CHAPTER 23

Setting Global Traffic Rules

This chapter describes denial-of-service detection, content screening, and other protection mechanisms applicable to traffic at the ACE XML Gateway. It covers these topics:

About Global Traffic Policies

Global traffic policies improve the security and reliability of a system deployment. By default, global policies apply to all service traffic at the ACE XML Gateway, regardless of the consumer-side service interface the traffic addresses or backend service destination.

The types of global policies you can implement include:

- Traffic throttling rules designed to protect against denial of service attacks
- Content screening and replacement rules
- Large message handling
- User restrictions and misuse monitoring

Preventing Denial-of-Service Attacks

A denial-of-service (DOS) attack is one in which an attacker attempts to interrupt or hamper service by overwhelming a host with requests. The most common types of attack include:

- Sending so many messages so quickly that the service cannot respond to legitimate users' requests
- Sending message content designed to keep the system too busy processing to respond to other requests
- Sending messages that take a long time to produce a response
- Sending messages that generate so many errors that the system becomes bogged down in error-handling
- Sending numerous messages that produce authentication errors

The ACE XML Gateway can detect and block each of these types of attack using configurable parameters, with limits for:

- **Overall request rate**. The number of request messages received by the ACE XML Gateway.
- **Authentication failures**. The number of failed attempts to authenticate the sending consumer.
- **CPU usage**. The number of milliseconds of Gateway CPU time needed to process a message.
- **Internal errors**. The number of Gateway errors caused by processing a message.
- **Service latency**. The time in seconds required for the destination service to process a message.
- **Service errors**. The number of errors reported by the destination service when processing a message.

For each parameter, you can configure both a maximum allowed rate and a maximum allowed burst. The maximum rate is a limit on the number of events counted per minute. The maximum burst is a limit on the number of events detected at the same time.

The ACE XML Gateway monitors statistics associated with each of these parameters. When one of the parameters exceeds a configured threshold, the ACE XML Gateway logs an attack. If attack protection is enabled, the ACE XML Gateway rejects messages from the source for a configurable length of time.

By default, DOS protection blocks traffic from offending sources automatically, but you can disable this option if you prefer to allow the traffic. In that case, the ACE XML Gateway still logs the traffic so that you can monitor it.

You can also adjust the minimum length of time that the ACE XML Gateway blocks the source. The minimum blocking interval instructs the ACE XML Gateway to block attacking sources for no less than the configured number of seconds. The ACE XML Gateway may block a source for a longer time if it calculates that the severity of the attack calls for a longer time to recover from over-utilization.

The ACE XML Gateway applies DOS limits separately to each originating IP address. This means that if you set a limit of twenty messages per second, the ACE XML Gateway limits each individual IP address to twenty messages per second. In other words, if two different IP addresses send fifteen messages per second each, the ACE XML Gateway allows them, even though the aggregate rate is thirty per second, because no individual source has exceeded the limit. If one of the sources sends twenty-one messages per second, however, DOS protection blocks the traffic.

Understanding Traffic Rate Thresholds

Several parts of an ACE XML Gateway policy provide for traffic threshold settings, including the global denial-of-service configuration, user-identity settings, and throttling settings for HTTP servers.

In each case, the ACE XML Gateway determines the type of traffic activity that produces an excessive usage event from two configurable factors: the traffic-rate level and the maximum allowed burst level.

Generally, the traffic-rate level specifies the number of requests that are permitted over the span of a minute, while the burst level determines the number of requests that are permitted in a brief span of time. A burst can be thought of as a brief surge in traffic capacity. The increased capacity is, in effect, borrowed against the capacity that is available across a broader, encompassing time period.

What this means in terms of the actual traffic activity that triggers a threshold event and the length of time that constitutes a burst period is determined from the combination of the burst level and traffic rate.

To take a closer look at how this works, consider the traffic rate; its value determines how long the ACE XML Gateway “remembers” a message, for purposes of counting it against the available capacity. That is, if the request rate is set to 60 requests per second, the ACE XML Gateway remembers each message event for one second over the course of a minute. If two messages arrive at the same time, it takes two seconds forget both messages. With a traffic rate of 60, if messages are actually received exactly one second apart and the burst rate is two, the threshold is never reached.

Another way to understand this concept is through the "leaky bucket" analogy. Consider a bucket with a hole at the bottom. Each time a request is processed, water is added to the bucket. Water leaks from the bucket at a constant rate. If the bucket overflows, no more water can be added—a throttling event is triggered. The size of the bucket is the burst, and the rate of the leak is the rate value.

Figure 23-1 Healthy Message Rate

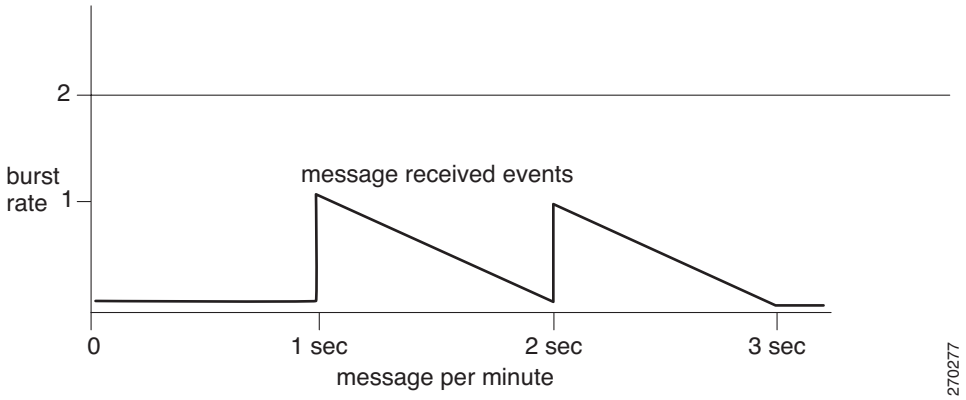

Notice that the ACE XML Gateway's memory of a message in load accounting degrades gradually, over the span of time determined by the traffic rate setting.

With the example burst level of 60, if another message arrives within half a second of the first and the threshold is set to 1, the threshold will be exceeded.

Figure 23-2 Burst Level Threshold Exceeded

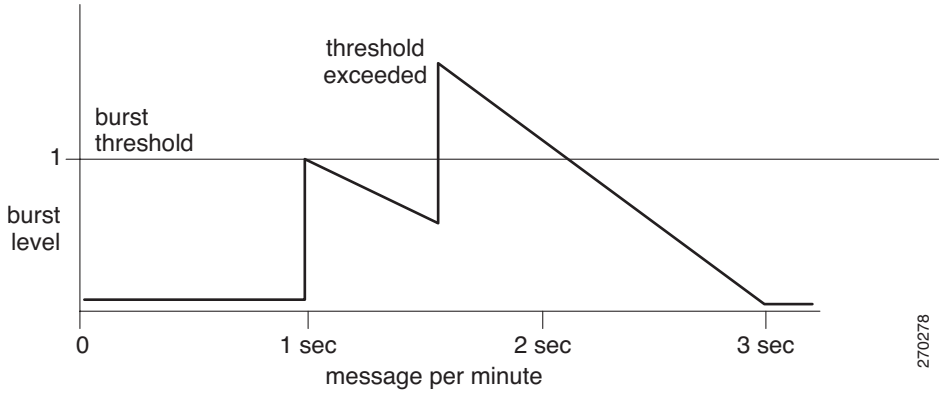

As implied by this scenario, the traffic rate per minute setting only indirectly specifies the number of messages that exceed a threshold; in literal terms, it specifies the amount of time the ACE XML Gateway considers the message in its accounting of the message load against the threshold. Therefore, a traffic per minute setting of two and a burst level of five does not reach the threshold until at least six messages are received—receiving five messages within a minute does not alone exceed the threshold. Since the message rate is set to a low level, until the limit is reached, traffic received is counted against the permissible load for a correspondingly longer period—over several minutes, in effect.

Setting the message rate and burst rate, therefore, is a matter of determining the expected traffic patterns given the types of services you are protecting. If protecting services that may naturally incur usage spikes, such as consumer-oriented services or services that collaborate in some way (for example, multiple services that are invoked by the same application event), you need to set the burst-level to a value that accommodates usage spikes.

For example, if an end-user application generates five service requests to the ACE XML Gateway, the burst level in your configuration should be at least five. How often the application is invoked by the user determines the traffic rate. If the application should only be run once a minute, the request rate can be set to five, so that it takes a minute for the load count from the event to return to zero, with each message taking about 12 seconds to decay.

Figure 23-3 Accommodating Traffic Bursts

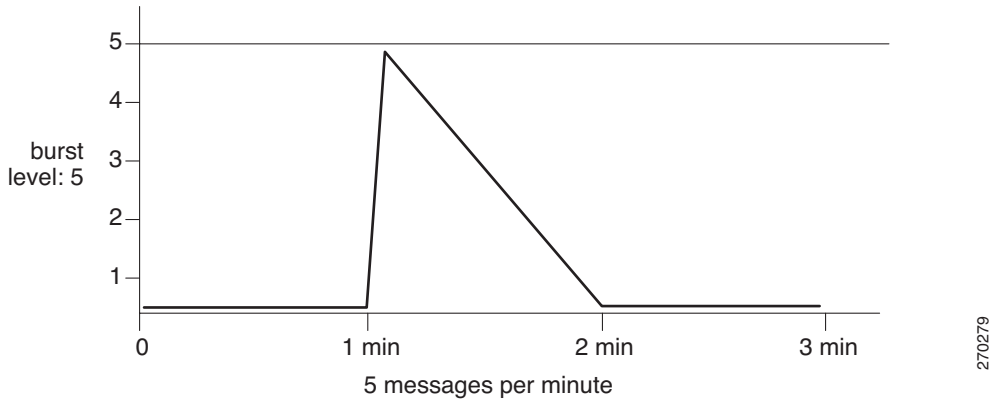

Alternatively, if a high level of concurrent activity would indicate suspicious or undesirable activity given the context environment, you would set the burst level to a relatively low number with the traffic rate at a relatively higher value (if appropriate), such as the earlier example in which the burst level is 2 and a rate per minute is 60.

Of course, the example values described in this section are for illustration purposes only. For most applications, these numbers would be unrealistically low. In an actual configuration, the actual rates and levels would be in the thousands or hundreds.

Configuring DOS Protection

To configure denial-of-service protection:

Step 1 As an `Administrator` user or `Privileged` user with the `Operations` role in the ACE XML Manager web console, click the **Denial-of-Service Protection** link in the **Policy** section of the console navigation menu.

Step 2 Select the check box labelled **Enable Denial-of-Service Protection**.

Step 3 If desired, use the controls on the **Denial-Of-Service Protection Settings** page to configure protection settings other than the defaults.

You can use the **Whitelist** option to specify zero or more IP addresses/IP ranges that are excluded from attack detection. Clients with the specified IP address or whose address is included in the specified IP address range are exempt from blocking even if the rate of their requests exceeds one of the configured threshold.

Note Localhost (127.0.0.1) is always considered to be in the whitelist, even if not specified.

Step 4 Click the **Save Changes** button. The ACE XML Manager commits the denial-of-service settings to the working policy and displays the **Changes Saved** page.

Denial-of-service protection is now configured. Deploy the policy to have the changes take effect at the ACE XML Gateway.

Using Content-Screening Rules

The emergence of dynamic, script-driven web sites has brought with it a form of attack that attempts to take advantage of the dynamic processing capabilities of web infrastructure. In a command-injection attack, an attacker sends a message that contains harmful commands inside the message, with the intent of causing the server that processes the request to execute the harmful command.

For example, an attacker can insert a SQL command, such as `DROP TABLE`, into a message in such a way that certain database-driven backend systems execute the command, resulting in data loss.

Ideally, the services on a protected network would not be vulnerable to such attacks, but in fact it's often much easier and much less expensive to recognize and reject such messages than to find and fix every vulnerability in affected services. The ACE XML Gateway provides an extensible facility for doing that.

The ACE XML Gateway matches incoming message contents against a list of rules that describe message data that the ACE XML Gateway should reject. The ACE XML Gateway includes built-in rules that protect against common attacks, such as the SQL-injection attack. Other rules exist for preventing private or sensitive information from passing through the ACE XML Gateway. For example, it can mask numbers it recognizes to be social security numbers (sometimes used in identity theft), phone numbers or email addresses.

When the ACE XML Gateway encounters a message that matches a content-screening rule, it can:

- Drop the message and note the event in the log
- Allow the message to pass, but note the event in the log
- Modify the matched content before allowing the message to pass

You can supplement the prebuilt content-screening rules with your own custom rules. The label “Cisco signature verified” distinguishes the rules provided for you from your own custom rules.

Figure 23-4 Cisco-Provided Content Screening Rules

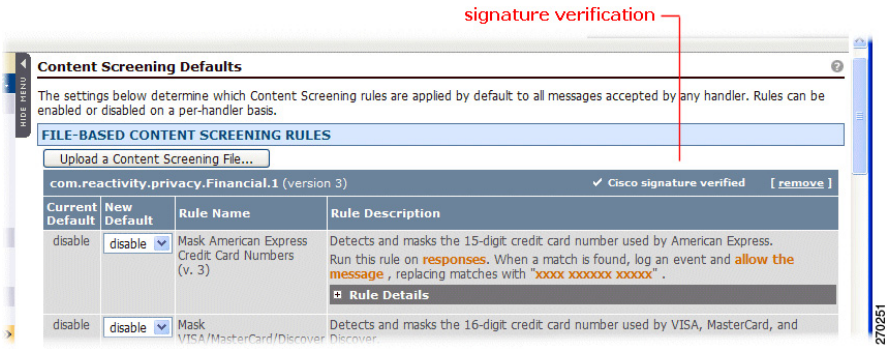

Updates to the built-in content screening rules may be issued independently of the software release cycle. In this case, the update will be provided as a content screening rule file that you can upload using the **Upload a Content Screening File** button.

You should not attempt to create or load a content screening rule file other than as directed by a Cisco ACE XML Gateway support representative. To add individual custom rules, see "Creating Content-Screening Rules" section on page 23-224.

Content screening rules can be applied to either the request or the response. It is important to note that when enabled for a basic virtual service object, content screening rules are applied to both the request and response. If the policy specifies content replacement, this behavior may cause unintended effects. Content replacement rules are typically intended to be applied in a particular context, only to request or only to response branch of the message transaction. To have content screening or replacement rules applied only in one message branch, you need to convert the basic virtual service object to an advanced virtual service and then disable the content screening rule in either the handler (to disable it for requests) or the service descriptor (to disable it for responses).

Applying Content-Screening Rules

A setting for a content screening rule can be applied globally or only for a particular virtual service.

Global Content-Screening Rules

Global content screening rules apply to traffic for services that do not contain more specific settings. To view and enable global content screening rules:

Step 1 In the navigation menu, click the **Content Screening Defaults** link.

Using the controls in the **Content Screening Defaults** page, you can enable or disable content-screening rules that are applied globally (that is, to all messages), upload new content-screening files, and create custom content-screening rules.

Step 2 To get more information on a rule, click the expand button next to the **Rule Details** heading under the rule description.

Figure 23-5 Rule Details

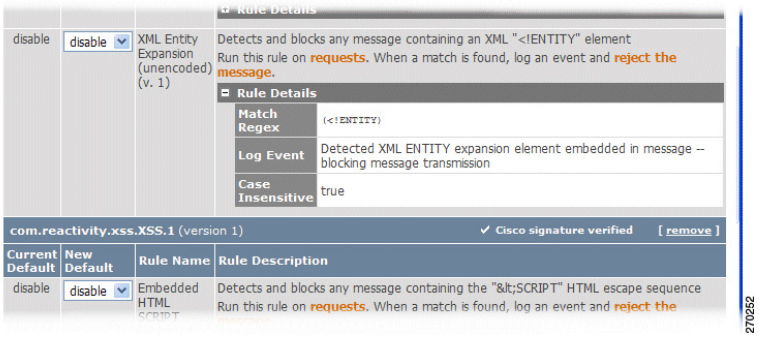

Step 3 To enable a content-screening rule, choose `enable` from the **New Default** menu. When enabled, the ACE XML Gateway actively checks messages for the regular expression specified by the rule.

Alternatively, turn content screening off for that particular rule by choosing `disable`.

Step 4 When finished, click **Save Changes to Default Settings** and deploy the policy to have your changes enforced by the ACE XML Gateway.

Overriding Content-Screening Rule Applicability

For a particular service, you can enable, disable, or accept the default setting for a particular content-screening rule.

To set the service-level applicability of a content-screening rule:

Step 1 In the **Virtual Services** browser, click on the name of the virtual service object for which you would like to view or change service-specific content-screening rules.

Step 2 Click the **Edit** link next to the **Content Screening** settings heading.

The list of content-screening rules appears.

Step 3 Choose from these options:

Option	Description
use default	Use the enabled/disabled state that the **Content Screening Defaults** page specifies for this rule.
always enable	Always enable this rule, even if it is enabled in the **Content Screening Defaults** page specifies for this rule. When you choose this item, this handler always applies this rule to its messages, even if the **Content Screening Defaults** page disables the rule.
always disable	Never enable this rule, even if enabled in the **Content Screening Defaults** page. When you choose this item, this handler never applies this rule to its messages, regardless of any other settings.

Step 4 When finished, click **Save Changes** and deploy the policy to have your changes enforced by the ACE XML Gateway.

Creating Content-Screening Rules

You can create your own content-screening rules. At the implementation level, a content-screening rule is a regular expression that defines the content to be matched in the outgoing or incoming message.

You can specify the results of a match, whether the message is blocked, passed through, or passed through with the matched text modified.

To create a content-screening rule:

Step 1 In the **Policy** portion of the navigation menu, click the **Content Screening Defaults** link.

The **Content Screening Defaults** page appears.

Step 2 Click the **Define a New Rule** button below the **Custom Content Screening Rule** heading, towards the bottom of the page.

The **Custom Content Screening Rule** page appears.

Step 3 Type a unique, descriptive name for the rule in the **Rule Name** field. You can use any name that helps to identify the rule to other policy developers.

Step 4 Add a description for the rule in the **Description** field.

The description helps to document the rule for other console users. It appears in the **Rule Description** column of the rule table on both the global and service-specific content-screening rules page.

Step 5 Type the regular expression that matches the content you want to screen for in the **Regular Expression** field. (For more information on regular expressions, see "Regular Expressions in the Policy" section on page 2-11.)

Step 6 To make the alphabetical characters in the regular expression case-insensitive (that is, the ACE XML Gateway checks both uppercase and lowercase matches of the character), select the **Use case-insensitive regular expression matching** checkbox.

Step 7 In **Rule Actions**, choose whether you want to run the rule on message requests, responses, or HTTP header of either inbound requests or responses. The inbound HTTP headers option allows you to check GET variable content, for example.

Step 8 In the **Log this Warning event** field, type the text that you want to appear in the log when the content screening rule is matched.

Step 9 If you do not want the ACE XML Gateway to immediately reject the message, choose **Allow the message to continue being processed** option.

Step 10 If you allow the message to continue, you can specify replacement text. To do so, select the **Replace any matching string with** checkbox, and type the text to substitute the matched text.

The configuration page should appear similar to the following.

Figure 23-6 Completed Content Screening Rule Configuration

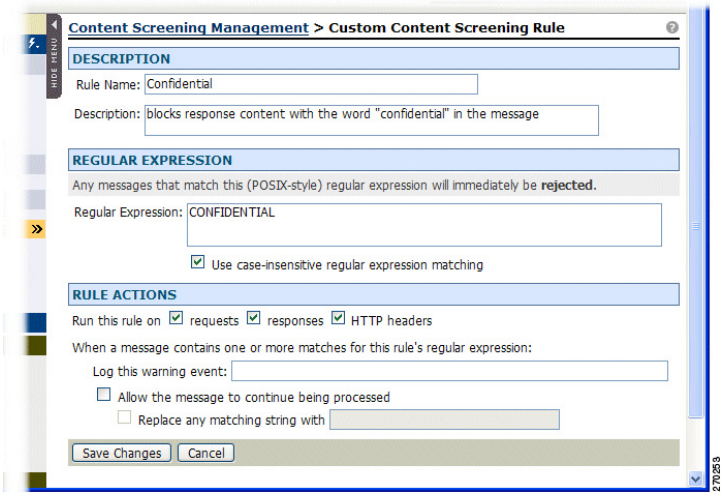

Step 11 Click **Save Changes** when finished.

A new rule is disabled by default. You can enable it globally or on a service-specific basis as described in sections "Global Content-Screening Rules" section on page 23-222 and "Overriding Content-Screening Rule Applicability" section on page 23-223.