being run.

Scheduling backups

Related topics:

- *Running immediate backups*, page 270
- *Maintaining the backup files*, page 271
- *Restoring your Websense data*, page 271
- *Discontinuing scheduled backups*, page 272
- *Command reference*, page 273

To schedule backups, open a command shell and navigate to the Websense bin directory (**C:\Program Files\Websense\bin** or **opt/Websense/bin**, by default). Enter the following command.

```
wsbackup -s -t "<m> <h> <day_of_month> <month>
   <day_of_week>" -d <directory>
```

Note that the time information uses **crontab** format, and the quotation marks and spaces are required.

In place of the variables shown in the example, provide the following information:

Variable	Information
<m>	0 - 59 Specify the precise minute to start the backup.
<h>	0 - 23 Specify the general hour of the day to start the backup.
<day_of_month>	1 - 31 Specify the date to perform the backup. If you schedule a backup for days 29 - 31, the utility uses the standard substitution procedure for the operating system in months that do not include that date.

Variable	Information
<month>	1 - 12 Specify the month to perform the backup.
<day_of_week>	0 - 6 Specify a day of the week. 0 represents Sunday.

Each field can take a number, an asterisk, or a list of parameters. Refer to any **crontab** reference for details.

Running immediate backups

Related topics:

- *Scheduling backups*, page 269
- *Maintaining the backup files*, page 271
- *Restoring your Websense data*, page 271
- *Discontinuing scheduled backups*, page 272
- *Command reference*, page 273

To launch an immediate backup, open a command shell and navigate to the Websense bin directory (**C:\Program Files\Websense\bin** or **opt/Websense/bin**, by default). Enter the following command.

```
wsbackup -b -d <directory>
```

Here, *directory* indicates the destination directory for the backup archive.

Warning

Do not store backup files in the Websense **bin** directory. This directory is deleted if you uninstall your Websense software.

When you initiate an immediate backup, any error messages and notifications are displayed on the console of the machine running the backup.

Maintaining the backup files

Related topics:

- *Scheduling backups*, page 269
- *Running immediate backups*, page 270
- *Restoring your Websense data*, page 271
- *Discontinuing scheduled backups*, page 272
- *Command reference*, page 273

When you perform a backup, a configuration file (**WebsenseBackup.cfg**) is created and stored with the backup archive. This configuration file specifies:

- How long to keep the backup archive in the backup directory
- The maximum amount of disk space that may be consumed by all backup files in the directory

Edit the **WebsenseBackup.cfg** file in any text editor to change either of these parameters:

Parameter	Value
KeepDays	Number of days archive files should remain in the backup directory. The default is 365.
KeepSize	Number of bytes allotted for backup files. The default is 10857600.

Any files older than the **KeepDays** value are deleted from the backup directory. If the amount of alloted disk space is exceeded, the oldest files are deleted from the backup directory to make room for newer files.

Restoring your Websense data

Related topics:

- *Scheduling backups*, page 269
- *Running immediate backups*, page 270
- *Maintaining the backup files*, page 271
- *Discontinuing scheduled backups*, page 272
- *Command reference*, page 273

When you restore Websense configuration data, make sure that you are restoring data for the components that exist on the current machine.

To initiate the restore process, open a command shell and navigate to the Websense bin directory (**C:\Program Files\Websense\bin** or **opt/Websense/bin**, by default). Enter the following command.

```
wsbackup -r -f archive_file.tar.gz
```

> **Important**
> The restore process may take several minutes. Do not stop the process while restoration is underway.

During the restore process, the Backup Utility stops all Websense services. If the utility is unable to stop the services, it sends a message asking the user to manually stop them. Services must be stopped in the order described in *Stopping and starting Websense services*, page 258.

The Backup Utility saves some files used for communication with third-party integration products. Because these files reside outside the Websense directory structure, you must restore them manually, by copying each file to the correct directory.

Files that must be restored manually include:

File name	Restore to
isa_ignore.txt	**Windows\system32**
ignore.txt	**Windows\system32\bin**
wsSquid.ini	**/etc/wsLib**

Discontinuing scheduled backups

Related topics:

- *Scheduling backups*, page 269
- *Running immediate backups*, page 270
- *Maintaining the backup files*, page 271
- *Restoring your Websense data*, page 271
- *Command reference*, page 273

To clear the backup schedule and stop running currently scheduled backups, open a command shell and navigate to the Websense bin directory (**C:\Program Files\Websense\bin** or **opt/Websense/bin**, by default). Enter the following command.:

```
wsbackup -u
```

Command reference

Related topics:

- *Scheduling backups*, page 269
- *Running immediate backups*, page 270
- *Maintaining the backup files*, page 271
- *Restoring your Websense data*, page 271
- *Discontinuing scheduled backups*, page 272

Only root (Linux) or a member of the Administrators group (Windows) can run the Backup Utility.

To see a complete list of Backup Utility command options at any time, enter:

```
wsbackup -h
```

or

```
wsbackup --help
```

The **wsbackup** command takes the following options:

- `-b` *or* `--backup`
- `-d` *`directory_path`* *or* `--dir` *`directory_path`*
- `-f` *`full_file_name`* *or* `--file` *`full_file_name`*
- `-h` *or* `--help` *or* `-?`
- `-r` *or* `--restore`
- `-s` *or* `--schedule`
- `-t` *or* `--time`
- `-u` *or* `--unschedule`
- `-v` *or* `--verbose [0...3]`

13 Reporting Administration

Related topics:

- *Planning your configuration*, page 276
- *Managing access to reporting tools*, page 276
- *Basic configuration*, page 277
- *Log Server Configuration utility*, page 281
- *Administering the Log Database*, page 294
- *Configuring investigative reports*, page 303
- *Self-reporting*, page 307

To use Websense presentation reports and investigative reports, you must install both Websense Manager and the reporting components on a Windows server. You also must configure Websense software to log Internet filtering activity.

Logging sends records to Websense Log Server, which processes them into a Log Database that must be installed on a supported database engine: Microsoft SQL Server Desktop Engine (commonly referred to as MSDE in this document) or Microsoft SQL Server Enterprise or Standard Editions (both commonly referred to as Microsoft SQL Server). See the Websense *Installation Guide* for more information on installing these reporting components.

When you generate a report, Websense Manager displays information from the Log Database according to the filter you define for the report.

Organizations that install Websense Manager on a Linux server, or that prefer to use Linux for their reporting needs can install the separate Websense Explorer for Linux product to generate reports. This product operates independently of Websense Manager. Refer to the Websense *Explorer for Linux Administrator's Guide* for instructions on installing and using that program.

Planning your configuration

Depending on the volume of Internet traffic in your network, the Log Database can become very large. To help determine an effective logging and reporting strategy for your organization, consider these questions:

- When is the network traffic busiest?

  Consider scheduling resource intensive database jobs and reporting jobs at times when the traffic volume is lower. This improves logging and reporting performance during peak periods. See *Configuring Internet browse time options*, page 297, and *Configuring Log Database maintenance options*, page 298.

- How long should log data be kept to support historical reporting?

  Consider automatically deleting partitions after they reach this age. This reduces the amount of disk space required for the Log Database. See *Configuring Log Database maintenance options*, page 298.

- How much detail is really needed?

  Consider which logging options to activate: logging full URLs and hits increase the Log Database size. To decrease Log Database size, consider:

  - disabling full URL logging (see *Configuring full URL logging*, page 296)
  - logging visits instead of hits (see *Configuring log cache files*, page 286)
  - enabling consolidation (see *Configuring consolidation options*, page 287)
  - enabling selective category logging (see *Configuring Filtering Service for logging*, page 280)

Successful reporting implementations are deployed on hardware that matches or exceeds the requirements for expected load and for historical data retention.

Managing access to reporting tools

When the Websense Manager and reporting components are installed on Windows servers, reporting options appear within Websense Manager and the Log Server Configuration utility.

When you install the reporting components, Log Server is connected to a specific Policy Server. You must select that Policy Server during Websense Manager logon to access reporting features. If you log on to a different Policy Server, you cannot access Presentation Reports or Investigative Reports on the Main tab, or the entire Reporting section of the Settings tab.

In organizations that use only the WebsenseAdministrator logon account, everyone who uses Websense Manager has access to all reporting options within Websense Manager, including presentation reports, investigative reports, and settings for the reporting tools.

In organizations that use delegated administration, access to reporting tools within Websense Manager is controlled by WebsenseAdministrator and members of the Super Administrator role. When creating a role, the Super Administrator designates whether that role has access to specific reporting options.

See *Editing roles*, page 233, for information on configuring access to reporting tools.

The Log Server Configuration utility is accessed from the Windows Start menu. Only those with access to the installation machine can open this utility and modify Log Server settings. See *Log Server Configuration utility*, page 281.

If your organization has installed Websense Manager on a Linux server, or chooses the Websense Explorer for Linux reporting program instead of the reporting components that run on Windows, reporting options do not appear in Websense Manager. No Internet filtering charts can be shown on the Today and History pages. See the *Explorer for Linux Administrator's Guide* for information on installing that program and using it to run reports.

Basic configuration

Related topics:

- *Configuring Filtering Service for logging*, page 280
- *Assigning categories to risk classes*, page 278
- *Configuring reporting preferences*, page 279
- *Log Server Configuration utility*, page 281
- *Administering the Log Database*, page 294

You can use a variety of configuration options to customize reporting for your environment.

The Websense Master Database organizes categories into **risk classes**. Risk classes suggest possible types or levels of vulnerability posed by sites in those categories. Use the General > Risk Classes page, accessed from the Settings tab, to customize the risk classes for your organization. See *Assigning categories to risk classes*, page 278.

Use the Reporting > Preferences page, accessed from the Settings tab, to configure the email server used to distribute reports, and to activate the self-reporting feature. See *Configuring reporting preferences*, page 279.

Logging is the process of storing information about Websense filtering activities in a Log Database so that you can generate reports.

Use the General > Logging page, accessed from the Settings tab, to enable logging, select the categories to be logged, and determine what user information is logged. See *Configuring Filtering Service for logging*, page 280, for more information.

Use the Log Server Configuration utility to manage the way the log records are processed and connections to the Log Database. See *Log Server Configuration utility*, page 281, for more information.

Use the Reporting > Log Database page, accessed from the Settings tab, to administer the Log Database, including Internet browse time controls, database partition options, and error logs. See *Administering the Log Database*, page 294, for more information.

Assigning categories to risk classes

Related topics:

- *Risk classes*, page 37
- *Block Pages*, page 77
- *Using Reports to Evaluate Filtering Policies*, page 85

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

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

Unconditional Super Administrators can view or change which categories comprise each risk class on the **Settings > Risk Classes** page. For example, some businesses may consider user-posted video sites to fall under the risk classes of legal liability, network bandwidth loss, and productivity loss. However, if your company does market research on a certain demographic, you might consider these part of the Business Usage risk class.

Note

The security block page appears for blocked sites in the Security Risk class default categories. Changes to the categories in the Security Risk class affect reporting, but do not affect block pages. See *Block Pages*, page 77.

Risk class information in Websense reports reflects the assignments you make on this page.

1. Select an entry in the **Risk Classes** list.
2. Review the **Categories** list to see which categories are currently included in that risk class.

   A check mark shows that the category is currently assigned to the selected risk class. The blue W icon indicates categories that are included in the risk class by default.
3. Mark or clear entries in the category tree to include or exclude a category from the selected risk class. Categories can belong to more than one risk class.

Other choices include:

Option	Description
Select All	Selects all categories in the tree.
Clear All	Deselects all categories in the tree.
Restore Defaults	Resets the category choices for the selected risk class to those provided by the Websense software. A blue W icon indicates a default category.

4. Repeat this process for each risk class.
5. Click **OK** to cache your changes. Changes are not implemented until you click **Save All**.

Configuring reporting preferences

Related topics:

- *Self-reporting*, page 307
- *Scheduling presentation reports*, page 99
- *Scheduling investigative reports*, page 125

When you schedule either presentation or investigative reports to be run at a later time or on a repeating cycle, the reports are distributed via email to specified recipients. Use the **Reporting > Preferences** page, accessed from the Settings tab, to supply key information for these email messages.

This page is also used to enable self-reporting, where individuals can generate investigative reports on their own Internet activity.

1. Enter the **Email address** to appear in the from field when scheduled reports are distributed via email.
2. Enter the **SMTP server IP or name** for the email server used to distribute scheduled reports via email.
3. Mark the **Allow self-reporting** check box to permit end users in your organization to access Websense Manager and run investigative reports on their personal Internet activity. See *Self-reporting*, page 307.
4. Click **Save Now** to implement your changes.

Configuring Filtering Service for logging

Related topics:

- *Introducing the Log Database*, page 292
- *Log Server Configuration utility*, page 281

Use the **General > Logging** page on the Settings tab to provide the IP address and port for sending the log records to Log Server. This page also lets you select what user information and URL categories Websense Filtering Service should send to Log Server and make available for reports and category usage alerts (see *Configuring category usage alerts*, page 263).

In an environment with multiple Policy Servers, configure the General > Logging page separately for each one. All Filtering Services associated with the active Policy Server send their log records to the Log Server identified on this page.

Keep the following facts in mind when working with multiple Policy Servers:

- If the Log Server IP address and port are blank for any Policy Server, the Filtering Services associated with that Policy Server cannot log any traffic for reporting or alerts.
- Each Filtering Service logs traffic according to the settings for the Policy Server it is connected to. If you change the user information or category logging selections for different Policy Servers, reports generated for users associated with different Policy Servers may appear inconsistent.

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. To log identifying information for machines accessing the Internet, mark **Log IP addresses**.
2. 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**.

3. Enter the IP address or machine name where Log Server is installed in the **Log Server IP address or name** field.

   Important

   If Log Server is installed on a separate machine from Policy Server, this entry may default to localhost. If this happens, enter the correct IP address of the Log Server machine to enable display of charts on the Today and History pages, as well as other reporting features.

4. Enter the **Port** number for sending log records to Log Server.
5. Click **Check Status** to determine whether Websense 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. Click the **Selective Category Logging** button to open the area for indicating which URL categories to log.

   Selections you make here apply to all category filters in all active policies.

   Notes

   If you disable logging for categories that have usage alerts set up (see *Configuring category usage alerts*, page 263), no usage alerts can be sent.

   Reports cannot include information on categories that are not logged.

   a. Expand or collapse the parent categories as needed to see the categories of interest.
   b. Select each category to be logged by marking its check box.

      You must select or deselect each category separately. Selecting a parent category does not automatically select its subcategories. Use **Select All** and **Clear All** to assist with selections.
7. Click **OK** to cache your changes. Changes are not implemented until you click **Save All**.

Log Server Configuration utility

Related topics:

- *Managing access to reporting tools*, page 276
- *Basic configuration*, page 277
- *Stopping and starting Log Server*, page 292

During installation, you configure certain aspects of Log Server operation, including how Log Server interacts with Websense filtering components.

The Log Server Configuration utility lets you change these settings when needed, and configure other details about Log Server operation. This utility is installed on the same machine as Log Server.

1. From the Windows Start menu, select **Programs > Websense > Utilities > Log Server Configuration**.

   The Log Server Configuration utility opens.
2. Select a tab to display its options and make any changes. For detailed instructions, see:
   - *Configuring Log Server connections*, page 282
   - *Configuring Log Server database options*, page 283
   - *Configuring log cache files*, page 286
   - *Configuring consolidation options*, page 287
   - *Configuring WebCatcher*, page 289
3. Click **Apply** to save the changes.
4. Use the **Connection** tab to stop and restart Log Server for the changes to take effect.

IMPORTANT

After making changes to any Log Server Configuration tab, click **Apply**. Then, you **must** stop and restart Log Server for the changes to take effect. To avoid restarting Log Server multiple times, make all Log Server Configuration changes before restarting Log Server.

Configuring Log Server connections

Related topics:

- *Log Server Configuration utility*, page 281
- *Configuring Log Server database options*, page 283
- *Configuring log cache files*, page 286
- *Configuring consolidation options*, page 287
- *Configuring WebCatcher*, page 289
- *Stopping and starting Log Server*, page 292

The **Connection** tab of the Log Server Configuration utility contains options for creating and maintaining a connection between Log Server and Websense filtering components.

1. Accept the default **Log Server input port** (55805) or enter another available port.

This is the port over which the Log Server communicates with Filtering Service. The port entered here must match the port entered on the General > Logging page (Settings tab) in Websense Manager.

2. Enter a number of hours as the **User/Group update interval** to specify how often Log Server should contact the directory service for updates.

   Log Server contacts the directory service to obtain updated information, such as full user name and group assignments, about the users with records in the Log Database.

   Activity for a user whose group has changed continues to be reported with the previous group 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 than the default of 12 hours.

3. Click **Apply** to save any changes.
4. Use the button in the Service Status area to **Start** or **Stop** Log Server. The label of the button changes to reflect the action that will occur when you click it.

   Note

   No Internet access activity can be logged when Log Server is stopped.

Changes made in the Log Server Configuration utility do not take effect until you stop and restart Log Server.

Configuring Log Server database options

Related topics:

- *Log Server Configuration utility*, page 281
- *Configuring Log Server connections*, page 282
- *Setting up the database connection*, page 285
- *Configuring log cache files*, page 286
- *Configuring consolidation options*, page 287
- *Configuring WebCatcher*, page 289
- *Stopping and starting Log Server*, page 292

Open the **Database** tab of the Log Server Configuration utility to configure how Log Server works with the Log Database.

1. Choose a **Log Insertion Method** from the following options.
   - Open Database Connectivity (ODBC): Inserts records into the database individually, using a database driver to manage data between Log Server and Log Database.

- Bulk Copy Program (BCP) *(recommended)*: Inserts records into the Log Database in groups called batches. This option is recommended because it offers better efficiency than ODBC insertion.

  Note
  The BCP option is available only if you install the SQL Server Client Tools on the Log Server machine.

2. Click the **Connection** button to select the Log Database for storing new Internet access information from Websense. See *Setting up the database connection*, page 285.

   ODBC Data Source Name (DSN) and **ODBC Login Name** display the settings established for the database connection.

3. If you chose BCP as the log insertion method in step 1, set the following options. If you chose ODBC as the log insertion method, skip this step.

Option	Description
BCP file path location	Directory path for storing BCP files. This must be a path where Log Server has read and write access. This option is available only if Log Server is installed on the Log Database machine, or if the SQL Server Client Tools are installed on the Log Server machine.
BCP 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.
BCP 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.

4. Set the **Maximum connections allowed** to indicate how many internal connections can be made between Log Server and the database engine. The options available depend on the database engine being used.

   - **MSDE**: This value is preset to 4, and cannot be changed.

- **SQL Server**: Set to a number from 4 to 50, as appropriate for your SQL Server license. The minimum number of connections depends on the selected log insertion method.

  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.

5. Check or uncheck **Enhanced logging** to enable or disable this option, which controls how Log Server resumes logging after it has been stopped.

   When this option is deselected (the default), Log Server begins processing at the beginning of the oldest log cache file after a stop. This could result in some duplicate entries in the Log Database, but speeds Log Server processing.

   When this option is checked, Log Server tracks its location in the active log cache file. After a restart, Log Server resumes processing where it stopped. Enhanced logging can slow Log Server processing.

6. Click **Apply** to save any changes, then stop and restart Log Server (see *Stopping and starting Log Server*, page 292).

Setting up the database connection

Related topics:

- *Configuring Log Server connections*, page 282
- *Configuring Log Server database options*, page 283

The **Connection** button in the Database tab of Log Server Configuration utility lets you select the Log Database for storing incoming Internet access information from Websense. This is configured automatically during installation, but can be changed whenever you need to change the database for logging. (The database must already exist to establish a connection.)

1. In the Data Source dialog box, select the **Machine Data Source** tab.
2. Select the ODBC connection for the database into which new information will be logged.
3. Click **OK** to display the SQL Server Logon dialog box.
4. If the **Use Trusted Connection** option is available, make sure it is set properly for your environment.

   MSDE users: Uncheck the Trusted Connection option.

SQL Server users: Contact your Database Administrator for assistance.

> **Note**
> If you use a trusted connection for communications with SQL Server, you may need to configure several Websense services with the trusted user name and password. See the Websense *Installation Guide* for details.

5. Enter the **Logon ID** and **Password** established when the database was created. Usually this is the same logon ID and password entered during Log Server installation and database creation.
6. Stop and restart Log Server via the **Connection** tab after making this and any other changes in the Log Server Configuration utility.

Configuring log cache files

Related topics:

- *Log Server Configuration utility*, page 281
- *Configuring Log Server connections*, page 282
- *Configuring Log Server database options*, page 283
- *Configuring consolidation options*, page 287
- *Configuring WebCatcher*, page 289
- *Stopping and starting Log Server*, page 292

The **Settings** tab of the Log Server Configuration utility lets you manage the log cache file creation options, and specify whether Log Server tracks the individual files that make up each Web site requested, or just the Web site.

1. Enter the path for storing log cache files in the **Log file path location** field. The default path is **<installation directory>\bin\Cache**. (The default installation directory is C:\Program Files\Websense\).
2. For **Cache file creation rate**, indicate the maximum number of minutes Log Server should spend sending Internet access information to a log cache file (**log*n*.tmp**) before closing it and creating a new file.

   This setting works in combination with the size setting: Log Server creates a new log cache file as soon as either limit is reached.
3. For **Cache file creation size**, specify how large a log cache file should be before Log Server closes it and creates a new one.

   This setting works in combination with the creation rate setting: Log Server creates a new log cache file as soon as either limit is reached.

4. Check **Enable visits** to create a log record for each Web site visited.

   Note
   Managing Log Database size is an important concern in high-volume networks. Enabling visits logging is one way to control database size and growth.

   When this option is deselected, a separate log record is created for each HTTP request generated to display the different page elements, such as graphics and advertisements. Also known as logging hits, this option 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.

   If you have installed Websense Web Security Gateway, real-time scanning activity is always reported in hits on the reports that are specific to real-time scanning, even when visits logging is enabled. In this situation, the numbers shown on Web filtering reports that include traffic blocked by real-time scanning will be lower than the numbers shown on real-time scanning reports.

   Note
   It is best to create a new database partition prior to changing the method of logging between visits and hits. See the Reporting > Log Database page (Settings tab) in Websense Manager to create a new database partition.

5. Click **Apply** to save any changes, then stop and restart Log Server (see *Stopping and starting Log Server*, page 292).

Configuring consolidation options

Related topics:

- *Log Server Configuration utility*, page 281
- *Configuring Log Server connections*, page 282
- *Configuring Log Server database options*, page 283
- *Configuring log cache files*, page 286
- *Configuring WebCatcher*, page 289
- *Stopping and starting Log Server*, page 292

Use the **Consolidation** tab of the Log Server Configuration utility to enable consolidation and set consolidation preferences.

> **Note**
> Managing Log Database size is an important concern in high-volume networks. Enabling consolidation is one way to control database size and growth.

Consolidation decreases the size of your Log Database by combining Internet requests that share the following elements:

- Domain name (for example: www.websense.com)
- Category
- Keyword
- Action (for example: Category Blocked)
- User/workstation

Reports run faster when the Log Database is smaller. However, consolidating log data may decrease the accuracy of some detail reports, as separate records for the same domain name may be lost.

> **Important**
> Enabling consolidation may skew the accuracy of some report data, such as Internet Browse Time calculations.

1. Check **Consolidate Log Records** to enable consolidation, which combines multiple similar Internet requests into a single log record.

   When this option is deselected, the default, the Log Database retains full hits or visits detail for each Internet request (depending on your selection on the Settings tab, see *Configuring log cache files*, page 286). This provides greater reporting detail, but a larger Log Database.

   Selecting this option creates a smaller Log Database with less reporting detail.

   > **Important**
   > To assure consistent reports, consider creating a new database partition whenever you enable or disable consolidation. Also, be sure to generate reports from partitions with the same consolidation setting.

   If you have installed Websense Web Security Gateway, real-time scanning activity is always reported as separate hits on the reports that are specific to real-time scanning, even when consolidation is enabled. In this situation, the numbers shown on Web filtering reports that include traffic blocked by real-time scanning will be lower than the numbers shown on real-time scanning reports.

2. For the **Consolidation time interval**, specify the maximum time between the first and last records to be combined.

This represents the greatest 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 enabled the full URL option on the Reporting > Log Database page (Settings tab) in Websense Manager, the consolidated log record will 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 active, consolidation would create a single log entry under the URL www.domain.com/shoeshopping.

3. Click **Apply** to save any changes, then stop and restart Log Server (see *Stopping and starting Log Server*, page 292).

Configuring WebCatcher

Related topics:

- *Log Server Configuration utility*, page 281
- *Configuring Log Server connections*, page 282
- *Configuring Log Server database options*, page 283
- *Configuring log cache files*, page 286
- *Configuring consolidation options*, page 287
- *Configuring WebCatcher*, page 289
- *WebCatcher Authentication*, page 291
- *Stopping and starting Log Server*, page 292

WebCatcher is an optional feature that collects unrecognized URLs and security URLs, and submits them to Websense, Inc., where they are analyzed for potential security and liability risks, and for categorization. (Full URL logging is not required for WebCatcher processing.) Websense, Inc., reviews information and updates the Master Database with newly categorized URLs, resulting in improved filtering.

Choose the types of URLs to send, and set the file size and processing time on the **WebCatcher** tab of the Log Server Configuration utility.

> **Note**
> In an environment with multiple Log Servers, WebCatcher is enabled for only one Log Server. Once it has been enabled, this tab is unavailable when running the Log Server Configuration tool for other Log Server instances.

The information sent to Websense, Inc., contains only URLs and does not include user information.

The following example illustrates the information that would be sent if you activated WebCatcher. The IP address in this example reflects the address of the machine hosting the URL, not the requestor's IP address.

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

WebCatcher data is sent to Websense, Inc., via HTTP Post. You may need to create roles or make other changes on your proxy server or firewall to permit the outgoing HTTP traffic. Refer to the proxy server or firewall documentation for instructions.

1. Select one of the following options:
   - **Yes, send only specified URLs to Websense** activates WebCatcher processing. You must indicate which URLs to send. Continue to step 2.
   - **No, do not send information to Websense** deactivates WebCatcher processing. No further entries are required if you choose this option.
2. Check **Send uncategorized URLs** to send a list of all uncategorized URLs found in your Log Database.

   Websense, Inc., analyzes uncategorized URLs it receives, and adds them to the Master Database categories, as appropriate. This improves filtering accuracy for all organizations.

   > **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.

3. Check **Send security URLs** to send a list of security URLs found in your Log Database.

   Security URLs received are analyzed by Websense, Inc., to determine the activity of sites in the Keyloggers, Malicious Web Sites, Phishing and Other Fraud, and Spyware categories.
4. Under **Select the country which best reflects your location**, select the country where the majority of activity is being logged.
5. Check the **Save a copy of the data transmitted to Websense** option to save a copy of the data that is being sent to Websense, Inc.

When this option is enabled, WebCatcher saves the data as unencrypted XML files in the Websense\Reporter directory. These files are date and time stamped.

6. Under **Maximum upload file size**, indicate how large the file can grow (from 4096 KB to 8192 KB) before sending it to Websense.

   Ensure that your system can post a file of this size via HTTP Post.

7. For **Minimum Daily start time**, set the start time for WebCatcher to send the file if the size threshold has not been reached that day.

   This assures that the information is submitted and cleared from your system at least once each day.

8. Click the **Authentication** button if the Log Server machine must authenticate to access the Internet.

   See *WebCatcher Authentication*, page 291, for information about the **Authentication** dialog box that appears.

9. Click **Apply** to save any changes, then stop and restart Log Server (see *Stopping and starting Log Server*, page 292).

WebCatcher Authentication

Related topics:

- *Log Server Configuration utility*, page 281
- *Configuring WebCatcher*, page 289
- *Stopping and starting Log Server*, page 292

The Authentication dialog box appears after you click **Authentication** on the WebCatcher tab.

1. Check the **Use a proxy server** option if the Log Server machine accesses the Internet through a proxy server, and then provide the requested information.

Field	Description
Proxy server name	Enter the machine name or IP address of the proxy server through which Log Server accesses the Internet.
Proxy server port	Enter the port number over which the proxy server communicates.

2. Check the **Use Basic Authentication** option if the Log Server machine must authenticate to access the Internet, and then enter the user name and password for authentication.
3. Click **OK** to save any changes and return to the WebCatcher tab.

Stopping and starting Log Server

Related topics:

- *Log Server Configuration utility*, page 281
- *Configuring Log Server connections*, page 282

Log Server receives information from Filtering Service and saves it in the Log Database for use when generating reports. It runs as a Windows service, typically started during installation, and starts any time you restart the machine.

Changes you make in the Log Server Configuration utility take effect only after you stop and restart Log Server. This can be done easily through the Connection tab in the Log Server Configuration utility.

1. From the Windows Start menu, select **Programs > Websense > Utilities > Log Server Configuration**.
2. In the **Connections** tab, click **Stop**.
3. Wait several seconds, and then click **Start** to restart the Log Server service.
4. Click **OK** to close the Log Server Configuration utility.

Note

Websense cannot log Internet access that occurs while the Log Server is stopped.

Introducing the Log Database

Related topics:

- *Database jobs*, page 293
- *Administering the Log Database*, page 294

The Log Database stores the records of Internet activity and the associated Websense filtering actions. Installation creates the Log Database with a catalog database and one database partition.

The **catalog database** provides a single connection point for the various Websense components that need to access the Log Database: Status pages, Log Server, presentation reports, and investigative reports. It contains supporting information for the database partitions, including the list of category names, risk class definitions, the mapping of users to groups, database jobs, and so forth. The catalog database also maintains a list of all the available database partitions.

Database partitions store the individual log records of Internet activity. For MSDE users, new partitions are created based on size rollover rules established by Websense software. Microsoft SQL Server users can configure the Log Database to start a new partition based on partition size or a date interval (see *Configuring rollover options*, page 295, for more information).

> **Note**
> Date-based partitions are available only when Websense software uses Microsoft SQL Server as the database engine.

When partitions are based on size, all incoming log records are inserted into the most recent active partition that satisfies the size rule. When the partition reaches the designated maximum size, a new partition is created for inserting new log records.

When the partitions are based on date, new partitions are created according to the established cycle. For example, if the rollover option is monthly, a new partition is created as soon as any records are received for the new month. Incoming log records are inserted into the appropriate partition based on date.

Database partitions provide flexibility and performance advantages. For example, you can generate reports from a single partition to limit the scope of data that must be analyzed to locate the requested information.

Database jobs

The following database jobs are installed along with the Log Database. The SQL Server Agent must be running on the machine running the database engine (MSDE or Microsoft SQL Server).

- The Extract, Transform, and Load (ETL) job runs continuously, receiving data from Log Server, processing it, and then inserting it into the partition database. The ETL job must be running to process log records into the Log Database.
- The database maintenance job performs database maintenance tasks and preserves optimal performance. This job runs nightly, by default.
- The Internet browse time (IBT) job analyzes the data received and calculates browse time for each client. The IBT database job is resource intensive, affecting most database resources. This job runs nightly, by default.

Certain aspects of these database jobs can be configured on the Settings > Log Database page. See *Log Database administration settings*, page 294, 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.

Administering the Log Database

Related topics:

- *Log Database administration settings*, page 294
- *Configuring rollover options*, page 295
- *Configuring Internet browse time options*, page 297
- *Configuring full URL logging*, page 296
- *Configuring Log Database maintenance options*, page 298
- *Configuring Log Database partition creation*, page 300
- *Configuring available partitions*, page 301
- *Viewing error logs*, page 302

Administering the Log Database involves controlling many aspects of database operations, including:

- What operations the database jobs perform, and when they run.
- The conditions for creating new database partitions.
- Which partitions are available for reporting.

These and other options give the person who administers the Log Database significant control. See *Log Database administration settings*, page 294.

The Super Administrator designates who can administer the Log Database when creating roles. See *Editing roles*, page 233.

Note

It is advisable to limit the number of administrators who have the permission to change Log Database settings.

Log Database administration settings

Related topics:

- *Administering the Log Database*, page 294

The **Reporting > Log Database** page, accessed from the Settings tab, allows you to manage various aspects of Log Database operations. The options are grouped into logical sections that are described separately.

You must click the Save Now button within a section to activate changes in that section. Clicking **Save Now** records the changes in that section immediately. (It is not necessary to also click Save All.)

The top of the page shows the name of the active Log Database and a **Refresh** link. This Refresh link redisplays the information currently in the Log Database page. Any changes that have not been applied with the appropriate Save Now button are lost.

For detailed instructions on using each section, click the appropriate link, below.

- Database Rollover Options: *Configuring rollover options*, page 295.
- Full URL Logging: *Configuring full URL logging*, page 296.
- Internet Browse Time Configuration: *Configuring Internet browse time options*, page 297.
- Maintenance Configuration: *Configuring Log Database maintenance options*, page 298.
- Database Partition Creation: *Configuring Log Database partition creation*, page 300.
- Available Partitions: *Configuring available partitions*, page 301.
- Error Log Activity: *Viewing error logs*, page 302.

Configuring rollover options

Related topics:

- *Log Database administration settings*, page 294
- *Configuring Internet browse time options*, page 297
- *Configuring full URL logging*, page 296
- *Configuring Log Database maintenance options*, page 298
- *Configuring Log Database partition creation*, page 300
- *Configuring available partitions*, page 301
- *Viewing error logs*, page 302

Use the **Database Rollover Options** section of the Reporting > Log Database page (Settings tab) to specify when you want the Log Database to create a new database partition (roll over).

1. Use the **Rollover every** options to indicate whether database partitions should roll over based on size (MB) or date (weeks or months), depending on the database engine being used.

   MSDE users must use the size rollover option. Microsoft SQL Server users can choose either size or date.

   - For date-based rollovers, select either **weeks** or **months** as the unit of measure, and specify how many full calendar weeks or months to keep in a database partition before a new one is created.
   - For size-based rollovers, select **MB** and specify the number of megabytes the database must reach for the rollover to begin.

     Microsoft SQL Server users may set a size up to 204800 MB.

MSDE users must set a size between 100 MB and 1536 MB.

Note

If the rollover begins during a busy part of the day, performance may slow during the rollover process.

To avoid this possibility, some environments choose to set the automatic rollover to a long time period or large maximum size. Then, they perform regular manual rollovers to prevent the automatic rollover from occurring. See *Configuring Log Database partition creation*, page 300, for information on manual rollovers.

Keep in mind that extremely large individual partitions are not recommended. Reporting performance can slow if data is not divided into multiple, smaller partitions.

When a new partition database is created based reporting is automatically enabled for the partition (see *Configuring available partitions*, page 301).

2. Click **Save Now** to activate changes to the database rollover options.

Configuring full URL logging

Related topics:

- *Log Database administration settings*, page 294
- *Configuring rollover options*, page 295
- *Configuring Internet browse time options*, page 297
- *Configuring Log Database maintenance options*, page 298
- *Configuring Log Database partition creation*, page 300
- *Configuring available partitions*, page 301
- *Viewing error logs*, page 302

The **Full URL Logging** section of the Reporting > Log Database page (Settings tab) lets you decide what portion of the URL is logged for each Internet request.

Note

Managing Log Database size is an important concern in high-volume networks. Disabling the Full URL Logging option is one way to control database size and growth.

1. Mark **Record full URL of each site requested** to log the entire URL, including the domain (www.domain.com) and the path to the particular page (/products/productA.html).

   Important

   Enable full URL logging if you plan to generate reports of real-time scanning activity (see *Reporting on real-time scanning activity*, page 139). Otherwise, reports can display only the domain (www.domain.com) of the site categorized, even though individual pages within the site may fall into different categories, or contain different threats.

   If this option is not checked, only domain names are logged. This choice results in a smaller database, but provides less detail.

   Logging full URLs produces a larger Log Database, but provides greater detail.

   If you activate full URL logging when consolidation is active, the consolidated record contains the full URL from the first record in the consolidation group. See *Configuring consolidation options*, page 287, for more information.

2. Click **Save Now** to activate changes to the full URL logging options.

Configuring Internet browse time options

Related topics:

- *Log Database administration settings*, page 294
- *Configuring rollover options*, page 295
- *Configuring full URL logging*, page 296
- *Configuring Log Database maintenance options*, page 298
- *Configuring Log Database partition creation*, page 300
- *Configuring available partitions*, page 301
- *Viewing error logs*, page 302

Internet browse time (IBT) reports give a view into the amount of time users spend on the Internet. A nightly database job calculates browse time for each client based on the new log records received that day. Set browse time options in the **Internet Browse Time Configuration** section of the Settings > Log Database page.

1. Choose a **Job start time** for the IBT database job.

   The time and system resources required by this job vary depending on the volume of data logged each day. It is best to run this job at a different time than the nightly maintenance job (see *Configuring Log Database maintenance options*, page 298), and to select a slow time on the network to minimize any impact on generating reports.

The IBT database job is resource intensive, affecting most database resources. If you enable this job, set the start time so that it does not interfere with the database system's ability to process scheduled reports and other important operations. Also, monitor the job to determine whether more robust hardware is needed to accommodate all processing needs.

2. For **Read time threshold**, set an average number of minutes for reading a specific Web site.

   The read time threshold defines browse sessions for the purpose of Internet browse time reports. Opening a browser generates HTTP traffic. This represents the beginning of a browse session. The session is open as long as HTTP traffic is continually generated within the time set here. The browse session is considered closed once this amount of time passes with no HTTP traffic. A new browse session begins as soon as HTTP traffic is generated again.

   Note

   It is best to change the Read Time Threshold as seldom as possible, and to start a new database partition whenever you do make a change.

   To avoid inconsistent data on the reports, generate IBT reports from database partitions that use the same Read Time Threshold value.

   Be aware that some Web sites use an automatic refresh technique to update information frequently. One example is a news site that rotates a display of the latest news stories. This refresh generates new HTTP traffic. Therefore, when this kind of site is left open, new log records are generated each time the site refreshes. There is no gap in HTTP traffic, so the browser session is not closed.

3. Set a **Last read time** value to account for time spent reading the last Web site before the end of a browse session.

   When the time gap of HTTP traffic is longer than the read time threshold, the session is ended and the value of the Last Read Time is added to the session time.

4. Click **Save Now** to activate changes to the Internet browse time configuration.

Configuring Log Database maintenance options

Related topics:

- *Log Database administration settings*, page 294
- *Configuring rollover options*, page 295
- *Configuring Internet browse time options*, page 297
- *Configuring full URL logging*, page 296
- *Configuring Log Database partition creation*, page 300
- *Configuring available partitions*, page 301
- *Viewing error logs*, page 302

Use the **Maintenance Configuration** section of the Reporting > Log Database page (Settings tab) to control certain aspects of database processing, such as the time for running the database maintenance job, some of the tasks it performs, and deleting of database partitions and error logs.

1. For **Maintenance start time**, select the time of day for running the database maintenance job.

   The time and system resources required by this job vary depending on the tasks you select in this area. To minimize any impact on other activities and systems, it is best to run this job during a slow time on the network, different from the time set for the IBT job (see *Configuring Internet browse time options*, page 297).

2. Check **Automatically delete partitions**, and then specify the number of days (from 2 to 365) after which partitions should be deleted.

   

   Warning

   After a partition has been deleted, the data cannot be recovered. See *Configuring available partitions*, page 301, for an alternative way to delete partitions.

3. Check **Enable automatic reindexing**, and then select a day of the week to have this processing performed automatically each week.

   Reindexing the database is important to maintain database integrity and to optimize reporting speed.

   

   Important

   It is best to perform this processing during a quiet time on the network. Reindexing database partitions is resource intensive and time-consuming. Reports should not be run during the process.

4. Check **Number of days before deleting failed batches** and then enter a number of days (from 0 to 90) after which to delete any failed batches.

   If this option is not checked, failed batches are retained indefinitely for future processing.

   If there is insufficient disk space or inadequate database permissions to insert log records into the database, the records are marked as a **failed batch**. Typically, these batches are successfully reprocessed and inserted into the database during the nightly database maintenance job.

   However, this reprocessing cannot be successful if the disk space or permission problem has not been resolved. Additionally, if **Process the unprocessed batches** is not selected, failed batches are never reprocessed. They are deleted after the time specified here.

5. Check **Process the unprocessed batches** to have the nightly database maintenance job reprocess any failed batches.

   If this option is unchecked, failed batches are never reprocessed. They are deleted after the time specified above, if any.

6. Check **Number of days before deleting error log**, and then enter a number of days (0 to 90) after which to delete database error records from the catalog database.

   If this option is not checked, error logs are retained indefinitely.

7. Click **Save Now** to activate changes to the maintenance configuration options.

Configuring Log Database partition creation

Related topics:

- *Log Database administration settings*, page 294
- *Configuring rollover options*, page 295
- *Configuring Internet browse time options*, page 297
- *Configuring full URL logging*, page 296
- *Configuring Log Database maintenance options*, page 298
- *Configuring available partitions*, page 301
- *Viewing error logs*, page 302

Use the **Database Partition Creation** section of the Reporting > Log Database page (Settings tab) to define characteristics for new database partitions, such as location and size options. This area also lets you create a new partition right away, rather than waiting for the planned rollover (see *Configuring rollover options*, page 295).

1. Enter the **File Path** for creating both the **Data** and **Log** files for new database partitions.
2. Under **Init Size** set the initial file size (from 100 to 204800 MB) for both the **Data** and **Log** files for new database partitions.

   Microsoft SQL Server users: Acceptable range is 100 - 204800

   MSDE users: Acceptable range is 100 - 1500

   

   Note

   Best practice recommends calculating the average partition size over a period of time. Then, update the initial size to that value. This approach minimizes the number of times the partition must be expanded, and frees resources to process data into the partitions.

3. Under **Growth** set the increment by which to increase the size, in megabytes (MB), of a partition's **Data** and **Log** files when additional space is required.

   Microsoft SQL Server users: Acceptable range is 1 - 999999

   MSDE users: Acceptable range is 1 - 450

4. Click **Save Now** to implement the path, size, and growth changes entered.

   Database partitions created after these changes use the new settings

5. Click **Create Now** to create a new partition the next time the ETL job runs (see *Database jobs*, page 293), regardless of the automatic rollover settings. This process usually takes a few minutes.

   To have the new partition use the changes made in this section, be sure to click **Save Now** before you click **Create Now**.

   Click the Refresh link in the content pane periodically. The Available Partitions area will show the new partition when the creation process is complete.

Configuring available partitions

Related topics:

- *Log Database administration settings*, page 294
- *Configuring rollover options*, page 295
- *Configuring Internet browse time options*, page 297
- *Configuring full URL logging*, page 296
- *Configuring Log Database maintenance options*, page 298
- *Configuring Log Database partition creation*, page 300
- *Viewing error logs*, page 302

The **Available Partitions** section of the Reporting > Log Database page (Settings tab) lists all the database partitions available for reporting. The list shows the dates covered, as well as the size and name of each partition.

Use this list to control what database partitions are included in reports, and to select individual partitions to be deleted.

1. Check **Enable** beside each partition to be included in reports.

   Use the **Select all** and **Select none** options above the list, as appropriate.

   You must enable at least one partition for reporting. Use the **Select none** option to disable all partitions at one time so that you can enable just a few.

   Use these options to manage how much data must be analyzed when generating reports and speed report processing. For example, if you plan to generate a series of reports for June, deselect all partitions except those with dates in June.

   

   Important

   This selection affects scheduled reports as well as reports run interactively. To avoid generating reports with no data, make sure the relevant partitions are enabled when reports are scheduled to run.

2. Click the **Delete** option beside a partition name if that partition is no longer needed. The partition is actually deleted the next time the nightly database maintenance job runs.

   Warning
   Use this option with care. You cannot recover deleted partitions.

   Deleting obsolete partitions minimizes the number of partitions in the Log Database, which improves database and reporting performance. Use this Delete option to delete individual partitions as needed. See *Configuring Log Database maintenance options*, page 298, if you prefer to delete older partitions according to a schedule.

3. Click **Save Now** to activate changes to the available partitions options.

Viewing error logs

Related topics:

- *Log Database administration settings*, page 294
- *Configuring rollover options*, page 295
- *Configuring Internet browse time options*, page 297
- *Configuring full URL logging*, page 296
- *Configuring Log Database maintenance options*, page 298
- *Configuring Log Database partition creation*, page 300
- *Configuring available partitions*, page 301

Use the **Error Log Activity** section of the Reporting > Log Database page (Settings tab) to view records of errors that occurred during the jobs run on the Websense Log Database (see *Database jobs*, page 293). This information may be helpful in troubleshooting.

Choose one of the following options.

- Choose a number from the drop-down list to display that many error log entries.
- Choose **View All** to display all error log entries.
- Choose **View None** to hide all error log entries.

Configuring investigative reports

Related topics:

- *Database connection and report defaults*, page 303
- *Display and output options*, page 305

Investigative reports let you interactively delve into the information about your organization's Internet usage. See *Investigative reports*, page 105.

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 303.

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 305.

Database connection and report defaults

Related topics:

- *Configuring investigative reports*, page 303
- *Display and output options*, page 305
- *Summary reports*, page 107
- *Multi-level summary reports*, page 111

Use the **Investigative Reports > Options** page to connect to the desired Log Database, and to control defaults for investigative reports detail view.

Changes made to this page affect your reports. Other administrators, or even users logging on for self-reporting, can change these values for their own reporting activities.

1. Choose the Log Database to use for investigative reports.
   - Check **View the catalog database** to connect to the Log Database where Log Server is logging. Proceed to step 2.
   - To access a different Log Database:
     a. Uncheck the **View the catalog database** option.

b. Enter the following information to identify the desired Log Database. (Investigative reports can be generated from a v6.3.x or v7.0 database.)

Field	Description
Server	Enter the machine name or IP address where the Log Database is located.
Database	Enter the name of the Log Database.
User ID	Enter the user ID for an account that has permission to access the database. Leave this blank if Log Server was installed to use a trusted connection to access the Log Database. If you are uncertain, enter **sa**. That is the default user ID for MSDE and the default administrator ID in Microsoft SQL Server.
Password	Enter the password for the specified User ID. 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.

Field	Description
Select report type	Choose whether to open detail reports initially showing: • **Detail**: each record appears on a separate row; time can be displayed. • **Summary**: combines into a single entry all records that share a common element. The specific element varies, according to the information reported. Typically, the right-most column before the measure shows the summarized element. Time cannot be displayed.
Available Columns / Current Report	Select a column name in the Available Columns list and click the appropriate arrow to move it to the Current Report list. Up to 7 columns can be on the Current Report list. After the Current Report list contains all the columns for initial detail reports, set the order of the columns. Select an entry in the list, and use the up and down arrow buttons to change its position.

3. Click **Save Options** to immediately save all changes.

Display and output options

Related topics:

- *Configuring investigative reports*, page 303
- *Database connection and report defaults*, page 303
- *Output to file*, page 129

You can make adjustments to the way certain report choices and report results are displayed in summary and multi-level investigative reports, and specify the default page size when reports are output to PDF format.

These investigative reports configuration options are set in the **wse.ini** file. The default location is:

```
C:\Program Files\Websense\webroot\Explorer\wse.ini
```

The following table lists the parameters that affect display and output of investigative reports, what each controls, and its default value. (Do NOT modify any other settings in the wse.ini file.)

Parameter	Description
maxUsersMenu	The database must have fewer users than this value (by default, 5000) to show User as a report choice in the Internet Use by list.
maxGroupsMenu	The database must have fewer groups than this value (by default, 3000) to show Group as a report choice in the Internet Use by list. **Note:** There must be 2 or more groups for Group to appear in the Internet Use by list. There also must be 2 or more domains for Domain to appear in the Internet Use by list. There is no maximum value for domains.
maxUsersDrilldown	This works with the warnTooManyHits parameter to control when the User option displays in red. The red lettering indicates that selecting User will produce a very large report, which could be slow to generate. If there are more users than this value (by default, 5000), and more hits than the warnTooManyHits value, the User option displays red in various drop-down lists and values lists. If there are more users than this value, but fewer hits than the warnTooManyHits value, the User option displays in normal color, as the resulting report will be a more reasonable size.
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	This works with the maxUsersDrilldown parameter to control when the User option displays in red. If there are more users than the maxUsersDrilldown value, but fewer hits than this value (by default, 10000), the User option does *not* display in red. 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.
hitsPerPage	This determines the maximum number of items (by default, 100) displayed per page. (This does not affect printed reports.)

Parameter	Description
maxOutputBufferSize	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. 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.
sendMulti	This option is disabled (0) by default. Set it to 1 (enabled) to divide very large, scheduled detail reports into multiple files of 10,000 rows each. The files that represent one report are zipped and sent to the email recipients. The report files can be extracted with most common file compression utilities.
maxSlices	This is the maximum number of distinct slices (by default, 6) in a pie chart, including an Other slice, which combines all values that do not have individual slices.
timelineCompressionThreshold	This option is used only for User Activity by Day or Month, when the Group Similar Hits/View All Hits option is available. The report collapses all hits with the same category that occur within the number of seconds set here (by default, 10).
PageSize	Investigative report results can be output to Portable Document Format (PDF) for easy distribution or printing. The page size (by default, Letter) can be: • A4 (8.27 X 11.69 inches) • Letter (8.5 X 11 inches)

Self-reporting

Related topics:

- *Configuring reporting preferences*, page 279
- *Accessing self-reporting*, page 130
- *Investigative reports*, page 105

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.

> **Note**
> Self-reporting is available only when Websense Manager and reporting components are installed on a Windows operating system. See the *Deployment Guide* for additional information.

To enable self-reporting:

1. Go to **Settings >General > Directory Services**, and configure the directory service used to authenticate users who access Websense Manager with their network credentials. This may have been done previously to enable filtering by user and group names. See *Directory services*, page 56.

   If your installation includes multiple Policy Servers, you must log on to each one and configure the Directory Services page with information for the appropriate directory service.
2. Go to the **Settings > Reporting > Preferences**, and mark the **Allow self-reporting** check box. See *Configuring reporting preferences*, page 279.

After enabling the option, be sure to give users the information they need to run the reports:

- The URL for accessing the self-reporting interface. Remind users that they can save the URL as a favorite or bookmark for future use.

  Read on for detailed information about the URL.
- Which Policy Server to select during logon.

  In networks with only one Policy Server, this is not needed. If your network includes multiple Policy Servers, give users the IP address of the Policy Server configured to communicate with the directory service that authenticates their network logon. This is also the Policy Server specified when you installed Log Server.
- What user name and password to use during logon.

  Self-reporting users must enter their network user name and password during logon.

The **URL** for accessing the self-reporting interface is:

```
https://<ServerIP>:9443/mng/login/pages/
   selfReportingLogin.jsf
```

In place of <ServerIP> use the IP address of the machine running Websense Manager.

Administrators and users can also access the self-reporting logon page by opening the Websense Manager logon page and clicking the Self-Reporting link.

If your network includes **multiple Policy Servers**, you must inform users which one to choose during self-reporting logon.

14 Network Configuration

Related topics:

- *Hardware configuration*, page 312
- *Network Agent configuration*, page 313
- *Verifying Network Agent configuration*, page 319

When you run Websense software in stand-alone mode (not integrated with a proxy or firewall product), Websense Network Agent enables:

- Internet content filtering
- Network protocol and Internet application management
- Bandwidth management
- Logging of bytes transferred

In an integrated Websense software deployment, a third-party product may handle the task of routing user requests to Websense software for filtering, and routing block pages back to the client. In this environment, Network Agent may still be used to filter non-HTTP requests, provide enhanced logging detail, or both.

Network Agent continually monitors overall network usage, including bytes transferred over the network. The agent sends usage summaries to Websense software at predefined intervals. Each summary includes start time and end time, overall bytes used, and bytes used per protocol.

By default, Network Agent also provides bandwidth usage data to Policy Server, and filtering log data to Filtering Service.

Network Agent is typically configured to see all traffic in your network. The agent distinguishes between:

- Requests sent from internal machines to internal machines (hits to an intranet server, for example)
- Requests sent from internal machines to external machines such as Web servers (user Internet requests, for example)

The latter is the primary concern in monitoring employee Internet usage.

Hardware configuration

Each Network Agent instance monitors traffic **from** the machines you identify as belonging to your network. By default, it monitors traffic **to** only those internal machines that you specify (for example, internal Web servers).

You can customize which internal machines (network segments) are monitored by each Network Agent instance, or even by each network interface card (NIC) on a Network Agent machine.

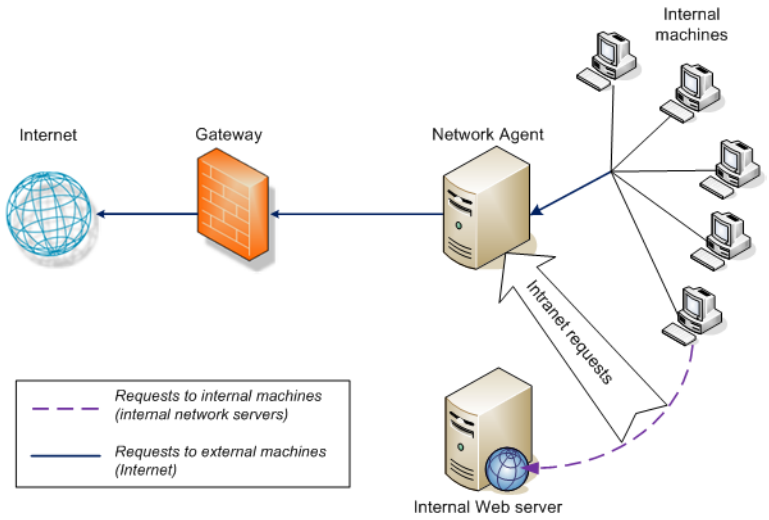

Monitoring requests to internal machines

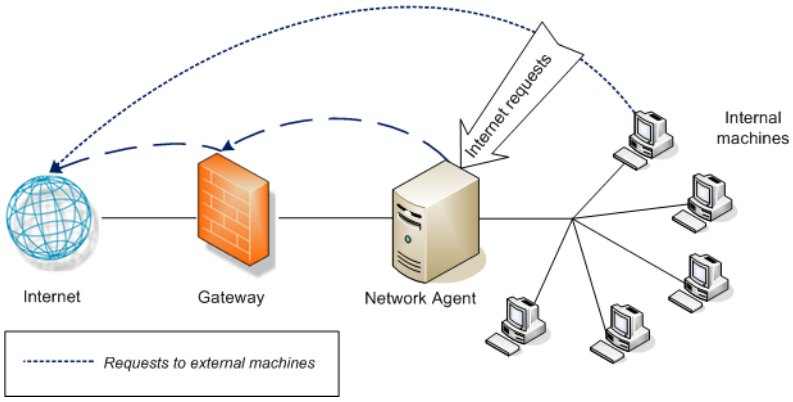

Monitoring requests to external machines

Each Network Agent instance must:

- Be positioned appropriately in the network to detect traffic to and from all monitored machines.
- Have at least 1 NIC dedicated to monitoring traffic.

Network Agent can be installed on a machine with multiple NICs, and can use multiple NICs for both monitoring requests and sending block pages. If you add a new

NIC to the Network Agent machine, restart the Network Agent service, and then configure the new NIC (see *Configuring NIC settings*, page 316).

> **Note**
> To determine whether Network Agent can see traffic in a network segment, use the Network Traffic Detector utility. See *Verifying Network Agent configuration*, page 319.

More information about Network Agent placement and NIC requirements can be found in the *Deployment Guide*.

For information about configuring Network Agent to monitor internal network requests, use specific NICs, and perform enhanced logging, see *Network Agent configuration*, page 313.

Network Agent configuration

Related topics:

- *Hardware configuration*, page 312
- *Configuring global settings*, page 314
- *Configuring local settings*, page 315
- *Configuring NIC settings*, page 316
- *Adding or editing IP addresses*, page 319

After installing Network Agent, use Websense Manager to configure its network monitoring behavior. Network Agent settings are divided into two main areas:

- **Global settings** affect all Network Agent instances. Use these settings to:
  - Identify the machines in your network.
  - List machines in your network that Network Agent should monitor for **incoming** requests (for example, internal Web servers).
  - Specify bandwidth calculation and protocol logging behavior.
- **Local settings** apply only to the selected Network Agent instance. Use these settings to:
  - Identify which Filtering Service instance is associated with each Network Agent.
  - Note proxies and caches used by the machines that this Network Agent monitors.
  - Configure how each network card (NIC) in the Network Agent machine is used (to monitor requests, send block pages, or both).

Network card settings also determine which segment of the network each Network Agent instance monitors.

Configuring global settings

Related topics:

- *Hardware configuration*, page 312
- *Configuring local settings*, page 315
- *Configuring NIC settings*, page 316
- *Adding or editing IP addresses*, page 319

Use the **Settings > Network Agent > Global** page to define basic monitoring and logging behavior for all instances of Network Agent.

The **Internal Network Definition** list identifies the machines that are part of your network. By default, Network Agent does not monitor the traffic (internal network communications) sent between these machines.

An initial set of entries is provided by default. You can add additional entries, or edit or delete existing entries.

The **Internal Traffic to Monitor** list includes any machines included within the Internal Network Definition for which you **do** want Network Agent to monitor traffic. This might include internal Web servers, for example, to help you to track internal connections.

Any requests sent from anywhere in the network to the specified internal machines is monitored. By default, this list is blank.

- Click **Add** to add an IP address or range to the appropriate list. See *Adding or editing IP addresses*, page 319, for more information.
- To edit an entry in the list, click the IP address or range. See *Adding or editing IP addresses*, page 319, for more information.
- To remove an entry from the list, mark the check box next to an IP address or range, and then click **Delete**.

The **Additional Settings** options allow you to determine how often Network Agent calculates bandwidth usage, and whether and how often protocol traffic is logged:

Field	What to do
Bandwidth calculation interval	Enter a number between 1 and 300 to specify how frequently, in seconds, Network Agent should calculate bandwidth usage. An entry of 300, for example, indicates that Network Agent will calculate bandwidth every 5 minutes. The default is 10 seconds.
Log protocol traffic periodically	Mark this option to enable the Logging interval field.
Logging interval	Enter a number between 1 and 300 to specify how frequently, in minutes, Network Agent logs protocols. An entry of 60, for example, indicates that Network Agent will write to the log file every hour. The default is 1 minute.

When you are finished making changes, click **OK** to cache the changes. Changes are not implemented until you click **Save All**.

Configuring local settings

Related topics:

- *Hardware configuration*, page 312
- *Configuring global settings*, page 314
- *Configuring NIC settings*, page 316

Use the **Settings > Network Agent > Local Settings** page to configure filtering behavior, proxy information, and other settings for the selected instance of Network Agent. The IP address of the selected Network Agent instance appears in the title bar of the content pane, and is highlighted in the left navigation 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, filtered, and logged correctly, use the **Proxies and Caches** list to specify the IP address of any proxy or cache server that communicates with Network Agent.

- Click **Add** to add an IP address or range to the list. See *Adding or editing IP addresses*, page 319, for more information.
- To edit an entry in the list, click the IP address or range.
- To remove an entry from the list, mark the check box next to an IP address or range, and then click **Delete**.

Use the **Network Interface Cards** list to configure individual NICs. Click on a NIC in the **Name** column, and then see *Configuring NIC settings*, page 316, for further instructions.

If HTTP requests in your network are passed through a non-standard port, click **Advanced Network Agent Settings** to provide the correct ports for Network Agent to monitor. By default the **Ports used for HTTP traffic** are **8080, 80**.

The other settings in this section should not be changed unless you are instructed to do so by Websense Technical Support.

Field	Description
Mode	• None (default) • General • Error • Detail • Bandwidth
Output	• File (default) • Window
Port	55870 (default)

When you are finished making changes to your Network Agent settings, click **OK** to cache the changes. Changes are not implemented until you click **Save All**.

Configuring NIC settings

Related topics:

- *Hardware configuration*, page 312
- *Network Agent configuration*, page 313
- *Configuring monitoring settings for a NIC*, page 318
- *Adding or editing IP addresses*, page 319

Use the **Network Agent > Local Settings > NIC Configuration** page to specify how Network Agent uses each available network interface card (NIC) to monitor and manage network usage.

The **NIC Information** area provides the context for the changes that you make, showing the **IP address**, brief NIC **Description**, and card **Name**. Use this information to ensure that you are configuring the correct NIC.

Monitoring

In a multiple-NIC configuration, you can identify one NIC to monitor network traffic, and another NIC to serve block pages. At least one NIC must be used for monitoring, and more than one NIC can monitor traffic.

Use the **Monitoring** section to indicate whether or not to **Use this NIC to monitor traffic**.

- If this NIC is not used for monitoring, deselect the check box and then continue with the next section.
- If the NIC is used for monitoring, select the check box, and then click **Configure**. You are taken to the Configure Monitoring Behavior page. See *Configuring monitoring settings for a NIC*, page 318, for instructions.

Other NIC options

In addition to configuring monitoring options, you can also determine other NIC behaviors:

1. Under Blocking, make sure that the appropriate NIC is listed in the **Blocking NIC** field. If you are configuring multiple NICs, the settings for each NIC should show the same value in this field. In other words, only one NIC is used for blocking.
2. If you are running Websense software in **Stand-Alone** mode, **Filter and log HTTP requests** is selected, and cannot be changed.
3. If you have integrated Websense software with a third-party device or application, use the **Integrations** options to indicate how this Network Agent should filter and log HTTP requests. Options that do not apply to your environment are disabled.
   - Select **Log HTTP requests** to improve accuracy in Websense reports.
   - Select **Filter all requests not sent over HTTP ports** to use Network Agent to filter only those HTTP requests not sent through the integration product.
4. Under Protocol Management, indicate whether Network Agent should use this NIC to filter non-HTTP protocols:
   - Check **Filter non-HTTP protocol requests** to activate the protocol management feature. This allows Websense software to filter Internet applications and data transfer methods, such as those used for instant messaging, streaming media, file sharing, Internet mail, and so on. See *Filtering categories and protocols*, page 34, and *Working with protocols*, page 168, 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 174, for more information.

Configuring monitoring settings for a NIC

Use the **Local Settings > NIC Configuration > Monitor List** page to specify which machines Network Agent monitors via the selected network interface card (NIC).

1. Under Monitor List, specify which requests Network Agent monitors:
   - **All**: Network Agent monitors requests from all machines it sees using the selected NIC. Typically, this includes all machines in the same network segment as the current Network Agent machine or NIC.
   - **None**: Network Agent does not monitor any requests.
   - **Specific**: Network Agent monitors only the network segments included in the Monitor List.
2. If you selected Specific, click **Add**, and then specify the IP addresses of the machines Network Agent should monitor. See *Adding or editing IP addresses*, page 319, for more information.

   Note

   You cannot enter overlapping IP address ranges. If ranges overlap, network bandwidth measurements may not be accurate, and bandwidth-based filtering may not be applied correctly.

   To remove an IP address or network range from the list, check the appropriate list item, and then click **Delete**.
3. Under Monitor List Exceptions, identify any internal machines Network Agent should exclude from monitoring.

   For example, Network Agent could ignore requests made by CPM Server. This way, CPM Server requests will not clutter Websense log data or any of the status monitors output.

   a. To identify a machine, click **Add**, and then enter its IP address.

   b. Repeat the process to identify additional machines.
4. Click **OK** to cache your changes and return to the NIC Configuration page. Changes are not implemented until you click **Save All**.

Adding or editing IP addresses

Related topics:

- *Configuring global settings*, page 314
- *Configuring local settings*, page 315
- *Configuring NIC settings*, page 316

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.

- When you add or edit an IP address range, make sure that it does not overlap any existing entry (single IP address or range) in the list.
- When you add or edit a single IP address, make sure that it does not fall into a range that already appears in the list.

To add a new IP address or range:

1. Select the **IP address** or **IP address range** radio button.
2. Enter a valid IP address or range.
3. Click **OK** to return to the previous Network Agent Settings page. The new IP address or range appears in the appropriate table.

   To return to the previous page without caching your changes, click **Cancel**.
4. Repeat this process for additional IP addresses, as needed.

When you edit an existing IP address or range, the Edit IP Addresses page displays the selected item with the correct radio button already selected. Make any necessary changes, and then click **OK** to return to the previous page.

When you are finished adding or editing IP addresses, click **OK** on the Network Agent Settings page. Changes are not implemented until you click **Save All**.

Verifying Network Agent configuration

After configuring Network Agent in Websense Manager, use the Network Traffic Detector to ensure that computers on your network are visible to Websense software.

1. Go to **Start > Programs > Websense > Utilities > Network Traffic Detector** to launch the tool.
2. Select a network card from the **Network Adapter** drop-down list.
3. Check the addresses that appear in the **Monitored Network Ranges** list to verify that all appropriate subnetworks are listed.
4. Use the **Add Subnetwork** and **Remove Subnetwork** buttons to change which parts of the network are tested.

5. Click **Start Monitoring**.

   The Network Traffic Detector detects computers in the network by monitoring the information they send across the network. The **Number of Computers Detected** list shows a running count of computers detected.

6. To see specific information about the computers detected by the tool, select a subnetwork in the Monitored Network Ranges list, and then click **View Detected Computers**.

   If a specific computer is not listed, verify that it is generating network traffic. To do this, go to the machine, launch a browser, and navigate to a Web site. Then return to the Network Traffic Detector and see if the computer appears in the **Detected Computers** dialog box.

7. When you have finished testing network traffic visibility, click **Stop Monitoring**.

If some computers are not visible:

- Review the network configuration and NIC placement requirements (see *Hardware configuration*, page 312).
- Review the more detailed network configuration information in the *Installation Guide* for your Websense software.
- Verify that you have properly configured the monitoring NIC (*Configuring NIC settings*, page 316).

15 Troubleshooting

Use this section to find solutions to common issues before contacting Technical Support.

The Websense Web site features an extensive Knowledge Base, available at www.websense.com/global/en/SupportAndKB/. Search for topics by keyword or reference number, or browse the most popular articles.

Troubleshooting instructions are grouped into the following sections:

Installation and subscription issues

Websense Status shows a subscription problem

A valid subscription key is needed to download the Websense Master Database and perform Internet filtering. When your subscription expires or is invalid, and when the Master Database has not been downloaded for more than 2 weeks, the Websense Health monitor displays a warning.

- Verify that you have entered your subscription key exactly as you received it. The key is case sensitive.
- Make sure that your subscription has not expired. See *Subscription key*, page 324.
- Ensure that the Master Database has been downloaded successfully within the last 2 weeks. You can check download status in Websense Manager: click **Database Download** on the Status > Today page.

  See *The Master Database does not download*, page 323, for help troubleshooting database download problems.

If you have entered the key correctly, but continue to receive a status error, or if your subscription has expired, contact Websense, Inc., or your authorized reseller.

When your subscription expires, Websense Manager settings determine whether all users are given unfiltered Internet access or all Internet requests are blocked. See *Your subscription*, page 25, for more information.

After upgrade, users are missing from Websense Manager

If you have defined Active Directory as your directory service, after an upgrade to Websense software, user names may not appear in Websense Manager. This occurs when user names include characters that are not part of the UTF-8 character set.

To support LDAP 3.0, the Websense installer changes the character set from MBCS to UTF-8 during upgrade. As a result, user names that include non-UTF-8 characters are not properly recognized.

To fix this problem, manually change the character set to MBCS:

1. In Websense Manager, go to **Settings** > **Directory Services**.
2. Make sure that **Active Directory (Native Mode)** is selected under Directories, near the top of the page.
3. Click **Advanced Directory Settings**.
4. Under Character Set, click **MBCS**. You may have to scroll down to see this option.
5. Click **OK** to cache the change. Changes are not implemented until you click **Save All**.

Master Database issues

The initial filtering database is being used

The Websense Master Database houses the category and protocol definitions that provide the basis for filtering Internet content.

A partial version of the Master Database is installed with your Websense software on each Filtering Service machine. This partial database is used to enable basic filtering functionality from the time you enter your subscription key.

You must download the full database for complete filtering to occur. See *The Websense Master Database*, page 28, for more information.

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

The Master Database is more than 1 week old

The Websense Master Database houses the category and protocol definitions that provide the basis for filtering Internet content. Websense software downloads changes to the Master Database according to the schedule defined in Websense Manager. By default, download is scheduled to occur once a day.

To manually initiate a database download:

1. In Websense Manager, go to the **Status > Today** page, and then click **Database Download**.
2. Click **Update** next to the appropriate Filtering Service instance to start the database download, or click **Update All** to start the download on all Filtering Service machines.

   Note

   After downloading updates to the Master Database, CPU usage can be 90% or more for a short time while the database is loaded into local memory. It is a good idea to perform the download at off-peak times.

3. To continue working while the database is downloaded, click **Close**.

   Click the **Database Download** button at any time to view download status.

If a new version of the Master Database adds or removes categories or protocols, administrators performing category- or protocol-related policy management tasks (like editing a category set) at the time of the download may receive errors. Although such updates are somewhat rare, as a best practice, try to avoid making changes to category- and protocol-related while a database is being updated.

The Master Database does not download

If you are unable to download the Websense Master Database successfully:

- Make sure that you have entered your subscription key correctly in Websense Manager, and that the key has not expired (*Subscription key*, page 324).
- Verify that the Filtering Service machine is able to access the Internet (*Internet access*, page 324).
- Check firewall or proxy server settings to make sure that Filtering Service can connect to the Websense download server (*Verify firewall or proxy server settings*, page 325).
- Make sure that there is enough disk space (*Insufficient disk space*, page 326) and memory (*Insufficient memory*, page 327) on the download machine.
- Look for any application or appliance in the network, such as anti-virus software, that might prevent the download connection (*Restriction applications*, page 327).

Subscription key

To verify that the subscription key is entered correctly and has not expired:

1. In Websense Manager, go to **Settings > Account**.
2. Compare the key that you received from Websense, Inc., or your reseller to the **Subscription key** field. The key must use the same capitalization as in your key document.
3. Check the date next to **Key expires**. If the date has passed, contact your reseller or Websense, Inc., to renew your subscription.
4. If you have made changes to the key in the Settings dialog box, click **OK** to activate the key and enable database download.

To manually initiate a database download, or to check the status of the most recent database download, click **Database Download** in the toolbar at the top of the Status > Today page.

Internet access

To download the Master Database, the Filtering Service machine sends an **HTTP post** command to the download servers at the following URLs:

download.websense.com
ddsdom.websense.com
ddsint.websense.com
portal.websense.com
my.websense.com

To verify that Filtering Service has the Internet access necessary to communicate with the download server:

1. Open a browser on the machine running Filtering Service.
2. Enter the following URL:

   http://download.websense.com/

   If the machine is able to open an HTTP connection to the site, a redirect page is displayed, and then the browser displays the Websense home page.

If this does not happen, ensure that the machine:

- Can communicate over port 80, or the port designated in your network for HTTP traffic
- Is configured to properly perform DNS lookups
- Is configured to use any necessary proxy servers (see *Verify firewall or proxy server settings*, page 325)

Also make sure that your gateway does not include any rules that block HTTP traffic from the Filtering Service machine.

3. Use one of the following methods to confirm that the machine can communicate with the download site:
   - From the command prompt, enter the following command:

     ```
     ping download.websense.com
     ```

     Verify that the ping receives a reply from the download server.
   - Use telnet to connect to **download.websense.com 80**. If you see a cursor and no error message, you can connect to the download server.

Verify firewall or proxy server settings

If the Master Database is downloaded through a firewall or proxy server that requires authentication, ensure that a browser on the Filtering Service machine can load Web pages properly. If pages open normally, but the Master Database does not download, check the proxy server settings in the Web browser.

Microsoft Internet Explorer:

1. Select **Tools > Internet Options**.
2. Open the **Connections** tab.
3. Click **LAN Settings**. Proxy server configuration information appears under **Proxy server**.

   Make a note of the proxy settings.

Mozilla Firefox:

1. Select **Tools > Options > Advanced**.
2. Select the **Network** tab.
3. Click **Settings**. The Connection Settings dialog box shows whether the browser is configured to connect to a proxy server.

   Make a note of the proxy settings.

Next, make sure that Websense software is configured to use the same proxy server to perform the download.

1. In Websense Manager, go to **Settings > Database Download**.
2. Verify that **Use proxy server or firewall** is selected, and that the correct server and port are listed.

3. Make sure that the **Authentication** settings are correct. Verify the user name and password, checking spelling and capitalization.

   If Websense software must provide authentication information, the firewall or proxy server must be configured to accept clear text or basic authentication. Information about enabling basic authentication is available from the Websense Knowledge Base.

If a firewall restricts Internet access at the time Websense software normally downloads the database, or restricts the size of a file that can be transferred via HTTP, Websense software cannot download the database. To determine if the firewall is causing the download failure, search for a rule on the firewall that might be blocking the download, and change the download times in Websense Manager (*Configuring database downloads*, page 29), if necessary.

Insufficient disk space

The Websense Master Database is stored in the Websense **bin** directory (/opt/Websense/bin or C:\Program Files\Websense\bin, by default). The drive containing this directory must have enough space to download the compressed database, and enough room for the database to be decompressed.

The machine should have at least 2 times the size of the Master Database in free disk space. As the entries in the Master Database increase, the size required for a successful download increases. As a general rule, Websense, Inc., recommends at least 3 GB of free disk space on the download drive.

In Windows, use Windows Explorer to check the disk space available:

1. Open **My Computer** in Windows Explorer (not Internet Explorer).
2. Select the drive on which Websense software is installed. By default, Websense software is located on the C drive.
3. Right-click and select **Properties** from the pop-up menu.
4. On the General tab, verify that at least 3 GB of free space is available. If there is insufficient free space on the drive, delete any unnecessary files to free up the required space.

On Linux systems, use the **df** command to verify the amount of available space in the file system in which Websense software is installed:

1. Open a terminal session.
2. At the prompt, enter:

   ```
   df -h /opt
   ```

   Websense software is usually installed in the /opt/Websense/bin directory. If it is installed elsewhere, use that path.
3. Make sure that at least 3 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 you verify that there is sufficient disk space, but still have download problems, try stopping all Websense services (see *Stopping and starting Websense services*, page

258), deleting the **Websense.xfr** and **Websense** (no extension) files, starting the services, and then manually downloading a new database.

Insufficient memory

The memory required to run Websense software and download the Master Database varies, depending on the size of the network. For example, in a small network, 2 GB of memory is recommended for all platforms.

Refer to the *Deployment Guide* for system recommendations.

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.
4. If less than 2 GB is installed, upgrade the RAM in the machine.

You also can select **Control Panel > Administrative Tools > Performance** to capture information.

To check the memory in a Linux system:

1. Open a terminal session.
2. At the prompt, enter:

   ```
   top
   ```
3. Compute the total memory available by adding **Mem: av** and **Swap: av**.
4. If less than 2 GB is installed, upgrade the RAM in the machine.

Restriction applications

Some restriction applications or appliances, such as virus scanners, size-limiting applications, or intrusion detection systems can interfere with database downloads. Ideally, configure Websense software to go straight to the last gateway so that it does not connect to these applications or appliances. Alternatively:

1. Disable the restrictions relating to the Filtering Service machine and to the Master Database download location.

   See the appliance or software documentation for instructions on changing the device's configuration.
2. Attempt to download the Master Database.

If this change has no effect, reconfigure the application or appliance to include the machine running Filtering Service.

Master Database download does not occur at the correct time

The system date and time may not be set correctly on the Filtering Service machine. Websense software uses the system clock to determine the proper time for downloading the Master Database.

If the download is not occurring at all, see *The Master Database does not download*, page 323.

Contacting Technical Support for database download issues

If you are still experiencing Master Database download problems after completing the troubleshooting steps in this Help section, send the following information to Websense Technical Support:

1. The exact error message that appears in the Database Download dialog box
2. External IP addresses of the machines attempting to download the database
3. Your Websense subscription key
4. Date and time of the last attempt
5. Number of bytes transferred, if any
6. Open a command prompt and perform an **nslookup** on **download.websense.com**. If connection to the download server is made, send the IP addresses returned to Technical Support.
7. Open a command prompt and perform a **tracert** to **download.websense.com**. If connection to the download server is made, send the route trace to Technical Support.
8. A packet trace or packet capture performed on the Websense download server during an attempted download.
9. A packet trace or packet capture performed on the network gateway during the same attempted download.
10. The following files from the Websense **bin** directory: **websense.ini**, **eimserver.ini** and **config.xml**.

Go to www.websense.com/SupportPortal/default.aspx for Technical Support contact information.

Filtering issues

- *An FTP request is not blocked as expected*, page 331
- *Websense software is not applying user or group policies*, page 332
- *Remote users are not filtered by correct policy*, page 332

Filtering service is not running

When Filtering Service is not running, Internet requests cannot be filtered and logged.

Filtering Service may stop running if:

- There is insufficient disk space on the Filtering Service machine.
- A Master Database download failed due to lack of disk space (see *The Master Database does not download*, page 323).
- 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\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 for error messages from Filtering Service (see *Troubleshooting tools*, page 361).
- Log off of Websense Manager, restart Websense Policy Server, and then restart Websense Filtering Service (see *Stopping and starting Websense services*, page 258).

  Wait 1 minute before logging on to Websense Manager again.

User Service is not available

When User Service is not running, or when Policy Server cannot communicate with User Service, Websense software cannot correctly apply user-based filtering policies.

User Service may appear to have stopped if you restarted Policy Server after restarting other Websense Services. To correct this issue:

1. Restart the Websense Policy Server service (see *Stopping and starting Websense services*, page 258).
2. Start or restart Websense User Service.
3. Close Websense Manager.

Wait 1 minute before logging on to Websense Manager again.

If the previous steps do not fix the problem:

- Check the Windows Event Viewer or **websense.log** file for error messages from User Service (see *Troubleshooting tools*, page 361).
- Navigate to the Websense **bin** directory (C:\Program Files\Websense\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.

Sites are incorrectly categorized as Information Technology

Internet Explorer versions 4.0 and later have the ability to accept searches from the Address bar. When this option is enabled, if a user enters only a domain name in the Address bar (**websense** instead of **http://www.websense.com**, for example), Internet Explorer considers the entry a search request, not a site request. It displays the most likely site the user is looking for, along with a list of closely matching sites.

As a result, Websense software permits, blocks, or limits the request based on the status of the Information Technology/Search Engines and Portals category in the active policy—not on the category of the requested site. For Websense software to filter based on the category of the requested site, you must turn off searching from the Address bar:

1. Go to **Tools > Internet Options**.
2. Go to the **Advanced** tab.
3. Under Search from the Address bar, select **Do not search from the Address bar**.
4. Click **OK**.

> **Note**
> These steps are valid for Internet Explorer versions 5, 6, and 7.

Keywords are not being blocked

There are 2 possible reasons for this problem: **Disable keyword blocking** is selected, or the site whose URL contains the keyword uses **post** to send data to your Web server.

To ensure that keyword blocking is enabled:

1. In Websense Manager, go to **Settings > Filtering**.
2. Under General Filtering, check the **Keyword search options** list. If **Disable keyword blocking** is shown, select another option from the list. See *Configuring Websense filtering settings*, page 50, for more information about the available options.
3. Click **OK** to cache the change. Changes are not implemented until you click **Save All**.

If a site uses **post** to send data to your Web server, Websense software does not recognize keyword filtering settings for that URL. Unless your integration product recognizes data sent via post, users can still access URLs containing blocked keywords.

To see whether a site uses a post command, view the site's source from within your browser. If the source code contains a string like **<method=post>**, then post is used to load that site.

Custom or limited access filter URLs are not filtered as expected

If an HTTPS URL in a limited access filter or custom URL list (recategorized or unfiltered) is not filtered as expected, an integration product may be transforming the URL into a format that Filtering Service cannot recognize.

Non-proxy integration products translate URLs from domain format into IP format. For example, the URL **https://<domain>** is read as **https://<IP address>:443**. When this occurs, Filtering Service cannot match the URL received from the integration product with a custom URL or limited access filter, and does not filter the site appropriately.

To work around this problem, add both the IP addresses and URLs for sites you want to filter using custom URLs or limited access filters.

A user cannot access a protocol or application as expected

If your network includes Microsoft ISA Server, certain authentication method configurations may result in dropped connections to messaging applications.

If any method other than Anonymous Authentication is active, the proxy server attempts to identify data packets received when users request application connections. The proxy server fails to identify the data packet, and the connection is dropped. This may skew Websense protocol filtering activity.

An inability to access a protocol or Internet application might also occur if the port used by the application is blocked. This could occur if:

- The port is blocked by a firewall.
- A blocked custom protocol includes the port (as a single port or as part of a port range) in any of its identifiers.

An FTP request is not blocked as expected

When integrated with Check Point® firewalls, Websense software requires **folder view** to be enabled in the client's browser to recognize and filter FTP requests.

When folder view is not enabled, FTP requests sent to the FireWall-1 proxy are sent to Websense software with an "http://" prefix. As a result, Websense software filters these requests as HTTP requests, rather than FTP requests.

Websense software is not applying user or group policies

If Websense software is applying computer or network policies, or the **Default** policy, even after user or group policies were assigned, see *User identification issues*, page 334. Additional information is available via the Knowledge Base.

Remote users are not filtered by correct policy

If a remote user accesses the network by logging on using cached domain credentials (network logon information), Websense software applies the policy assigned to that user, or to the user's group or domain, if appropriate. If there is no policy assigned to the user, group, or domain, or if the user logs on to the computer with a local user account, Websense software applies the Default policy.

Occasionally, a user is not filtered by a user or group policy or the Default policy. This occurs when the user logs on to the remote computer with a local user account, and the last portion of the remote computer's Media Access Control (MAC) address overlaps with an in-network IP address to which a policy has been assigned. In this case, the policy assigned to that particular IP address is applied to the remote user.

Network Agent issues

Network Agent is not installed

Network Agent is required to enable protocol filtering. With some integrations, Network Agent is also used to provide more accurate logging.

If you are running with an integration product, and do not require Network Agent protocol filtering or logging, you can hide the "No Network Agent is installed" status message. See *Reviewing current system status*, page 266, for instructions.

For stand-alone installations, Network Agent must be installed for network traffic to be monitored and filtered. See the *Installation Guide* for installation instructions, and then see *Network Agent configuration*, page 313.

Network Agent is not running

Network Agent is required to enable protocol filtering. With some integrations, Network Agent is also used to provide more accurate logging.

For stand-alone installations, Network Agent must be running to monitor and filter network traffic.

To troubleshoot this problem:

1. Check the Windows Services dialog box (see *The Windows Services dialog box*, page 361) to see if the **Websense Network Agent** service has started.
2. Restart the **Websense Policy Broker** and **Websense Policy Server** services (see *Stopping and starting Websense services*, page 258).
3. Start or restart the **Websense Network Agent** service.
4. Close Websense Manager.
5. Wait 1 minute, and then log on to Websense Manager again.

If that does not fix the problem:

- Check the **Windows Event Viewer** for error messages from Network Agent (see *The Windows Event Viewer*, page 361).
- Check the **Websense.log** file for error messages from Network Agent (see *The Websense log file*, page 362).

Network Agent is not monitoring any NICs

Network Agent must be associated with at least one network interface card (NIC) to monitor network traffic.

If you add or remove network cards from the Network Agent machine, you must update your Network Agent configuration.

1. In Websense Manager, go to **Settings**.
2. In the left navigation pane, under Network Agent, select the IP address of the Network Agent machine.
3. Verify that all NICs for the selected machine are listed.
4. Verify that at least one NIC is set to monitor network traffic.

See *Network Agent configuration*, page 313, for more information.

Network Agent can't communicate with Filtering Service

Network Agent must be able to communicate with Filtering Service to enforce your Internet usage policies.

- Did you change the IP address of Filtering Service machine or reinstall Filtering Service?

  If so, see *Update Filtering Service IP address or UID information*, page 334.
- Do you have more than 2 network interface cards (NICs) on the Network Agent machine?

  If so, see *Network Configuration*, page 311, to verify your Websense software settings.

- Have you reconfigured the switch connected to the Network Agent?

  If so, refer to the *Installation Guide* to verify your hardware setup, and see *Network Agent configuration*, page 313, to verify your Websense settings.

If none of these apply, see *Configuring local settings*, page 315, for information about associating Network Agent and Filtering Service.

Update Filtering Service IP address or UID information

When Filtering Service has been uninstalled and reinstalled, Network Agent does not automatically update the internal identifier (UID) for the Filtering Service. Websense Manager attempts to query Filtering Service using the old UID, which no longer exists.

Likewise, when you change the IP address of the Filtering Service machine, this change is not automatically registered.

To re-establish connection to the Filtering Service:

1. Open Websense Manager.

   A status message indicates that a Network Agent instance is unable to connect to Filtering Service.
2. Click **Settings** at the top of the left navigation pane.
3. In the left navigation pane, under Network Agent, select the IP address of the Network Agent machine.
4. At the top of the page, under Filtering Service Definition, expand the **Server IP address** list, and then select the IP address of the Filtering Service machine.
5. Click **OK** at the bottom of the page to cache the update. Changes are not implemented until you click **Save All**.

User identification issues

Related topics:

- *Filtering issues*, page 328
- *Remote users are not prompted for manual authentication*, page 343
- *Remote users are not being filtered correctly*, page 344

If Websense software is using computer or network policies, or the **Default** policy, to filter Internet requests, even after you have assigned user or group-based policies, or if the wrong user or group-based policy is being applied, use the following steps to pinpoint the problem:

- If you are using Microsoft ISA Server, and changed its authentication method, ensure that the Web Proxy Service was restarted.
- If you are using nested groups in Windows Active Directory, policies assigned to a parent group are applied to users belonging to a sub-group, and not directly to the parent group. For information on user and group hierarchies, see your directory service documentation.
- The User Service cache may be outdated. User Service caches user name to IP address mappings for 3 hours. You can force the User Service cache to update by caching any change in Websense Manager, and then clicking **Save All**.
- If the user being filtered incorrectly is on a machine running Windows XP SP2, the problem could be due to the Windows Internet Connection Firewall (ICF), included and enabled by default in Windows XP SP2. For more information about the Windows ICF, see Microsoft Knowledge Base Article #320855.

  For DC Agent or Logon Agent to get user logon information from a machine running Windows XP SP2:

  1. From the Windows **Start** menu on the client machine, select **Settings** > **Control Panel** > **Security Center** > **Windows Firewall**.
  2. Go to the **Exceptions** tab.
  3. Check **File and Printer Sharing**.
  4. Click **OK** to close the ICF dialog box, and then close any other open windows.

If you are using a Websense transparent identification agent, consult the appropriate troubleshooting section:

- *Troubleshooting DC Agent*, page 335.
- *Troubleshooting Logon Agent*, page 337.
- *Troubleshooting eDirectory Agent*, page 340.
- *Troubleshooting RADIUS Agent*, page 342.

Troubleshooting DC Agent

To troubleshoot user identification problems with DC Agent:

1. Check all network connections.
2. Check the Windows Event Viewer for error messages (see *The Windows Event Viewer*, page 361).
3. Check the Websense log file (Websense.log) for detailed error information (see *The Websense log file*, page 362).

Common causes for DC Agent user identification problems include:

- Network or Windows services are communicating with the domain controller in a way that makes DC Agent see the service as a new user, to whom no policy has been defined. See *Users are incorrectly filtered by the Default policy*, page 336.
- DC Agent or User Service may have been installed as a service 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, DC Agent is not allowed to download the list. See *Changing DC Agent and User Service permissions manually*, page 336.

- The User Service cache is outdated. User Service caches user-name-to-IP-address mappings for 3 hours, by default. The cache is also updated each time you make changes and click **Save All** in Websense Manager.

Users are incorrectly filtered by the Default policy

When some network or Microsoft Windows 200x contact the domain controller, the account name they use can cause Websense software to believe that an unidentified user is accessing the Internet from the filtered machine. Because no user or group-based policy has been assigned to this user, the computer or network policy, or the Default policy, is applied.

- Network services may require domain privileges to access data on the network, and use the domain user name under which they are running to contact the domain controller.

  To address this issue, see *Configuring an agent to ignore certain user names*, page 213.

- Windows 200x services contact the domain controller periodically with a user name made up of the computer name followed by a dollar sign (jdoe-computer$). DC Agent interprets the service as a new user, to whom no policy has been assigned.

  To address this issue, configure DC Agent to ignore any logon of the form **computer$**.

  1. On the DC Agent machine, navigate to the Websense **bin** directory (by default, **C:\Program Files\Websense\bin**).
  2. Open the **transid.ini** file in a text editor.
  3. Add the following entry to the file:

     ```
     IgnoreDollarSign=true
     ```

  4. Save and close the file.
  5. Restart DC Agent (see *Stopping and starting Websense services*, page 258).

Changing DC Agent and User Service permissions manually

On the machine running the domain controller:

1. 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. Open the Windows Services dialog box on each Websense DC Agent machine (go to **Start > Programs > Administrative Tools > Services**).
3. Select the **Websense DC Agent** entry, and then click **Stop**.
4. Double-click the **Websense DC Agent** entry.
5. On the **Log On** tab, select the **This account** option.
6. Enter the user name of the Websense DC Agent 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 DC Agent** entry in the Services dialog box, and then click **Start**.
10. Repeat this procedure for each instance of the Websense User Service.

Troubleshooting Logon Agent

If some users in your network are filtered by the **Default** policy because Logon Agent is not able to identify them:

- Make sure that Windows Group Policy Objects (GPO) are being applied correctly to these users' machines (see *Group Policy Objects*, page 337).
- If User Service is installed on a Linux machine and you are using Windows Active Directory (Native Mode), check your directory service configuration (see *User Service running on Linux*, page 338).
- Verify that the client machine can communicate with the domain controller from which the logon script is being run (see *Domain controller visibility*, page 338).
- Ensure that NetBIOS is enabled on the client machine (see *NetBIOS*, page 338).
- Make sure that the user profile on the client machine has not become corrupt (see *User profile issues*, page 339).

Group Policy Objects

After verifying that your environment meets the prerequisites described in the *Installation Guide* for your Websense software, make sure that Group Policy Objects are being applied correctly:

1. On the Active Directory machine, open the Windows Control Panel and go to **Administrative Tools > Active Directory Users and Computers**.
2. Right-click the domain entry, and then select **Properties**.
3. Click the **Group Policy** tab, and then select the domain policy from the Group Domain Policy Objects Links list.
4. Click **Edit**, and then expand the User Configuration node in the directory tree.
5. Expand the Windows Settings node, and then select **Scripts**.
6. In the right pane, double-click **Logon**, and then verify that **logon.bat** is listed in the Logon Properties dialog box.

This script is required by the client Logon Application.

- If **logon.bat** is not in the script, refer to the *Initial Setup* chapter of your Websense software *Installation Guide*.
- If **logon.bat** does appear in the script, but Logon Agent is not working, use the additional troubleshooting steps in this section to verify that there is not a network connectivity problem, or refer to the Websense Knowledge Base.

User Service running on Linux

When you use Logon Agent for transparent identification of users, and User Service is installed on a Linux machine, you must temporarily configure Websense software to communicate with Active Directory in Mixed Mode.

1. In Websense Manager, go to **Settings > Directory Services**.
2. Make a note of your current directory settings.
3. Under Directories, select **Windows NT Directory / Active Directory (Mixed Mode)**.
4. Click **OK** to cache changes, and then click **Save All**.
5. Under Directories, select **Active Directory (Native Mode)**. If your original configuration does not appear, use the notes recorded in step 2 to recreate your directory settings. See *Windows Active Directory (Native Mode)*, page 57, for detailed instructions.
6. When you are finished making configuration changes, click **OK**, and then click **Save All**.

Domain controller visibility

To verify that the client machine can communicate with the domain controller:

1. Attempt to map a drive on the client machine to the domain controller's root shared drive. This is where the logon script normally runs, and where **LogonApp.exe** resides.
2. On the client machine, open a Windows command prompt and execute the following command:

   ```
   net view /domain:<domain name>
   ```

If either of these tests fails, see your Windows operating system documentation for possible solutions. There is a network connectivity problem not related to Websense software.

NetBIOS

NetBIOS for TCP/IP must be enabled and the TCP/IP NetBIOS Helper service must be running for the Websense logon script to execute on the user's machine.

To make sure that NetBIOS for TCP/IP is enabled on the client machine.

1. Right-click **My Network Places**, and then select **Properties**.

2. Right-click **Local Area Connection**, and then select **Properties**.
3. Select **Internet Protocol (TCP/IP)**, and then click **Properties**.
4. Click **Advanced**.
5. Select the **WINS** tab, and then verify that the correct NetBIOS option is set.
6. If you make a change, click **OK**, then click **OK** twice more to close the different Properties dialog boxes and save your changes.

   If no change was needed, click **Cancel** to close each dialog box without making changes.

Use the Windows Services dialog box to verify that the **TCP/IP NetBIOS Helper** service is running on the client machine (see *The Windows Services dialog box*, page 361). The TCP/IP NetBIOS Helper service runs on Windows 2000, Windows XP, Windows Server 2003, and Windows NT.

User profile issues

If the user profile on the client machine is corrupt, the Websense logon script (and Windows GPO settings) cannot run. This problem can be resolved by recreating the user profile.

When you recreate a user profile, the user's existing My Documents folder, Favorites, and other custom data and settings are not automatically transferred to the new profile. Do not delete the existing, corrupted profile until you have verified that the new profile has solved the problem and copied the user's existing data to the new profile.

To recreate the user profile:

1. Log on to the client machine as a local administrator.
2. Rename the directory that contains the user profile:

   ```
   C:\Documents and Settings\<user name>
   ```

3. Restart the machine.
4. Log on to the machine as the filtered user. A new user profile is created automatically.
5. Check to make sure the user is filtered as expected.
6. Copy the custom data (such as the contents of the My Documents folder) from the old profile to the new one. Do not use the File and Settings Transfer Wizard, which may transfer the corruption to the new profile.

Troubleshooting eDirectory Agent

Related topics:

- *Enabling eDirectory Agent diagnostics*, page 341
- *eDirectory Agent miscounts eDirectory Server connections*, page 341
- *Running eDirectory Agent in console mode*, page 342

A user may not be filtered properly if the user name is not being passed to eDirectory Agent. If a user does not log on to Novell eDirectory server, eDirectory Agent cannot detect the logon. This happens because:

- A user logs on to a domain that is not included in the default root context for eDirectory user logon sessions. This root context is specified during installation, and should match the root context specified for Novell eDirectory on the **Settings** > **Directory Services** page.
- A user tries to bypass a logon prompt to circumvent Websense filtering.
- A user does not have an account set up in eDirectory server.

If a user does not log on to eDirectory server, user-specific policies cannot be applied to that user. Instead, the **Default** policy takes effect. If there are shared workstations in your network where users log on anonymously, set up a filtering policy for those particular machines.

To determine whether eDirectory Agent is receiving a user name and identifying that user:

1. Activate eDirectory Agent logging, as described under *Enabling eDirectory Agent diagnostics*, page 341.
2. Open the log file you have specified in a text editor.
3. Search for an entry corresponding to the user who is not being filtered properly.
4. An entry like the following indicates that eDirectory Agent has identified a user:

   ```
   WsUserData::WsUserData()
   User: cn=Admin,o=novell (10.202.4.78)
   WsUserData::~WsUserData()
   ```

   In the example above, the user **Admin** logged on to eDirectory server, and was identified successfully.
5. If a user is being identified, but is still not being filtered as expected, check your policy configuration to verify that the appropriate policy is applied to that user, and that the user name in Websense Manager corresponds to the user name in Novell eDirectory.

   If the user is *not* being identified, verify that:

   - The user has a Novell eDirectory account.
   - The user is logging on to a domain that is included in the default root context for eDirectory user logons.

- The user is not bypassing a logon prompt.

Enabling eDirectory Agent diagnostics

eDirectory Agent has built-in diagnostic capabilities, but these are not activated by default. You can enable logging and debugging during installation, or at any other time.

1. Stop eDirectory Agent (see *Stopping and starting Websense services*, page 258).
2. On the eDirectory Agent machine, go to the eDirectory Agent installation directory.
3. Open the file **wsedir.ini** in a text editor.
4. Locate the **[eDirAgent]** section.
5. To enable logging and debugging, change the value of **DebugMode** to **On**:

   ```
   DebugMode=On
   ```

6. To specify the log detail level, modify the following line:

   ```
   DebugLevel=<N>
   ```

   N can be a value from 0-3, where 3 indicates the most detail.
7. Modify the **LogFile** line to specify the name of the log output file:

   ```
   LogFile=filename.txt
   ```

   By default, log output is sent to the eDirectory Agent console. If you are running the agent in console mode (see *Running eDirectory Agent in console mode*, page 342), you can keep the default value.
8. Save and close the **wsedir.ini** file.
9. Start the eDirectory Agent service (see *Stopping and starting Websense services*, page 258).

eDirectory Agent miscounts eDirectory Server connections

If eDirectory Agent is monitoring more than 1000 users in your network, but shows only 1000 connections to the Novell eDirectory server, it may be due to a limitation of the Windows API that conveys information from the eDirectory server to the Websense eDirectory Agent. This is occurs very rarely.

To work around this limitation, add a parameter to the **wsedir.ini** file that counts server connections accurately (Windows only):

1. Stop the Websense eDirectory Agent service (see *Stopping and starting Websense services*, page 258).
2. Go to the Websense **bin** directory (by default, **C:\Program Files\Websense\bin**).
3. Open the **wsedir.ini** file in a text editor.
4. Insert a blank line, and then enter:

   ```
   MaxConnNumber = <NNNN>
   ```

Here, <*NNNN*> is the maximum number of possible connections to the Novell eDirectory server. For example, if your network has 1,950 users, you might enter 2000 as the maximum number.

5. Save the file.
6. Restart eDirectory Agent.

Running eDirectory Agent in console mode

1. Do one of the following:
   - At the Windows command prompt (**Start > Run > cmd**), enter the command:

     ```
     eDirectoryAgent.exe -c
     ```

   - At the Linux command shell, enter the command:

     ```
     eDirectoryAgent -c
     ```

2. When you are ready to stop the agent, press **Enter**. It may take a few seconds for the agent to stop running.

Troubleshooting RADIUS Agent

RADIUS Agent has built-in diagnostic capabilities, but these are not activated by default. To activate RADIUS Agent logging and debugging:

1. Stop the RADIUS Agent service (see *Stopping and starting Websense services*, page 258).
2. On the RADIUS Agent machine, go to the agent installation directory (by default, **Websense\bin**).
3. Open the **wsradius.ini** file in a text editor.
4. Locate the **[RADIUSAgent]** section.
5. To enable logging and debugging, change the value of **DebugMode** to **On**:

   ```
   DebugMode=On
   ```

6. To specify the log detail level, modify the following line:

   ```
   DebugLevel=<N>
   ```

   N can be a value from 0-3, where 3 indicates the most detail.

7. Modify the **LogFile** line to indicate the name of the output file:

   ```
   LogFile=filename.txt
   ```

   By default, log output is sent to the RADIUS Agent console. If you are running the agent in console mode (see *Running RADIUS Agent in console mode*, page 343), you can optionally keep the default value.

8. Save and close the **wsradius.ini** file.
9. Start the RADIUS Agent service (see *Stopping and starting Websense services*, page 258).

If remote users are not being identified and filtered as expected, the likely cause is communication problems between RADIUS Agent and your RADIUS server. Check your RADIUS Agent logs for errors to determine the cause.

Running RADIUS Agent in console mode

To start RADIUS Agent in console mode (as an application), enter the following:

- At the Windows command prompt:

  ```
  RadiusAgent.exe -c
  ```

- At the Linux shell prompt:

  ```
  ./RadiusAgent -c
  ```

To stop the agent at any time, press **Enter** again. It may take a couple of seconds for the agent to stop running.

RADIUS Agent accepts the following command line parameters:

> **Note**
> On Linux, Websense, Inc., recommends using the script provided to start or stop Websense RADIUS Agent (**WsRADIUSAgent start|stop**), instead of the -r and -s parameters.

Parameter	Description
-i	Installs RADIUS Agent service/daemon.
-r	Runs RADIUS Agent service/daemon.
-s	Stops RADIUS Agent service/daemon.
-c	Runs RADIUS Agent as an application process instead of as a service or daemon. When in console mode, RADIUS Agent can be configured to send log output to the console or to a text file.
-v	Displays the version number of RADIUS Agent.
-? -h -help <*no option*>	Displays usage information on the command line. Lists and describes all possible command line parameters.

Remote users are not prompted for manual authentication

If you have configured remote users to manually authenticate when accessing the Internet, there may be some occasions when individual users are not prompted for the authentication. This can occur in situations where some in-network IP addresses have been configured to bypass manual authentication.

When a remote user accesses the network, Websense software reads the last portion of the computer's Media Access Control (MAC) address. If this matches an in-network

IP address that has been configured to bypass manual authentication, the remote user will not be prompted to authenticate manually when accessing the Internet.

One solution is to reconfigure the in-network IP address to use manual authentication. An alternative solution is to disable the manual authentication requirement for the affected remote user.

Remote users are not being filtered correctly

If remote users are not being filtered, or are not being filtered by particular policies assigned to them, check the RADIUS Agent logs for the message **Error receiving from server: 10060** (Windows) or **Error receiving from server: 0** (Linux).

This usually occurs when the RADIUS server does not recognize RADIUS Agent as a client (source of RADIUS requests). Ensure that your RADIUS server is configured properly (see *Configuring the RADIUS environment*, page 200).

You can use RADIUS Agent's built-in diagnostic tool to troubleshoot filtering problems (see *Troubleshooting RADIUS Agent*, page 342).

If you have implemented the Remote Filtering feature (see *Filter Remote Clients*, page 143), remote users cannot be filtered if the Remote Filtering Client cannot communicate with the Remote Filtering Server within the network.

For instructions on setting up Remote Filtering, see the *Remote Filtering* technical paper.

Block message issues

No block page appears for a blocked file type

When file type blocking is used, the block message may not always be visible to the user. For example, when a downloadable file is contained within an internal frame (IFRAME) on a permitted site, the block message sent to that frame is not visible because the frame size is zero.

This is only a display problem; users cannot access or download the blocked file.

Users receive a browser error instead of a block page

If users receive an error message instead of a block page, the 2 most likely causes are:

- The user's browser is configured to use an external proxy. In most browsers, there is a setting that enables use of an external proxy. Verify that the browser is not set to use an external proxy.
- There is a problem identifying or communicating with the Filtering Service machine.

If the user's browser settings are correct, make sure that the IP address of the Filtering Service machine is listed correctly in the **eimserver.ini** file.

1. Stop **Websense Filtering Service** (see *Stopping and starting Websense services*, page 258).
2. Navigate to the Websense **bin** directory (C:\Program Files\Websense\bin or /opt/Websense/bin, by default).
3. Open the **eimserver.ini** file in a text editor.
4. Under [WebsenseServer], add a blank line, and enter the following:

   ```
   BlockMsgServerName = <Filtering Service IP address>
   ```

   For example, if the Filtering Service IP address is 10.201.72.15, enter:

   ```
   BlockMsgServerName = 10.201.72.15
   ```

5. Save and close the file.
6. Restart Filtering Service.

If the Filtering Service machine has more than one NIC, and the block page still does not display correctly after editing the **eimserver.ini** file, try the IP addresses of the other NICs in the **BlockMsgServerName** parameter.

If the block page still does not appear, make sure that users have read access to the files in the Websense block page directories:

- Websense\BlockPages\en\Default
- Websense\BlockPages\en\Custom

If the block page problem persists, see the Websense Knowledge Base for additional troubleshooting hints.

A blank white page appears instead of a block page

When advertisements are blocked, or when a browser does not correctly detect the encoding associated with a block page, users may receive a blank white page instead of a block page. The reasons for this behavior are as follows:

- When the Advertisements category is blocked, Websense software sometimes interprets a request for a graphic file as an advertisement request, and displays a blank image instead of a block message (the normal method for blocking advertisements). If the requested URL ends in .gif or similar, have the user reenter the URL, leaving off the *.gif portion.
- Some older browsers may not detect the encoding of block pages. To enable proper character detection, configure your browser to display the appropriate character set (UTF-8 for French, German, Italian, Spanish, Brazilian Portuguese,

Simplified Chinese, Traditional Chinese, or Korean; and Shift_JIS for Japanese). See your browser's documentation for instructions, or upgrade the browser to a newer version.

Protocol block messages don't appear as expected

Protocol block messages may not appear, or appear only after a delay, for any of the following reasons:

- User Service must be installed on a Windows machine in order for protocol block messages to display properly. For more information, see the *Installation Guide*.
- Protocol block messages may not reach client machines if Network Agent is installed on a machine with multiple network interface cards (NICs), and a NIC is monitoring a different network segment from Filtering Service. Ensure that the Filtering Service machine has NetBIOS and Server Message Block protocol access to client machines, and that port 15871 is not blocked.
- A protocol block message may be slightly delayed, or appear on an internal machine where the requested protocol data originated (instead of on the client machine), when Network Agent is configured to monitor requests **sent to** internal machines.
- If the filtered client or the Websense filtering machine is running Windows 200x, the Windows **Messenger** service must be running for the protocol block message to display. Use the Windows Services dialog box on the client or server machine to see if the Messenger service is running (see *The Windows Services dialog box*, page 361). Even though the block message does not appear, protocol requests are still blocked.

A protocol block message appears instead of a block page

If your integration product does not send HTTPS information to Websense software, or if Websense software is running in stand-alone mode, Network Agent may interpret an HTTPS site request that is blocked via category settings as a protocol request. As a result, a protocol block message is displayed. The HTTPS request is also logged as a protocol request.

Log, status message, and alert issues

Where do I find error messages for Websense components?

When there are errors or warnings related to core Websense components, short alert messages are displayed in the **Health Alert Summary** list at the top of the **Status > Today** page in Websense Manager (see *Websense Health alerts*, page 347).

- Click an alert message to see more detailed information on the **Status > Alerts** page.
- Click **Solutions** next to a message on the Status > Alerts page for troubleshooting assistance.

Errors, warnings, and messages from Websense software components, as well as database download status messages, are recorded in the **websense.log** file in the Websense **bin** directory (C:\Program Files\Websense\bin or /opt/Websense/bin, by default). See *The Websense log file*, page 362.

For Websense software components installed on Windows machines, you can also check the Windows Event viewer. See *The Windows Event Viewer*, page 361.

Websense Health alerts

The Websense Health Alert Summary lists any potential concerns encountered by monitored components of your Websense software. These include:

- A Filtering Service is not running
- User Service is not available
- A Log Server is not running
- There is no Log Server configured for a Policy Server
- The Log Database is not available
- Network Agent is not running
- There is no Network Agent configured for a Policy Server
- No monitoring NIC has been configured for a Network Agent
- No Filtering Service has been configured for a Network Agent
- The initial filtering database is in use
- The Master Database is downloading for the first time
- The Master Database is being updated
- The Master Database is more than 1 week old
- The Master Database did not download successfully
- WebCatcher is not enabled
- There is a subscription problem
- The subscription key is about to expire
- No subscription key has been entered

The Alerts page provides basic information about any error or warning condition. Click **Solutions** for information about addressing the problem.

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 266, for more information.

Two log records are generated for a single request

When Windows QoS Packet Scheduler is installed on the same machine as Network Agent, 2 requests are logged for each single HTTP or protocol request made from the Network Agent machine. (This duplication does not occur with requests made by client machines within your network.)

To fix the problem, disable Windows QoS Packet Scheduler on the Network Agent machine.

This problem does not occur if you use Network Agent for all logging. See *Configuring NIC settings*, page 316, for details.

Policy Server and Policy Database issues

I forgot my password

If you are a Super Administrator or delegated administrator using a Websense user account to log on to Policy Server via Websense Manager, any unconditional Super Administrator can reset the password.

- The WebsenseAdministrator password is set on the **Settings > Account** page.
- Other administrator account passwords are set on the **Delegated Administration > Manage Websense User Accounts** page.

If you are not using delegated administration, and have forgotten the WebsenseAdministrator password, log on to MyWebsense to reset the password.

- The subscription key associated with the MyWebsense account must match your current Websense Web Security or Websense Web Filter subscription key.
- If you have multiple subscription keys, you must select the appropriate Websense Web Security or Websense Web Filter key for the password reset process to succeed.
- You must have access to the Websense Manager machine to complete the reset process.

I cannot log on to Policy Server

Verify that the selected Policy Server IP address is correct. If the address of the Policy Server machine has changed since the Policy Server was added to Websense Manager, you will need to log on to a different Policy Server, remove the old IP address from Websense Manager, and then add the new Policy Server IP address. See *Adding and editing Policy Server instances*, page 251.

If Websense Manager has stopped suddenly, or has been stopped via the kill (Linux) or End Task (Windows) commands, wait a few minutes before logging on again. Websense software detects and closes the terminated session within 3 minutes.

The Websense Policy Database service fails to start

The Websense Policy Database runs as a special account: **WebsenseDBUser**. If this account experiences logon problems, the Policy Database is unable to start.

To address this issue, change the WebsenseDBUser password.

1. Log on to the Policy Database machine as a local administrator.
2. Go to **Start > Programs > Administrative Tools > Computer Management**.
3. In the navigation pane, under System Tools, expand **Local Users and Groups**, and then select **Users**. User information is displayed in the content pane.
4. Right-click **WebsenseDBUser** and select **Set Password**.
5. Enter and confirm the new password for this user account, and then click **OK**.
6. Close the Computer Management dialog box.
7. Go to **Start > Programs > Administrative Tools > Services**.
8. Right-click **Websense Policy Database** and select **Properties**.
9. On the Log On tab of the Properties dialog box, enter the new WebsenseDBUser password information, and then click **OK**.
10. Right-click Websense Policy Database again, and then select **Start**.

    When the service has started, close the Services dialog box.

Delegated administration issues

- *Managed clients cannot be deleted from role*, page 350
- *Logon error says someone else is logged on at my machine*, page 350
- *Some users cannot access a site in the Unfiltered URLs list*, page 350
- *Recategorized sites are filtered according to the wrong category*, page 350
- *I cannot create a custom protocol*, page 351

Managed clients cannot be deleted from role

Clients cannot be deleted directly from the managed clients list on the Delegated Administration >Edit Role page if:

- the administrator has applied a policy to the client
- the administrator has applied a policy to one or more members of a network, group, domain, or organizational unit

There may also be problems if, during Websense Manager logon, the Super Administrator chooses a different Policy Server than the one that communicates with the directory service containing the clients to be deleted. In this situation, the current Policy Server and directory service do not recognize the clients.

For assistance deleting managed clients, see *Deleting managed clients*, page 240.

Logon error says someone else is logged on at my machine

When you attempt to log on to Websense Manager you may sometimes receive the error “Logon failed. The role <role name> has been in use by <user name>, since <date, time>, on computer 127.0.0.1.” The IP address 127.0.0.1 is also called the loopback address, and typically indicates the local machine.

This message means that someone is logged on at the Websense Manager installation machine, in the same role you are requesting. You can select a different role (if you administer multiple roles), log on for reporting only, or wait until the other administrator logs off.

Some users cannot access a site in the Unfiltered URLs list

Unfiltered URLs affect only the clients managed by role in which the URLs are added. For example, if a Super Administrator adds unfiltered URLs, clients managed by delegated administration roles are not granted access to those sites.

To make the site available to clients in other roles, the Super Administrator can switch to each role and add the relevant sites to that role’s unfiltered URLs list.

Recategorized sites are filtered according to the wrong category

Recategorized URLs affect only the clients managed by role in which the URLs are added. For example, when a Super Administrator recategorizes URLs, clients managed by delegated administration roles continue to be filtered according to the Master Database category for those sites.

To apply the recategorization to clients in other roles, the Super Administrator can switch to each role and recategorize the sites for that role.

I cannot create a custom protocol

Only Super Administrators are able to create custom protocols. However, delegated administrators can set filtering actions for custom protocols.

When Super Administrators create custom protocols, they should set the appropriate default action for most clients. Then, inform delegated administrators of the new protocol so they can update the filters for their role, as appropriate.

Reporting issues

Log Server is not running

If Log Server is not running, or if other Websense components are unable to communicate with Log Server, Internet usage information is not stored, and you may not be able to generate Internet usage reports.

Log Server may be unavailable if:

- There is insufficient disk space on the Log Server machine.
- You changed the Microsoft SQL Server or MSDE password without updating the ODBC or Log Server configuration.

- It has been more than 14 days since the Master Database was downloaded successfully.
- The logserver.ini file is missing or corrupted.
- You stopped Log Server to avoid logging Internet usage information.

To troubleshoot the problem:

- Verify the amount of free disk space, and remove extraneous files, as needed.
- If you believe that a password change is the source of the problem, see *Updating the Log Server connection password*, page 355.
- Navigate to the Websense **bin** directory (C:\Program Files\Websense\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.
- Check the Windows Services dialog box to verify that Log Server has started, and restart the service if necessary (see *Stopping and starting Websense services*, page 258).
- Check the Windows Event Viewer and **websense.log** file for error messages from Log Server (see *Troubleshooting tools*, page 361).

No Log Server is installed for a Policy Server

Websense Log Server collects Internet usage information and stores it in the Log Database for use in investigative reports, presentation reports, and the charts and summaries on the Today and History pages in Websense 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 Websense Manager.
- Log Server is installed on a Linux machine.
- You are not using Websense reporting tools.

To verify that the correct Log Server IP address is set in Websense Manager:

1. Select the **Settings** tab of the left navigation pane, and then go to **General > Logging**.
2. Enter the IP address of the Log Server machine in the **Log Server IP address or name** field.
3. Click **OK** to cache your change, and then click **Save All**.

If Log Server is installed on a Linux machine, or if you are not using Websense reporting tools, you can hide the alert message in Websense Manager.

1. On the Main tab of the left navigation pane, go to **Status > Alerts**.
2. Under Active Alerts, click **Advanced**.
3. Mark **Hide this alert** for the “No Log Server installed” message.

4. Click **Save Now**. The change is implemented immediately.

Log Database was not created

Sometimes the installer cannot create the Log Database. The following list describes the most common causes and solutions.

Problem:	A file or files exist that use the names Websense software uses for the Log Database (wslogdb70 and wslogdb70_1), but the files are not properly connected to the database engine, so cannot be used by the Websense installer.
Solution:	Remove or rename the existing files, and then run the installer again.
Problem:	The account used to log on for installation has inadequate permissions on the drive where the database is being installed.
Solution:	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.
Problem:	There is insufficient disk space available to create and maintain the Log Database at the specified location.
Solution:	Clear enough space on the selected disk to install and maintain the Log Database. Then, run the installer again. Alternatively, choose another location.
Problem:	The account used to log on for installation has inadequate SQL Server permissions to create a database.
Solution:	Update the logon account or log on with an account that already has the required permissions. Then, run the installer again. The required permissions depend on the version of Microsoft SQL Server: ■ SQL Server 2000 or MSDE: **dbo** (database owner) permissions required ■ SQL Server 2005: **dbo** and **SQLServerAgentReader** permissions required

Log Database is not available

The Websense Log Database stores Internet usage information for use in presentation reports, investigative reports, and the charts and summaries on the Today and History pages in Websense 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 Desktop Engine [MSDE]) is running on the Log Database machine.

1. Open the Windows Services dialog box (see *The Windows Services dialog box*, page 361) and verify that the following services are running:

- Microsoft SQL Server:
  - MSSQLSERVER
  - SQLSERVERAGENT
- Microsoft SQL Desktop Engine (MSDE):
  - MSSQL$WEBSENSE (if you obtained MSDE from Websense, Inc.)
  - SQLAgent$WEBSENSE

2. If a service has stopped, right-click the service name and click **Start**.

   If the service does not restart, check the Windows Event Viewer (see *The Windows Event Viewer*, page 361) for Microsoft SQL Server or MSDE errors and warnings.

If the database engine is running:

- Make sure SQL Server Agent is running on the machine running the database engine.
- Use the Windows Services dialog box to make sure that the **Websense Log Server** service is running.
- If Log Server and the Log Database are on different machines, make sure that both machines are running, and that the network connection between the machines is not impaired.
- Make sure that there is enough disk space on the Log Database machine, and that the Log Database has a sufficient quantity of allocated disk space (see *Log Server is not recording data in the Log Database*, page 354).
- Make sure that the Microsoft SQL Server or MSDE password has not been changed. If the password changes, you must update the password information Log Server uses to connect to the database. See *Updating the Log Server connection password*, page 355.

Log Database size

Log Database size is always a concern. If you have been successfully generating Websense reports and notice the reports are now taking much longer to display, or you begin receiving timeout messages from your Web browser, consider disabling some database partitions.

1. In Websense Manager, go to **Settings > Reporting >Log Database**.
2. Locate the **Available Partitions** section of the page.
3. Clear the **Enable** check box for any partitions that are not required for current reporting operations.
4. Click **Save Now** to implement the change.

Log Server is not recording data in the Log Database

Usually, when Log Server is unable to write data to the Log Database, the database has run out of allocated disk space. This can occur either when the disk drive is full, or

in the case of Microsoft SQL Server, if there is a maximum size set for how large the database can grow.

If the disk drive that houses the Log Database is full, you must add disk space to the machine to restore logging.

If your SQL Server Database Administrator has set a maximum size for how large an individual database within Microsoft SQL Server can grow, do one of the following:

- Contact your SQL Server Database Administrator to increase the maximum.
- Find out the maximum size, and go to **Settings > Reporting >Log Database** to configure the Log Database to roll over when it reaches approximate 90% of the maximum size. See *Configuring rollover options*, page 295.

If your Information Technology department has established a maximum amount of disk space for SQL Server operations, contact them for assistance.

Updating the Log Server connection password

If you change the password for the account that Websense software uses to connect to the Log Database, you must also update Log Server to use the new password.

1. On the Log Server machine, go to **Start > Programs > Websense >Utilities > Log Server Configuration**. The Log Server Configuration utility opens.
2. Click the **Database** tab, and then verify that correct database (by default, **wslogdb70**) appears in the ODBC Data Source Name (DSN) field.
3. Click **Connection**. The Select Data Source dialog box opens.
4. Click the **Machine Data Source** tab, and then double-click **wslogdb70** (or your Log Database name). The SQL Server Login dialog box opens.
5. Make sure that the LoginID field contains the correct account name (usually **sa**), and then enter the new password.
6. Click **OK**, and then, in the Log Server Configuration dialog box, click **Apply**.
7. Click the **Connection** tab, and then stop and restart Log Server.
8. When Log Server is running again, click **OK** to close the utility.

Configuring user permissions for Microsoft SQL Server 2005

Microsoft SQL Server 2005 defines SQL Server Agent roles that govern accessibility of the job framework. The SQL Server Agent jobs for SQL Server 2005 are stored in the SQL Server msdb database.

To install Websense Log Server successfully, the user account that owns the Websense database must have membership in one of the following roles in the msdb database:

- SQLAgentUserRole
- SQLAgentReader Role

- SQLAgentOperator Role

> **Note**
> The SQL user account must also be a member of the **DBCreator** fixed server role.

Go to Microsoft SQL Server 2005 to grant the SQL Server user account the necessary permissions to successfully install the Websense reporting components.

1. On the SQL Server machine, go to **Start > Programs > Microsoft SQL Server 2005 > Microsoft SQL Server Management Studio**.
2. Select the **Object Explorer** tree.
3. Select **Security > Logins**.
4. Select the login account to be used during the installation.
5. Right-click the login account and select **Properties** for this user.
6. Select **User Mapping** and do the following:
   a. Select **msdb** in database mapping.
   b. Grant membership to one of these roles:
      - SQLAgentUserRole
      - SQLAgentReader Role
      - SQLAgentOperator Role
   c. Click **OK** to save.
7. Select **Server Roles**, and then select **dbcreator**. The dbcreator role is created.
8. Click **OK** to save.

Log Server cannot connect to the directory service

If either of errors listed below occurs, Log Server is unable to access the directory service, which is necessary for updating user-to-group mappings for reports. These errors appear in the Windows Event Viewer (see *The Windows Event Viewer*, page 361).

- 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 Websense Log Server and Websense User Service are on different sides of a firewall that is limiting access.

To resolve this problem, configure the firewall to permit access over the ports used for communication between these components.

Data on Internet browse time reports is skewed

Be aware that consolidation may skew the data for Internet browse time reports. These reports show the time users spend accessing the Internet and can include details about the time spent at each site. Internet browse time is calculated using a special algorithm, and enabling consolidation may skew the accuracy of the calculations for these reports.

Bandwidth is larger than expected

Many, but not all, Websense integrations provide bandwidth information. If your integration does not provide bandwidth information, you can configure Network Agent to perform logging so that bandwidth data is included.

When a user requests a permitted file download, the integration product or Network Agent sends the full file size, which Websense software logs as bytes received.

If the user subsequently cancels the actual download, or the file does not download completely, the bytes received value in the Log Database still represents the full file size. In these circumstances, the reported bytes received will be larger than the actual number of bytes received.

This also affects reported bandwidth values, which represent a combination of bytes received and bytes sent.

Some protocol requests are not being logged

A few protocols, such as those used by ICQ and AOL, prompt users to log into a server using one IP address, and then send a different identifying IP address and port number to the client for messaging purposes. In this case, all messages sent and received may not be monitored and logged by the Websense Network Agent, because the messaging server is not known at the time messages are exchanged.

As a result, the number of requests logged may not match the number of requests actually sent. This affects the accuracy of reports produced by Websense reporting tools.

All reports are empty

If there is no data for any of your reports, make sure that:

- The active database partitions include information for the dates included in the reports. See *Database partitions*, page 358.
- The SQL Server Agent job is active in Microsoft SQL Server or MSDE. See *SQL Server Agent job*, page 358.
- Log Server is correctly set up to receive log information from Filtering Service. See *Log Server configuration*, page 358.

Database partitions

Websense log records are stored in partitions within the database. New partitions may be created based on size or date, depending on your database engine and configuration.

You can activate or deactivate individual partitions in Websense Manager. If you attempt to generate report based on information stored in deactivated partitions, no information is found and the report is empty.

To make sure the appropriate database partitions are active:

1. Go to **Settings > Reporting > Log Database**.
2. Scroll down to the **Available Partitions** section.
3. Mark the **Enable** check box for each partition that contains data to be included on the reports.
4. Click **Save Now** to implement the change.

SQL Server Agent job

It is possible that the SQL Server Agent database job has been disabled. This job must be running for the log records to be processed into the database by the ETL database job.

If you are running with MSDE:

1. Go to **Start > Administrative Tools > Services**.
2. Make sure that both the SQL Server and SQL Server Agent services are started. If you obtained MSDE from Websense, Inc., these services are called MSSQL$WEBSENSE and SQLAgent$WEBSENSE.

If you are running full Microsoft SQL Server, ask your Database Administrator to make sure the SQL Server Agent job is running.

Log Server configuration

Configuration settings must be correct in both Websense Manager and Log Server to make sure that Log Server receives log information from Filtering Service. Otherwise, log data is never processed into the Log Database.

First, verify that Websense Manager is connecting to the Log Server successfully.

1. Log on to Websense Manager with unconditional Super Administrator permissions.
2. Go to **Settings > General > Logging**.
3. Enter the **machine name** or **IP address** where the Log Server is located.
4. Enter the **port** that Log Server is listening on (the default is 55805).
5. Click **Check Status** to determine whether Websense 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 All**.

Next, verify the settings in the Log Server Configuration utility.

1. On the machine where Log Server is running, go to **Start > Programs > Websense > Utilities > Log Server Configuration**.
2. On the **Connections** tab, verify that the Port matches the value entered in Websense Manager.
3. Click **OK** to save any changes.
4. Use the button on the **Connections** tab to stop and then start Log Server.
5. Click **Quit** to close the Log Server Configuration utility.

No charts appear on Today or History pages

In organizations that use delegated administration, review the reporting permissions for the delegated administrator's role. If **View reports on Today and History pages** is not selected, these chart do not appear for delegated administrators in that role.

In environments that use multiple Policy Servers, the Log Server is installed to communicate with only one Policy Server. You must log on to that Policy Server to view charts on the Today and History pages, or to access other reporting features.

Cannot access certain reporting features

If your Web browser has pop-up blocking at a very strict setting, it may block certain reporting features. To use those features, you must decrease the blocking level or disable pop-up blocking entirely.

Microsoft Excel output is missing some report data

The largest number of rows that can be opened in a Microsoft Excel worksheet is 65,536. If you export a report with more than 65,536 records to Microsoft Excel format, the 65,537th and all following records are not available in the worksheet.

To assure access to all information in the exported report, do one of the following:

- For presentation reports, edit the report filter to define a smaller report, perhaps by setting a shorter date range, selecting fewer users and groups, or selecting fewer actions.
- For investigative reports, drill down into the data to define a smaller report.
- Select a different export format.

Saving presentation reports output to HTML

If you generate a report directly from the Reporting > Presentation Reports page, you can choose from 3 display formats: HTML, PDF, and XLS. If you choose the HTML display format, you can view the report in the Websense Manager window.

Printing and saving presentation reports from the browser are not recommended. The printed output includes the entire browser window, and opening a saved file launches Websense Manager.

To print or save reports more effectively, choose PDF or XLS as the output format. You can open these file types immediately if the viewing software (Adobe Reader or Microsoft Excel) is installed on the local machine. You also can save the file to disk (the only option if the proper viewing software is not available).

After you open a report in Adobe Reader or Microsoft Excel, use that program's print and save options to produce the desired final output.

Investigative reports search issues

There are two potential concerns related to searching investigative reports.

- Extended ASCII characters cannot be entered
- Search pattern may not be found

Extended ASCII characters

The Search fields above the bar chart on the main investigative reports page let you search for a specific term or text string in the chart element that you select.

If you are using Mozilla Firefox on a Linux server to access Websense Manager, you cannot enter extended ASCII characters in these 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 Websense Manager from a Windows server, using any supported browser.

Search pattern not found

Sometimes, investigative reports is unable to find URLs associated with a pattern entered in the Search fields on the main investigative reports page. If this occurs, and you are reasonably certain that the pattern exists within the URLs reported, try entering a different pattern that would also find the URLs of interest.

General investigative reports issues

- Some queries take a very long time. You may see a blank screen or get a message saying that your query has timed out. This can happen for the following reasons:
  - Web server times out

- MSDE or 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**.

Troubleshooting tools

- *The Windows Services dialog box*, page 361
- *The Windows Event Viewer*, page 361
- *The Websense log file*, page 362

The Windows Services dialog box

On Microsoft Windows machines, Filtering Service, Network Agent, Policy Server, User Service, and all Websense transparent identification agents run as services. You can use the Windows Services dialog box to check the status of these services.

1. In the Windows Control Panel, open the **Administrative Tools** folder.
2. Double-click **Services**.
3. Scroll through the list of services to find the service you are troubleshooting.

   The service entry includes the service name, a brief service description, the service status (started or stopped), how the service starts, and what account the service uses to perform its tasks.
4. Double-click a service name to open a properties dialog box with more detailed information about the service.

The Windows Event Viewer

The Windows Event Viewer records error messages about Windows events, including service activities. These messages can help you identify network or service errors that may be causing Internet filtering or user identification problems.

1. In the Windows Control panel, open the **Administrative Tools** folder.
2. Double-click **Event Viewer**.
3. In the Event Viewer, click **Application** for a list of error messages, warnings, and informational messages.
4. Scroll through the list to identify errors or warnings from Websense services.

The Websense log file

Websense software writes error messages to the **websense.log** file, located in the Websense **bin** directory (C:\Program Files\Websense\bin or /opt/Websense/bin, by default).

The information in this file is comparable to that found in the Windows Event Viewer. In Windows environments, the Event Viewer presents messages in a more user-friendly format. The **websense.log** file, however, is available on Linux systems, and can be sent to Websense Technical Support if you need help troubleshooting a problem.

Index

B

C

D

E

F

M

T

U

V

W

X