Web Protection Frequently Asked Questions

TRITON AP-WEB and Web Filter & Security | v8.2.x, v8.3.x | 16-Dec-2016

This collection includes answers to the following frequently asked questions about TRITON AP-WEB and Web Filter & Security:

For answers to common reporting questions, see the Web Protection Reporting FAQ.

How is a policy or exception assigned to a request?

Web protection **policies** are used to determine how to respond to a user request for a website or Internet application. **Exceptions** allow administrators to identify specific websites that should be handled in a different manner than is defined in policies.

Exceptions and policies can be applied to directory clients (user, group, or OU) or computer and network clients (individual IP addresses or IP address ranges). It is therefore possible to have a policy for all of the following:

- The user making the request
- The group or groups to which the user is assigned
- The specific IP address from which a request originates
- The IP address range from which the request originates

In situations where there are multiple exceptions or policies as described above, the service handling the request uses a **precedence order** to figure out the most applicable exception or policy:

- When Filtering Service (an on-premises component) is used to respond to a request, one of two precedence orders is used:

  User > Computer > Network > Group > OU (*default*)

  User > Group > OU > Computer > Network

  For more information about configuring the Filtering Service precedence order, see "Prioritizing group and domain policies" (version 8.3 or version 8.2).
- The hybrid service always uses the following policy precedence order:

  User > Group > OU > IP address (filtered location)

Exceptions take precedence over policies. The general rules for determining which exception to apply are:

- Super Administrator exceptions take precedence over exceptions created by delegated administrators, *unless* the Super Administrator has configured an option to allow delegated administrator exceptions take precedence.
- Exceptions that apply to one or more individual clients take precedence over exceptions applied to an entire delegated administrator role.
- If multiple equivalent exceptions could be applied (for example, 2 Super Administrator exceptions applied to the same group):
  - Blocked takes precedence over permit.
  - If there are multiple blocked exceptions, the first one found is applied.
  - If there are multiple permitted exceptions and no blocked exceptions, the first permitted exception found is applied.

If no applicable exceptions are found, the service determines which policy to apply:

- When on-premises components respond to a request, by default, a computer or network policy takes precedence over a group policy.
- When the hybrid service enforces policy, a group policy takes precedence over a computer or network policy.
- A policy assigned to a computer (single IP address) takes precedence over a policy assigned to a network (IP address range).
- If multiple group policies apply to the same user, and no higher-priority policy applies, precedence is applied based on the **Use most restrictive group policy** setting set on the **Settings > General > Filtering** page in the Web module of the TRITON Manager.
  - If the option is selected, the request is blocked if **any** of the applicable policies blocks the URL category.
  - If the option is **not** selected, the request is permitted if **any** of the applicable policies permits the URL category.
  - If all groups have the same policy, that policy is used.
- Custom protocols take precedence over pre-defined protocols.
- Custom categorization take precedence over pre-defined categories.
- The **Manage Role Priority** option in Delegated Administration sets precedence when a user is in multiple groups managed by different delegated administrator roles.
- If no other policy is found, the Default policy is applied.

How do I know which policy is being applied to a client's requests?

When multiple exceptions or policies might apply to a client's request, web protection software uses the rules described in *How is a policy or exception assigned to a request?*, page 2, to determine which to apply.

If you aren't sure which exception or policy is currently being applied to requests from a specific client, you can use the Check Policy tool to find out.

1. Log on to the Web module of the TRITON Manager.
2. Click **Check Policy** in the Toolbox in the right navigation pane.
3. To identify a client, enter one of the following:
   - A fully qualified user name

     If your organization uses an LDAP-based directory service, you can also click **Find User** to search the directory.
   - An IP address
4. Click **Go**.

The tool displays the name of one or more policies. The tool returns multiple policies when all of the following are true:

- The user belongs to multiple groups or OUs
- Different policies are assigned to each group or OU
- No policy is assigned specifically to the user

This can also occur when an IP address is included in more than one network range.

If the Check Policy tool returns an unexpected result, and client requests are being blocked, you can use information provided in a block page to determine what policy is being applied and how the user has been identified.

If you use either the default block page or a custom block page that includes the More Information option, you can find out which policy was applied to a request, and get category and client information.

1. Have the user browse to a blocked URL.
2. Click the **More Information** link or button.
3. When the category information displays at the top of the block page, right-click within the section containing the category information and select **View Source** (Internet Explorer) or T**his Frame > View Frame Source** (Mozilla Firefox).
4. Scroll down to the bottom of the HTML output. The information includes:
   - User name and IP address

     Use this information to verify the user was identified correctly.
   - The policy that was applied

- The delegated administration role associated with the policy
- How the URL was assigned to a category

If a customized block page does not contain the **More Information** option, you can still generate a blocked request, then change the block URL to retrieve the additional information. The block page URL looks something like this:

```
http://<ipaddress>:15871/cgi-bin/blockpage.cgi?ws-
session=.....
```

Replace **blockpage** with **moreBlockInfo** as follows

```
http://<ipaddress>:15871/cgi-bin/moreBlockInfo.cgi?ws-
session=.....
```

Follow steps 3 and 4 in the procedure above to find user and policy information.

If the user is not correctly identified and there is no policy assigned to the user's IP address, the Default policy is used.

If the user is not being correctly identified, verify the directory path for the user or group object, especially if modifications have recently been made to directory service settings:

1. In the Web module of the TRITON console, go to **Main > Policy Management > Clients** page.
2. Expand the Directory tree and verify that the LDAP path displayed matches what is currently configured on the **Settings > General > Directory Service** page.
3. If necessary, delete, re-add, and save the user or group objects. When re-adding clients, make sure to assign them the correct policy.

What do I do when the wrong policy is being applied to requests?

When a user's requests are not managed by the expected policy, and you have confirmed that the user is being identified correctly, use these steps to help identify and resolve the problem.

1. Verify your subscription key.
   a. In the Web module of the TRITON console, go to the **Main > Status > Alerts** page and make sure there are no subscription-related alerts in the Health Alerts Summary.
   b. Navigate to the **Settings > General > Account** page and verify that your subscription key appears, the expiration date has not passed, and the number of subscribed network users is greater than 0.
2. Make sure the Master Database has downloaded successfully.
   a. Check for alerts on the **Status > Alerts** page.
   b. If no alerts appear, click **Database Download** in the toolbar at the top of the dashboard, and make sure all Filtering Service instances show a successful last download, and that all downloads happened within the last 2 weeks (14 days).

      If there are any messages, or if the database is outdated, click **Update** to initiate a manual update.
3. Check for Filtering Service errors.
   - Look for Filtering Service alerts on the **Status > Alerts** page. Click the **Solutions** link next to any alert for troubleshooting steps.
   - If Filtering Service resides on a Windows machine, check the Event Viewer (Start > Administrative Tools > Event Viewer or Server Manager > Tools > Event Viewer) for **Websense Filtering Service** errors.
   - Check the **websense.log** file in the **bin** directory (C:\Program Files\Websense\Web Security\bin or /opt/Websense/bin/, by default) for **EIMServer** (Filtering Service) errors.
4. Use the TestLogServer utility to verify that Filtering Service is receiving URL requests. For instructions, see Using TestLogServer for Web Protection Troubleshooting.
   - If Filtering Service is not receiving Internet traffic, verify that Content Gateway, Network Agent, or your third-party integration product has been properly configured to communicate with Filtering Service.
   - If you have a standalone Web Filter & Security deployment, verify that Network Agent is able to see all traffic (incoming and outgoing) and that port spanning is configured.
5. Run the **WebsensePing** utility to see what happens when a user requests a site.

a. Open a command prompt on the Filtering Service machine and navigate to the appropriate directory (C:\Program Files\Websense\Web Security\bin or /opt/Websense/, by default).

b. Enter one of the following commands:

Windows (TRITON AP-WEB):

```
websenseping -m 18 -user <username> -url <URL>
websenseping -m 18 -uip <IPaddress> -url <URL>
```

Windows (Web Filter & Security):

```
websenseping -m 8 -user <username> -url <URL>
websenseping -m 8 -uip <IPaddress> -url <URL>
```

Linux (TRITON AP-WEB):

```
./WebsenseTools -p -m 18 -user <username> -url <URL>
./WebsenseTools -p -m 18 -uip <IPaddress> -url <URL>
```

Linux (Web Filter & Security):

```
./WebsenseTools -p -m 8 -user <username> -url <URL>
./WebsenseTools -p -m 8 -uip <IPaddress> -url <URL>
```

Here, <username> is the name of the user and <IPaddress> is the client IP address, depending on whether the policy is user-based or IP address-based.

A user name can be entered in Windows NT format (winNT://Test/jdoe) or LDAP format (LDAP://GC OU=Technical Support,OU=US Technical Services,DC=Test,DC=com/John Doe).

Both user name and client IP address can be entered in the same command to help make sure the information provided by WebsensePing is based on the policy that would be applied.

c. Review the output of the command to determine what action (disposition) would be applied and confirm the category of the URL is what you expected.

6. Verify that connections to the client and origin server are being closed by running a packet capture on the Filtering Service machine and on the client.

7. Refresh the Filtering Service user/group cache. By default, Filtering Service caches user and group information for 2 hours. The cache needs to be updated when any changes are made to users or groups.

   To update the cache, go to the **Settings > General > Directory Services** page in the Web module of the TRITON console, then click **Clear Cache**.

8. Make sure the client machine can communicate with the Filtering Service machine.

   a. On the client machine, open a Command Prompt and **ping** the Filtering Service machine.

   b. If the ping succeeds, on the Filtering Service machine, make sure that Filtering Service (EIMServer) is listening on port **15871**.

      ```
      netstat -an |find "15871"
      ```

c. From the client, open a **telnet** session to the Filtering Service machine on port 15871.

```
telnet <ip_address> 15871
```

If telnet fails, ensure there are no local firewalls or devices between the client and Filtering Service that are blocking the port.

9. To make sure the client machine can receive a block page, go to the client machine and enter the following URL:

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

- If you see an **Invalid Request** message, Filtering Service is active and listening. This means that the client can reach Filtering Service but there may be DNS issues.
- If you see a **Page Cannot be Displayed** message, there are connectivity issues between the machines.

How do keywords and regular expressions work?

Keywords and regular expressions are ways to identify types or groups of related URLs for special handling.

- A **keyword** is a string of characters (a word, phrase, number, or acronym) that might be found in a URL.

  When you create a keyword, you associate it with a category. When your software finds that keyword in a URL, it assigns the URL to the category that you have selected, and blocks the URL.

  Keywords are used only to recategorize and block URLs; they cannot be used to permit URLs.

- A **regular expression** is a template or pattern used to match multiple strings, or groups of characters. Regular expressions can be used:

  - To identify groups of related URLs in limited access filters and exceptions
  - As a more flexible form of keyword, to assign URLs to categories for blocking

  When you use regular expressions, Filtering Service tries to match the general pattern, rather than a specific, single URL or keyword. Regular expressions match precise terms and are case sensitive.

Use keywords and regular expression with caution to avoid unintended over blocking or under blocking.

More about keywords

There are two steps involved in using keywords:

1. Define specific keywords and associate them with categories.
2. Turn on keyword-based blocking in some or all of your policies.

Once keyword-based blocking is enabled, web protection software tries to match the keyword against each requested URL as follows:

- If the keyword contains only ASCII characters, the keyword is matched against the domain, path, and query (CGI) portions of a URL. The match is case independent.

  For example, if you associated the keyword “nba” with the permitted Sports category, the following URLs are blocked:

  - sports.espn.go.com/**nba**/
  - moder**nba**kery.com
  - fashio**nba**r.com

- If the keyword contains characters outside the ASCII character set, the keyword is matched against only the path and query (CGI) portions of the string. The match is case independent.

  For example, if you associated the keyword "fútbol" with the permitted Sports category:

  - "www.**fútbol**.com" is **permitted** (the domain portion of the URL is not matched).
  - "es.wikipedia.org/wiki/**Fútbol**" is **blocked** (the path portion of the URL is matched).

When web protection software identifies a keyword in a URL:

- The URL is recategorized according to the keyword match.
- Reports show the keyword category, rather than the Master Database category, for the URL.
- The block page the user receives shows that the URL was blocked by keyword.

For more information on using and defining keywords see your *Administrator Help* (version 8.3 or version 8.2).

More about regular expressions

Regular expressions can be as simple or complex as your environment requires, but **use them with care**. Poorly constructed regular expressions can result in excessive filtering overhead.

As with keywords, if non-ASCII characters appear in a regular expression, the expression is matched against only the path and query (CGI) strings in the URL. No matching is done against the domain.

If only ASCII characters appear in the regular expression, the match is performed against the entire URL.

Most Perl regular expression syntax is supported.

For further help with regular expressions, see:

en.wikipedia.org/wiki/Regular_expression

www.regular-expressions.info/

Can I exclude specific traffic from logging?

With web protection solutions, you can prevent requests for specified categories or protocols from being logged. You can also configure Network Agent so that internal (intranet) traffic is not logged.

When you omit categories and protocols from logging:

- Requests for the categories or protocols are stilled managed by policies (permitted, blocked, and so on).
- No record of any requests for the categories is included in the Log Database, so the traffic does not appear in reports.
- Usage alerts can **not** be generated for the non-logged categories and protocols.

Excluding categories from logging

Category logging settings are global, affecting your entire software deployment.

1. Go to the **Settings** > **General** > **Logging** page in the Web module of the TRITON console.
2. Use the Selective Category Logging list to identify all categories that should not be logged. By default, requests are logged for all categories.
   - Expand parent categories to configure subcategories.
   - Clear the check box next to a category name to stop logging the category.
   - You must select or deselect each category separately. Selecting a parent category does not automatically select its subcategories. Use **Select All** and **Clear All** to assist with selections.
3. Click **OK** to cache your changes. Changes are not implemented until you click **Save and Deploy**.

To exclude a specific URL from logging, add it to a custom category, then omit that category from logging as described above. See “Editing categories and their attributes” (version 8.3 or version 8.2) in the *Administrator Help*.

Disabling a protocol from being logged

When you omit a protocol from being logged, as described below, the change initially affects all policies and filters. You can, however, override the default configuration in individual protocol filters. When you do this, requests for the protocol are logged only when the active filter specifies that logging is enabled.

To disable logging for a protocol by default:

1. Go to the **Main > Policy Management** > **Filter Components** page in the Web module of the TRITON console and click **Edit Protocols**.
2. Select a protocol in the list.
3. Click **Override Action**.
4. If the protocol is currently set for logging by default, click **Change Settings**, then clear the **Log protocol data** check box.
5. Click **OK** to return to the **Edit Protocols** page. When you are finished making changes, click **OK** again to cache the changes, then click **Save and Deploy** to implement them.

The changes that you made are applied to all active protocol filters in your delegated administration role. Administrators can override the change in individual protocol filters.

See "Editing custom protocols" (version 8.3 or version 8.2) in the *Administrator Help* for more information.

Ensuring that internal traffic is not logged

Network Agent is generally configured to ignore traffic between machines in your network, and monitor only traffic that leaves your network (Internet requests). To do this, Network Agent needs to know which machines are part of your network.

1. Go to the **Settings > Network Agent > Global** page in the Web module of the TRITON console.
2. Check the IP address ranges listed in the **Ignore Internal Traffic** list.
   - Click an IP address or range to edit it.
   - Click **Add** to add missing IP addresses or ranges to the list.
   - Click **Delete** to remove entries.
   - If the client machines whose traffic you don't want logged do not have a static IP address, ensure that they resides in a DHCP range that can be added to this list.

   Network Agent ignores traffic between these machines, monitoring only traffic that leaves the defined network.
3. If you want to monitor, block, or permit traffic to some internal machines, add those IP addresses to the **Internal Traffic to Monitor** list.
   - By default, this traffic is both monitored and logged.
   - If you want to be able to block traffic to these machines, but don't want the blocked requests logged, you can configure that later.
4. When you are finished making changes, click **OK**, and then click **Save and Deploy** to implement the change.

If you are managing traffic to one or more of your internal machines, add the IP addresses of those machines (or the URLs used to access the machines) to a custom category. Then, use the **Settings > General > Logging** page to exclude the custom categories from logging (see *Excluding categories from logging*, page 11).

How do I create exceptions and how do they work?

Exceptions override policies to permit or block access to specific URLs for specific clients. Use an exception to:

- Permit access for all employees to a specific URL that would otherwise be blocked based on its category.
- Block all clients in a specific role from a suspicious uncategorized URL while the site is investigated.
- Permit temporary access to a URL for certain members of a group while continuing to block the site for the rest of the group.
- Permit access to URLs only when accessed from specific sites (referers). See "What is a Referer?" (version 8.3 or version 8.2) in the *Administrator Help* for additional information.

Exceptions allow a flexible and rapid response to user requests, changes in company policies, spikes in Internet activity, or other changes in circumstance.

Use the **Policy Management > Exceptions** page to add, review, edit, or delete exceptions. The page shows the following information for each exception:

- The **Name** given to the exception when it was created.
- The **URL** or regular expression permitted or blocked by the exception, or a link to the complete list of URLs or expressions (if the exception includes more than one).

  If the exception is defined with referer information and no URLs or regular expressions were specified in the exception, text in this column explains that access to all URLs is permitted if accessed from a referer site.
- The URL of the approved **Referer** or a link to a complete list of referer URLs assigned to the exception (if one or more approved referers is included in the exception).
- The name of the **Client**, if the exception applies to one client, or the role name (if the exception applies to an entire delegated administration role), Global (if all clients are affected), or a link to a complete list of affected clients.
- The exception **Type**, an icon indicating whether URLs in the exception are blocked, permitted, or permitted with security override disabled.

  Additional hover over text is provided ("Permitted by referer" or "Permitted by referer, even when security risks are found") when approved referers are included in the exception.
- When the exception **Expires**, either an expiration date, or Never.
- Whether the exception is **Active** (used to block or permit requests) or Inactive (either expired, or just not currently in use).
- The date the exception was **Last Modified**.

Click **Add** below the list to create an exception.

Click the link that is the name of the exception or mark the check box next to an exception and click **Edit** to edit an existing exception.

See the *Administrator Help* (version 8.3 or version 8.2) for instructions for creating exceptions.

If multiple exceptions could apply to a request, how is the right one selected?

To apply exceptions, Filtering Service uses the following rules:

- By default, Super Administrator exceptions take precedence over exceptions created by delegated administrators.
- A delegated administrator exception takes precedence when the Super Administrator exception has been defined to allow delegated administrator override.
- If multiple equivalent exceptions could be applied:
  - Blocked takes precedence over permit.
  - If there are multiple blocked exceptions, the first one found is applied.
  - If there are multiple permitted exceptions and no blocked exceptions, the first permitted exception found is applied.
  - If there are multiple referer exceptions and no blocked exception, and one of the referer exceptions includes no specific URLs or regular expressions, the referer exception that lists the URL is applied.
- Client exceptions (that apply to one or more individual clients) take precedence over role exceptions.

Use the Test Filtering tool in (under Toolbox in the Web module of the TRITON Manager) to verify that client requests will be blocked or permitted as expected.

How do I create policies?

Policies govern Internet access and are made up of:

- Category filters, used to apply actions to URL categories
- Limited access filters, used to permit access to only a list of URLs
- Protocol filters, used to apply actions to Internet protocols
- A schedule used to determine when each filter is enforced.

Each new software installation includes a **Default** policy that controls Internet access for all clients not governed by another policy. It begins monitoring Internet usage as soon as you enter your subscription key.

As a best practice, edit the Default policy first, to set the baseline for Internet access at your organization. Next, create custom policies as needed to provide the levels of access needed for different groups in your organization.

To create a new policy:

1. Go to the **Policy Management > Policies** page in the Web module of the TRITON console and click **Add**.
2. 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.
3. 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 (,).
4. 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.
5. Click **OK** to cache your changes and go to the Edit Policy page.

In order for the policy to take effect, you must both apply it to clients (to determine who is governed by the policy) and click **Save and Deploy** to implement your changes.

For a quick but thorough introduction to creating and editing policies and applying them to clients, see the “Policy Management” section of the New Admin Quick Start tutorial (version 8.3 or version 8.2).

How do block pages work?

A block page is displayed in a user's browser when your software prevents access to a URL. The block page has 3 sections:

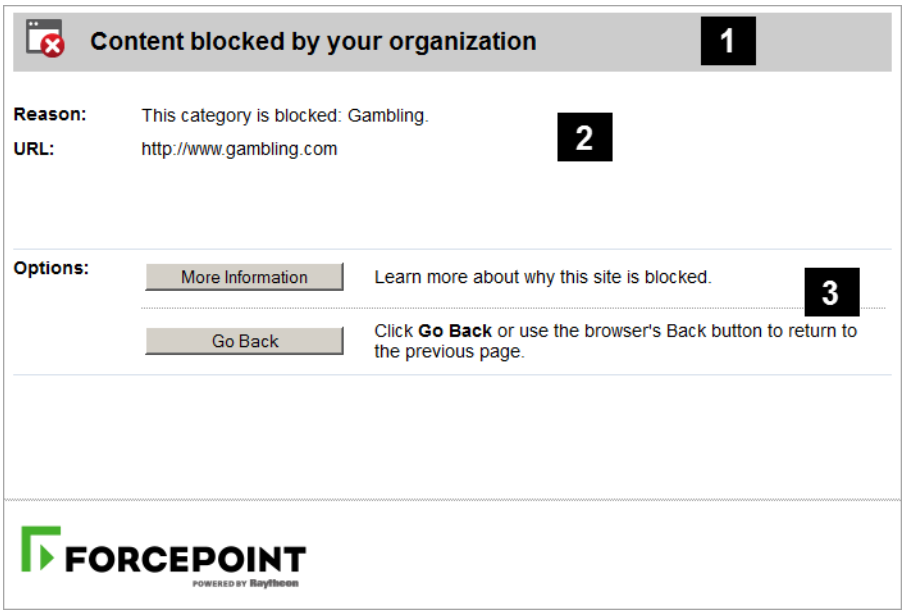

1. The **header** explains that the site has been blocked.
2. The **top frame** contains the actual block message, including the requested URL and the reason it was blocked.
3. The **bottom frame** provides any options that are available to the user, based on the way the site has been blocked. The user may be presented with an option to **Go Back** to the previous page, **Continue** to the site, or **Use Quota Time** to view the site.

A slightly different block page is presented if the URL is blocked because its category is in the Security Risk class. In this case, the header information indicates that a security risk has been blocked, with an explanation that the URL may pose a security threat.

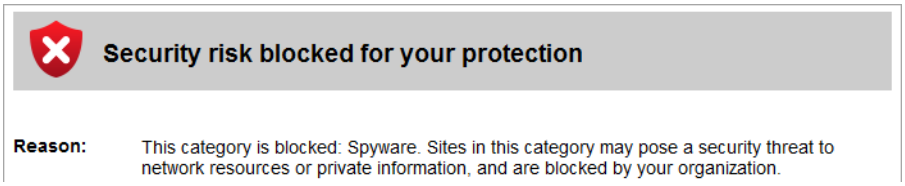

When a user clicks the **More Information** button, additional information is displayed on the block page to explain why the request was blocked.

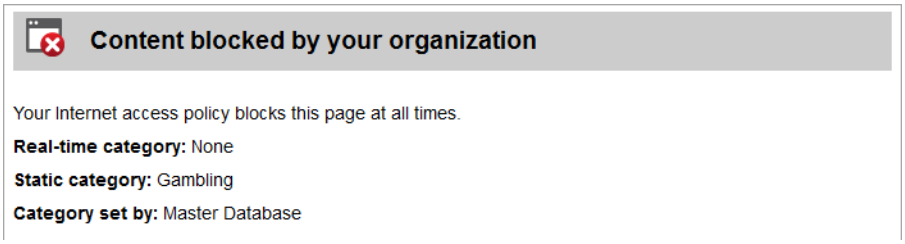

In addition to the visible details displayed on the More Information page, hidden information is added to the block page source code. Administrators can use this data to help with troubleshooting if a user's requests are being blocked unexpectedly.

To access the hidden information on the More Information page:

1. Right-click inside the top frame and select **View Source** (Internet Explorer) or **This Frame > View Frame Source** (Mozilla Firefox).
2. Scroll down to the bottom of the resulting HTML output. The information includes:
   - User name and IP address
   - The policy that was applied
   - The delegated administrator role associated with the policy
   - How the categorization was done

   For example:

   ```
   User name: LDAP://10.203.128.200 OU=UA,DC=ua-
   ux,DC=forcepoint,DC=com/Chinua Achebe Source IP address:
   10.203.128.45 Current time: 10:42
   ```

   ```
   This user receives policy: role-8**Security Only. The policy
   includes a category or limited access filter for the current
   time.
   ```

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

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

If a customized block page does not contain the **More Information** option, the URL that generates the page can be edited. A block page URL looks something like this:

```
http://<ipaddress>:15871/cgi-bin/blockpage.cgi?ws-
session=...
```

Replace **blockpage** with **moreBlockInfo**:

```
http://<ipaddress>:15871/cgi-bin/moreBlockInfo.cgi?ws-
session=...
```

The page that displays contains the information described in the procedure above.

Note that this information is viewable only if your browser supports iframes.

For information about customizing block pages, see Creating Custom Block Pages.

Why is the block page sometimes blank?

In some cases, a very small, blank image file (BlockImage.gif) is displayed instead of a standard or security block page. This happens when the Advertisements category is blocked, and a site tries to display an image (like a GIF or JPG file) hosted at a URL in the Advertisements category.

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

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

Why is only part of the block page visible?

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

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

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

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

How is quota time configured and used?

When you create category filters, you can apply the **Limit by Quota** action to categories to let users access URLs in those categories for a limited amount of time each day.

You can assign quota time in two ways:

- Define a global period of quota time (60 minutes, by default) on the **Settings > General > Filtering** page in the Web module of the TRITON console.

  When clients request a URL in a quota-limited category, they receive this amount of quota time, unless otherwise specified (see the next bullet).
- Assign a custom period of quota time to specific clients on the **Main > Policy Management > Clients > Edit Clients** page.

When users request a URL in a quota-limited category, they see a block page that offers a **Use Quota Time** button. If they click this button to continue to the website, the amount of time they spend browsing the site is monitored. If they request another URL in a quota-limited category:

- And their quota session has not ended, they see the new website.
- And their quota session has ended, they receive a new block page, either offering them the option to use additional quota time, or telling them that they have used all available quota time for the day.

The amount of time available in a quota session (10 minutes, by default) is configured on the **Settings > General > Filtering** page.

If your organization uses multiple Filtering Service instances, an additional component—State Server—may be required for correct application of quota time. See "Policy Server, Filtering Service, and State Server" in the *Administrator Help* (version 8.3 or version 8.2) for details.

What are filters, and how do they work?

Filters are one building block of web protection policies. The 3 types of filters define which categories, URLs, and protocols users can access. By adding filters to a policy, you define how clients at your organization can access the Internet.

- **Category filters** list:
  - All of the categories available to your organization, including both Master Database categories and custom categories.
  - The action (like permit, block, confirm, or quota) assigned to each category.

  Each category filter can assign a different action to each category. Your software includes 5 sample category filters that can be customized and used in policies, as well as a set of templates that you can use to create new filters.
- **Protocol filters** list:
  - All non-HTTP protocols, including both Master Database protocols and custom protocols.
  - The action (like permit, block, or limit by bandwidth) assigned to each protocol.

  Each filter can assign a different action to each protocol. Your software includes 3 sample protocol filters that can be customized and used in policies, as well as a set of templates that you can use to create new filters.

  The components required for protocol management vary based on your subscription level:
  - (*TRITON AP-WEB*) Content Gateway offers protocol management for protocols that tunnel over HTTP. It can be used in conjunction with Network Agent to provide full protocol management.

    The hybrid service does not enforce protocol filters.
  - (*Web Filter & Security*) Network Agent is required to enable protocol management.

  Web protection software can block TCP-based protocol requests, but not UDP-based protocol requests. If an application uses both TCP- and UDP-based messages, and the original network request is made via TCP, any subsequent data sent using UDP is blocked since the initial TCP request is blocked.
- **Limited access filters** are a restrictive list of permitted URLs that can be used in place of a category filter in web protection policies.

  When a limited access filter is in effect, users can visit only the URLs in the list. All other sites are blocked.

  If a URL that is permitted by a limited access filter becomes infected with malicious code, user requests to that URL are blocked as long as Security Risk categories are blocked in the Default policy. Check the category filter currently used by the Default policy to verify that all security-related categories are blocked.

See the *Administrator Help* for more information.

What are custom URLs, and how do they work?

With **custom URLs**, you can:

- Assign URLs that don't appear in the Master Database to a category.
- Change the category assigned to a URL.

The category that you assign to the URL is used for policy enforcement, and displayed in reports.

The Master Database is not consulted for a URL that has been recategorized, **unless** a permitted site has been infected with malicious code. In this case, by default, the URL is filtered according to the action applied to the security category that best describes the type of security risk. If the Default policy blocks that security category, the URL is blocked.

Create and manage custom URLs on the **Policy Management > Filtering Components > Edit Categories** page in the Web module of the TRITON console. You can:

- Add a new category and enter one or more URLs in the **Recategorized URLs** list as part of the category definition.
- Add one or more URLs to an existing category, using the **Recategorized URLs** list.

Enter each URL on a separate line.

To ensure that sites are managed correctly, enter both their URL and their IP address to the Recategorized URLs list. In addition:

- If a site can be accessed via multiple URLs, define each as a custom URL to ensure it is filtered as intended.
- Include the protocol for any non-HTTP site.
- If an HTTP redirect is used to send users to a new URL when a site is moved to a new domain, the new URL is not filtered the same way as the redirecting site. Create a new custom URL to ensure the site is filtered appropriately at its new address.
- Include the port number for HTTPS sites.
- Filtering Service performs a string match, recognizing URLs exactly as they are entered.

  For example, if the Search Engines and Portals category is blocked, but you recategorize www.yahoo.com in a permitted category, a user who accesses the site using images.search.yahoo.com or just yahoo.com will be blocked. If however, you recategorize yahoo.com, those sites will be permitted.

An implicit wildcard is assumed at the end of all recategorized URLs. That is, if www.domain2.com is added to a blocked category, the following will also be blocked:

- www.domain2.com/product
- www.domain2.com/services

- Regular expressions can be used to define a recategorized URL. See *How do keywords and regular expressions work?*, page 9.

Use the **URL Category** tool in the Toolbox on the Web tab of the TRITON Manager to verify that recategorized sites are being assigned to the correct category.

To see a list of custom URLs, you can:

- Use the **View all Custom URLs/Keywords** option on the **Edit Categories** page to navigate to a page that shows all custom URLs and their assigned category.
  1. Navigate to the **Policy Management > Filter Components > Edit Categories** page.
  2. At the top of the page, click **View All Custom URLs/Keywords**.
- Use the **Print Policies to File** option on the **Policies** page to generate a Microsoft Excel spreadsheet that includes custom URLs. This option gives a comprehensive list of all of your policy-related information, so the resulting file may be quite large.
  1. Navigate to the **Policy Management > Policies** page.
  2. At the top of the page, click **Print Policies to File**.
  3. Scroll to the bottom of the spreadsheet to see a list of recategorized URLs.
- Use the **See custom URLs in this category** option on the **Edit Category Filter** page to open a pop-up window that lists the recategorized URLs assigned to a selected category.
  1. Navigate to the **Policy Management > Filters** page.
  2. Select a category filter to open the **Edit Category Filter** page.
  3. Select a category that has been assigned to recategorized URLs and click the **See custom URLs in this category** link in the right column.

How do I back up and restore my policy and configuration data?

Your web protection software includes a command-line **Backup Utility** that makes it easy to back up your settings and policy data and to revert to a previous configuration. The backed up data can also be used to import configuration data after an upgrade.

The Backup Utility saves global configuration information stored in the Policy Database, local configuration information stored by each Policy Server, and specific component initialization and configuration files. The backup can be run manually at any time or scheduled to run at a regular interval.

Files created by the Backup Utility can be used to restore configuration data or revert to an earlier configuration. When restoring configuration data:

- Make sure you restore data for the components that exist on the machine.
- Do not restore file backups from a machine with one operating system/platform to a machine with a different operating system/platform.

See the *Backup and Restore FAQ* (version 8.3 or version 8.2) for full instructions for backing up and restoring web, data, and email protection solutions.

What is my subscription level and what happens if it's exceeded?

TRITON AP-WEB and Web Filter & Security are offered on a subscription basis.

- The subscription type (TRITON AP-WEB or Web Filter & Security) determines what features are available to configure and use.
- Subscriptions are issued on a per-client (IP address) basis.

After installation, the first time you log on to the management console (the TRITON Manager), you are prompted to enter your subscription key. This prompts your web protection solution to verify the key, confirm your subscription type, and start downloading the Master Database. This verification and download process enables policy enforcement.

Your software maintains a subscription table that keeps track of the number of clients managed 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.

The client count includes:

- Any client who makes an Internet (HTTP/HTTPS/FTP) request that is passed to your web protection software by Content Gateway, Network Agent, or a third-party integration product
- If Network Agent is used for protocol management, any client who makes a non-browser-based Internet request (for example, an IM connection)
- Any client machine running software that connects to the Internet

To help ensure that each computer is counted only once per day, use static IP addresses.

- If you have a DHCP environment, consider setting your IP address leases to last more than 1 day. An increase to the user count may occur if your current configuration assigns multiple IP addresses to a machine in a given day.
- Laptop computers using wireless connections may be counted multiple times.

When the number of clients listed in the table reaches the subscribed maximum, any previously-unlisted client that requests Internet access exceeds the subscription.

- In TRITON AP-WEB deployments, there is no change in policy enforcement. Full security protection capabilities are maintained even after the licensed IP levels are exceeded.
- In Web Filter & Security deployments, when the number of subscribed users is exceeded, requests from users who exceed the subscription count are permitted or blocked based on the setting **Block users when subscription expires**, found on the **Settings > General > Account** page in the Web module of the TRITON Manager.

In all deployments, if your subscription were to expire, all requests are permitted or blocked, depending on the same configurable setting. Whether requests are permitted or blocked depends on the **Block users when subscription expires** selection, configured on the **Settings > General > Account** page in the TRITON Manager. Note that expiration notices are provided in advance of a possible subscription expiration.

Occasionally, the subscription count may be set to zero. If that happens, try one of the following:

- Initiate a manual database update.
  1. Navigate to the **Status > Dashboard** page and click the **Database Download** button in the toolbar.
  2. Click **Update** for the appropriate Filtering Service IP address.
- Remove and re-add your subscription key on the **Settings > General > Accounts** page. This also initiates a database download.

For subscription-related troubleshooting information, see "Installation and Subscription Issues" in the *Administrator Help* (version 8.3 or version 8.2).