



CHAPTER 14

Analyzing Guard Mitigation

This chapter describes how to analyze the Cisco Guard (Guard) mitigation and the zone traffic, and it shows how to identify configuration problems. It provides a brief explanation of how to identify the type of attack. This chapter contains the following sections:

- [Analyzing Zone Traffic Patterns](#)
- [Verifying Attack Mitigation](#)

Analyzing Zone Traffic Patterns

We recommend that once the current attack ends, you allow the Guard to learn the zone traffic patterns if the zone is an on-demand zone configuration using default parameters or if the zone traffic characteristics have changed since the last time the Guard learned the zone traffic.

Use the **show rates** command to display the current zone traffic rates. See the [“Using Counters to Analyze Traffic”](#) section on page 12-3 for more information.

View the received traffic rate and follow these guidelines:

- If the received rate is zero, this rate indicates a diversion problem. See the [“Recognizing a Traffic Diversion Problem”](#) section on page 14-2 for more information.
- If the received rate is greater than the legitimate rate, this rate indicates that the Guard mitigation is functioning, and the following problems might exist:
 - If the legitimate traffic rate for the zone is a lot higher than the zone traffic rate under normal traffic conditions, see the [“Blocking Flows to the Zone Based on Flow Characteristics”](#) section on page 14-2.
 - If the legitimate traffic for the zone is a lot lower than the zone traffic rate under normal traffic conditions, see the [“Verifying Traffic Blocking Criteria”](#) section on page 14-3.

This section contains the following topics:

- [Recognizing a Traffic Diversion Problem](#)
- [Blocking Flows to the Zone Based on Flow Characteristics](#)
- [Verifying Traffic Blocking Criteria](#)

Recognizing a Traffic Diversion Problem

If the Guard does not receive any packets, this condition may indicate a traffic diversion problem where the Guard does not receive the traffic that is sent to the zone.

See [Chapter 4, “Configuring Traffic Diversion”](#) and [Appendix B, “Troubleshooting Diversion,”](#) for more information.

Blocking Flows to the Zone Based on Flow Characteristics

If the legitimate traffic rate for the zone is a lot higher than the zone traffic rate under normal traffic conditions, the Guard might not be blocking all attack traffic. A high rate of legitimate traffic can occur if you did not allow the Guard to learn the traffic patterns of a zone, such as when you use an on-demand zone configuration for zone protection (see the [“Activating On-Demand Protection”](#) section on [page 9-2](#)). If the Guard does not know the zone traffic patterns, the policy thresholds may be too high for the specific zone.

To prevent the Guard from forwarding unwanted flows to the zone, we recommend that you perform the following tasks:

- Lower the threshold of policies that measure traffic according to source IP addresses.
- View the legitimate traffic rate. If the legitimate traffic rate still seems too high, this condition could indicate a sophisticated, large-scale zombie or client attack. Such attacks consist of many flows that do not differ in the rate or in the number of connections from a regular flow. You can configure a flex-content filter to block such anomaly traffic flows. See the [“Configuring Flex-Content Filters”](#) section on [page 6-3](#) for more information.

To lower the policy thresholds, perform the following steps:

Step 1 Display the current policy thresholds by entering the following command in zone configuration mode:

```
show policies
```

See the [“Displaying Policies”](#) section on [page 7-21](#) for more information about the policies.

Step 2 Examine the zone global traffic by entering the following command in zone configuration mode:

```
show policies */*/*/global statistics
```

The Guard displays the traffic flows with the highest rates that are forwarded to the zone, as measured by the protection policies. Determine whether or not the type of services and the volume represent the zone traffic. See the [“Displaying Policy Statistics”](#) section on [page 7-22](#) for more information about policy statistics.

Step 3 Examine the traffic of single users that are represented by a source IP address and determine which policies have a high threshold that should be decreased by entering the following command in zone configuration mode:

```
show policies */*/*/src_ip statistics
```

The Guard displays the traffic flows with the highest rates forwarded to the zone, as measured by the protection policies. See the [“Displaying Policy Statistics”](#) section on [page 7-22](#) for more information about policy statistics.

Step 4 If the traffic volume does not represent the zone traffic, decrease the threshold of the source IP address policies by entering the following command in zone configuration mode:

```
policy */*/*/src_ip thresh-mult threshold-multiply-factor
```

The *threshold-multiply-factor* argument specifies the number by which to multiply the policy threshold. Enter a number less than 1 to decrease the policy threshold. For example, enter 0.5 to decrease the threshold by half. See the “[Multiplying a Threshold by a Factor](#)” section on page 7-16 for more information.

Verifying Traffic Blocking Criteria

If the legitimate traffic rate seems too low, it could indicate that the Guard is blocking access to the zone by legitimate clients. This condition could occur if the learning process was performed some time ago and the policy thresholds no longer fit the zone traffic pattern. The result is that the policy thresholds are not tuned properly and are set too low for current normal traffic patterns.

To verify and change the Guard blocking criteria, perform the following steps:

-
- Step 1** If you suspect that the Guard is blocking access to the zone by legitimate clients, verify that dynamic filters are not blocking access from these clients by entering the following command in zone configuration mode:
- ```
show dynamic-filters [details]
```
- The dynamic filters provide details on the policy that caused the production of the dynamic filters. See the “[Displaying Dynamic Filters](#)” section on page 6-19 for more information.
- Step 2** Identify the policies that caused the production of the dynamic filters and display the statistics of these policies. For example, examine the traffic of single users, represented by a source IP address. Determine which policies have a low threshold that should be increased by entering the following command in zone configuration mode:
- ```
show policies */*/*/src_ip statistics
```
- The Guard displays the traffic flows with the highest rates forwarded to the zone, as measured by the zone policies. See the “[Displaying Policy Statistics](#)” section on page 7-22 for further information about policy statistics.
- Step 3** If the traffic volume does not represent the zone traffic, increase the threshold by entering the following command in zone configuration mode:
- ```
policy */*/*/src_ip thresh-mult threshold-multiply-factor
```
- The *threshold-multiply-factor* argument specifies the number by which to multiply the policy threshold. Enter a number greater than 1 to increase the policy threshold. For example, enter 2 to increase the threshold by 2. See the “[Multiplying a Threshold by a Factor](#)” section on page 7-16 for more information.
- Step 4** Display the list of dynamic filters. (See [Step 1](#).) If the list of dynamic filters includes dynamic filters for IP addresses of legitimate clients with an action of drop, remove those dynamic filters by entering the following command in zone configuration mode:
- ```
no dynamic-filter filter-id
```
- See the “[Configuring Dynamic Filters](#)” section on page 6-18 for more information about dynamic filters.
- Step 5** If the Guard continues to produce drop-action dynamic filters from specific policies, deactivate these policies by entering the following command in policy configuration mode:
- ```
state inactive
```

See the [“Changing the Policy State”](#) section on page 7-13 for more information.


**Tip**

If several policies that are part of the same policy branch produce drop-action dynamic filters, you can deactivate the policy branch by changing the policy state at the higher-level policy sections (such as policy template or service sections).

**Step 6**

Configure known client IP addresses that are crucial to proper zone functioning to bypass the Guard protection functions so that the Guard forwards these traffic flows directly to the zone. Create a bypass filter with the IP address of these clients by entering the following command in zone configuration mode:

```
bypass-filter row-num ip-address protocol dest-port fragments-flag
```

See the [“Configuring Bypass Filters”](#) section on page 6-11 for more information.

## Verifying Attack Mitigation

After you identify an attack on the zone, you can verify that the Guard is mitigating the attack. This action is especially important if you are not familiar with the zone traffic patterns or if the zone is using an on-demand protection configuration (see the [“Activating On-Demand Protection”](#) section on page 9-2) and the Guard did not learn the zone traffic patterns.

To verify attack mitigation, perform the following actions:

- Display the zone current attack report to analyze the attack statistical information. See the [“Displaying the Zone Current Attack Report”](#) section on page 14-4 for more information.
- View the Guard filters, counters, and statistics.

This section contains the following topics:

- [Displaying the Zone Current Attack Report](#)
- [Displaying the Guard Advanced Statistics](#)
- [Displaying Dropped Traffic Statistics](#)

## Displaying the Zone Current Attack Report

You can display the report of an ongoing attack to learn more about the attack characteristics and the measures that the Guard took to mitigate the attack by entering the **show reports current** command. See the [“Displaying Attack Reports”](#) section on page 11-8 for more information.

The report provides you with details about the attack. The information includes when the attack started, a general analysis of the zone traffic flow, an analysis of the packets that were dropped and replied, details of the traffic anomalies that the Guard detected in the zone traffic, and the steps that the Guard took to mitigate the attack. See the [“Understanding the Report Layout”](#) section on page 11-1 for more information.

The report provides you with details about the two main classes of DDoS attack classifications as follows:

- Bandwidth depletion—Attacks designed to flood the zone with unwanted traffic that prevents legitimate traffic from reaching the zone. These attacks include spoofed attacks and malformed packets.
- Resource depletion—Attacks that are designed to tie up the resources of the zone.

See the “[Mitigated Attacks](#)” section on page 11-4 for more information about the types of mitigated attack.

## Displaying the Guard Advanced Statistics

You can view the Guard filters, counters, and diagnostics to learn about the attack characteristics and the measures that the Guard is taking to mitigate the attack. The Guard advanced statistics include the following information:

- Dynamic filters—Provides details about how the Guard is handling the attack. To view the dynamic filters, use the **show dynamic-filters** command. See the “[Displaying Dynamic Filters](#)” section on page 6-19 for more information.
- User filters—Defines how to handle traffic flows that are suspected as DDoS attacks. The zone configuration includes a default set of user filters. You can add or delete user filters. To display the user filters, use the **show** command or the **show running-config** command. The Guard displays the current traffic rate measured for each user filter. See the “[Displaying User Filters](#)” section on page 6-17 for more information.
- Statistics on dropped packets—Provides a list that details the distribution of dropped packets for the ongoing attack. To display the statistics about dropped packets, use the **show drop-statistics** command. See the “[Displaying Dropped Traffic Statistics](#)” section on page 14-5 for more information.
- Zone rate history—Provides the rate that the Guard measured for each counter in the past 24 hours and the details about the attack involvement. To display the zone rate history, use the **show rates history** command. See the “[Using Counters to Analyze Traffic](#)” section on page 12-3 for more information.
- Zone counters—Provides the number of packets that the Guard measured for each counter and enables you to analyze how the Guard has handled the zone traffic since the attack was initiated. See the “[Using Counters to Analyze Traffic](#)” section on page 12-3 for more information.

## Displaying Dropped Traffic Statistics

You can view the distribution of dropped packets for an ongoing attack by entering the following command in configuration mode:

```
show drop-statistics
```

The Guard displays the packets dropped by its protection functions by rate, packet, and bit units.

Table 14-1 provides the drop statistics.

**Table 14-1 Drop Statistics**

| <b>Drop Statistics</b>                    | <b>Description</b>                                                                                                                                                                                                                      |
|-------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Total dropped                             | Total amount of dropped traffic.                                                                                                                                                                                                        |
| Dynamic filters                           | Amount of traffic dropped by the dynamic filters.                                                                                                                                                                                       |
| User filters                              | Amount of traffic dropped by the user filters.                                                                                                                                                                                          |
| Flex-Content filters                      | Amount of traffic dropped by the flex-content filters.                                                                                                                                                                                  |
| Rate limit                                | Packets that are defined by the rate limit parameter of the user filters and the zone <b>rate-limit</b> command that were dropped.                                                                                                      |
| Incoming TCP unauthenticated basic        | Traffic that the TCP basic anti-spoofing functions could not authenticate and dropped. In the attack reports, these packets are counted under the spoofed packets in the Malicious Packets Statistics table.                            |
| Incoming TCP unauthenticated-strong       | Traffic that the TCP strong anti-spoofing functions could not authenticate and dropped. In the attack reports, these packets are counted under the spoofed packets in the Malicious Packets Statistics table.                           |
| Outgoing TCP unauthenticated              | Zone-initiated-connections traffic that the TCP anti-spoofing functions could not authenticate and dropped. In the attack reports, these packets are counted under the spoofed packets in the Malicious Packets Statistics table.       |
| UDP unauthenticated-basic                 | UDP traffic that the basic anti-spoofing functions could not authenticate and dropped. In the attack reports, these packets are counted under the spoofed packets in the Malicious Packets Statistics table.                            |
| UDP unauthenticated-strong                | UDP traffic that the strong anti-spoofing functions could not authenticate and dropped. In the attack reports, these packets are counted under the spoofed packets in the Malicious Packets Statistics table.                           |
| Other protocols unauthenticated           | Non-TCP and non-UDP traffic that the Guard anti-spoofing functions could not authenticate and dropped. In the attack reports, these packets are counted under the spoofed packets in the Malicious Packets Statistics table.            |
| TCP fragments unauthenticated             | TCP fragmented packets that the Guard anti-spoofing functions could not authenticate and dropped. In the attack reports, these packets are counted under the spoofed packets in the Malicious Packets Statistics table.                 |
| UDP fragments unauthenticated             | UDP fragmented packets that the Guard anti-spoofing functions could not authenticate and dropped. In the attack reports, these packets are counted under the spoofed packets in the Malicious Packets Statistics table.                 |
| Other protocols fragments unauthenticated | Non-TCP and non-UDP fragmented packets that the Guard anti-spoofing functions could not authenticate and dropped. In the attack reports, these packets are counted under the spoofed packets in the Malicious Packets Statistics table. |
| DNS malformed replies                     | Malformed DNS replies that the Guard protection functions dropped. In the attack reports, these packets are counted under the malformed packets in the Malicious Packets Statistics table.                                              |

**Table 14-1 Drop Statistics (continued)**

| <b>Drop Statistics</b>                  | <b>Description</b>                                                                                                                                                                                                                                 |
|-----------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| DNS spoofed replies                     | DNS packets that are in response to zone-initiated connections that the Guard anti-spoofing functions dropped. In the attack reports, these packets are counted under the spoofed packets in the Malicious Packets Statistics table.               |
| DNS short queries                       | Short (malformed) DNS queries that the Guard protection functions dropped. In the attack reports, these packets are counted under the malformed packets in the Malicious Packets Statistics table.                                                 |
| Non DNS packets to/from DNS port        | Non-DNS traffic destined to a DNS port or from a DNS port that the Guard protection functions dropped. In the attack reports, these packets are counted under the malformed packets in the Malicious Packets Statistics table.                     |
| Bad packets to proxy addresses          | Malformed traffic destined to the Guard proxy IP address that the Guard protection functions dropped.                                                                                                                                              |
| TCP anti-spoofing features related pkts | Number of dropped packets due to side operations that the Guard TCP anti-spoofing functions performed. In the attack reports, these packets are counted under the malformed packets in the Malicious Packets Statistics table.                     |
| DNS anti-spoofing features related pkts | Number of dropped packets due to side operations that the Guard DNS anti-spoofing functions performed. In the attack reports, these packets are counted under the malformed packets in the Malicious Packets Statistics table.                     |
| Anti-spoofing internal errors           | Number of packets dropped due to the Guard anti-spoofing function errors. In the attack reports, these packets are counted under the Packets table.                                                                                                |
| SIP anti-spoofing features related pkts | Number of SIP <sup>1</sup> over UDP packets that the Guard dropped due to side operations. In the attack reports, these packets are counted under the spoofed packets in the Malicious Packets Statistics table.                                   |
| SIP malformed packets                   | SIP over UDP packets that the Guard protection functions dropped because they were malformed. In the attack reports, these packets are counted under the malformed packets in the Malicious Packets Statistics table.                              |
| Land attack                             | Number of packets dropped because they had identical source and destination IP addresses. In the attack reports, these packets are counted under the malformed packets in the Malicious Packets Statistics table.                                  |
| Malformed packets                       | Number of packets dropped due to a malformed header in which the port, protocol or IP field in the header equals zero (0). In the attack reports, these packets are counted under the malformed packets in the Malicious Packets Statistics table. |

1. SIP = Session Initiation Protocol

The following example shows how to display the drop statistics:

```
user@GUARD-conf-zone-scanner# show drop-statistics
```

