FORTINET

High Performance Network Security

FortiDDoS Handbook

Version 4.3.1

FORTINET DOCUMENT LIBRARY

http://docs.fortinet.com

FORTINET VIDEO GUIDE

http://video.fortinet.com

FORTINET BLOG

https://blog.fortinet.com

CUSTOMER SERVICE & SUPPORT

https://support.fortinet.com

http://cookbook.fortinet.com/how-to-work-with-fortinet-support/

FORTIGATE COOKBOOK

http://cookbook.fortinet.com

FORTINET TRAINING SERVICES

http://www.fortinet.com/training

FORTIGUARD CENTER

http://www.fortiguard.com

END USER LICENSE AGREEMENT

http://www.fortinet.com/doc/legal/EULA.pdf

FEEDBACK

Email: techdocs@fortinet.com

FORTINET®

Wednesday, May 3, 2017

FortiDDoS Handbook

Version 4.3.1

TABLE OF CONTENTS

Change Log

Date	Change Description
2017-03-05	FortiDDoS 4.3.1 initial release

What’s New

This chapter lists features introduced in recent releases:

FortiDDoS 4.3.1

Cloud Monitoring feature is removed from FortiDDoS 4.3.1 Release. Cloud Monitoring option is removed from Global Settings > Settings.

Refer to FortiDDoS 4.3.1 Release notes for the more information about the resolved issues in this release.

FortiDDoS 4.3.0

- Larger Console: A new console window that’s larger than earlier versions allows the administrator to use the command line interface in a better way.
- Color scheme of the GUI is different compared to 4.2.x version.
- Cloud Monitoring: Users with paid subscription to FortiDDoS cloud monitoring can now visualize traffic statistics of FortiDDoS along with data statistics of access to protected resources from multiple vantage points on the Internet.
- Better looking reports: You will now see reports which are very similar to existing GUI display of tables and charts. Text based reports are not available now. Event Subtype for reports is limited to successful and failed logins.
- FortiDDoS 4.3 allows you to generate Certificate Signing Requests that you can send to a CA to give you a signed certificate. **Generate** and **Import** tabs are added under Certificate for this feature.
- **Configured status** and **Linked status** columns are available under System > Network > Interface table.
- **Dashboard** updates:
  - Dashboard is available as a separate tab aligned with System, Global Settings, Protection Profiles, Monitor, and Log & Report.
  - Host name can be more conveniently changed in System Information window by entering the new name and clicking Update.
  - Shortcut to change system time is removed from system information window. System time and Time zone is moved to separate tabs under System>Maintenance.
  - SPP type and time period is generalize on top right and not included separately in each graphs.
- **Restrict DNS Queries to Specific Subnets** is a new option under Protection Profile > SPP Settings which allows you to restrict DNS queries from unwanted sources from the Internet, if enabled.
- **Blacklisted Domains** and **Blacklisted IPv4 Addresses** options are available under Global settings which helps you to deny a large set of blacklist Domains/IPv4 Addresses.
- **Mbps** unit is available as a new option along with PPS under Global Settings > Settings > Settings > General tab > SPP Switching Threshold Measurement Unit.
- Monitor graphs show ingress and egress max packet rates for Layer 3, Layer 4 and Layer 7 graphs. Traffic Statistics uses the egress packet rates. Hence, the system recommendation is also be based on egress rates.
- **Save As CSV** option is available under the following tabs that exports information in a format that is suitable for printing and sharing.:
  - Protection Profiles > Traffic Statistics > Generate
  - Log & Report > Log Access > Logs > DDoS Attack log/Event Log
- **Save As PDF** option is available for all graphs under Monitor.
- **HTTP partial request per source per second threshold** and **HTTP partial request to response observation period** options are removed from Global Settings > Settings > Settings > Slow Connection.
- Destination IP in events will be renamed/reported as Protected IP.

- Invalid ICMP Anomaly is newly added under Global setting > Settings > Settings > General tab.
- Event log and DDoS log details can be check by clicking on preview button on right.
- Logs and report settings are reorganized under different tabs:
    - Separate tab for 'Report Purge Setting' which was under Report configuration tab in earlier releases.
    - New tab for Executive summary (instead of attack graph) which includes DDOS Attack log, DDoS attack Graph and Event log.
- HTTP Anomaly options under Global Settings > Settings > Settings > General tab is renamed as follows:
    - Known Method Anomaly
    - Unknown Method Anomaly
    - Invalid HTTP Version Anomaly
    - Do Not Parse HTTP 0.9
- Any configuration changes saved from FortiDDoS 4.3 GUI, by clicking **Save** button, gets updated without a dialog box with saved successfully message compared to earlier versions of FortiDDoS.
- Some of the tree elements on the left panel of the FortiDDoS 4.3 interface is moved as tabs on the main page. For example, Administrator, Access Profiles and Settings are changes to tabs on the Admin page.
- Some of the DNS Feature control settings are changed.
- Aggressive Aging due to slow connection is newly added under Graphs.
- A new scalar - Methods per Source is added under Protection Profiles > Thresholds > Scalars. High volumes of any HTTP Method from single sources will trigger this Threshold. You can see the thresholds under Protection Profiles > Thresholds > Thresholds > Scalars and the related graph under Monitor > Layer 7 > HTTP > Methods per source.
- DNS RD bit is no more included as DNS query anomaly.

FortiDDoS 4.2.1, 4.2.2 and 4.2.3

This release includes bug fixes only. No new features. See the release notes.

FortiDDoS 4.2.0

- DNS attack mitigation—Massive-scale DNS attack mitigation using ACLs, anomaly detection, and patented DNS flood mitigation methods to enable your business to continue to serve legitimate client purposes during floods. Start with Understanding FortiDDoS DNS attack mitigation.
- DNS monitor graphs—New graphs to monitor DNS traffic and DNS attack mitigation mechanisms. See Chapter 5: Monitor Graphs.
- DNS reports—New reports to monitor DNS attacks. See Chapter 6: Logs and Reports.
- LDAP—Support for administrator authentication against LDAP servers. See Configuring LDAP authentication.
- RADIUS—Changes to the RADIUS configuration to match the LDAP implementation. See Configuring RADIUS authentication.
- Cloud Signaling—REST API and configuration that enables small/medium businesses and enterprises to work with participating service providers and cloud providers to route traffic through a "scrubbing station" in the service provider network (SPN) before it is forwarded through the WAN link to the customer premises network (CPN). See Chapter 11: Service Provider Signaling Deployments.
- Global distress ACL—Configuration has been added to the Web UI. See Configuring a global distress ACL for protocol traffic.
- New Monitor graphs—New SPP Statistics graph shows throughput per SPP, and the Layer 4 Flood graph shows drops due to Slow Connection detection (slow TCP or HTTP). See Chapter 5: Monitor Graphs.
- Monitor graph enhancements—Links from aggregate graphs to detailed graphs. See Chapter 5: Monitor Graphs.
- HA sync changes—Refined HA sync so that network configurations are not synchronized. See HA synchronization.

- Web UI changes—The Global Settings and SPP Settings pages use tab groups to make it easier to navigate to the options you want to configure. See Configuring global settings and Configuring SPP settings.

FortiDDoS 4.1.11

This release includes bug fixes only. No new features. See the release notes.

FortiDDoS 4.1.10

This release includes bug fixes only. No new features. See the release notes.

FortiDDoS 4.1.9

This release includes bug fixes only. No new features. See the release notes.

FortiDDoS 4.1.8

- Source blocking for slow connection attacks was removed from Global Settings in 4.1.7. In 4.1.8, it has been added to the SPP settings configuration.

FortiDDoS 4.1.7

New features:

- Option to generate reports per SPP policy (subnet).
- Option to generate reports hourly.
- Option to purge reports automatically by size or manually by date range, similar to the functionality for DDoS Attack logs. The Report Configuration > Purge Settings displays Log Disk Status information giving total, used, and available space.
- Option to back up and restore a single SPP configuration.
- Built-in bypass for copper ports and FDD-2000B fiber bypass ports 17-20 is now configurable as fail-open or fail-closed.
- The system now supports TAP mode, compatible with FortiBridge-3000-series and other external Bridge/TAP products. In TAP mode, FortiDDoS monitors ingress traffic on both WAN- and LAN-side ports but does not pass traffic to the egress port. TAP mode also does not pass external bridge heartbeats. Contact your sales representative for details on interoperation with FortiBridge.

Changed features:

- Slow connection detection is automatically disabled in Detection Mode.
- Source blocking for slow connection attacks has been removed from Global Settings.
- The Monitor > Specific Graphs section has been removed. The graphs formerly included in this section have been moved to Monitor > Layer 3, Monitor > Layer 4, or Monitor > Layer 7, as appropriate.
- Beginning with release 4.1.6, the UDP service is identified when either the source or destination port is a well known port (the IANA assigned ports 0-1023).

FortiDDoS 4.1.6

Key bug fixes:

- Software upgrade with the web UI.
- HA active-passive configuration synchronization.

New features:

- New anti-spoofing ACL that drops traffic that matches local addresses when the addresses appear to be spoofed.
- New table to track up to 2^32 IPv4 ACLs. The table includes rules from the Local Address Anti-Spoofing, IP address, Geolocation, and IP Reputation lists. You can use a new Monitor graph called Address Denied to monitor drops.
- New HTTP header options for detecting proxy IP addresses: X-Real-IP, X-True-Client-IP.
- New option to drop sessions when packets contain the HTTP Range header.
- Asymmetric mode configurable through the web UI. New configuration options are available to ease setup in networks with asymmetric traffic.
- Tap Mode. System now supports Tap Mode as a Beta feature. Tap Mode is designed to work with FortiBridge-300xS/L in Bypass/Tap mode to allow continuous offline monitoring of network traffic.
- FortiDDoS now supports attack logging to FortiAnalyzer.
- The dashboard System Status portlet now displays LAN/WAN port labels.
- On the dashboard System Status portlet, unconfigured SPPs are now represented by a gray circle.
- Improved Monitor graph workflow. New Aggregate Drops graph showing Flood, ACL, Anomaly, Hash, and Memory Drops all on one graph. Organization of the graphs below this is more logical.
- Tooltips on all graphs now show more granular time information.
- Syslog and SNMP traps now contain subnet ID.
- Added an event log and SNMP trap to notify when a link goes up or down.
- Added an event log for FortiGuard IP Reputation updates.
- Added a CLI command to back up Event logs and other diagnostic data.
- Added the `diagnose debug RRD` commands to verify the integrity of RRD (reporting) files.

Changed features:

- The SPP ACL Drops portlet is now called Top SPPs with Denied Packets; and the SPP Attacks portlet is now called Top Attacked SPPs.
- Source penalty factor is now called Source multiplier; and Application penalty factor is now called Layer 7 multiplier.
- The Aggressive Aging TCP Feature Control URL-Flood option was mislabelled. It is now the layer7-flood option.
- The system recommended threshold values for TCP/UDP ports and ICMP types/codes are based on a new and improved heuristic algorithm.
- Separate Subnet attack logs and reports have been removed and subnet information integrated into the DDoS Attack Log.
- Syslog format changes to better interoperate with FortiAnalyzer.

Removed features:

- The DDoS Attack Log no longer includes the frequent Most Active Source and Most Active Destination notifications that were sent when the Most Active Source / Destination data points were recorded. These logs had been sent to give details on the recorded data point even when the effective rate limit was not met, causing confusion. When these thresholds are exceeded, however, the events are logged as a Source Flood and a Destination Flood, respectively.
- Syslog and SNMP traps also no longer include the Most Active Source and Most Active Destination notifications.
- Destination penalty factors have been removed to prevent rate-limiting of all users to a specific destination.
- Thresholds, logs, and graph plots for URL Scan events have been removed. URL Scan events included the HTTP request anomalies related to sequential requests and HTTP mandatory header counts.
- MyList and MyGraph functionality has been removed.

- Dark Address Scan graph has been removed.
- The System Dashboard system Reset button. You can use CLI commands to completely reset the system.

FortiDDoS 4.1.5

Bug fixes only.

FortiDDoS 4.1.4

- Support for administrator authentication against an external RADIUS server.

FortiDDoS 4.1.3

Bug fixes only.

FortiDDoS 4.1.2

- **Asymmetric Mode** — We recommend that you enable Asymmetric Mode if the FortiDDoS appliance is deployed in a network path where asymmetric routes are possible. An example of an asymmetric route is one in which the client request traverses the FortiDDoS system, but the server response takes a route that does not.

FortiDDoS 4.1.1

- **Enhanced IPv6 support** — FortiDDoS now supports both Internet Protocol version 4 (IPv4) and Internet Protocol version 6 (IPv6) by default.
  If your appliance is deployed on a network with IPv6 traffic, specify the IPv6 prefix settings before you configure SPPs that monitor and protect IPv6 subnets.

  SPPs can now monitor and protect IPv4 and IPv6 traffic simultaneously. Use either IPv4 or IPv6 address formats to specify the subnet in an SPP policy, and then apply policies that use either format to the same SPP

  In addition, you can now specify IPv6 prefixes to /32.

- **Enhanced slow connection configuration** — To simplify configuration, FortiDDoS now provides a global setting that allows you to switch from moderate to aggressive slow connection attack protection, as well as settings you apply to individual SPPs.
- **FortiDDoS 200B**— Support for this new hardware model.

FortiDDoS 4.1.0

- **Logging & report enhancements**
- **SNMP traps & MIBs for attack logs** — You can now configure FortiDDoS to send attack log information to SNMP managers.
- **DDoS Subnet Attack Log** — The new **DDoS Subnet Attack Log** displays events associated with a specific SPP policy, with counts updated every five minutes.
- **Subnet Executive Summary dashboard** — The new **Subnet Executive Summary** dashboard displays all attacks in the **Top Attacked Subnet** and **Top ACL Subnet Drops** report categories.
- **Destination tracking** — For all attack log event categories, FortiDDoS now provides the IP address of the first destination it identifies as the target of the attack activity. Information that is organized by this destination is available as a report type and widgets on the **Executive Summary** and **Attack Graphs** dashboards.
- **Filter report information by SPP or subnet** — When you create a report configuration, you can now restrict the information in the report to a specific SPP or subnet.
- **Enhanced blocking by geolocation** — The **Geo Location Policy** setting allows you to either permit traffic from all geographic locations and add exceptions or deny access to all locations with exceptions.
- **Access dropped packet and other statistics via API** — You can now use the FortiDDoS REST API to access dropped and blocked traffic statistics and traffic graph information. See the *FortiDDoS REST API Reference*.

- **MySQL access to DDoS attack log** — You can now access the DDoS attack log with read-only permission using a third-party tool such as the MySQL command-line tool or MySQL Workbench.
- **Alert email message for SPP switching** — You can now configure FortiDDoS to generate a system event log and send a corresponding email message whenever the appliance switches a subnet to its alternative SPP. (If you want FortiDDoS to notify you that the traffic level has exceeded the SPP switching threshold without switching the SPP, in the SPP policy settings, specify the same SPP for both **Service Protection Profile** and **Alternate Service Protection Profile**.)
- **Improved dual-stack IPv6 support** — Additional settings and functionality that make it easier to deploy FortiDDoS in networks with IPv6 traffic.
- **Double VLAN (DVLAN) detection** — FortiDDoS now tracks traffic with an additional 802.1Q tag (for example, VLAN Q-in-Q).

FortiDDoS 4.0.1

No new features. Bug fixes only.

FortiDDoS 4.0.0

- **Additional data ports** — 16 physical LAN and WAN ports are configured as linked pairs. Odd-numbered ports are LAN connections that have a corresponding even-numbered port, which is the associated WAN connection. That is, Port 1/Port 2 behaves as LAN 1/WAN 1, Port 3/Port 4 as LAN 2/WAN 2, Port 5/Port 6 as LAN 3/WAN 3, and so on. These port pairs enable you to protect up to 8 links with a single appliance.
- **Increased throughput**—Increased throughput on all hardware models. For details on throughput rates, see the product datasheet.
- **Configuration synchronization** — High Availability (HA) configuration allows you to synchronize configuration information between two FortiDDoS appliances to create a secondary appliance that always has an up-to-date configuration.
- **Automatic bypass for copper links** — For Ethernet links (copper, RJ-45), the FortiDDoS appliance automatically passes traffic through when the appliance is not powered up, its FortiASIC processor or integrated switch fabric fail, or it is booting up and all services are not yet available.
- **Link down synchronization** — The appliance has two options for Link Down Synchronization: Wire and Hub. When Wire is selected, FortiDDoS monitors the link state of both ports in a port pair. If the link goes down on either port, it disables the other port. The appliance re-enables the port when it detects that the link for other port in the pair is up again. When Hub is selected, FortiDDoS does not disable both ports in a port pair if the link goes down on one of the ports.
- **Redesigned web UI—** The graphical user interface is organized by component and tasks. Many of its system settings and options are shared with other Fortinet products.
- **Management via command-line interface** — You can perform all appliance configuration from a Secure Shell (SSH) or Telnet terminal or from the JavaScript CLI Console widget in the web UI.
- **RESTful web API configuration** — Use a web API that uses HTTP and REST principles to perform tasks such as allowing or denying sources, setting thresholds and changing SPP (formerly VIDs) configuration.
- For more information, contact Fortinet Technical Support.
- **BIOS-based signed appliance certificate** — The validation mechanism for the appliance’s identity is built into its hardware.
- **Faster threshold report generation** — FortiDDoS now takes less time to generate the traffic statistics it uses to calculate system-recommended thresholds.
- **Save reports as PDF** — Many FortiDDoS system events and attack activity reports and graphs have a **Save as PDF** option that exports information in a format that is suitable for printing and sharing.

- **Attack activity at a glance dashboard** — Access the most popular attack activity reports information on a single web page and in table format using **Log & Report > Report Browse > Executive Summary**. Use **Log & Report > Attack Graphs > Attack Graphs** to access the most popular attack activity graphs on a single web page.
- **Context-sensitive help** — Click **Help** to open the HTML help information for the current content pane.
- **Filter and sort log information** — For system event and DDoS attack logs, you can use the column headers to sort log information or arrange the columns. The filter feature allows you to select items to include or exclude based on date, category, or other criteria.
- **Enhanced reports** — New features include the ability to generate reports as HTML, text, PDF and customize reports with a logo.
- **Built-in DoS control** — FortiDDoS blocks packets with a pre-defined set of anomalies before they reach the appliance's processor. Traffic graphs and reports do not report the packets that this feature drops.
- **Block protocols on subnets (Distress ACL)** — The distress ACL feature helps to block brute force protocol attacks on a specified subnet or IP address. It allows you to block packets that can flood the pipe before they reach the appliance's processor. Traffic graphs and reports do not report the packets that this feature drops.
- **Bypass fiber ports for 2000B model** — Two physical port pairs on the FortiDDoS 2000B have built-in bypass capability. Built-in bypass works during a power failure, critical component failure and during startup and shutdown.
- **IP Reputation update using file upload** — You can update the addresses in the IP Reputation Service list by uploading a .pkg file.

Introduction

FortiDDoS is a network behavior anomaly (NBA) prevention system that detects and blocks attacks that intend to disrupt network service by overutilizing server resources. These attacks are commonly referred to as distributed denial of service (DDoS) attacks.

FortiDDoS uses a variety of schemes, including anomaly detection and statistical techniques, to provide nonstop protection, even against attacks that cannot yet be recognized by an attack signature. When FortiDDoS detects an attack, it immediately drops traffic from the offending source, thus protecting the systems it defends from floods.

Product features

The following features make FortiDDoS the best in its class:

- Purpose-built for low latency and rapid response
  The patented combination of purpose-built hardware and heuristics allow you to deploy the FortiDDoS appliance inline (for example, between the external network and a protected server), where it can receive, process, and transmit packets at a high rate, even when an attack is underway. FortiDDoS introduces a latency of approximately a few microseconds and has a response time of 2 seconds or less.
- Massive-scale SYN and DNS flood mitigation
  SYN flood mitigation and DNS flood mitigation techniques not only protect your network from DDoS attacks but, importantly, enable your business to continue to serve legitimate client purposes during floods.
- Initial learning period
  FortiDDoS learns based on inbound and outbound traffic patterns. You first deploy the system in Detection Mode. In Detection Mode, the system operates with high (factory default) thresholds and does not drop any packets.

  At the end of the initial learning period, you can adopt system-recommended thresholds (usually lower than the factory default) and continue to use Detection Mode to review logs for false positives and false negatives. As needed, you repeat the tuning: adjust thresholds and monitor the results.

  When you are satisfied with the system settings, change to Prevention Mode. In Prevention Mode, the appliance drops packets and blocks sources that violate ACL rules and DDoS attack detection thresholds.
- Continuous learning
  FortiDDoS begins learning traffic patterns as soon as it begins monitoring traffic, and it never stops learning. It continuously analyzes traffic rates and dynamically adjusts the thresholds that differentiate between legitimate traffic volume and attacks.
- Zero Day attack prevention
  FortiDDoS uses rate-based analysis, which protects against attacks that hackers have not yet imagined. Administrators do not need to intervene, and the appliance is “on guard” 24/7, automatically protecting your network systems and bandwidth.
- Granular attack detection thresholds
  The FortiDDoS specialized hardware is designed to monitor thresholds for all traffic it sees at Layers 3, 4, and 7. It tracks throughput, packet rate, new connections, TCP state transitions, fragments, checksums, flags, and so on. You can set thresholds on the appropriate traffic parameter to limit traffic for particular systems or applications.
- Deep packet inspection
  The FortiDDoS specialized hardware enables deep packet inspection. FortiDDoS can identify header fields in HTTP packets and maintain specific thresholds for specific URLs. This granularity enables the system to distinguish between attacks against a specific URL and legitimate traffic to other resources.
- Slow connection detection
  Many botnets have started using slow connection build up as a mechanism to confuse security appliances and thus effectively overload the servers. FortiDDoS can identify these types of attacks by monitoring thresholds for partial requests. When an attack is detected, the system “aggressively ages” the connections, recovering resources for protected servers.

- Known IP address matching
  A proprietary algorithm matches incoming connection requests with known IP addresses to mitigate SYN attacks without the overhead of connection proxies. Legitimate users can connect or remain connected, even during a SYN flood attack.
- Source tracking
  FortiDDoS tracks connection and rate behavior per source IP address, so it can identify the source of attacks and apply more stringent limits to the traffic they send to your servers.
- Service Protection Profiles (SPPs)
  FortiDDoS maintains up to 8 sets of counters and thresholds that you assign to a subnet as a group. Thus, a single appliance can protect up to 511 subnets, each identified by an IP address representing a server or group of networked servers. Each of these virtual protection zones—called Service Protection Profiles—learns traffic patterns and estimates adaptive thresholds independently. You can assign each profile an independent administrator, which is useful in multi-tenant environments such as an ISP.
- Cloud signaling
  REST API and configuration that enables small/medium businesses and enterprises to work with participating service providers and cloud providers to route traffic through a "scrubbing station" in the service provider network (SPN) before it is forwarded through the WAN link to the customer premises network (CPN).
- Intuitive analysis tools and reports
  The on-box reporting tools enable graphical analysis of network traffic history from five minutes to one year. You can analyze traffic profiles using a broad range of Layer 3, 4 or 7 parameters. With just a few clicks, you can create intuitive and useful reports such as top attackers, top attacks, top attack destinations, top connections, and so on.
- Viewing traffic monitor graphs
  Traffic monitor graphs display trends in throughput rates and drop counts due to threat prevention actions. In Detection Mode, the drop count is hypothetical, but useful as you tune detection thresholds.
- Configurable event monitoring
  You can monitor FortiDDoS events using the web UI, SNMP, or email event notification.
- Cloud Monitoring
  FortiDDoS can interact with FortiDDoS Cloud Monitoring and send statistics of ports, SPP and drops. This helps you to simultaneously visualize these statistics with the view of the traffic from multiple vantage points on the Internet during normal times and floods. The value is in being to correlate the statistics.
- Local Address Anti-spoofing
  FortiDDoS allows you to block inbound packets that have a source address inside the network, block inbound packets that do not have a destination in your network, block outbound source that are not local addresses and block outbound packets with a destination inside your local network.

Deployment topology

Figure 1 shows a simple deployment. The FortiDDoS appliance is deployed inline between the Internet and the local network to protect the local network servers from volume-based attacks like floods and attacks that send anomalous packets to exploit known vulnerabilities.

Figure 1: Basic deployment

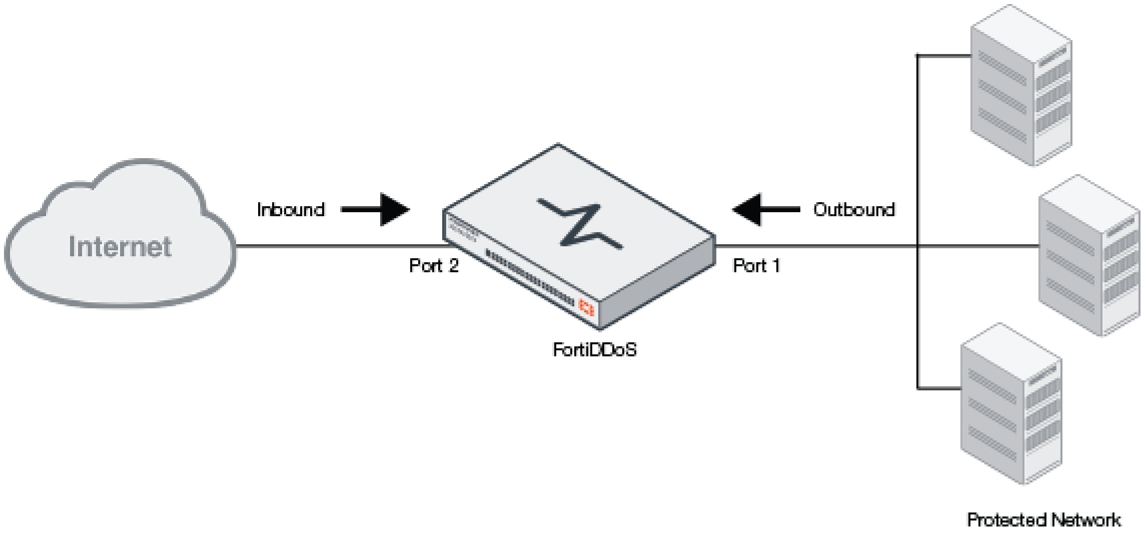

You can deploy FortiDDoS in more complex and specialized topologies. See Chapter 9: Deployment Topologies.

Document scope

The FortiDDoS QuickStart Guide for your appliance has details on the FortiDDoS hardware components, ports, and LEDs, and the FortiDDoS product datasheet has detailed specifications. The product datasheet also lists throughput per model. Please use those resources to size your deployment, select the appropriate hardware models, and install the hardware into an appropriate location and machine environment.

This handbook describes how to get started with the FortiDDoS system, how to modify and manage configurations, how to monitor traffic, and how to troubleshoot system issues.

Figure 2 shows the order in which FortiDDoS applies its rules and actions. It is provided in this introduction to give you an overview of the important features that you can learn about in this manual. Figure 2 shows that packets matching the Do Not Track policy are forwarded without inspection. Otherwise, the packets are evaluated against sets of built-in rules and user-defined rules. Legitimate traffic is forwarded with low latency.

1. Global ACL Deny—You can configure rules that deny traffic to/from local addresses geolocations to prevent spoofing, and from IPv4 address spaces and geolocations known to have no business requesting resources from any of the protected subnets. You can also block addresses maintained by the FortiGuard IP Reputation service. Packets that match deny rules are dropped. Packets that do not match the deny rules continue for further processing.
2. Protocol Anomalies—Drop packets identified by built-in protocol anomaly detection methods. No configuration is required. Layer 3 protocol anomaly detection is performed first. If none found, the traffic continues to Layer 4 protocol anomaly detection. If none found, the packets continue for further processing.
3. Global ACL Deny IP netmask—Rules configured to match an IPv4 netmask are consulted next. Packets that match deny rules are dropped. Otherwise, processing continues.
4. SPP ACL—Use SPP allow/deny rules to enforce nuanced policy decisions based on Layer 3, Layer 4, and Layer 7 parameters. An SPP administrator can create granular rules based on his or her knowledge of the IP addresses and services that have reason or no reason to travel inbound or outbound in its network. Layer 3 rules are processed first. Packets not dropped continue to Layer 4 rule processing. Packets not dropped continue to Layer 7 processing. Packets that are not dropped continue.
5. TCP State Anomalies—You can enable rules to drop packets identified by TCP state anomaly detection methods. Packets that have TCP state anomalies are either harmful or useless, so we recommend you use the TCP state anomalies detection options to drop these. Packets that are not dropped continue.
6. Source SPP Thresholds—Packets are evaluated against the source table. Packets from source IP addresses subject to a FortiDDoS blocking period are dropped. Packets that exceed per-source thresholds are dropped. Packets that are not dropped continue.
7. Destination SPP Thresholds—Packets are evaluated against the destination table. Packets that exceed per-destination thresholds are dropped. Packets that are not dropped continue.
8. Port rules—Packets are evaluated against the SPP ACL and SPP thresholds. Packets that are not dropped continue.
9. SPP Thresholds—Packets are evaluated against SPP rate limits. Layer 3 thresholds are processed first, then Layer 4, then Layer 7.
   - If a maximum rate limit is reached, such as packet rate for a specified protocol, FortiDDoS drops the packets.
   - If a slow connection threshold is reached, FortiDDoS sends a TCP reset to the server.
   - If a SYN flood threshold is reached, FortiDDoS challenges the client using the configured SYN Flood Mitigation Mode method.
   - Otherwise, processing continues.

10. HTTP Header rules—Packets are evaluated against the SPP ACL and SPP thresholds. Packets that are not dropped are forwarded toward their destination.

Figure 2: FortiDDoS drop precedence

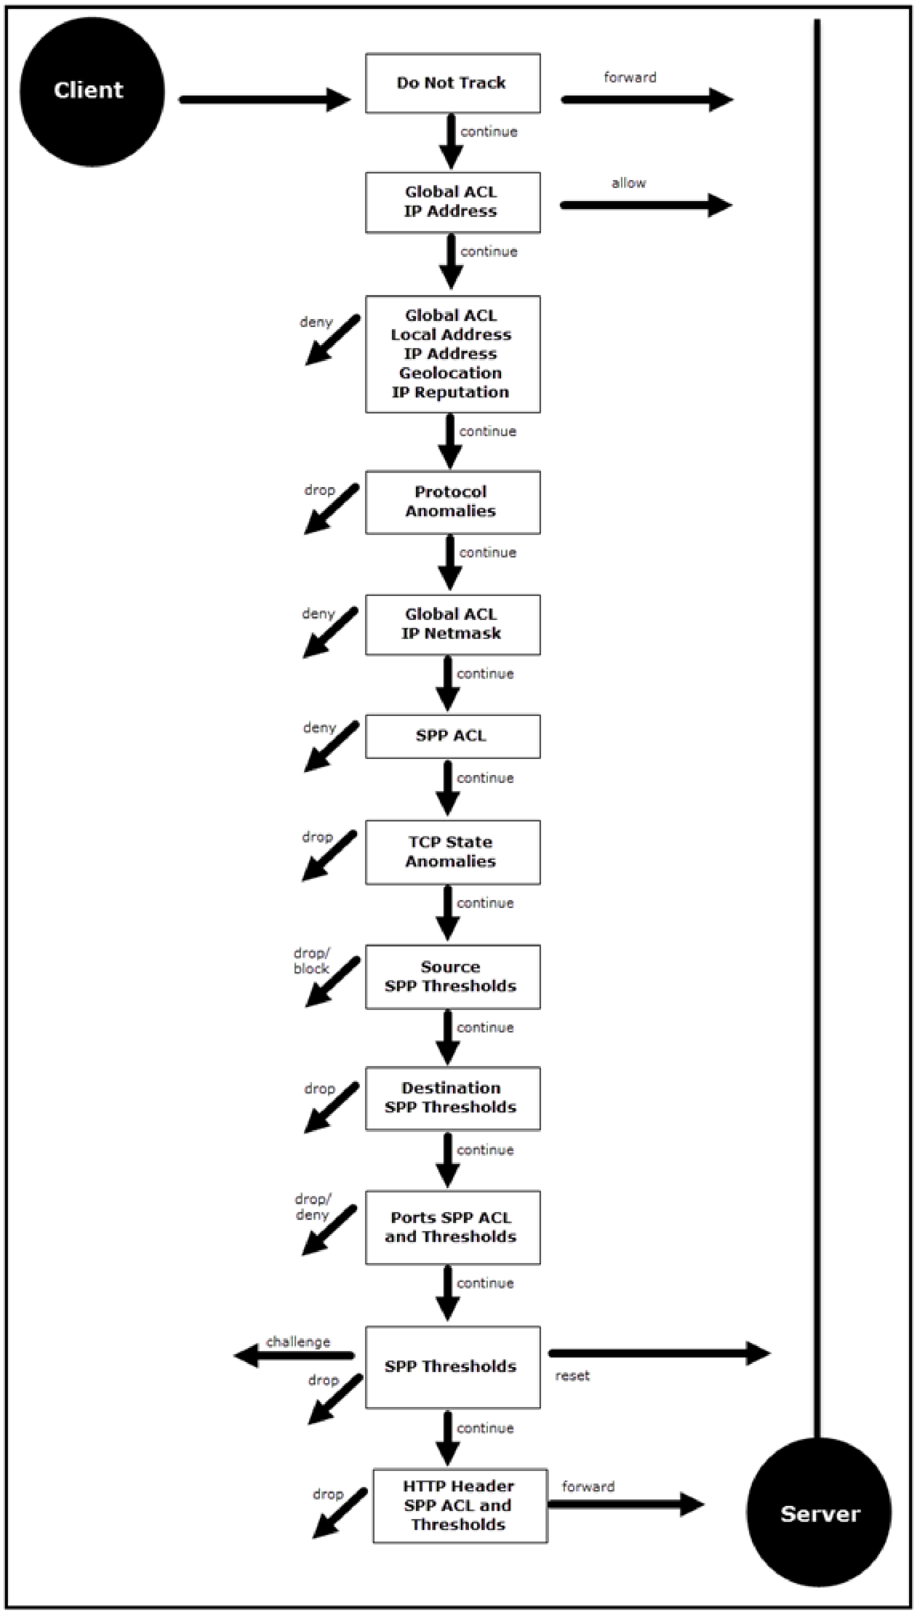

Chapter 1: Key Concepts

This chapter describes FortiDDoS concepts, terms, and features.

If you are new to FortiDDoS, or new to distributed denial of service (DDoS), this chapter can help you understand the problem and the mitigation techniques.

This chapter includes the following sections:

DDoS attack overview

Computer network security is a challenge as old as the Internet itself. The sophistication and infamy of network-based system attacks has kept pace with the security technology and hackers only feel more challenged by the latest heuristics designed to foil their efforts.

Some attackers exploit system weaknesses for political purposes, disgruntled about the state of software or hardware in the market today. Others target specific systems out of spite or a grudge against a specific company.

Yet others are simply in search of the infamy of bringing a high-traffic site to its knees with a denial of service (DoS) attack. In such an attack, the hacker attempts to consume all the resources of a networked system so that no other users can be served. The implications for victims range from a nuisance to millions of dollars in lost revenue.

In distributed denial of service (DDoS) attacks, attackers write a program that will covertly send itself to dozens, hundreds, or even thousands of other computers. These computers are known as 'agents' or 'zombies', because they act on behalf of the hackers to launch an attack against target systems. A network of these computers is called a botnet.

At a predetermined time, the worm will cause all of these zombies to attempt repeated connections to a target site. If the attack is successful, it will deplete all system or network resources, thereby denying service to legitimate users or customers.

E-commerce sites, domain name servers, web servers, and email servers are all vulnerable to these types of attacks. IT managers must take steps to protect their systems—and their businesses—from irreparable damage.

Any computer can be infected, and the consequences can range from a nuisance popup ad to thousands of dollars in costs for replacement or repair. For this reason, antivirus software for all PCs should be a mandatory element of any network security strategy. But whether you measure cost in terms of lost revenue, lost productivity, or actual repair/restore expenses, the cost of losing a server to an attack is far more severe than losing a laptop or desktop.

Servers that host hundreds or thousands of internal users, partners, and revenue-bearing services are usually the targets of hackers, because this is where the pain is felt most. Protecting these valuable assets appropriately is paramount. In early 2000, the industry saw a new kind of 'worm' attack, in which hundreds or thousands of (sometimes unsuspecting) systems were employed to simultaneously bombard a target host, paralyzing its productivity. Several high traffic sites such as Amazon.com, Buy.com, CNN, Yahoo, and eBAY were affected by these DDoS attacks.

To circumvent detection, attackers are increasingly mimicking the behavior of a large number of clients. The resulting attacks are hard to defend against with standard techniques, as the malicious requests differ from the legitimate ones in intent but not in content. Because each attacking system looks innocent, advanced techniques are required to separate the 'bad' traffic from the 'good' traffic.

DDoS mitigation techniques overview

The best security strategies encompass people, operations, and technology. The first two typically fall within an autonomous domain, e.g. within a company or IT department that can enforce procedures among employees, contractors, or partners. But since the Internet is a public resource, such policies cannot be applied to all potential users of a public website or email server. Thankfully, technology offers a range of security products to address the various vulnerabilities.

Firewalls

Firewalls can go a long way to solving some problems by restricting access to authorized users and blocking unwanted protocols. As such, they are a valuable part of a security strategy. But public websites and eCommerce servers cannot know in advance who will access them and cannot 'prescreen' users via an access list. Certain protocols can be blocked by firewalls, but most DoS attacks utilize authorized ports (e.g. TCP port 80 for a web server) that cannot be blocked by a firewall without effectively blocking all legitimate HTTP traffic to the site, thereby accomplishing the hacker’s objective.

Firewalls offer some security against a single user DoS attack by denying access to the offending connection (once it is known), but most DoS attacks today are distributed among hundreds or thousands of zombies, each of which could be sending legal packets that would pass firewall scrutiny. Firewalls perform a valuable service in an integrated security strategy, but firewalls alone are not enough.

Router access control lists

Likewise, access lists in the router can be used to block certain addresses, if such addresses can be known a priori. But websites open to the public are, by nature, open to connections from individual computers, which are exactly the agents hackers use to initiate attacks. In a DDoS attack, thousands of innocent looking connections are used in parallel. Although router access lists can be used to eliminate offending packets once they are identified, routers lack the processing power and profiling heuristics to make such identifications on their own.

In addition, complex access lists can cause processing bottlenecks in routers, whose main function is to route IP packets. Performing packet inspections at Layers 3, 4, and 7 taxes the resources of the router and can limit network throughput.

Antivirus software

End systems cannot be considered secure without antivirus software. Such software scans all inputs to the system for known viruses and worms, which can cause damage to the end system and any others they may infect. Even after a virus is known and characterized, instances of it are still circulating on the Internet, through email, on CDs and floppy disks. A good antivirus subscription that is frequently updated for the latest protection is invaluable to any corporate or individual computer user.

But even antivirus software is not enough to catch certain attacks that have been cleverly disguised. Once a system is infected with a new strain, the damage can be done before the virus or worm is detected and the system is disinfected.

Application protection

Such packages include software that watches for email anomalies, database access queries, or other behavior that may exploit vulnerability in the application. Because it must be very specific—and very close—to the application it is protecting, application protection is typically implemented as software on the host. Dedicated servers would benefit from well-designed application security software that will maintain the integrity of the code and detect anomalous behavior that could indicate an attack. Certain malicious code can attempt to overwrite registers on the end-system and thereby hijack the hardware for destructive purposes.

Intrusion detection systems

Intrusion Detection Systems (IDS) are designed to 'listen' to traffic and behavior and set an alarm if certain conditions are met. Some IDS implementations are implemented in the host, while others are deployed in the network. The IDS sensor monitors traffic, looking for protocol violations, traffic rate changes or matches to known attack 'signatures'. When a threat is detected, an alarm is sent to notify a (human) network administrator to intervene.

Host-based intrusion detection systems are designed as software running on general purpose computing platforms. Not to be confused with application security software (mentioned above), which runs on the end system and focuses primarily on Layers 5-7, software based intrusion systems must also focus on Layers 3 and 4 of the protocol stack. These packages rely on the CPU power of the host system to analyze traffic as it comes into the server. General purpose computers often lack the performance required to monitor real-time network traffic and perform their primary functions. Creating a bottleneck in the network or on the server actually helps the hacker accomplish his goal by restricting access to valuable resources.

End-systems provide the best environment for signature recognition because packets are fully reassembled and any necessary decryption has been performed. However, signature-based intrusion detection has its limitations, as described below.

The next step in the evolution of intrusion security was content-based Intrusion Prevention Systems (IPS). Unlike IDS, which require manual intervention from an administrator to stop an attack, a content-based IPS automatically takes action to prevent an attack once it is recognized. This can cut down response time to near zero, which is the ultimate goal of intrusion security.

IPS must be intelligent, however, or the remedy might actually accomplish the hacker's goal: denying resources to legitimate users.

Prevention mechanisms can also be harmful if detection is subject to false positives, or incorrect identification of intrusion. If the prevention action is to disable a port, protocol, or address, a false positive could result in denial of service to one or more legitimate users.

Network behavior analysis

An alternative to signature recognition is network behavior analysis (NBA). Rate-based systems must provide detailed analysis and/or control of traffic flow. A baseline of traffic patterns is established, usually during a learning mode in which the device only 'listens' without acting on any alarm conditions. A good system will have default parameters set to reasonable levels, but the 'listening' period is required to learn the traffic behavior on various systems. The listening period should be 'typical,' in the sense that no attacks or unusual traffic patterns should be present. For example, Saturday and Sunday are probably not good days to build a baseline for a corporate server that is much busier during the workweek. Periods of unusually high or low traffic also make bad listening intervals, such as Christmas vacation week, unusually high traffic due to external events (press releases, sales promotions, Super Bowl halftime shows, and so on).

Once a baseline is established, rate-based systems watch for deviations from the known traffic patterns to detect anomalies. Good systems will allow an administrator to override the baseline parameters if events causing traffic surges are foreseen, for example, a server backup scheduled overnight.

While signature-based systems are scrutinized for false-negatives, or failing to identify an attack, rate-based systems should be scrutinized for false positives, or misidentifying legitimate changes in traffic patterns as attacks. Whether setting alarms or taking preventative action, rate-based systems must be well-designed to avoid unnecessary overhead.

Equally important for rate-based systems are their analysis tools. Administrators should be able to view their traffic patterns on a variety of levels, and use this information to tune their network resources.

FortiDDoS compared with conventional firewalls

Conventional stateful firewalls drop packets or stateful connections, but they cannot correlate packets to a source. FortiDDoS has a unique feature that allows it to promptly correlate attacks and verify if they are initiated by a single host. If it can do that (in case it is a non-spoofed attack), it blocks the offending source for a longer period of time.

It is important to understand the differences between a stateful firewall and a stateful NBA system such as FortiDDoS. Here are the key differences: Conventional stateful firewalls have rules that allow or deny packets or individual connections based on their individual characteristics. They do not remember packets in an aggregate way.

FortiDDoS operates on an aggregate basis. It looks at packet rates—typically within one second, over a period of time. It measures packet rates for various Layer 3, 4, and 7 parameters and compares against thresholds set for them. If the rate exceeds the threshold, it blocks them for a configured period.

In a firewall, the administrator can set a rule that allows the UDP destination port 1434 regardless of the rate. A FortiDDoS administrator, on the other hand, can set a rule that allows UDP 1434 only if the rate is within 10 packets per second. Beyond this rate, the UDP packets destined to that port are dropped.

There are some features in FortiDDoS that are similar to a firewall. Like a firewall, FortiDDoS allows you to configure Layer 3, 4, and 7 blocking conditions. It is therefore important to learn how to migrate a firewall security policy to a FortiDDoS security policy.

FortiDDoS compared with conventional intrusion prevention systems

FortiDDoS is a rate-based IPS device that detects and blocks network attacks which are characterized by excessive use of network resources. It uses a variety of schemes, including anomaly detection and statistical techniques, to detect and block malicious network traffic. When it detects an intrusion, the FortiDDoS blocks traffic immediately, thus protecting the systems it is defending from being overwhelmed.

Unlike conventional content-based IPS, an NBA system does not rely on a predefined attack “signature” to recognize malicious traffic. An IPS is vulnerable to “zero-day” attacks, or attacks that cannot be recognized because no signature has been identified to match the attack traffic. In addition, attack traffic that is compressed, encrypted, or effectively fragmented can escape many pattern-matching algorithms in content-based IPS. And many rate-based attacks are based on genuine and compliant traffic being sent at high rates, effectively evading the IPS.

An NBA provides a network with unique protection capabilities. It delivers security services not available from traditional firewalls, IPS, or antivirus/spam detectors. The detection, prevention, and reporting of network attacks is based on traffic patterns rather than individual transaction or packet-based detection, which enables the FortiDDoS to serve a vital role in an effective security infrastructure. Rather than replacing these elements, an NBA complements their presence to form a defense-in-depth network security architecture.

FortiDDoS compared with conventional network behavior analysis

FortiDDoS is a hardware-based NBA solution. Unlike software-based solutions, it maintains normal levels of processing and data throughput during denial of service attacks.

FortiDDoS appliances are powered by one or more purpose-built FortiASIC-TP2 traffic processors that maintain massive connection tables and still perform with the lowest latency in the industry. Each FortiASIC-TP2 processor maintains the following resources:

- Source table with 1,000,000 rows. This table tracks the packet rate for every source IP address and is used for “per-source” thresholds.
- Destination table with 1,000,000 rows. This table tracks the packet rate for every destination IP address and is used for “per-destination” thresholds.
- Connection (session) table with 1,000,000 rows. This table tracks the status of every active TCP session. Connections are identified using the 4-tuple of Source IP Address, Source Port, Protected IP Address, and Destination Port. It is used for connection count and connection rate thresholds.
- Legitimate IP address table with 2,000,000 rows. This table tracks every IP address that has successfully created the TCP three-way handshake. Entries are timed-out in order to maintain the table as a source of recently validated source IP addresses.
- DNS query response match table with 1,900,000 rows. This table stores DNS queries so that it can match DNS responses. DNS responses that do not have a corresponding query are considered unsolicited response and are dropped. An entry is cleared when the matching response is received. Stale entries are periodically cleaned up.
- DNS TTL table with 1,500,000 rows. This table stores DNS query details correlated with the client IP address. During a flood, the system drops queries that have an entry in the table. It is not expected that a client would send the same query before the TTL expires.
- DNS legitimate query table that can store 128k unique queries. This table stores DNS query details for queries that have successful responses. An entry is cleared when the TTL expires. During a flood, the system drops queries that do not have an entry in the table.
- DNS cache that can store 64k responses. During a flood, the DNS response to valid queries can be served from the cache, reducing the load on the protected DNS server.

Figure 3 illustrates the number of FortiASIC-TP2 traffic processors for each FortiDDoS appliance model. Note the following:

- FortiDDoS 200B and FortiDDoS 400B—These models have 1 TP2.
- FortiDDoS 600B/800B—These models have 2 TP2s. Interfaces 1-8 are bound to one TP2, and interfaces 9-16 are bound to the other.
- FortiDDoS 900B/1000B and FortiDDoS 1200B/2000B—The FortiDDoS 900B/1000B has 3 TP2s, and the FortiDDoS 1200B/2000B has 6 TP2s. Sessions are distributed among the TP2s using a hash-based load balancing

algorithm. For TCP/UDP traffic, the hash includes Source IP / Source Port / Protected IP / Destination Port / Protocol. For non-TCP/non-UDP traffic, the hash includes Source IP / Protected IP / Protocol.

Figure 3: FortiASIC-TP2

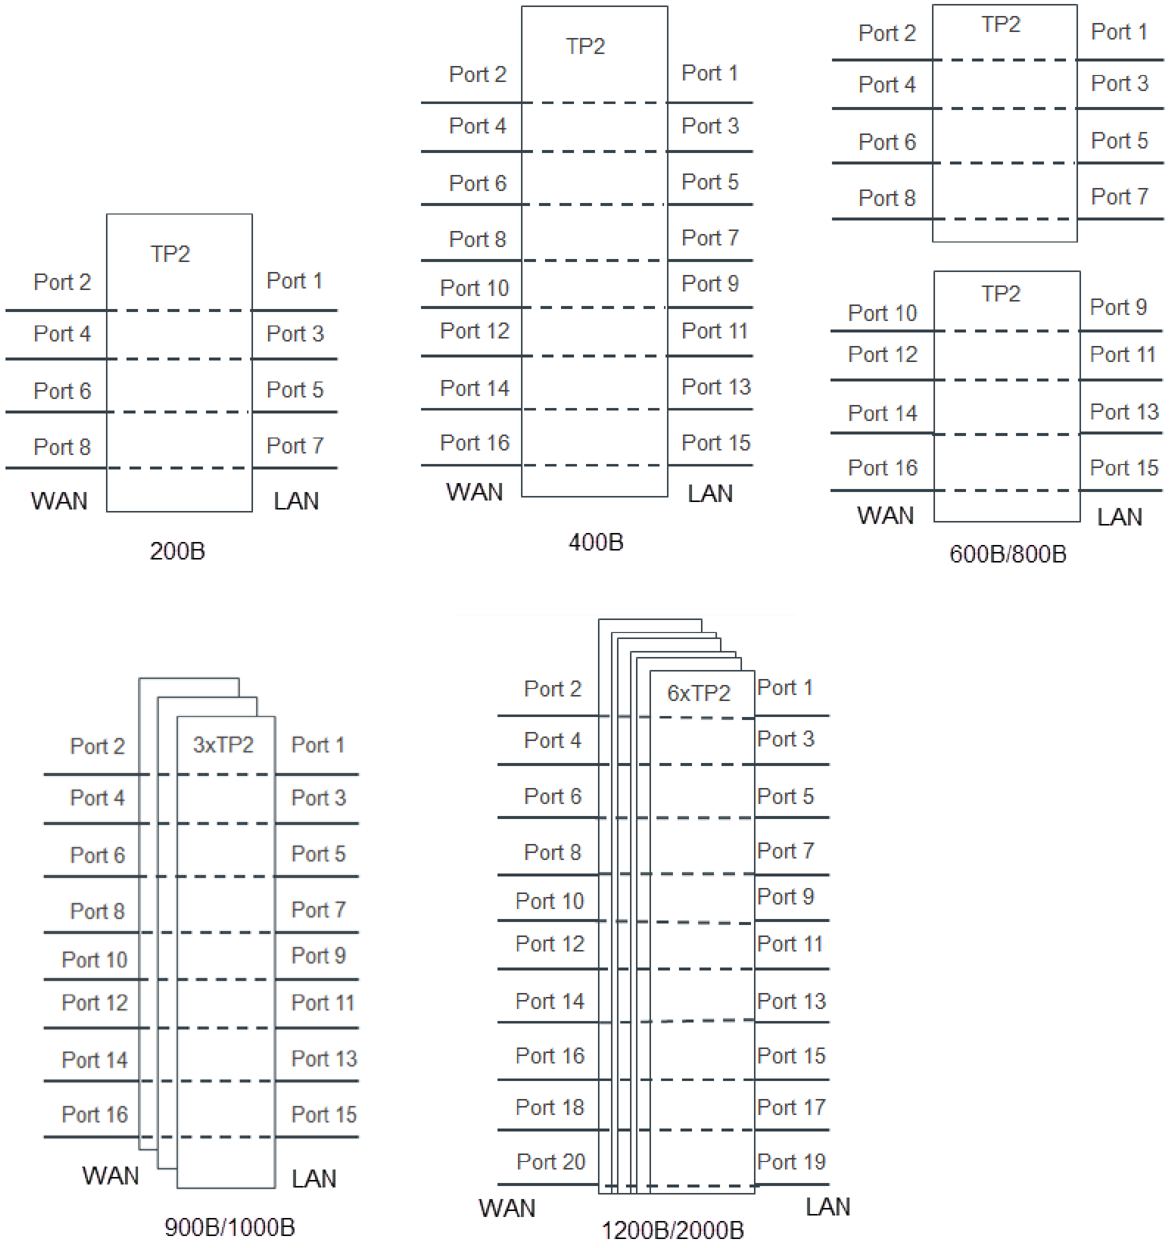

With its massive computing power, the FortiDDoS system maintains the magnitude of bidirectional traffic data that security administrators need to prevent DDoS attacks. The system uses counters, historical data, and predictive models to enforce intelligent rate limits based on granular Layer 3, Layer 4, and Layer 7 parameters and aggregations.

The result is excellent security, fewer false positives, and visibility into key trends.

Note: FortiDDoS 600B and 900B are not designed to support DNS ACLs, DNS anomaly detection, or DNS flood mitigation.

Understanding FortiDDoS rate limiting thresholds

This section includes the following information:

- Granular monitoring and rate limiting
- Source tracking table
- Destination tracking table
- Continuous learning and adaptive thresholds
- Hierarchical nature of protocols and implication on thresholds

Granular monitoring and rate limiting

Increasingly, instead of simple bandwidth attacks, attackers try to avoid detection by creating attacks that mimic the behavior of a large number of clients. Evading an NBA system is easy if attackers do coarse-grained rate-based control. Because the content of the malicious requests does not differ from that of legitimate ones, the resulting attacks are hard to defend against using standard techniques.

In contrast, FortiDDoS uses a combination of Layer 3, 4, and 7 counters and continuously adapts expected inbound and outbound rates for each of these traffic parameters.

Granular analytics also enable targeted mitigation responses. For example, if a few TCP connections are exceeding bandwidth, the system blocks those connections rather than all connections. If a single destination is under attack, FortiDDoS drops packets to that destination while others continue. During fragmented flood attacks, all non-fragmented packets continue as usual. During a port flood to a non-service port, the packets to other ports continue.

Granularity helps to increase the goodput (the throughput of useful data) of the system.

Table 1 lists the counters that FortiDDoS uses to detect subtle changes in the behavior of network traffic.

Table 1: FortiDDoS counters

Type	Counters
Layer 3	
Protocol flood	256 protocols per SPP per direction
Fragment flood	1 counter per SPP per direction
IP source flood & source tracking	1 million sources per TP2
IP destination flood	1 million destinations per TP2
Layer 4	
TCP port flood	65k per SPP per direction

Type	Counters
UDP port flood	65k per SPP per direction
ICMP type/code flood	256 types and 256 codes per SPP per direction
TCP connection flood	1 million connections per TP2
Legitimate IP table	2 million IP addresses per TP2
SYN flood	1 SYN counter per SPP per direction
SYN rate/source	1 million sources per TP2
SYN/destination	1 million destinations per TP2
HTTP Method/Source	1 million sources per TP2
Concurrent connections/source	1 million sources per TP2
Layer 7	
HTTP method	1 counter per SPP per direction
URLs	32,767 URLs per SPP per direction
Host	512 headers per SPP per direction
Referer	512 headers per SPP per direction
Cookie	512 headers per SPP per direction
User-Agent	512 headers per SPP per direction
DNS query rate	1 counter per SPP per direction
DNS query rate/source	1 million sources per TP2
DNS suspicious activity/source	1 million sources per TP2
DNS question count	1 counter per SPP per direction
DNS MX count	1 counter per SPP per direction
DNS All count	1 counter per SPP per direction
DNS zone transfer count	1 counter per SPP per direction
DNS fragment count	1 counter per SPP per direction

Source tracking table

FortiDDoS TP2 traffic processors maintain massive connection tables and still perform with low latency. Each TP2 has a source tracking table with 1,000,000 rows.

The source tracking table enables FortiDDoS to correlate sources with attack events and apply a more stringent blocking period to the sources that exceeded maximum rate limits.

The source tracking table also enables the special “per-source” thresholds described in Table 2.

Table 2: Per-source thresholds

Counter	Description
most-active-source	This counter establishes a maximum packet rate for any IP packet from a single source. A rate that exceeds the adjusted baseline is anomalous and treated as a Source Flood attack event. In conjunction with the Source Multiplier, the most-active-source threshold is useful in tracking and blocking non-spoofed sources that are participating in an attack. See Figure 17. How is the threshold determined? When it establishes baseline traffic statistics, FortiDDoS records the highest packet rate from a single source during the observation period. In a one hour observation period, FortiDDoS collects a data point for twelve five minute windows. The data point is the highest rate observed in any one second during the five minute window. If the packet rate data points for most-active-source are 1000, 2000, 1000, 2000, 1000, 2000, 1000, 2000, 3000, 2000, 1000, and 2000, the generated statistic is the highest one: 3000.
syn-per-src	This counter establishes a maximum packet rate for SYN packets from a single source. A rate that exceeds the adjusted baseline is anomalous and treated as a SYN Flood From Source attack event.
concurrent-connections-per-source	This counter establishes a maximum packet rate for concurrent connections from a single source. A count that exceeds the adjusted baseline is anomalous and treated as an Excessive Concurrent Connections Per Source attack event.
dns-query-per-src	This counter establishes the maximum rate of DNS queries from a single source. A count that exceeds the adjusted baseline is anomalous and treated as DNS Query Flood From Source attack event.
dns-packet-track-per-src	This counter is based on heuristics to detect suspicious actions from sources. The source tracking counter is incremented when a query is not found in the DQRM, when there are fragmented packets in the query or response, and when the response has an RCODE other than 0.
methods-per-source	Drops due to method per source threshold.

Destination tracking table

Each TP2 has a destination table with 1,000,000 rows. This table tracks the packet rate for every destination and is used for “per-destination” thresholds.

The destination tracking table enables FortiDDoS to prevent destination flood attacks and slow connection attacks that are targeted at individual destinations. The “per-destination” thresholds enable it to do so without affecting the rates for other destinations in the SPP.

Table 3 describes the per-destination thresholds.

Table 3: Per-destination thresholds

Counter	Description
most-active-destination	This counter establishes a maximum packet rate to any one destination. A rate that exceeds the adjusted baseline is anomalous and treated as a Destination Flood attack event. How is the threshold determined? When it establishes baseline traffic statistics, FortiDDoS records the highest packet rate to any single destination during the observation period. In a one hour observation period, FortiDDoS collects a data point for twelve five minute windows. The data point is the highest rate observed in any one second during the five minute window. If the packet rate data points for most-active-destination are 100000, 200000, 100000, 200000, 100000, 200000, 100000, 200000, 300000, 200000, 100000, and 2000, the generated statistic is the highest one: 300000.
syn-per-dst	This counter establishes a maximum packet rate for particular TCP packets to a single destination. A rate that exceeds the adjusted baseline is anomalous and treated as a Excessive TCP Packets Per Destination flood attack event. When the syn-per-dst limits are exceeded for a particular destination, the SYN flood mitigation mode tests are applied to all new connection requests to that particular destination. Traffic to other destinations is not subject to the tests.

Continuous learning and adaptive thresholds

Most NBA systems use fixed value thresholds. Traffic, however, is never static. It shows trends and seasonality (a predictable or expected variation).

FortiDDoS uses adaptive thresholds. Adaptive thresholds take into account the traffic’s average, trend, and seasonality (expected or predictable variations).

Traffic prediction

Unlike other network behavior analysis (NBA) systems, FortiDDoS never stops learning. It continuously models inbound and outbound traffic patterns for key Layer 3, Layer 4, and Layer 7 parameters.

FortiDDoS uses the following information to model normal and abnormal traffic:

- The historical base, or weighted average, of recent traffic (more weight is given to recent traffic)
- The trend, or slope, of the traffic
- The seasonality of traffic over historical time periods

Figure 4: Trend, slope, and base of traffic

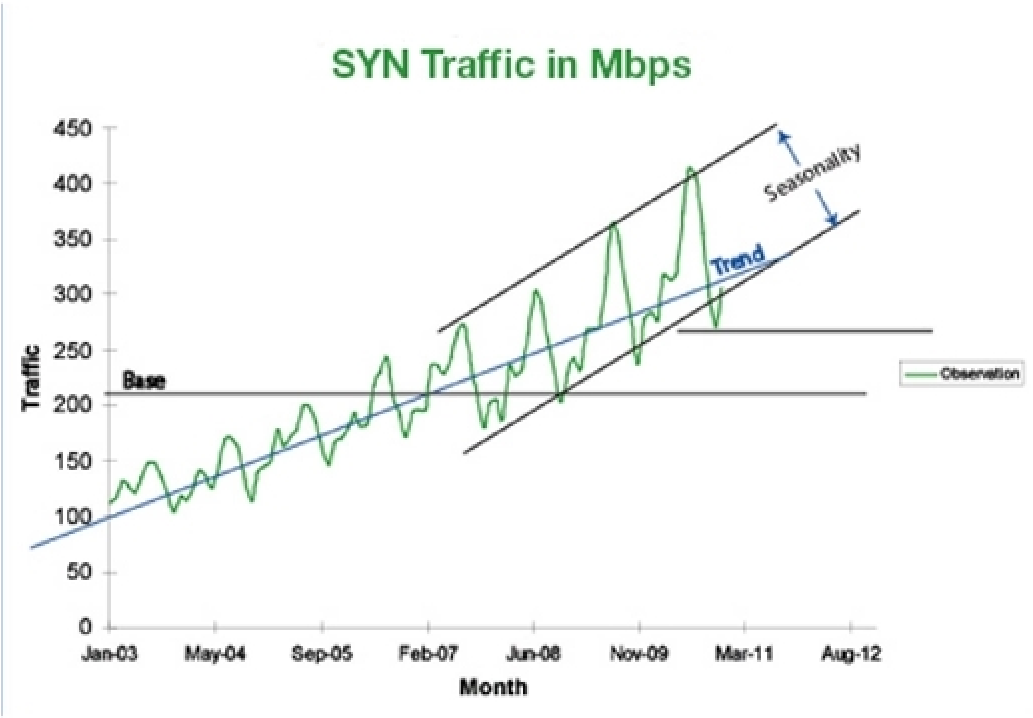

FortiDDoS uses these statistics to create a forecast for the next traffic period.

Figure 5: Forecast vs. actual traffic

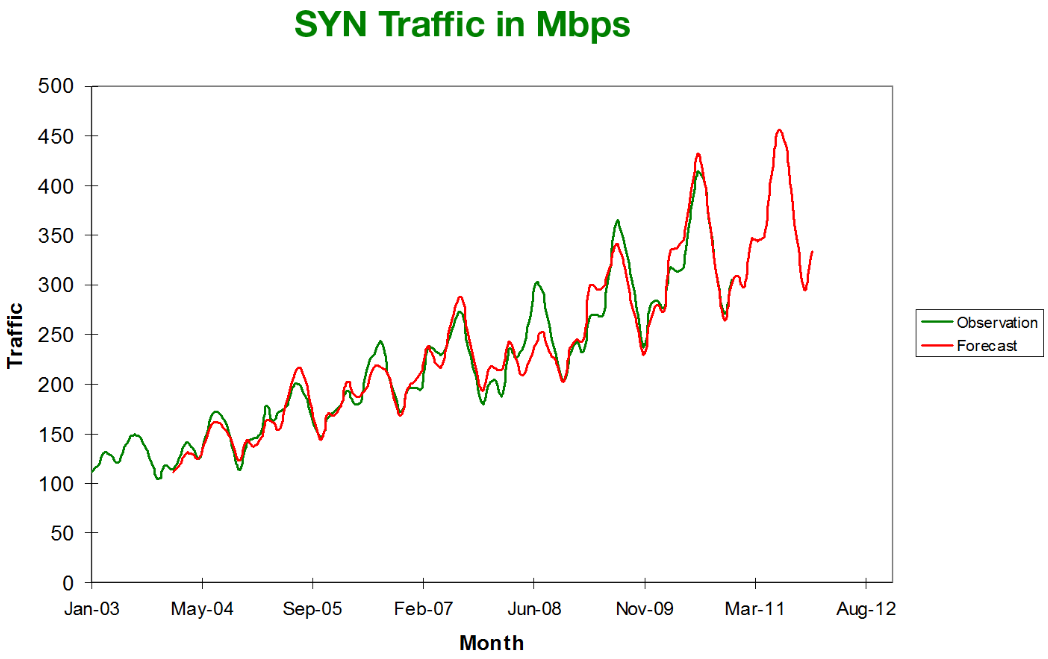

Traffic is nondeterministic; therefore, the forecast cannot be exact. The extent to which an observed traffic pattern is allowed to exceed its forecast is bounded by thresholds. Generally speaking, a threshold is a baseline rate that the system uses to compare observed traffic rates to determine whether a rate anomaly is occurring.

The FortiDDoS system maintains multiple thresholds for each key Layer 3, Layer 4, and Layer 7 parameter:

- Configured minimum threshold
- Estimated threshold
- Adaptive limit maximum threshold
- Adjustments for proxy IP addresses
- Packet count multipliers applied to traffic associated with an attack

Figure 6 illustrates how the system maintains multiple thresholds. The sections that follow explain the significance of each.

Figure 6: Adaptive, minimum and fixed

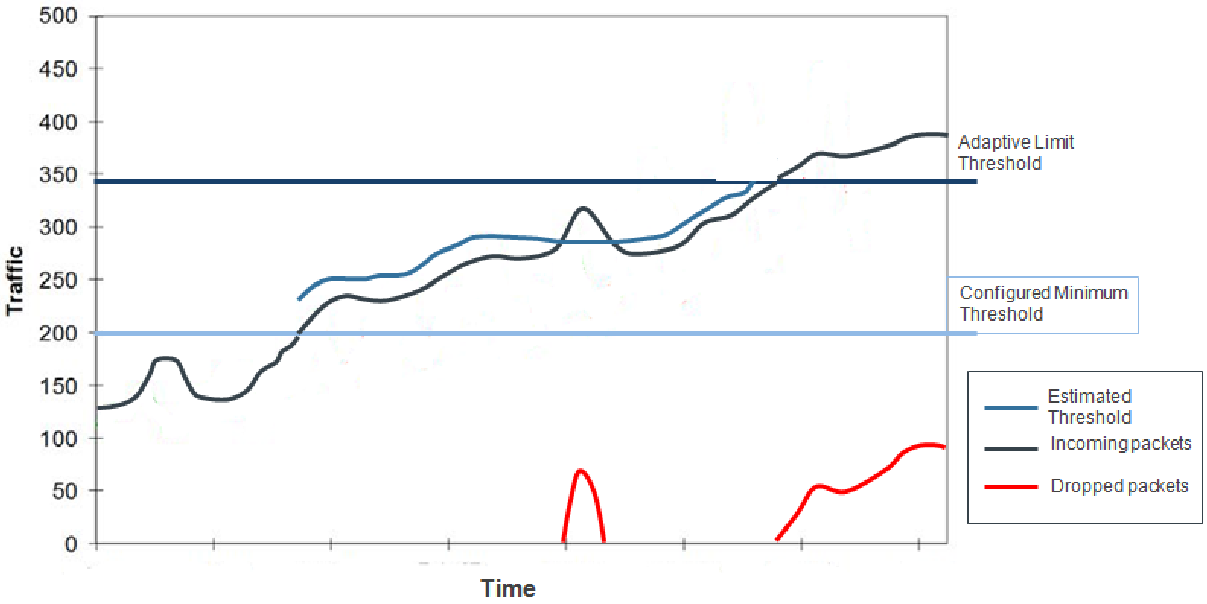

Configured minimum thresholds

The configured minimum threshold is a baseline of normal counts or rates. The baseline can be generated (based on statistics collected during the learning period) or stipulated (based on defaults or manually configured settings).

The configured minimum threshold is a factor in setting rate limits, but it is not itself the rate limit. Rate limits are set by the estimated threshold, a limit that is subject to heuristic adjustment based on average, trend, and seasonality.

Many of the graphs in the Monitor menu display the configured minimum threshold as a reference.

Table 4 summarizes the alternative methods for setting the configured minimum threshold.

Table 4: Setting the configured minimum threshold

Menu	Usage
Protection Profile > Thresholds > System Recommendation	The recommended method for setting the configured minimum thresholds. The configured minimum thresholds are a product of the observed rates adjusted by a percentage that you specify.
Protection Profile > Thresholds > Thresholds	The thresholds configuration is open. You can set user-defined thresholds and fine-tune them. You might be able to set reasonable values for port and protocol thresholds based on your knowledge of your network’s services and server capacity. Most likely, you must become a FortiDDoS expert before you will be able to set reasonable values for Scalar thresholds.
Protection Profile > Thresholds > Emergency Setup	Use if you do not have time to use Detection Mode to establish a baseline.
Protection Profile > Thresholds > Factory Default	Use to quickly restore the system to high values. The factory defaults are high to avoid possible traffic disruption when you first put the system inline. In general, you use these settings together with Detection Mode when you are setting an initial baseline or a new baseline.
Protection Profile > Thresholds > Percent Adjust	Use when you expect a spike in legitimate traffic due to an event that impacts business, like a news announcement or holiday shopping season.

Estimated thresholds

The estimated threshold is a calculated rate limit, based on heuristic adjustments.

The system models an adjusted normal baseline based on average, trend, and seasonality. It uses the heuristics to distinguish attack traffic from increases in traffic volume that is the result of legitimate users accessing protected resources.

The minimum value of an estimated threshold is the configured minimum threshold. In other words, if it is not predicting normal traffic becoming heavier than the baseline, it allows a rate at least as high as the configured minimum threshold.

The maximum value of an estimated threshold is the product of the configured minimum threshold and the adaptive limit. In other words, the system does enforce an absolute maximum rate limit.

Adaptive limit

The adaptive limit is a percentage of the configured minimum threshold.

An adaptive limit of 100% means no dynamic threshold estimation adjustment takes place once the configured minimum threshold is reached (that is, the threshold is a fixed value).

The product of the configured minimum threshold and adaptive limit is the absolute maximum rate limit. If the adaptive limit is 150% (the default), the system can increase the estimated threshold up to 150% of the value of the configured minimum threshold.

There are scenarios where FortiDDoS drops legitimate traffic because it cannot adapt quickly enough to a sudden change in traffic patterns. For example, when a news flash or other important announcement increases traffic to a company's website. In these situations, you can use the Protection Profiles > Thresholds > Percent Adjust configuration page to increase all configured thresholds by a specific percentage.

Adjustments for proxy IP addresses

FortiDDoS can take account of the possibility that a source IP address might be a proxy IP address, and adjust the threshold triggers accordingly. If a source IP address is determined to be a proxy IP address, the system adjusts thresholds for a few key parameters by a factor you specify on the Global Settings > Proxy IP page.

Packet count multipliers applied to traffic associated with an attack

Packet count multipliers are adjustments to counters that are applied to traffic associated with an attack so that the thresholds that control drop and block responses are triggered sooner. You can configure multipliers for the following types of traffic:

- Source floods—Traffic from a source that the system has identified as the source of a flood.
- Layer 7 floods—Traffic for attacks detected based on a URL or Host, Referer, Cookie, or User-Agent header field.

You can use the Protection Profiles > Settings page to specify packet count multipliers.

When both Source flood and Layer 7 flood conditions are met, the packet count multipliers are compounded. For example, when there is a User Agent flood attack, a source is sending a User-Agent that is overloaded. If the Source multiplier is 4 and the Layer 7 multiplier is 64, the total multiplier that is applied to such traffic is 4 x 64 = 264. In effect, each time the source sends a Layer 7 packet with that particular User-Agent header, FortiDDoS considers each packet the equivalent of 256 packets.

Hierarchical nature of protocols and implication on thresholds

An HTTP packet has Layer 7, Layer 4 (TCP), and Layer 3 (IP) properties. See Figure 7. A packet must be within the estimated thresholds of all these counters in order to pass through the FortiDDoS gateway. When it sets recommended thresholds, the system takes account of this complexity. If you set thresholds manually, you must also be sure that Layer 7 rates are consistent with Layer 4 and Layer 3 rates.

Figure 7: Protocol hierarchy for determining thresholds

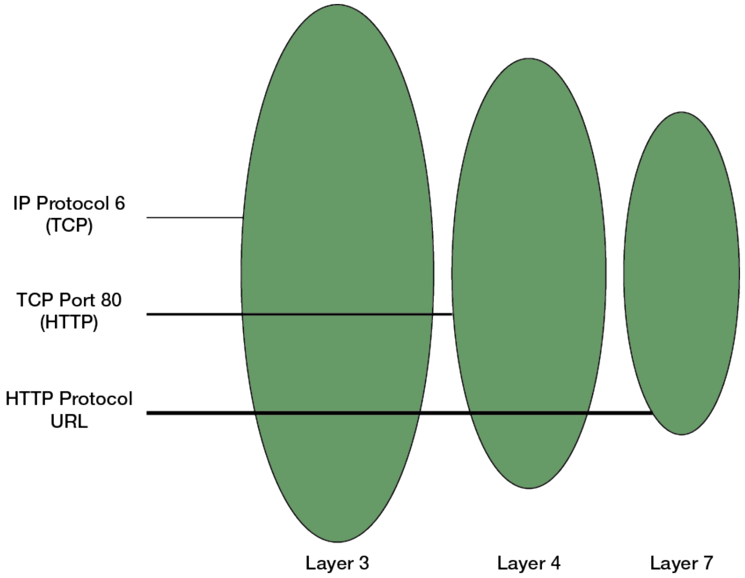

Figure 8 illustrates system processing for an HTTP packet.

Figure 8: HTTP packet properties

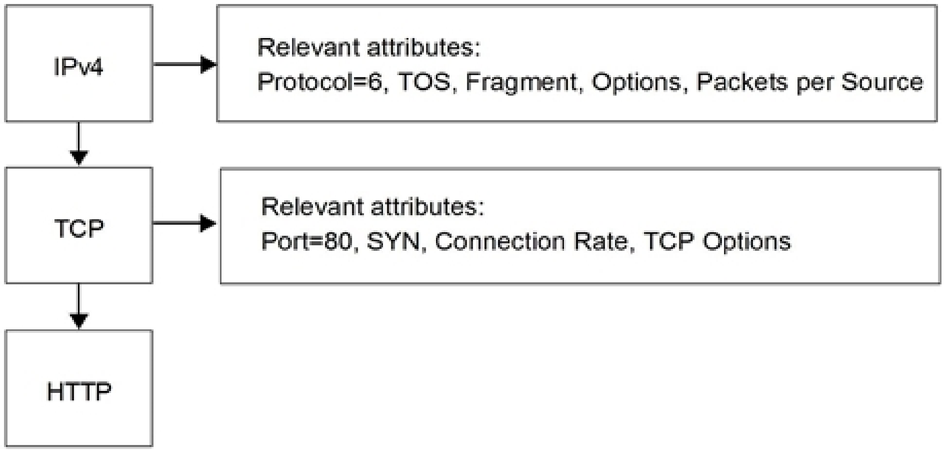

Relevant attributes:
Protocol=6, TOS, Fragment, Options, Packets per Source

Relevant attributes:
Port=80, SYN, Connection Rate, TCP Options

The following IPv4 packet properties are tracked:

- Protocol
- Fragment or not a fragment
- Source IP address (the system can monitor the packet rate from that specific source)
- The following TCP packet properties are tracked:
- Destination port
- SYN or not a SYN packet
- TCP connection tuple (the system can monitor the packet rate within that connection)

An HTTP packet has the following properties that can be tracked:

- Method (for example, GET)
- URL
- Headers

Figure 9 illustrates system processing for a UDP packet.

Figure 9: UDP packet properties

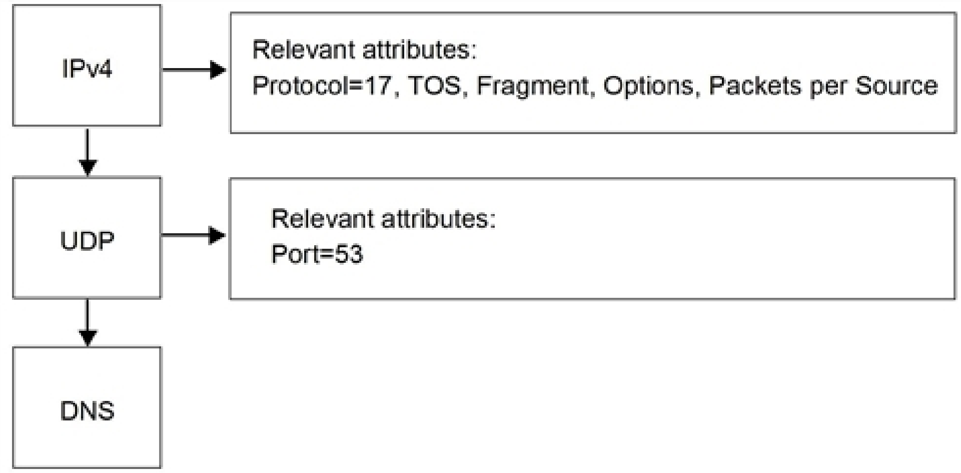

The following IPv4 packet properties are tracked:

- Protocol
- Fragment or not a fragment
- IP option values
- Source IP address, and hence packet rate from that specific source

The following UDP packet properties are tracked:

- Destination port

A DNS message has the following additional properties that can be tracked in queries and responses:

- QR code (query or response)
- QTYPE
- Question count
- RCODE

If a server supports TCP, UDP, and ICMP services, the rate of IP packets is an aggregate of rates for TCP, UDP, and ICMP packets. Similarly if the same server is a web server as well as an SMTP server, the TCP packet rate is the sum of packet rates for port 80 and port 25.

To summarize, because determining thresholds is a hierarchical process, avoid setting low thresholds on common conditions that can cause FortiDDoS to block legitimate traffic as well as attack traffic. The more specific you are about the type of traffic you want to allow as ‘normal’, the more effective the FortiDDoS is in blocking other traffic.

Using FortiDDoS ACLs

You can configure access control lists (ACLs) to deny known attacks and unwarranted traffic. For example, in a data center environment, you can use ACLs to protect the router from getting overloaded by floods from known attacks.

The ACLs are part of the core hardware architecture, so they do not add to latency through the device when you enable or disable them.

FortiDDoS enforces a Global ACL that applies to all traffic, and SPP ACLs that are applied after traffic has been sorted into an SPP.

The Global ACL features include:

- An anti-spoofing ACL based on the local address configuration
- An ACL based on source IP address
- An ACL based on Geolocation addresses
- An ACL based on the IP Reputation list provided through FortiGuard

It is possible for traffic to be denied based on multiple Global ACL rules, but only one deny reason-code is logged. The reason-code is based on the following order of precedence.

1. Anti-spoofing
2. Source IP address
3. GeoLocation
4. IP Reputation

You can configure additional ACLs per SPP. The SPP ACL rules can be based on source IP address, service, or Layer 7 parameter.

The following table summarizes the traffic parameters you can use to enforce an ACL.

Table 5: ACL parameters

Parameter	ACL
Layer 3	
Any protocol (up to 256)	SPP
Fragment	SPP
IP netmask or address (up to 4 billion)	Global, SPP
Geolocation (countries and regions), anonymous proxy, satellite provider	Global
IP-reputation (based on data from external public sources)	IP Reputation (subscription)

Parameter	ACL
Layer 4	
TCP port (up to 64k)	SPP
UDP port (up to 64k)	SPP
ICMP type/code (up to 64k)	SPP
Layer 7	
URLs (up to 32k)	SPP
Host (512)	SPP
Referer (512)	SPP
Cookie (512)	SPP
User-Agent (512)	SPP
DNS-All	SPP
DNS-Fragment	SPP
DNS-MX	SPP
DNS-Zone-Transfer	SPP
Restrict DNS Queries to specific subnets	SPP
DNS Blacklisted Domains	Global
DNS Blacklisted IPv4 Addresses	Global

Understanding FortiDDoS protocol anomaly protection

This section includes the following topics:

- TCP/IP anomalies
- TCP session state anomalies
- HTTP anomalies
- DNS anomalies

TCP/IP anomalies

Legitimate traffic conforms with standards set out in Internet Engineering Task Force (IETF) documents known as Requests for Comments (RFC). Traffic that does not conform with RFCs is anomalous. Often, anomalous traffic contains malicious components. In any case, it should be dropped to prevent resource issues.

The FortiDDoS system drops and logs the following Layer 3 anomalies:

- IP version other than 4 or 6
- Header length less than 5 words
- End of packet (EOP) before 20 bytes of IPv4 Data
- Total length less than 20 bytes
- EOP comes before the length specified by Total length
- End of Header before the data offset (while parsing options)
- Length field in LSRR/SSRR option is other than (3+(n*4)) where n takes value greater than or equal to 1
- Pointer in LSRR/SSRR is other than (n*4) where n takes value greater than or equal to 1
- For IP Options length less than 3
- Source and destination addresses are the same (LAND attack)
- Source or destination address is the same as the localhost (loopback address spoofing)

The FortiDDoS system drops and logs the following Layer 4 anomalies:

- Checksum errors
- Invalid flag combinations, such as SYN/RST
- Other header anomalies, such as incomplete packet
- Urgent flag is set then the urgent pointer must be non-zero
- SYN or FIN or RST is set for fragmented packets
- Data offset is less than 5 for a TCP packet
- End of packet is detected before the 20 bytes of TCP header
- EOP before the data offset indicated data offset
- Length field in Window scale option other than 3 in a TCP packet
- Missing UDP payload
- Missing ICMP payload
- SYN with payload

TCP session state anomalies

TCP session state anomalies are a symptom of an attack or invalid junk traffic, but they can also be seen as a by-product of traffic load tools used in test environments. You can use the Protection Profiles > SPP Settings configuration page to enable detection for TCP session state anomalies and to allow for the anomalies that are sometimes triggered by traffic load tools.

Table 6 summarizes recommended settings for TCP session state for the FortiDDoS deployment modes. In a typical Prevention Mode deployment where FortiDDoS receives both sides of the TCP connection, all settings are available and can be useful. Some settings are not appropriate when FortiDDoS is deployed in Detection Mode or Asymmetric Mode. See Understanding FortiDDoS Detection Mode or Understanding FortiDDoS Asymmetric Mode for TCP for additional information on the guidelines for those modes.

Table 6: TCP session state anomalies detection options

Setting	Detection Mode	Prevention - Symmetric	Prevention - Asymmetric
Sequence validation Drops packets with invalid TCP sequence numbers.	Do not enable	Recommended	Do not enable
SYN validation Drops SYNs during a flood if the source has not completed the TCP three-way handshake.	Do not enable	Recommended	Recommended
State transition anomalies validation Drops packets with TCP state transitions that are invalid. For example, if an ACK packet is received when FortiDDoS has not observed a SYN/ACK packet, it is a state transition anomaly.	Do not enable	Recommended	Do not enable
Foreign packet validation Drops TCP packets without an existing TCP connection and reports them as a foreign packet. In most cases, the foreign packets validation is useful for filtering out junk.	Do not enable	Recommended	Recommended
Allow tuple reuse Allows tuple reuse. Updates the TCP entry during the closed or close-wait, fin-wait, time-wait states, when the connection is just about to retire.	Recommended	Recommended	Recommended

Setting	Detection Mode	Prevention - Symmetric	Prevention - Asymmetric
Allow duplicate SYN-in-SYN-SENT Allows duplicate TCP SYN packets during the SYN-SENT state. It allows this type of packet even if the sequence numbers are different.	Recommended	Useful in some lab environments	Useful in some lab environments
Allow duplicate-SYN-in-SYN-RECV Allows duplicate TCP SYN packets during the SYN-RECV state. It allows this type of packet even if the sequence numbers are different.	Do not enable	Useful in some lab environments	Do not enable
Allow SYN anomaly, **Allow SYN-ACK anomaly**, **Allow ACK anomaly**, **Allow RST anomaly**, **Allow FIN anomaly** Allows duplicate TCP packets during any other state even if the sequence numbers are different from the existing connection entry. This is equivalent to allowing the packet without updating an existing connection entry with the new information from the allowed packet.	Do not enable	Seldom necessary but available in case these anomalies are false positives in legitimate traffic.	Do not enable

HTTP anomalies

You can use the Global Settings > Settings > General tab to enable detection for the following HTTP anomalies:

- Known Method Anomaly—Drops HTTP traffic that uses one of the eight known methods: GET, HEAD, OPTIONS, PUT, POST, CONNECT, DELETE, or TRACE. By default, all the methods are treated as valid and therefore no Monitor graphs are provisioned, even if there are drops.
- Unknown Method Anomaly—Drops HTTP traffic that uses a method other than one of the following: GET, HEAD, OPTIONS, PUT, POST, CONNECT, DELETE, or TRACE. For example, TEST or PROPFIND. The dropped packets will be shown in the Monitor Graphs as well as in the Attack Log.
- Invalid HTTP Version Anomaly—Drops HTTP traffic with an HTTP version other than one of the following: 0.9, 1.0, or 1.1. The dropped packets will be shown in the Monitor Graphs as well as in the Attack Log.
- Do Not Parse HTTP 0.9—Drops sessions when the HTTP request includes the HTTP Range header. The Range header can be abused by attackers to exhaust HTTP server resources. There are no drops associated with this feature and the default setting is to treat HTTP 0.9 packets as HTTP packets and further parse.

DNS anomalies

DNS anomalies are packet or session state irregularities known to be exploited by attackers.

Table 7 lists the types of DNS anomalies that can be detected.

Table 7: DNS anomaly detection

Group	Anomaly
DNS header anomaly	• Invalid op-code—Invalid value in the OpCode field. • Illegal flag combination—Invalid combination in the flags field. • SP, DP both 53—Normally, all DNS queries are sent from a high-numbered source port (49152 or above) to destination port 53, and responses are sent from source port 53 to a high-numbered destination port. If the header has port 53 for both, it is probably a crafted packet.
DNS query anomaly	• Query bit set—DNS query with the query reply (QR) bit set to 1. In a legitimate query, QR=0. • Null query—DNS query in which the question, answer, additional, and name server counts are 0. • RA bit set—DNS query with the recursion allowed (RA) bit set. The RA bit is set in responses, not queries. • QDCNT not 1 in query—Number of entries in the question section of the DNS packet is normally 1. Otherwise, it might be an exploit attempt.
DNS response anomaly	• QCLASS in reply—DNS response with a resource specifying a CLASS ID reserved for queries only (QCLASS). • QTYPE in reply—DNS response with a resource specifying a TYPE ID reserved for queries only (QTYPE). • Query bit not set—DNS response with the query reply (QR) bit set to 0. In a legitimate response, QR=1. • QDCNT not 1 in response—Number of entries in the question section of the DNS packet is normally 1. Otherwise, it might be an exploit attempt.
DNS buffer overflow anomaly	• TCP Message too long—TCP query or response message that exceeds the maximum length specified in the message header. • UDP message too long—UDP query or response message that exceeds the maximum length specified in the message header. • Label length too large—Query or response with a label that exceeds the maximum length (63) specified in the RFC. • Name too long—DNS name that exceeds 255 characters. This can cause problems for some DNS servers.

Group	Anomaly
DNS exploit anomaly	• Pointer loop—DNS message with a pointer that points beyond the end of data (RFC sec4.1.4). This is an exploit attempt. • Zone transfer—An asynchronous Transfer Full Range (AXFR) request (QTYPE=252) from untrusted networks is likely an exploit attempt. • Class is not IN—A query/response in which the question/resource address class is not IN (Internet Address). Although allowed by the RFC, this is rare and might indicate an exploit attempt. • Empty UDP message—An empty message might indicate an exploit attempt. • Message ends prematurely—A message that ends prematurely might indicate an exploit attempt. • TCP Buffer underflow—A query/response with less than two bytes of data specified in the two-byte prefix field.
DNS info anomaly	Type ALL used—Detects a DNS request with request type set to ALL (QTYPE-E=255). Typical user queries to not request ALL.
DNS data anomaly	• Invalid type, class—A query/response with TYPE or CLASS reserved values. • Extraneous data—A query/response with excess data in the packet after valid DNS data. • TTL too long—TTL value is greater than 7 days (or 604800 seconds). • Name length too short—A query/response with a null DNS name.

Understanding FortiDDoS Detection Mode

In Detection Mode, FortiDDoS logs events and builds traffic statistics for SPPs, but it does not take actions: it does not drop or block traffic, and it does not aggressively age connections. Packets are passed through the system to and from protected subnets. Any logs and reports that show drop or blocking activity are actually simulations of drop or block actions the system would have taken if it were deployed in Prevention Mode.

When you get started with FortiDDoS, you deploy it in Detection Mode for 2-14 days so that the FortiDDoS system can learn the baseline of normal inbound and outbound traffic. The length of the initial learning period depends upon the seasonality of traffic (its predictable or expected variations) and how representative of normal traffic conditions the learning period is. Ensure that there are no attacks during the initial learning period and that it is long enough to be a representative period of activity. If activity is heavier in one part of the week than another, ensure that your initial learning period includes periods of both high and low activity. Weekends alone are an insufficient learning period for businesses that have substantially different traffic during the week. Thus, it is better to start the learning period on a weekday. In most cases, 7 days is sufficient to capture the weekly seasonality in traffic.

At the end of the initial learning period, you can adopt system-recommended thresholds (usually lower than the factory default) and continue to use Detection Mode to review logs for false positives and false negatives. As needed, you repeat the tuning: adjust thresholds and monitor the results.

When you are satisfied with the system settings, change to Prevention Mode. In Prevention Mode, the appliance drops packets and blocks sources that violate ACL rules and DDoS attack detection thresholds.

Important: In Detection Mode, the FortiDDoS system forwards all packets, but a simulated drop might be recorded. TCP session control options depend on the true TCP state, and simulated drops when the appliance is in Detection Mode can lead to unexpected results. For example, if the system records a (simulated) drop for a TCP connection, when subsequent packets arrive for the connection, the system treats them as foreign packets because the state table entry indicates the session has already been closed.

Table 8 summarizes our guidelines for SYN flood mitigation and TCP session state settings in Detection Mode.

Table 8: SYN flood mitigation and TCP state anomaly detection settings

Settings	Guidelines
Global Settings > Settings	
SYN Flood Mitigation	The SYN flood mitigation mode settings are not applicable and disregarded. In Detection Mode, the FortiDDoS system does not drop packets, so it cannot test the legitimacy of source IP addresses.
Protection Profiles > SPP Settings > TCP session feature control	
SYN validation	Do not enable. This option enables SYN flood mitigation mode, which is not applicable in Detection Mode.
Sequence validation	Do not enable. Simulated “drops” in Detection Mode lead to incorrect window validations for subsequent session packets.

Settings	Guidelines
State transition anomalies validation	Do not enable. Simulated “drops” in Detection Mode lead to faulty tracking of session state.
Foreign packet validation	Do not enable. Simulated “drops” in Detection Mode lead to unexpected foreign packet violations.
Allow tuple reuse	Exception to the rule. Enabled by default to support standard test environments that reuse tuples in quick succession. The setting is valid in Detection Mode. Recommended to avoid unnecessary logging of the event when it is detected.
Allow duplicate SYN-in-SYN-SENT	Exception to the rule. Not enabled by default, but the setting is valid in Detection Mode. Recommended to avoid unnecessary logging.
Allow duplicate SYN-in-SYN-RECV Allow SYN anomaly Allow SYN-ACK anomaly Allow ACK anomaly Allow RST anomaly Allow FIN anomaly	Do not enable.

Understanding FortiDDoS Prevention Mode

This section includes the following information about attack mitigation features when you deploy FortiDDoS in Prevention Mode:

- SYN flood mitigation
- Aggressive aging
- Rate limiting
- Blocking
- Reducing false positives

SYN flood mitigation

This section includes the following information:

- Overview
- ACK Cookie
- SYN Cookie
- SYN Retransmission

Overview

When a client attempts to start a TCP connection to a server, the client and server perform a three-way handshake:

The client sends a SYN message to the server to request a connection.

The server creates an entry for the connection request in the Transmission Control Block (TCB) table with status SYN-RECEIVED, sends an acknowledgment (SYN-ACK) to the client, and waits for a response.

The client responds with an acknowledgment (ACK), the connection is established, and the server updates the entry in the TCB table to ESTABLISHED.

Figure 10: TCP Connection Three-Way Handshake

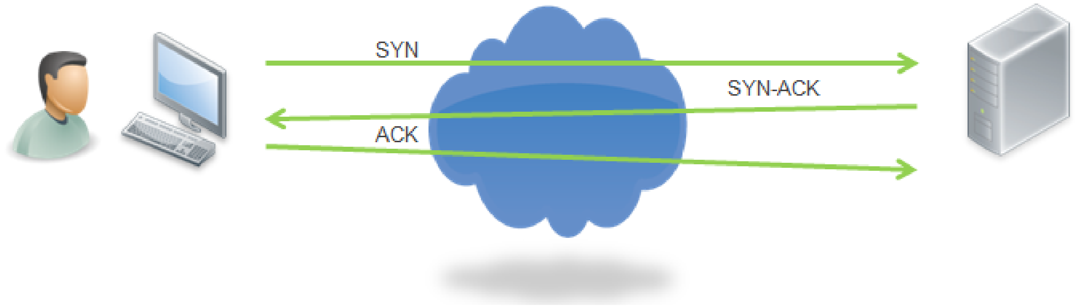

A SYN flood attack on a server exploits how the server maintains TCP connection state for the three-way handshake in the TCB table. In a spoofed attack, the attacker sends a large number of SYN packets from spoofed IP addresses to the server; or in a zombie attack, the attacker has used a virus to gain control of unwitting clients and sends a large number of SYN packets from legitimate IP addresses to the server. Each SYN packet that arrives creates an entry in the table. The spoofed addresses make it impossible to resolve the three-way handshake, and the TCP connection state in the TCB table remains ‘half-open’ instead of completing the cycle. It never transitions to ‘established’ and ultimately to ‘closed’. As a result, TCB table entries are not “cleaned up” by the expected life cycle, resources can be exhausted, and there can be system failure and outages.

Figure 11: Half-Open TCP Connection SYN Flood Attack

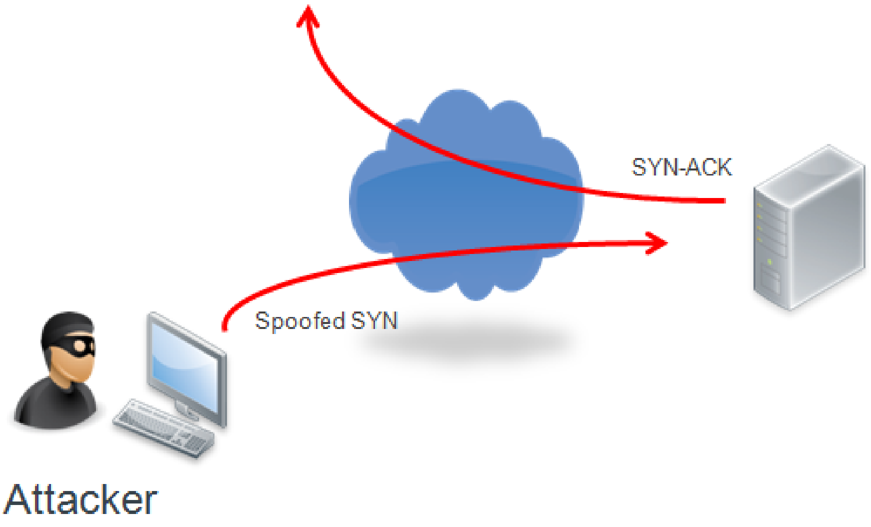

To prepare for SYN flood attacks, FortiDDoS maintains a table of IP addresses that have completed a three-way handshake. This is called the legitimate IP address (LIP) table. Entries in the LIP expire after 5 minutes.

When FortiDDoS detects a SYN flood attack, it enters SYN flood mitigation mode. In this mode, the system acts as a proxy for TCP connection requests and uses the LIP table to validate new connections:

- New SYN packets coming from addresses in the LIP table are presumed legitimate and are allowed
- FortiDDoS takes a guarded approach to other SYN packets, and they are processed according to the configured SYN flood mitigation mode option:
  - ACK Cookie
  - SYN Cookie
  - SYN Retransmission

The SYN flood mitigation mode behavior applies only when FortiDDoS has detected a SYN flood with either of the following thresholds:

- syn: When total SYNs to the subnet exceeds the threshold, the SYN flood mitigation mode tests are applied to all new connection requests.
- syn-per-dst: When the per-destination limits are exceeded for a particular destination, the SYN flood mitigation mode tests are applied to all new connection requests to that particular destination. Traffic to other destinations is not subject to the tests.

ACK Cookie

Figure 12 illustrates the ACK Cookie mitigation mode option. FortiDDoS sends the client two ACK packets: one with a correct ACK number and another with a wrong number. The system determines whether the source is not spoofed based on the client’s response. If the client’s response indicates that the source is not spoofed, FortiDDoS allows the connection and adds the source to the legitimate IP address table. Fortinet recommends this option if you have enough bandwidth in the reverse direction of the attack. (This method generates 2 responses for every SYN. Thus, a 1 Gbps SYN flood will generate 2 Gbps reverse traffic.)

Figure 12: SYN Flood Mitigation Mode—ACK Cookie

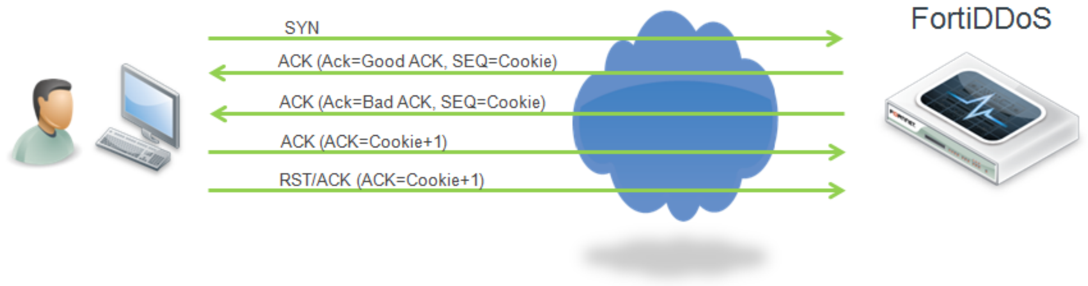

SYN Cookie

Figure 13 illustrates the SYN Cookie mitigation mode option. FortiDDoS sends a SYN/ACK with a cookie value in the TCP sequence field. If it receives an ACK back with the right cookie, a RST/ACK packet is sent and the IP address is added to the LIP table. If the client then retries, it succeeds in making a TCP connection. Fortinet recommends this option if you cannot use ACK Cookie and you anticipate high volume attacks.

Figure 13: SYN Flood Mitigation Mode—SYN Cookie

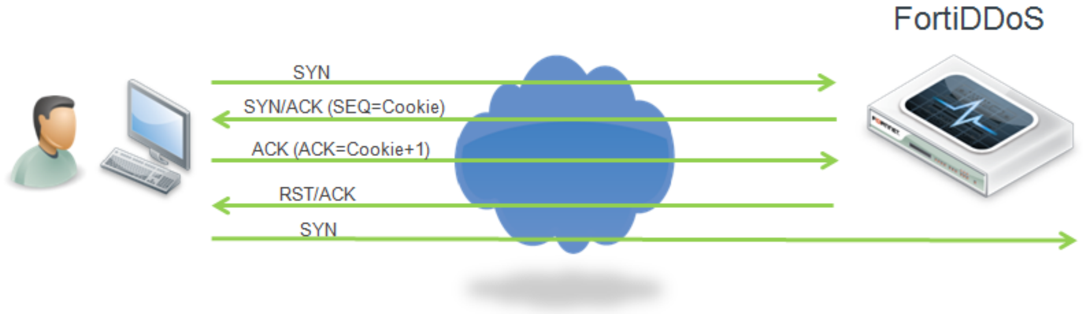

SYN Retransmission

Figure 14 illustrates the SYN Retransmission mitigation mode option. FortiDDoS drops the initial SYNs to force the client to send a SYN again. If a preconfigured number of retransmitted SYNs arrive within a predefined time period, the FortiDDoS considers the source to be legitimate. It allows the connection to go through and adds the source to the legitimate IP address table. Fortinet recommends this option if you cannot use ACK Cookie and you anticipate low volume attacks.

Figure 14: SYN Flood Mitigation Mode—SYN Retransmission

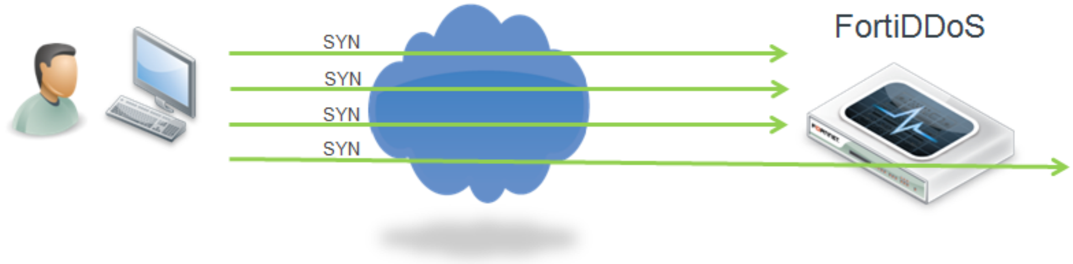

Aggressive aging

This section includes the following topics:

- Slow connection detection and aggressive aging
- Rate anomalies and aggressive aging
- Idle connections and aggressive aging

Slow connection detection and aggressive aging

Slow connection attacks are Layer 7 attacks that aim to make a service unavailable or increase latency to a service. These attacks are not detected by Layer 4 detection methods because these are legitimate TCP or UDP connections. With these attacks, distinguishing attackers from legitimate users is a complex task.

Variations of the Slowloris attack involve opening a legitimate TCP connection and not doing anything at all. Such idle connections fill up the connection tables in firewall and servers.

FortiDDoS can detect slow connection attacks and combat them by “aggressively aging” idle connections. When slow connection detection is enabled, the system monitors TCP ports 21, 22, 23, 25, 80, and 443, as well as user-configured HTTP service ports, for slow connection anomalies. If the traffic volume for a connection is below a specified byte threshold during an observation period, the connection is deemed a slow connection attack and the following actions can be taken:

- The session is dropped. The event is logged as a "Slow Connection: Source flood" event, and drops are reported on the Monitor > Flood Drops > Layer 4 page.
- The session entry in the FortiDDoS TCP state table is timed out.
- If the SPP aggressive aging track-slow-tcp-connections option is enabled, FortiDDoS sends a RST packet to the server so that the server can remove the connection from its connection table.
- If the SPP TCP state anomaly detection foreign-packet-validation option is enabled, subsequent packets for the connection are treated as foreign packets and dropped. The event is logged as a “State Anomalies: Foreign packet” event and drops are reported on the Monitor > Anomaly Drops > TCP State Anomalies page.
- If the SPP source blocking option is enabled, FortiDDoS applies the “Blocking Period for Identified Sources” configured on the Global Settings > Settings page. Drops based on this blocking period action are also logged as "Slow Connection: Source flood" events and reported on the Monitor > Flood Drops > Layer 4 page.

Figure 15 illustrates how FortiDDoS deployed between the client and server can monitor slow attack threats and take action to aggressively age them.

Figure 15: Slow connection detection and aggressive aging

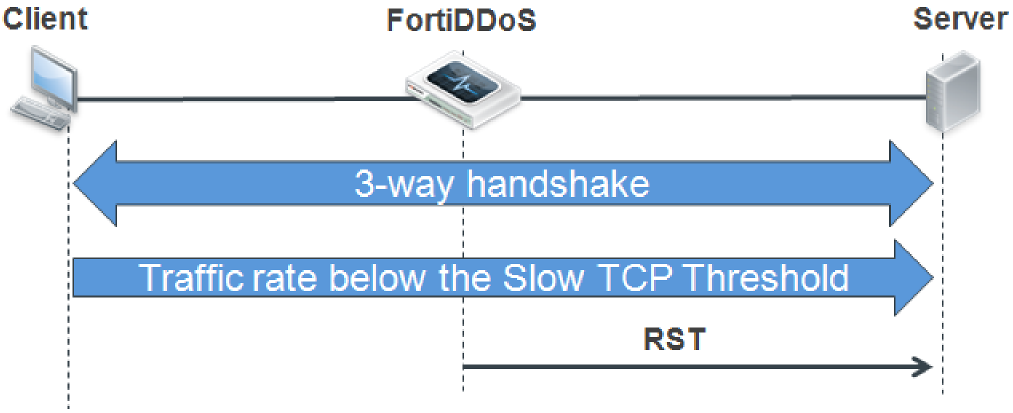

Note: By default, FortiDDoS uses the MAC address for the management interface (mgmt1) when it sends a TCP reset to aggressively age the connection. To configure a different MAC address for the resets, go to Global Settings > Settings > Settings.

Another slow connection attack, the R U Dead Yet? (RUDY) attack, injects one byte of information into an HTTP POST request. The partial request causes the targeted web server to hang while it waits for the rest of the request. When repeated, multiple simultaneous RUDY connections can fill up a web server’s connection table.

When deployed between clients and servers, FortiDDoS can detect HTTP connections that resemble RUDY attacks and “aggressively age” the connections in the same way it does for slow TCP connection attacks. When a partial request is sent from a client, the slow connection observation period starts. FortiDDoS monitors both the client and the server sides of the connection. A server cannot respond until it has received the complete HTTP request. If it has not responded to the source of the partial request by the end of the observation period, FortiDDoS considers it a slow connection. Not all slow connections are attacks, however, so FortiDDoS uses a client-side threshold to determine attack behavior. If the connection is deemed slow and the rate of partial HTTP requests from the identified source exceeds the configured threshold, it is considered a slow connection attack. The following actions can be taken (same actions as slow TCP connection):

- The session is dropped. The event is logged as a "Slow Connection: Source flood" event, and drops are reported on the Monitor > Flood Drops > Layer 4 page.
- The session entry in the FortiDDoS TCP state table is timed out.
- If the SPP aggressive aging track-slow-tcp-connections option is enabled, FortiDDoS sends a RST packet to the server so that the server can remove the connection from its connection table.
- If the SPP TCP state anomaly detection foreign-packet-validation option is enabled, subsequent packets for the connection are treated as foreign packets and dropped. The event is logged as a “State Anomalies: Foreign packet” event and drops are reported on the Monitor > Anomaly Drops > TCP State Anomalies page.
- If the SPP source blocking option is enabled, FortiDDoS applies the “Blocking Period for Identified Sources” configured on the Global Settings > Settings page. Drops based on this blocking period action are also logged as "Slow Connection: Source flood" events and reported on the Monitor > Flood Drops > Layer 4 page.

Table 9 summarizes the predefined thresholds for slow connection detection.

Table 9: Slow connection detection thresholds

Setting	Moderate	Aggressive
Slow TCP connection byte threshold	512 bytes	2048 bytes
Slow TCP connection observation period	30 seconds	15 seconds

To enable aggressive aging when these thresholds are reached, go to Protection Profiles > SPP Settings and select the Aggressive aging TCP connection feature control option **track-slow-tcp-connections**.

To enable the system to block traffic from the offending source IP address, go to Protection Profiles > SPP Settings and enable **Source blocking for slow connections**.

Caution: Source blocking for slow connection detection is disabled by default. Do not enable if it is typical for the SPP to receive traffic with source IP addresses that are proxy IP addresses (for example, a CDN proxy like Akamai). You want to avoid blocking a proxy IP address because the block potentially affects many users that are legitimately using the same proxy IP address.

Rate anomalies and aggressive aging

In addition to the slow connection detection, you can use the SPP aggressive aging TCP connection feature control options to reset the connection (instead of just dropping the packets) when the following rate anomalies are detected:

- high-concurrent-connection-per-source
- high-concurrent-connection-per-destination
- layer7-flood

Figure 16 illustrates aggressive aging when high concurrent connection or Layer 7 rate anomalies are detected.

Figure 16: Rate anomalies and aggressive aging

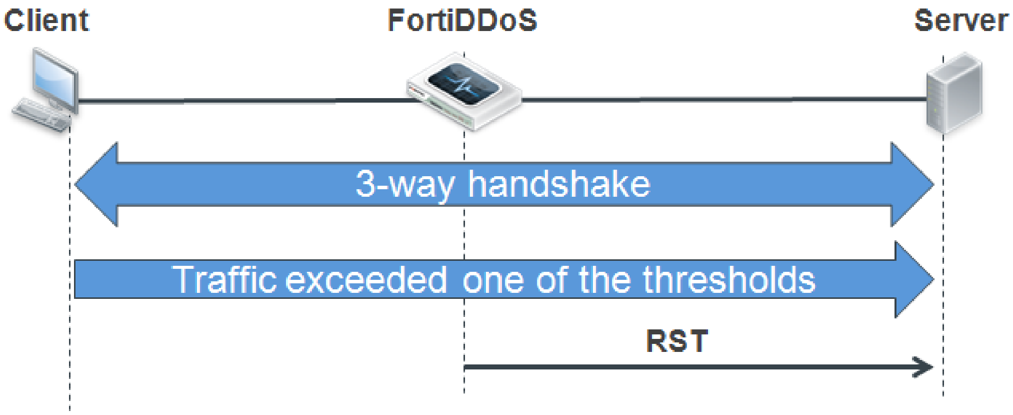

Note: The initial drops resulting from aggressive aging appear in logs and reports as SYN per Source flood drops or HTTP method flood drops, as appropriate. If the TCP session feature control option **foreign-packet-validation** option is also enabled, subsequent packets from these sources are dropped as foreign packet anomalies because the packets are correlated with a connection that has been reset.

Idle connections and aggressive aging

FortiDDoS maintains its own massive TCP connection table. To reserve space in this table for active traffic, FortiDDoS periodically uses aggressive aging to reset inactive connections. This behavior is not configurable, and it generates no logs.

Rate limiting

FortiDDoS maintains rate meters for packets, connections, and requests. It drops packets that exceed the maximum rates (which are based on history, heuristics, and a multiplier that you specify or based on an absolute limit that you specify).

Rate limiting thresholds are not only a good way to detect attacks, but also an effective method to protect servers. When deployed between client and server traffic, the rate limits ensure that a server is not inundated with more traffic than it can handle.

When FortiDDoS drops packets that exceed the maximum rates, the originating client retransmits the packets. Traffic originating from attackers is likely to be marked by extended blocking periods, while traffic originating from legitimate clients is likely to find itself within the acceptable rates as thresholds are reevaluated.

Blocking

In Prevention Mode, traffic that exceeds protection profile thresholds is blocked for the configured blocking period. When blocking period is over, the threshold is checked again.

The following examples assume that the blocking period has the default value of 15 seconds.

Example 1: Too many packets with a specified protocol

- The system drops incoming packets with the protocol that are destined for a specific network (specified as a subnet) for 15 seconds. It forwards all other packets.
- The system tracks the source of the packets to determine if this is a single-source attack.
- After 15 seconds, the system checks the rate of the packets against the threshold again.

Example 2: Too many mail messages to an SMTP server

- The system drops incoming TCP packets destined for port 25 on the mail server (or the mail server’s network) for 15 seconds. It forwards all other packets.
- The system tracks the source of the packets to determine if this is a single-source attack. If there is a single source, the appliance blocks all packets from that source for 15 seconds.
- After 15 seconds, the system checks the rate of the packets against the threshold again.
- Mail clients assume that the network is slowing down because TCP packets are lost. The clients start to send packets at a slower rate. No mail messages are lost.

Example 3: Too many SYN packets to a web server

- The system checks SYN packets destined for a web server. If they come from an IP address in the legitimate IP address table, the system permits them to continue to the web server. The appliance allows these packets as long as their rate is lower than the new-connections threshold (designed to indicate zombie floods). The system forwards all other SYN packets.

- If the IP address does not exist in the legitimate IP address table, and if the SYN flood mitigation method is SYN cookie, the system performs a proxy three-way handshake to validate the IP address.
- After 15 seconds, the system checks the packet rate against the threshold again.

Example 4: Too many concurrent connections from a single source

- If there are too many concurrent TCP connections from a single source, the system blocks new connections until the number of concurrent connections is less than the threshold.
- Once the concurrent connection count goes down, the system allows the source to establish new connections.
- The system tracks the source of the connections to determine if this is a single-source attack. If there is a single source, the appliance blocks all packets for 15 seconds.
- After 15 seconds, the system checks the connection rate against the threshold again.

Reducing false positives

When the FortiDDoS system blocks traffic because it exceeds the threshold of a specific traffic parameter, it blocks subsequent traffic with the offending characteristic. As a result, during the blocking period, the system might block traffic from legitimate sources in addition to traffic from a malicious source.

The system uses the following mechanisms to minimize the impact of these false positives:

- Because the blocking period is short (1 to 15 seconds), the system frequently checks to see if the traffic no longer exceeds the threshold that detected the attack.
- The system simultaneously attempts to determine whether the attack is not spoofed and can be attributed to one or a few sources. If it can identify these sources (called source attackers), it applies a “multiplier” to them. The multiplier makes traffic from these source attackers more likely to exceed the most active source threshold, which causes the system to apply a longer blocking period.
- If it identifies attackers, the system can stop blocking traffic from legitimate sources as soon as the standard, shorter blocking period is over, but continue to block traffic from source attackers for a longer period.

Figure 17 illustrates how the FortiDDoS system responds immediately to attacks but then adjusts its attack mitigation activity to packets from specific sources only.

In this example, the standard blocking period is 15 seconds and the blocking period for source attackers is 60 seconds (the default value). The multiplier for source attackers is 16, and the most active source threshold is 100 packets per second.

When Source B sends 90 fragmented packets, the calculated rate is 1440 packets per second, which exceeds the most active source threshold. But when Source C sends 2 fragmented packets per second, the calculated rate of 32 packets per second does not exceed the threshold. Thus, the system applies the longer blocking period to Source B only.

Source C, which sends an insignificant number of fragmented of packets, is blocked only for the length of the shorter, standard blocking period.

Figure 17: System attack response timeline

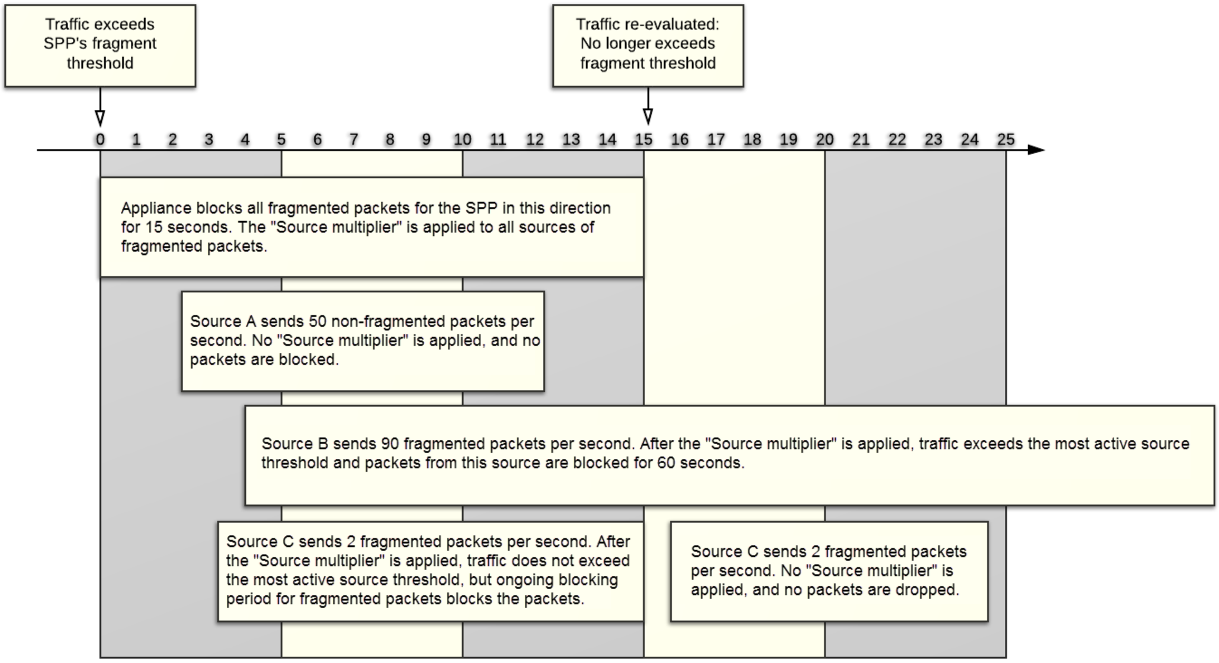
Traffic exceeds SPP's fragment threshold
Traffic re-evaluated: No longer exceeds fragment threshold
0 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25
Appliance blocks all fragmented packets for the SPP in this direction for 15 seconds. The "Source multiplier" is applied to all sources of fragmented packets.
Source A sends 50 non-fragmented packets per second. No "Source multiplier" is applied, and no packets are blocked.
Source B sends 90 fragmented packets per second. After the "Source multiplier" is applied, traffic exceeds the most active source threshold and packets from this source are blocked for 60 seconds.
Source C sends 2 fragmented packets per second. After the "Source multiplier" is applied, traffic does not exceed the most active source threshold, but ongoing blocking period for fragmented packets blocks the packets.
Source C sends 2 fragmented packets per second. No "Source multiplier" is applied, and no packets are dropped.

Understanding FortiDDoS Asymmetric Mode for TCP

FortiDDoS monitors TCP states. For TCP state monitoring to work fully and properly for all types of related mitigation, bidirectional traffic must pass through FortiDDoS.

When only one direction of traffic passes through the device, from FortiDDoS' perspective, we call it Asymmetric traffic and the appliance must be set in Asymmetric Mode.

Combinations of multiple links and BGP routing tables at the ISPs and the customer can result in inbound and outbound using any combination of the links.

Figure 23 shows an asymmetric route when an external client initiates the connection, such as a web server request. The initial TCP SYN traverses the network path where FortiDDoS has been deployed, but the SYN-ACK response takes a different route to the client.

Figure 18: Asymmetric route when an external client initiates the connection

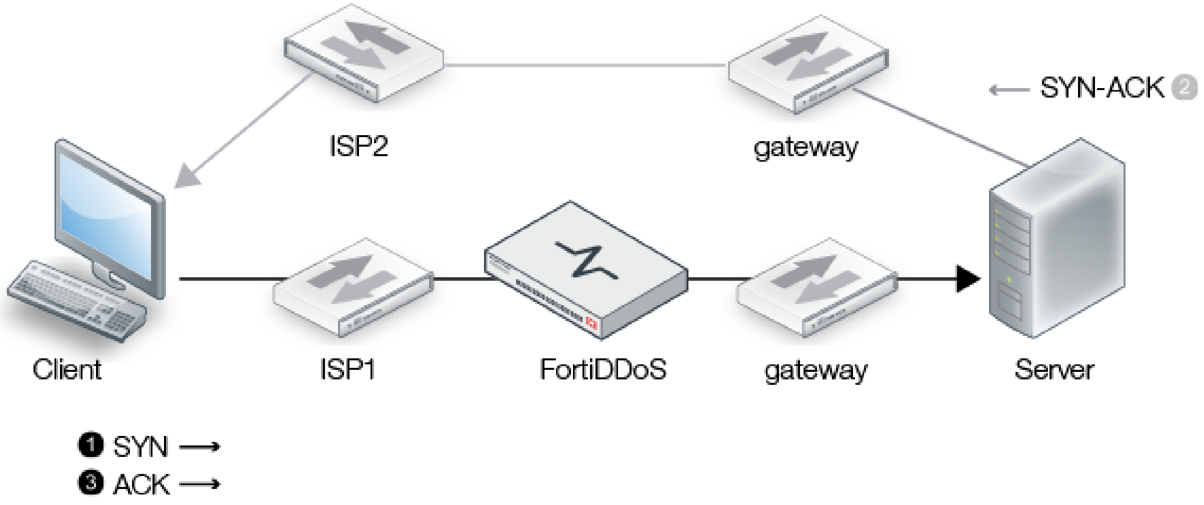

Figure 24 shows an asymmetric route when the internal resource initiates the connection, such as when a backup server initiates a scheduled job. The TCP SYN takes an out-of-path route, and the SYN-ACK packet is the first packet that FortiDDoS sees for the session.

Figure 19: Asymmetric route when an internal server initiates the connection

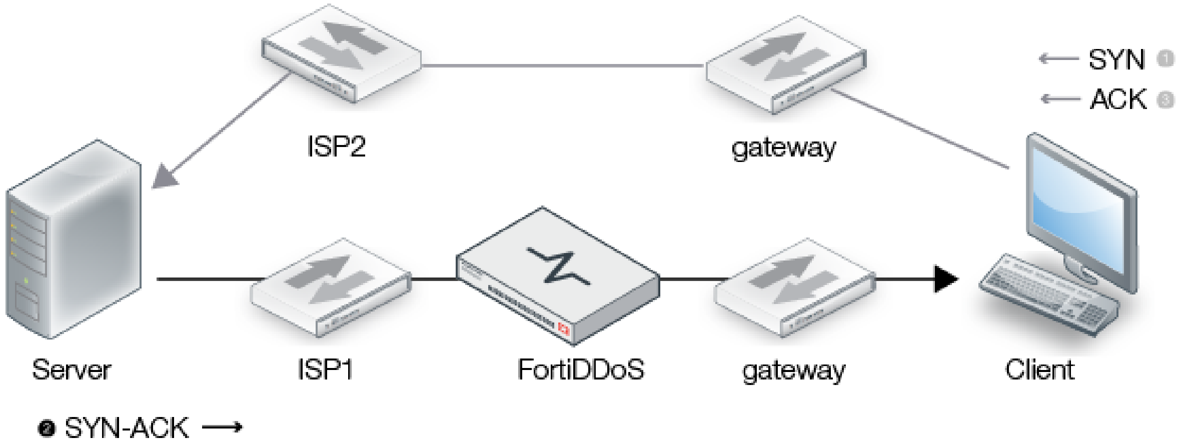

We have two key recommendations if you plan to deploy the FortiDDOS appliance in a network path where asymmetric routes are possible:

- When feasible, design the network routes so that FortiDDoS sees both sides of the client-server connection. You might be able to do this with the preferred routes, persistence, or active/active synchronization features of the routing devices in your deployment.
- If you cannot avoid asymmetric traffic, enable FortiDDoS Asymmetric Mode. In Asymmetric mode, FortiDDoS can use virtually 100% of its methods to detect abnormal network traffic, with the exception of the parameters noted below. Disabling these parameters results in a very small loss of attack detection capability.

In Asymmetric Mode, the system can parse Layer 4 and Layer 7 headers for most floods and URL-related features. If this feature is off, such floods are not detected when two-way session traffic is not completely seen by the appliance.

You must enable both **Asymmetric Mode** and the **Allow Inbound SYN-ACK** option so the system can properly handle asymmetric TCP traffic. When enabled, the system treats an inbound SYN-ACK as if a SYN, and it creates an entry for it in the TCP connection table. It does not increment the **syn** threshold counter, but it does track **syn-per-src** in order to protect against attacks that might attempt to exploit this behavior.

TCP state anomaly detection depends on tracking a two-way traffic flow, so some feature options on the Protection Profiles > SPP Settings page do not work in Asymmetric Mode. Table 10 summarizes the configuration guidelines for these feature options.

Table 10: Recommended TCP state anomaly detection settings in Asymmetric Mode

Settings	Guidelines
SYN validation	Recommended. This option enables SYN flood mitigation mode.
Sequence validation	Do not enable. Depends on tracking a two-way traffic flow.
State transition anomalies validation	Do not enable. Depends on tracking a two-way traffic flow.
Foreign packet validation	Recommended. In Asymmetric Mode, FortiDDoS can still track foreign packets.

Settings	Guidelines
Allow tuple reuse	Enabled by default to support standard test environments that reuse tuples in quick succession. The setting is valid in Asymmetric Mode. Recommended to avoid unnecessary logging of the event when it is detected.
Allow duplicate SYN-in-SYN-SENT	Not enabled by default, but the setting is valid in Asymmetric Mode. Recommended when FortiDDoS is in Detection Mode to avoid unnecessary logging of the event when it is detected.
Allow duplicate SYN-in-SYN-RECV Allow SYN anomaly Allow SYN-ACK anomaly Allow ACK anomaly Allow RST anomaly Allow FIN anomaly	Do not enable.

Workflow for getting started with Asymmetric Mode

1. Go to Global Settings > Settings > Settings > Deployment tab and enable the following settings:
   - Asymmetric Mode
   - Allow inbound SYN/ACK
2. Get started in Detection Mode:
   a. For each SPP, go to Protection Profiles > SPP Settings and ensure that the following TCP state anomaly options are enabled and no other:
      - Syn validation
      - Foreign packet validation
      - Allow tuple reuse
      - Allow duplicate SYN-in-SYN-SENT

   b. Enable Detection Mode.

   c. Establish a baseline of traffic statistics and set thresholds.
3. Change settings to the ones appropriate for Prevention Mode when there is asymmetric traffic:
   a. For each SPP, go to Protection Profiles > SPP Settings and ensure that the following TCP state anomaly options are enabled and no other:
      - SYN validation
      - Foreign packet validation
      - Allow tuple reuse

   b. Enable Prevention Mode.

Understanding Asymmetric Mode and DNS

An asymmetric route is one in which the traffic in one direction traverses FortiDDoS system, but traffic in the other direction takes a route that does not go via FortiDDoS. Combinations of multiple links and BGP routing tables at the ISPs and customer can result in inbound and outbound using any combination of the links.

If FortiDDoS is deployed in Asymmetric mode, then most DNS Feature controls must be disabled.

To mitigate many DNS DDoS attacks, FortiDDoS maintains tables for both DNS Queries and Responses. If both Queries and Responses do not go through FortiDDoS, then the functionality must be disabled for the device to work and report attacks properly.

Figure 20: Asymmetric Mode and DNS

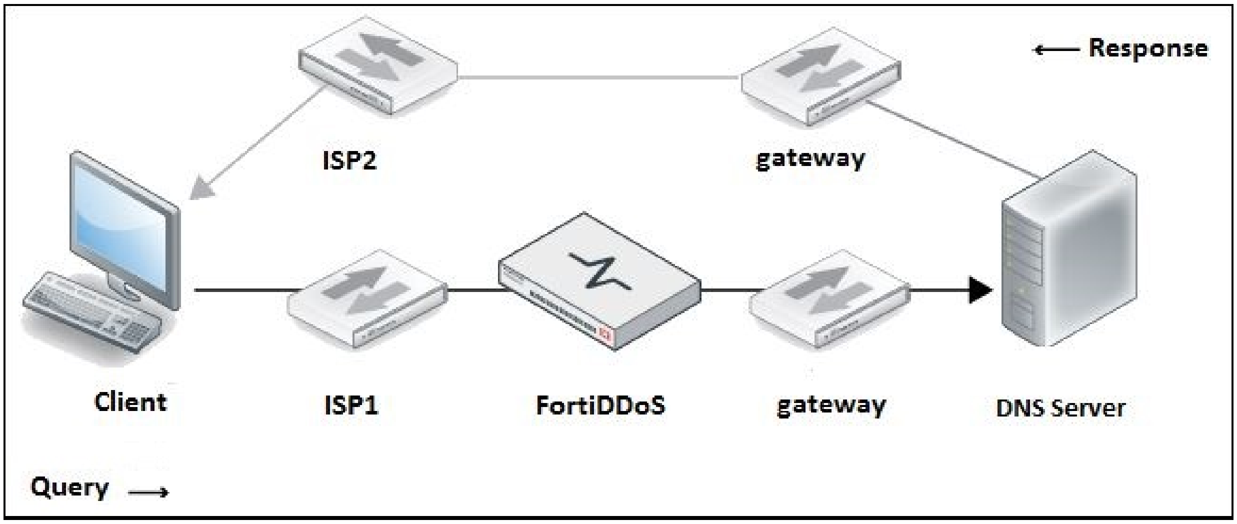

In asymmetric mode, the expectation is that FortiDDoS will only see half of most DNS Query/Response transactions. Since DNS is normally UDP and stateless, FortiDDoS cannot make assumptions of the state and may drop DNS Queries and/or Responses as anomalous.

DNS Anomaly Feature Controls are primarily DNS header anomalies. This can be enforced in Asymmetric mode because they work on packet by packet basis rather than maintaining state across packets.

Figure 25 shows the allowed DNS Feature Control configuration for use with asymmetric traffic:

Figure 21: DNS Feature Control configuration for asymmetric traffic

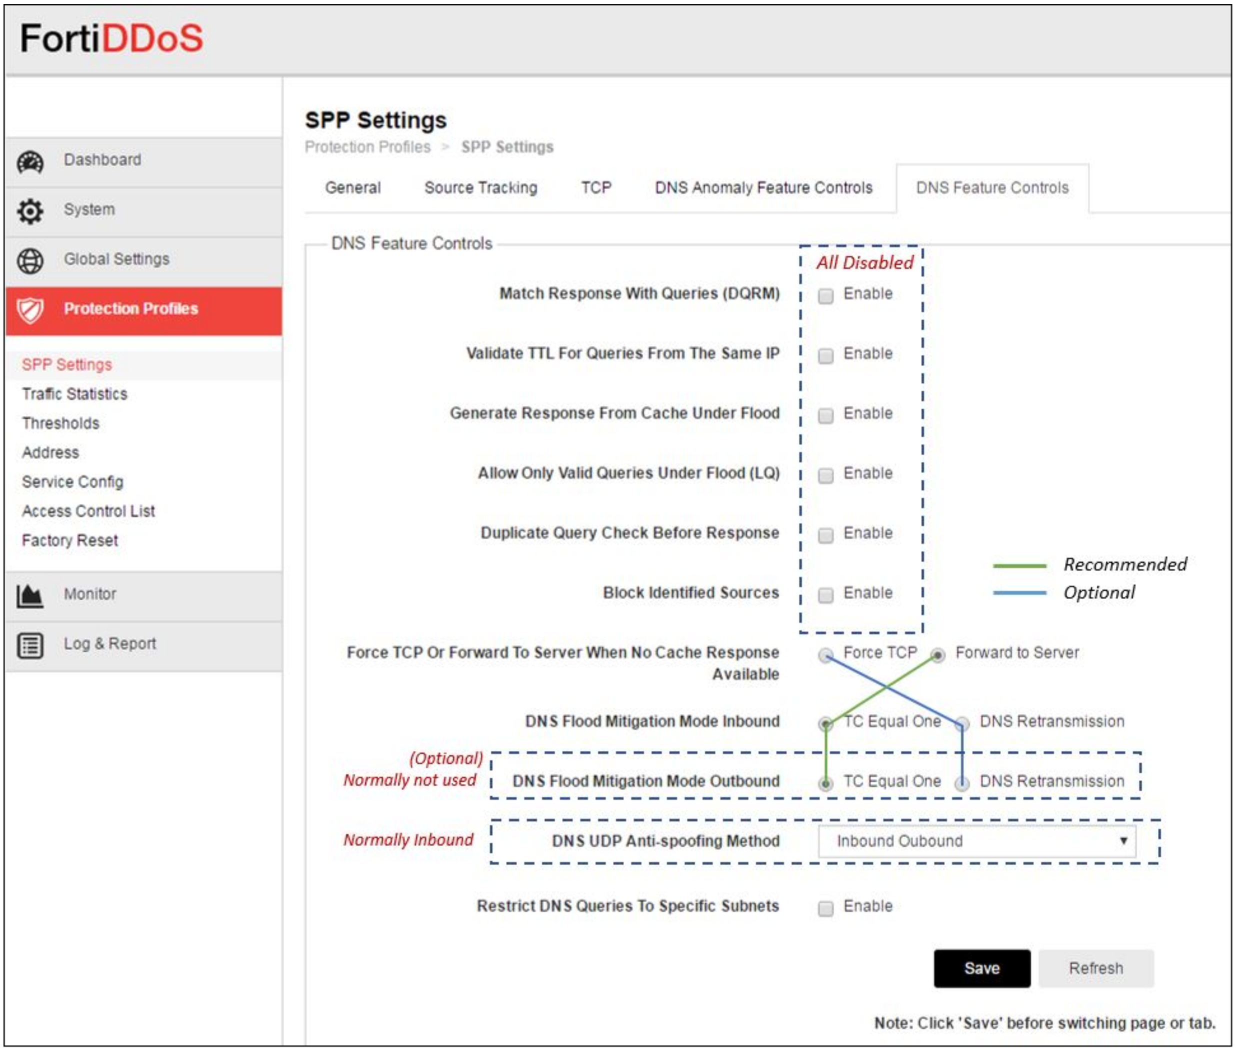

SPP Settings

Protection Profiles > SPP Settings

DNS Feature Controls

All Disabled

Match Response With Queries (DQRM) □ Enable

Validate TTL For Queries From The Same IP □ Enable

Generate Response From Cache Under Flood □ Enable

Allow Only Valid Queries Under Flood (LQ) □ Enable

Duplicate Query Check Before Response □ Enable

Block Identified Sources □ Enable

Force TCP Or Forward To Server When No Cache Response Available ○ Force TCP ● Forward to Server

DNS Flood Mitigation Mode Inbound ● TC Equal One ○ DNS Retransmission

(Optional) Normally not used DNS Flood Mitigation Mode Outbound ● TC Equal One ○ DNS Retransmission

Normally Inbound DNS UDP Anti-spoofing Method Inbound Oubound ▾

Restrict DNS Queries To Specific Subnets □ Enable

Note: Click 'Save' before switching page or tab.

Understanding FortiDDoS DNS attack mitigation

This section includes the following information:

- DNS attack vulnerability overview
- FortiDDoS DNS protection module summary
- FortiDDoS DNS flood mitigation overview
- FortiDDoS DNS flood types
- FortiDDoS DNS deployment topologies
- Getting started

Note: FortiDDoS 600B and 900B do not support DNS ACLs, DNS anomaly detection, or DNS flood mitigation.

DNS attack vulnerability overview

DNS was designed for robustness and reliability, not security. It is vulnerable to multiple types of attacks that can compromise or take down a network. Some of these attacks are described here.

DNS tunneling

DNS tunneling exploits the fact that firewall administrators must open port 53 in order for DNS authoritative name servers to respond to queries from the Internet. The attacker compromises a host in the internal network and runs a DNS tunnel server on it. A DNS tunnel client outside the internal network can then gain access to the internal network by sending a DNS query to the compromised host that sets up a DNS tunnel.

Figure 22: DNS tunneling attack

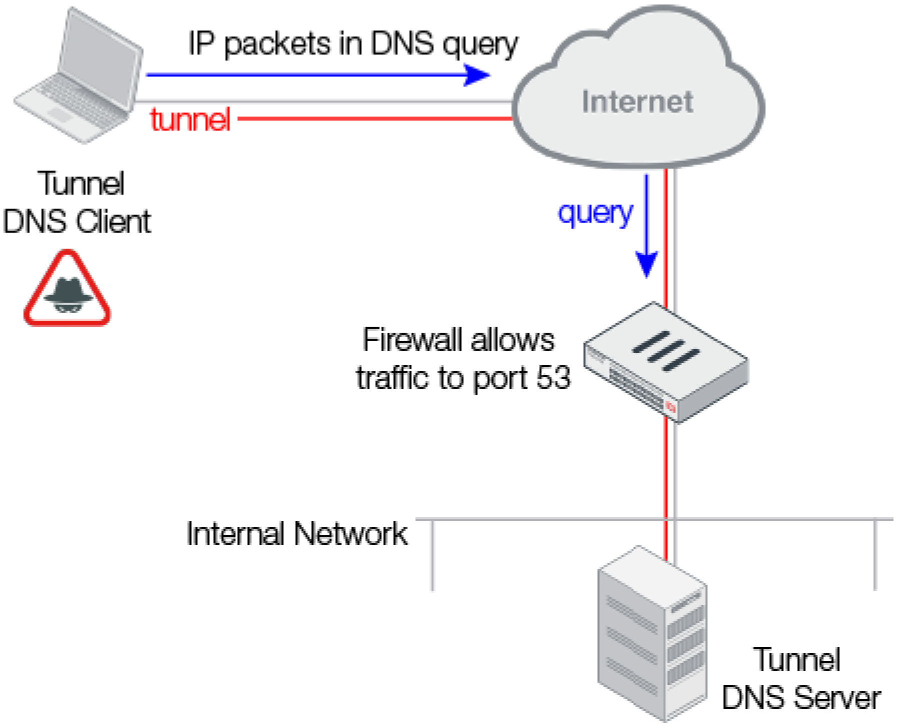

You can use FortiDDoS DNS anomaly detection to drop DNS tunneling attempts if the tunneling attempts do not conform to DNS header syntax.

DNS query floods

A DNS flood is an attempt to create a network outage by flooding critical DNS servers with excessive queries.

Some DNS floods target the authoritative name server for a domain. In these types of attacks, malware bots send a continuous flood of queries for random, nonexistent subdomains of a legitimate domain. All of the DNS servers in the recursive chain consume resources processing and responding to the bogus queries.

Figure 23: DNS slow-drip, random, non-existent subdomain attack

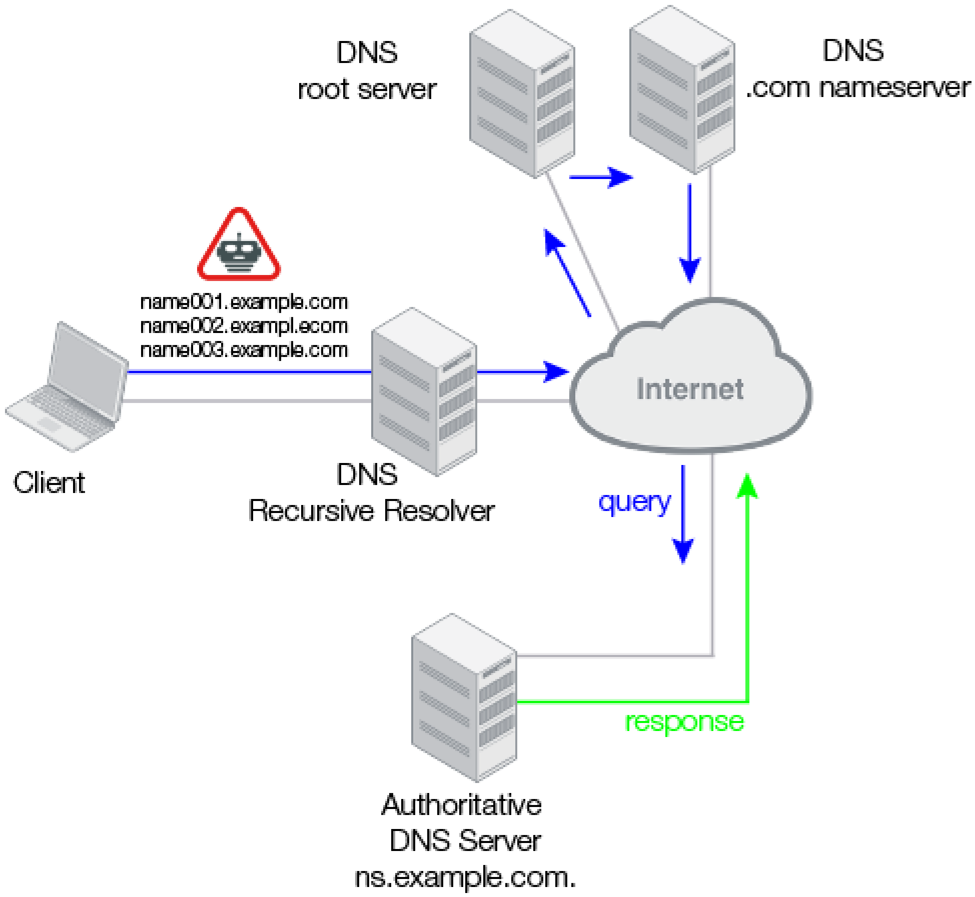

If clients in your internal network have been compromised by malware, your internal DNS resolvers could also be targets of query flood attacks.

In non-existent NX domain (NXDOMAIN) attacks, the clients that have been compromised send queries for domains that do not exist. This uses resources and can fill up the cache.

In phantom domain attacks, the clients that have been compromised send DNS queries for a phantom domain name—a domain server that exists, but it is controlled by an attacker. The attacker might have configured it to send no responses or slow responses. These illegitimate transactions waste resources, and a flood of them can take down the DNS resolver.

Figure 24: DNS NX domain and phantom domain attack

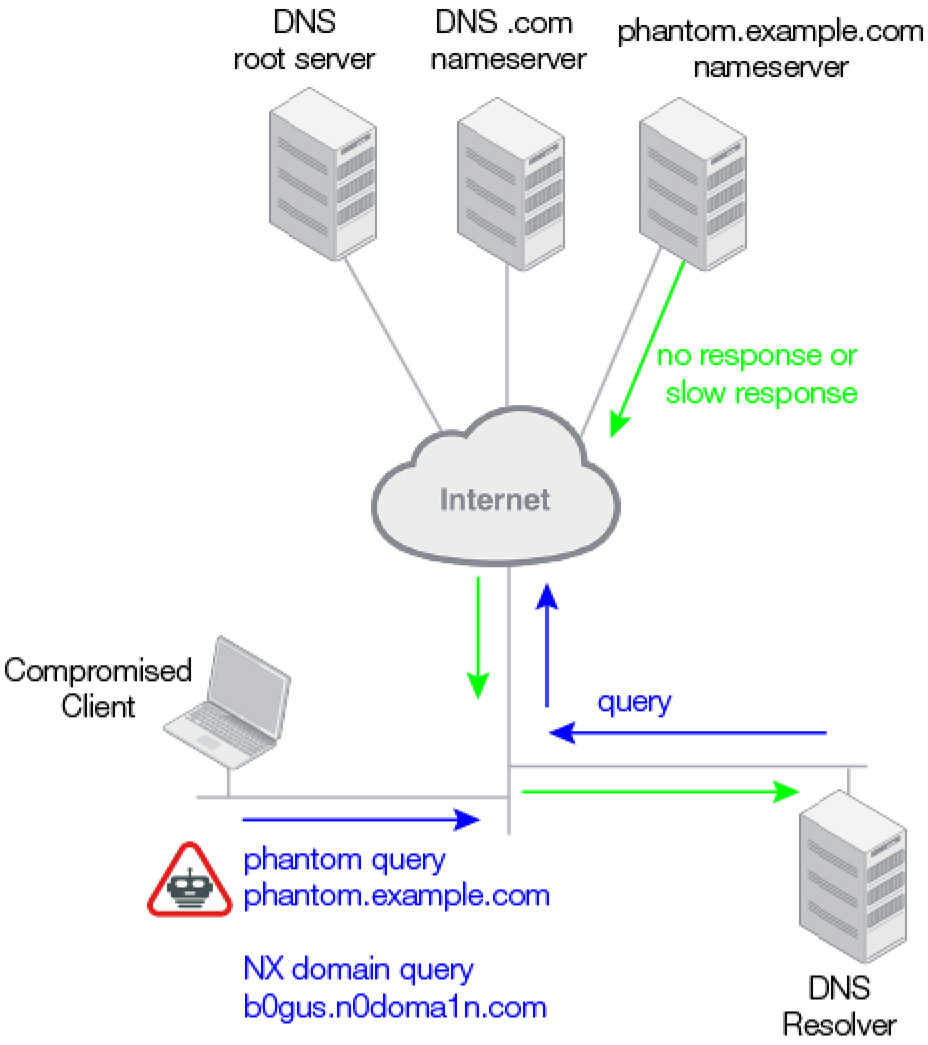

You can use FortiDDoS DNS flood mitigation features to prevent query floods.

DNS response exploits

There are also many attacks that use DNS responses to do damage. Unsolicited responses are a symptom of DNS Distributed Reflective Denial of Service attacks, DNS amplification attacks, and DNS cache poisoning.

Figure 25: DNS reflection attack

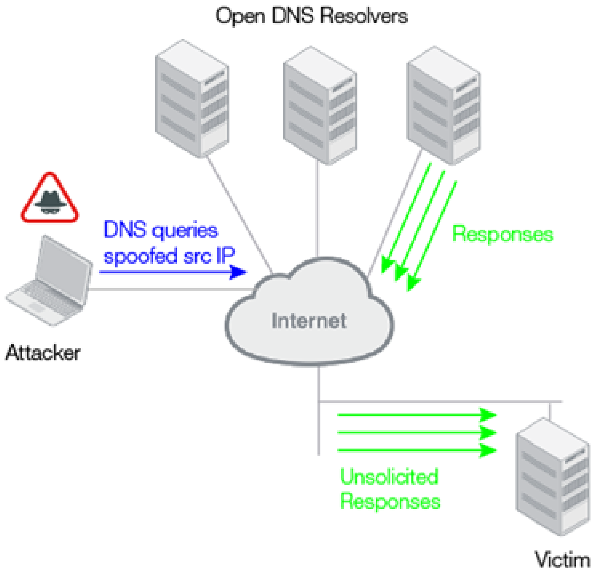

Figure 26: DNS amplification attack

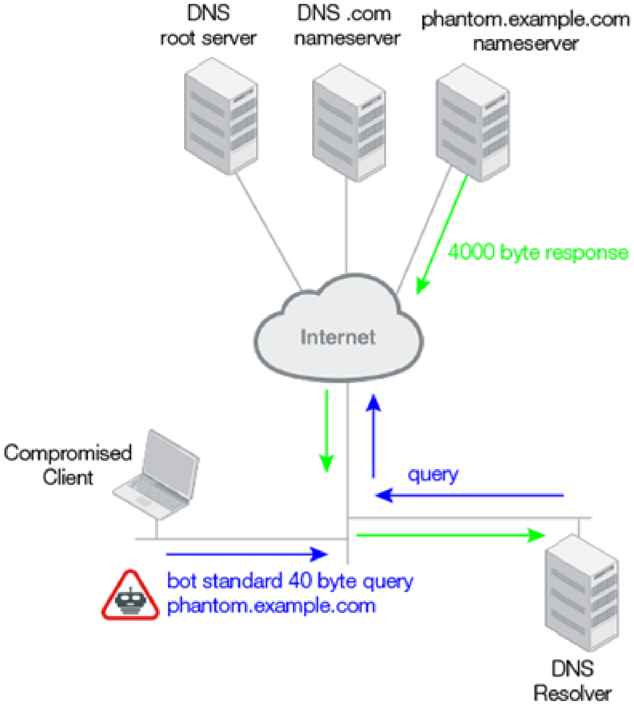

Figure 27: DNS cache poisoning

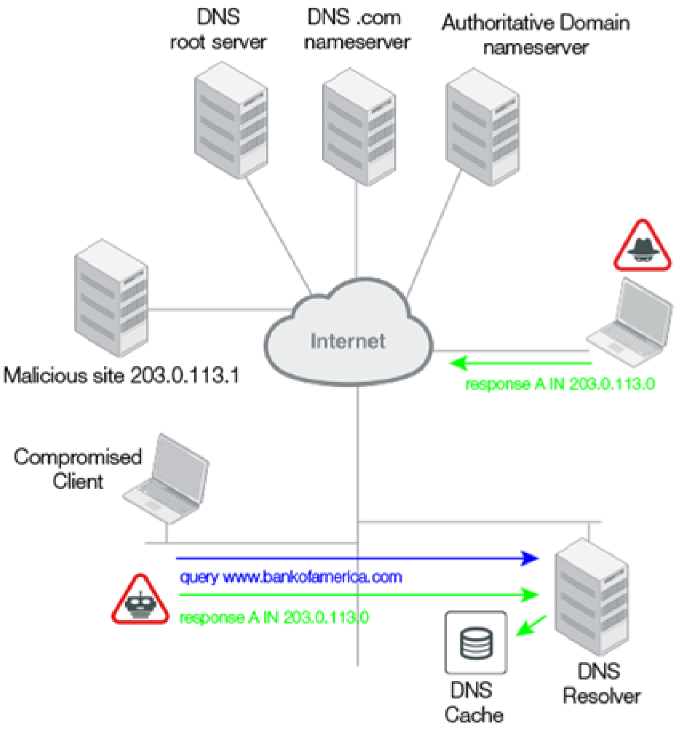

You can use the FortiDDoS DNS query response matching (DQRM) feature to prevent DNS response exploits.

FortiDDoS DNS protection module summary

FortiDDoS has the following protection modules for DNS (transport over TCP or UDP):

- ACL rules
  You can use the Do Not Track and Global ACL Allow policy to whitelist trusted IP addresses. For example, to permit DNS query type ALL or Zone Transfer from specified hosts, you can whitelist them and then create rules that deny those types of queries from all other sources. For an overview of ACLs, see Using FortiDDoS ACLs.
- Protocol anomaly rules
  Built-in and user-enabled rules filter malformed traffic and known protocol exploits. There is a special set of anomalies that can be detected in DNS traffic. For an overview of protocol anomalies, see Understanding FortiDDoS protocol anomaly protection.
- Rate meters and flood mitigation mechanisms
  For TCP, the DNS rate meters enforce rate limits (drops). For UDP, the DNS rate meters trigger flood mitigation responses that drop illegitimate queries but continue DNS services for legitimate user queries. For details, see FortiDDoS DNS flood mitigation overview.
- DNS Query Response Matching (DQRM)
  Blocks unsolicited responses and throttles duplicate queries (regardless of flood state). See FortiDDoS DNS flood mitigation overview.

Figure 28 and Figure 29 illustrate the order in which FortiDDoS applies its rules and actions for TCP and UDP DNS traffic, respectively.

Figure 28: TCP DNS Drop Precedence

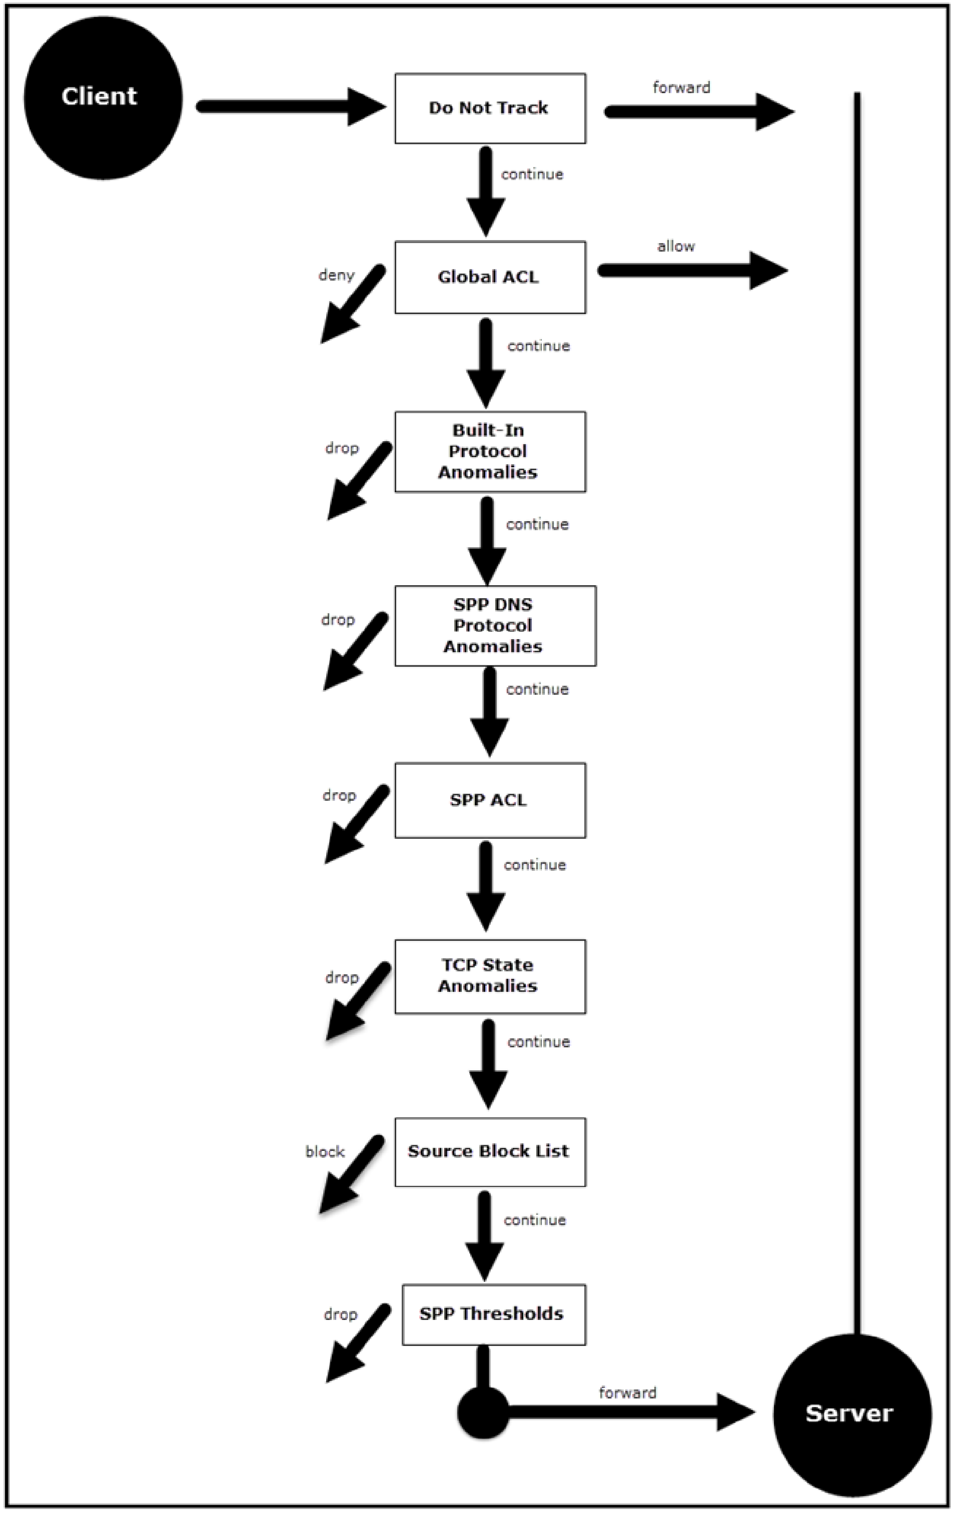

Figure 29: UDP DNS Drop Precedence

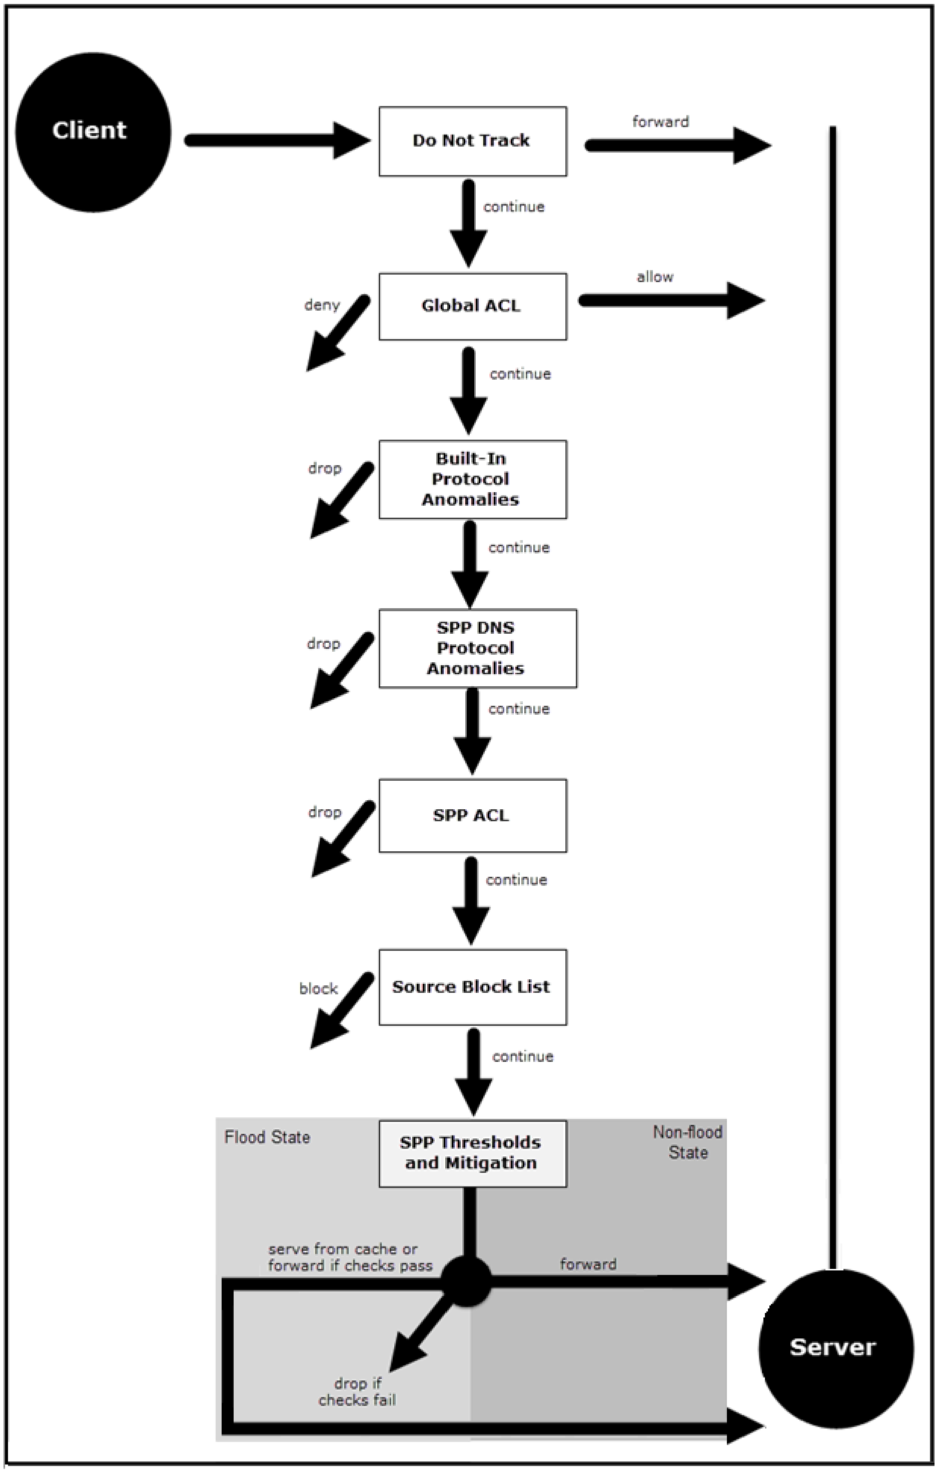

FortiDDoS DNS flood mitigation overview

FortiDDoS mitigates DNS threats by applying tests to determine whether queries and responses are legitimate. These methods minimize illegitimate traffic from reaching protected DNS servers and maximize the availability of DNS services for legitimate queries during a flood.

Under normal conditions (no floods), FortiDDoS builds a baseline of DNS traffic statistics and stores DNS query and response data in tables. At all times, the tables are used to validate response traffic. During UDP floods, the tables are used to test queries and responses.

Table 11 describes the system tables used for DNS attack mitigation.

Table 11: DNS-related system tables

System Table	DNS Flood Mitigation
DNS Query Response Match (DQRM) table	Used for all DNS traffic—UDP or TCP. When it receives a DNS query, the system stores the DNS transaction details in the DQRM table. It can store up to 1.9 million records. When it receives a response, it searches this table for a matching query. If the response has no matching query, FortiDDoS drops the unmatched response. Drops are reported on the Monitor > Layer 7 > DNS > Unsolicited Response graph. The table entry is cleared after the matching response is received. The DQRM table response validation prevents attacks that attempt to exploit DNS responses, such as DNS cache poisoning and DNS amplification attacks (also called Distributed Reflective Denial of Service attacks). The DQRM can also be used to throttle repeated queries that would otherwise result in unnecessary server activity. The "Duplicate query check before response" option drops identical queries (same transaction details) that are repeated at a rate of 3/second. Drops are reported on the Monitor > Layer 7 > DNS > Unexpected Query graph.
Legitimate Query (LQ) table	Used to mitigate UDP floods. When a valid response is received, the query details are stored in the table. It can store 128,000 records. Entries are cleared when the TTL expires. During a flood, the system drops queries that do not have entries in the table. Drops are reported on the Monitor > Layer 7 > DNS > LQ Drop graph.

System Table	DNS Flood Mitigation
TTL table	Used to mitigate UDP floods. When a valid response is received, the query details are correlated with the client IP address and stored in the table. It can store 1.5 million records. Entries are cleared when the TTL expires. Responses with TTL=0 are not added to the table. During a flood, the system drops queries that have an entry in the table. It is not expected that a client would send the same query before the TTL expires. Drops are reported on the Monitor > Layer 7 > DNS > TTL Drop graph.
Legitimate IP table	Used to mitigate UDP floods. During DNS query floods, you can leverage the legitimate IP (LIP) table to test whether the source IP address is spoofed. If the source IP address is found in the LIP table, processing continues; if there is no entry, the system can test source IP legitimacy by performing a UDP retransmission test or by sending a response with the TC flag set. The TC flag indicates to the client to retry the request over TCP. When the query is retried over TCP, other flood mitigation mechanisms may be available, such as SYN flood antispoofing features. Drops are reported on the Monitor > Layer 7 > DNS > Spoofed IP Drop graph.
DNS cache	Used to mitigate UDP floods. When a valid response is received, the system caches the response packets. It can store 64,000 records. Entries are cleared when the TTL expires. During a flood, if the query passes the LQ and TTL checks, the response is served from the cache and the query is not forwarded to the DNS server. This enables legitimate clients to get DNS results without adding load to the server that is being attacked. If there is not an entry in the cache, you can configure whether you want the query to be forwarded to the DNS server or have FortiDDoS send a response with the TC flag set. The TC flag indicates to the client to retry the request over TCP. Drops are reported on the Monitor > Layer 7 > DNS > Cache Drop graph.
Source tracking table	Used for source flood tracking—UDP or TCP. Tracks DNS queries per source and suspicious actions per source. It drops packets that exceed the maximum thresholds and applies the blocking period for identified sources. Drops are reported on the Monitor > Layer 7 > DNS Query Per Source and the Monitor > Layer 7 > Suspicious Sources graphs.

Source tracking thresholds and TCP thresholds are rate limits, resulting in drops when the flood rate thresholds are crossed. For UDP, rate thresholds trigger mitigation mechanisms. Drops are based on results of the mitigation checks. Figure 30 illustrates the packet flow through mitigation mechanisms during a UDP flood.

Figure 30: UDP mitigation process flow

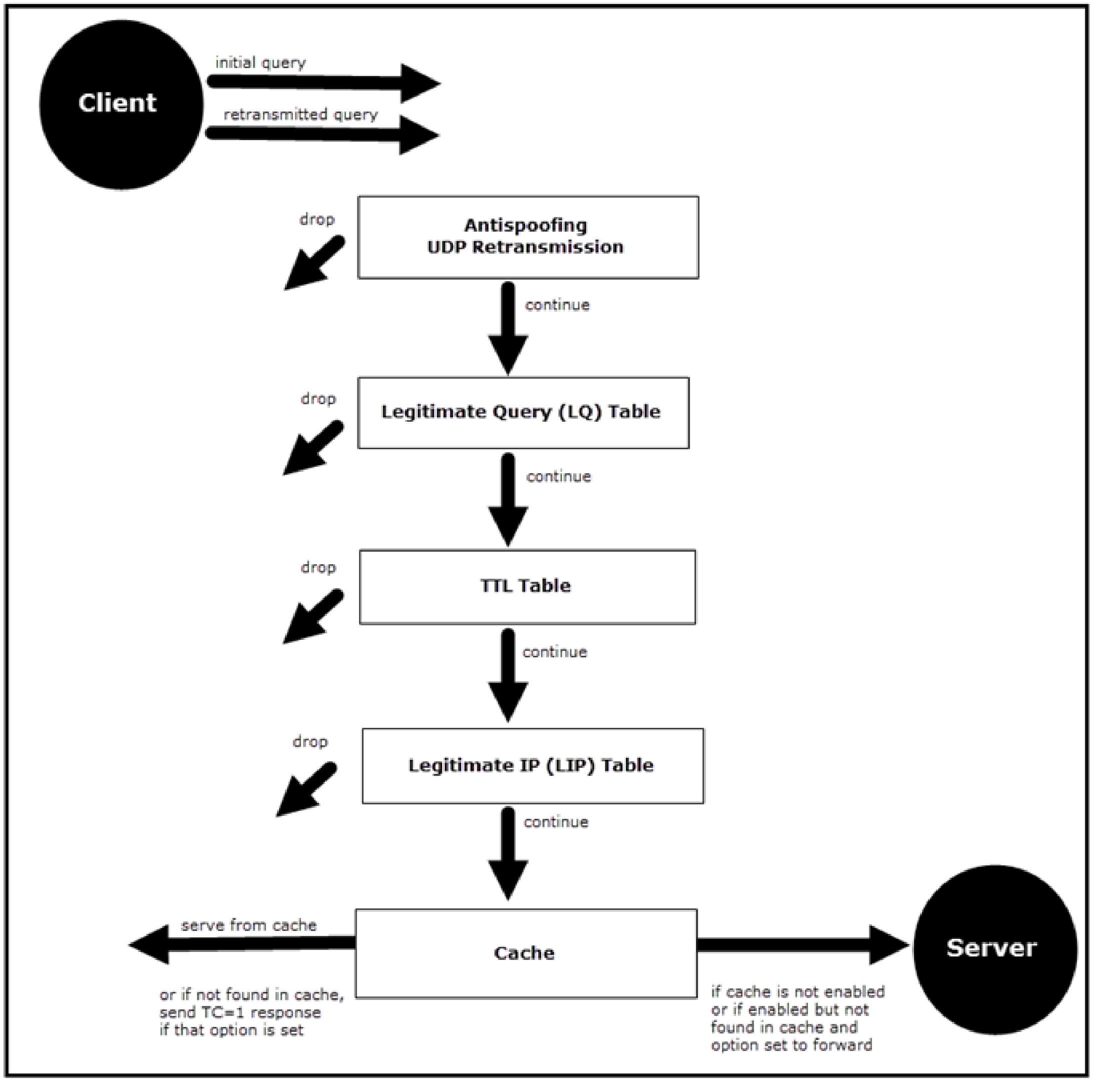

FortiDDoS DNS flood types

Table 12 summarizes the types of DNS floods mitigated by FortiDDoS.

Table 12: DNS flood types

DNS Flood Type	Thresholds
Query Flood	Abnormal rate of DNS queries or occurrences of query data. Spikes in DNS queries and fragmented queries are obvious symptoms of an attempt to take down the DNS server. Changes in norms for query data, such as question type and question count, are also symptoms of exploit attempts. Detected by the dns-query, dns-fragment, dns-question-count, dns-mx - count, dns-all-count, and dns-zone-xfer-count thresholds. The Monitor > Layer 7 graphs include packet rate graphs for each key threshold, and the Layer 7 drops graphs show which thresholds were at a flood state when the packets were dropped.
Per Source Flood	Rate limit for DNS queries from a single source. Detected by the dns-query-per-source threshold. The system applies the blocking period for identified sources. The Monitor > Layer 7 graphs include a Query Per Source graph.
Suspicious Sources	Heuristics to track other abnormal activity from a single source. Detected by the dns-packet-track-per-src threshold. This counter is incremented when a query is not found in the DQRM, when there are fragmented packets in the query or response, and when the response has an RCODE other than 0. The system applies the blocking period for identified sources. The Monitor > Layer 7 graphs include a Suspicious Sources graph.

FortiDDoS DNS deployment topologies

FortiDDoS can be deployed to protect:

- Authoritative DNS servers that receive queries from the Internet.
- DNS recursive resolvers that send queries to and receive responses from Internet DNS authorities.

Figure 31 shows a topology where FortiDDoS is deployed primarily to protect the authoritative DNS server for a domain. Under normal traffic rates, FortiDDoS builds a baseline of DNS traffic statistics and stores DNS query and response data in tables. The tables are used to validate response traffic.

Figure 31: DNS no flood: inbound queries

	Query	Response
UDP	• Adds an entry to the DQRM table. • Performs a duplicate query check to prevent unnecessary queries to the server.	• Validates the response against the DQRM table. If there is an entry, the traffic is forwarded; otherwise, it is dropped. • Updates the LQ table, the TTL table, and the DNS cache.
TCP	• Adds an entry to the DQRM table. • Performs a duplicate query check to prevent unnecessary queries to the server.	Validates the response against the DQRM table. If there is an entry, the traffic is forwarded; otherwise, it is dropped.

Figure 32 shows a topology where FortiDDoS is deployed in front of an internal DNS resolver that sends queries to and receives responses from the Internet. This type of deployment is useful for open resolvers where the DNS resolver is protected primarily from Internet-originating inbound reflection attacks.

Figure 32: DNS no flood: inbound response traffic

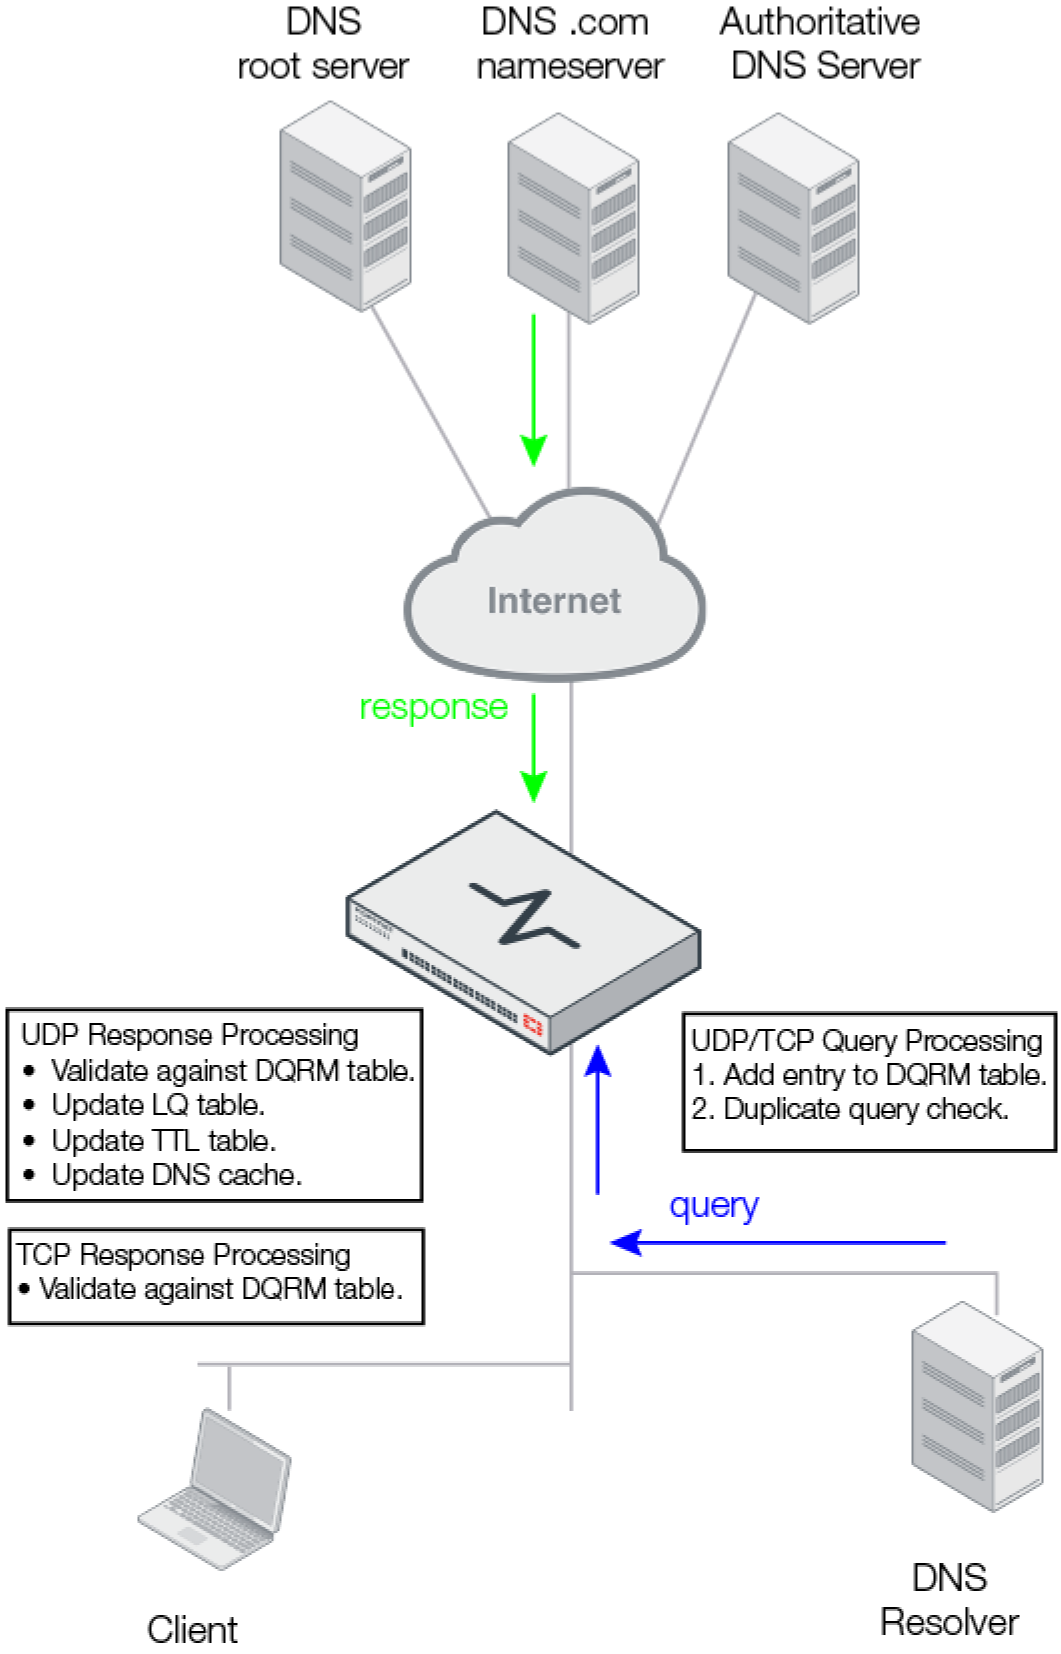
DNS
root server
DNS .com
nameserver
Authoritative
DNS Server
Internet
response
UDP Response Processing
• Validate against DQRM table.
• Update LQ table.
• Update TTL table.
• Update DNS cache.
TCP Response Processing
• Validate against DQRM table.
UDP/TCP Query Processing
1. Add entry to DQRM table.
2. Duplicate query check.
query
Client
DNS
Resolver

FortiDDoS collects data and validates the inbound responses and outbound requests the same as when queries are inbound. This deployment protects your network against different threats, such as DNS amplification attacks that result in unsolicited DNS response floods to targeted victims and DNS cache poisoning attacks, in which attackers send responses with malicious records to DNS recursive resolvers. In a deployment like this, the unsolicited responses would fail the DQRM check and be dropped.

DNS query flood mitigation

Figure 33 shows how FortiDDoS mitigates a DNS query flood. It uses the DNS tables and LIP table to validate queries and responses.

Figure 33: DNS Query Flood

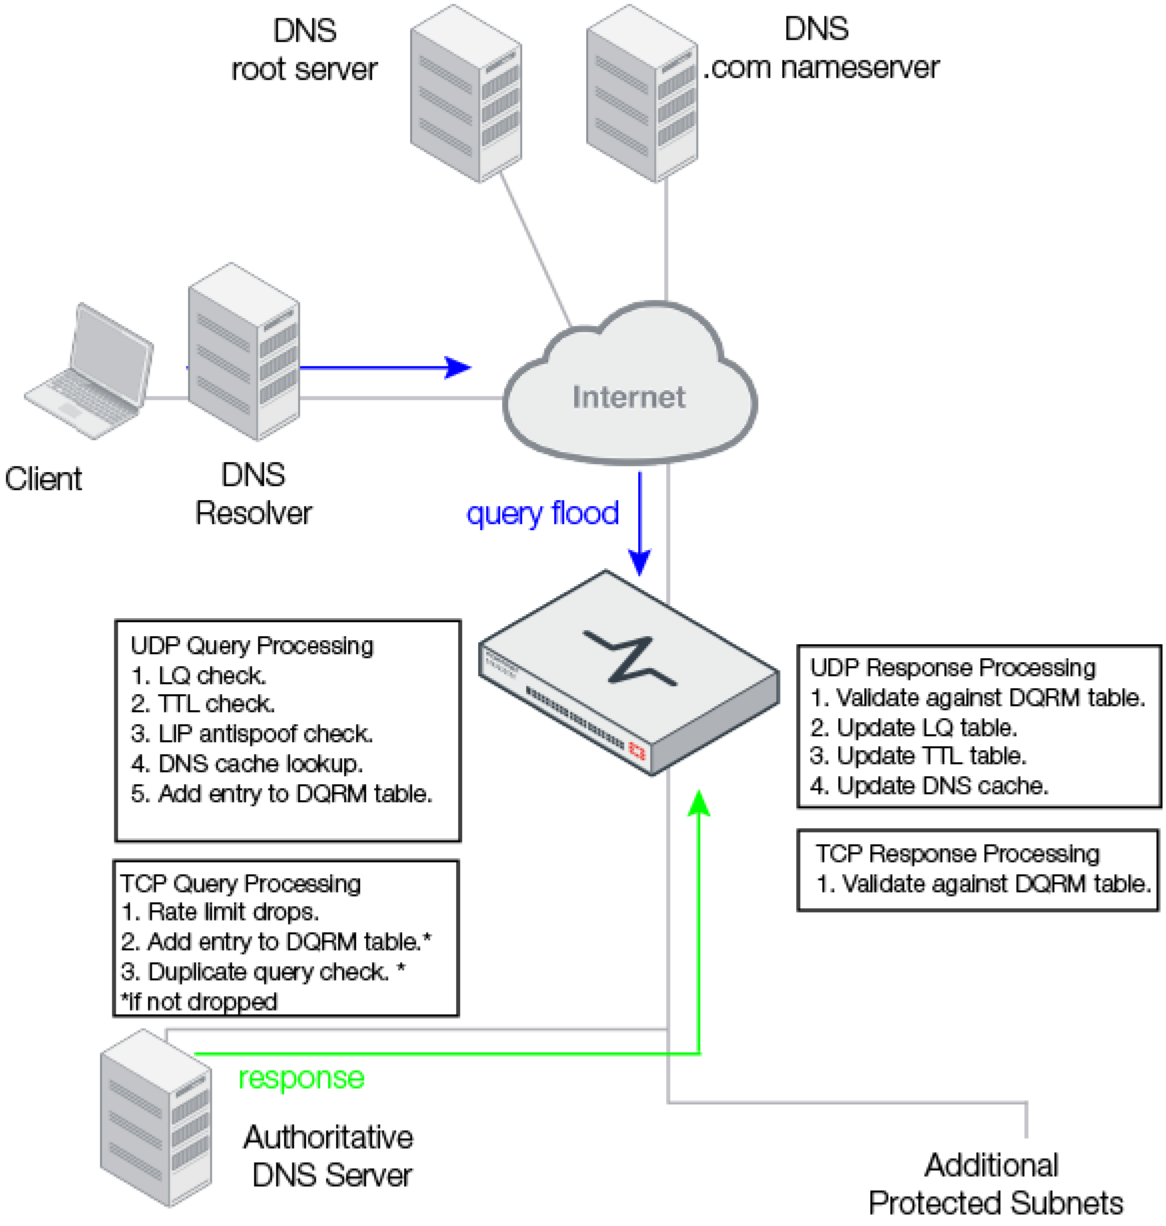

UDP Query Processing
1. LQ check.
2. TTL check.
3. LIP antispoof check.
4. DNS cache lookup.
5. Add entry to DQRM table.

TCP Query Processing
1. Rate limit drops.
2. Add entry to DQRM table.*
3. Duplicate query check. *
*If not dropped

UDP Response Processing
1. Validate against DQRM table.
2. Update LQ table.
3. Update TTL table.
4. Update DNS cache.

TCP Response Processing
1. Validate against DQRM table.

	Query	Response
UDP	• Validates against the LQ table. Under flood conditions, a query must have an entry in the LQ table or it is dropped. • Validates against the TTL table. If a match is found, the TTL check fails and the packets are dropped. It is not expected that a client would send the same query before the TTL expires. • Perform a lookup in the LIP table. If an entry exists, processing continues; otherwise, FortiDDoS drops the packets and tests the legitimacy of the source IP address. You can configure FortiDDoS to do so by performing a UDP retransmission challenge or by sending the requestor a response with the TC flag set. The TC flag indicates to the client to retry the request over TCP. • Performs a lookup in the DNS cache. If found, the response to the query is sent from the cache and the query is not forwarded to the protected server. If not found, you can configure whether to forward the query to the server or to send a TC=1 response to force the client to retry using TCP. • Adds an entry to the DQRM table.	• Validates the response against the DQRM table. If there is an entry, the traffic is forwarded; otherwise, it is dropped. • Updates the LQ table, the TTL table, and the DNS cache.
TCP	• Drops packets according to thresholds. • Adds an entry to the DQRM table. • Performs a duplicate query check to avoid unnecessary queries to the server.	Validates the response against the DQRM table. If there is an entry, the traffic is forwarded; otherwise, it is dropped.

Getting started

We recommend you allocate an SPP exclusively for DNS traffic. It takes a week to establish a baseline of traffic statistics for the SPP.

Getting started workflow

1. Go to Global Settings > Service Protection Profiles and create an SPP configuration exclusively for DNS traffic.
2. Go to Protection Profiles > SPP Settings and click the **General** tab. Ensure the SPP is in **Detection** mode.
3. Create ACL rules (if desired):
   a. Go to Protection Profiles > Service and create service configuration objects for DNS QTYPE or fragment.
   b. Go to Protection Profiles > ACL and create deny rules for those services.
4. Go to Protection Profiles > SPP Settings and click the **DNS Protocol Anomalies** tab. We recommend you enable detection for all anomalies and disable only if you encounter false positives (not expected).
5. Go to Protection Profiles > SPP Settings and click the **DNS Feature Controls** tab. We recommend you enable all features and leave disabled only features that are not suitable for your deployment.
6. Go to Monitor Graphs > Layer 7 > DNS and observe the accumulation of traffic statistics for the SPP's DNS meters.

   After you have established a baseline (a week's worth of traffic), take the following steps to prepare for and switch to prevention mode.

7. Configure thresholds. A threshold applies to both UDP and TCP rates, but there are separate counters for each protocol:
   a. Go to Protection Profiles > Traffic Statistics and generate baseline statistics.
   b. Go to Protection Profiles > Thresholds > System Recommendation and generate thresholds.
   c. Go to Protection Profiles > Thresholds > Thresholds, review them, and make manual changes (if any).
8. Go to Protection Profiles > SPP Settings and click the **TCP** tab. Enable the following options:
   - **Layer-7 flood**—This setting is useful when there is DNS traffic over TCP and FortiDDoS drops connections due to rate limits. If this setting is enabled, FortiDDoS sends a reset to the server to clear the connection from its TCP table.
   - **Foreign packet validation**—This setting is useful when there is DNS traffic over TCP and FortiDDoS drops connections due to rate limits. When this settings is enabled, subsequent packets from the same connection are dropped.
9. Go to Protection Profiles > SPP Settings and click the **General** tab. Change to **Prevention** mode.

Using FortiDDoS SPPs

Service Protection Profiles (SPP) are used to enable a single FortiDDoS appliance to protect multiple network zones with thresholds appropriate for the traffic in each of those zones.

One SPP (SPP-0) is designated as the default SPP. You allocate the remaining seven SPPs to subnets.

In an enterprise deployment, you can configure SPPs for specific departments, geographic locations, or functions within an organization.

In a multi-tenant deployment, you can use SPPs to separate a single physical FortiDDoS device into up to 8 logical devices. Each SPP has its own configuration and traffic database. The configuration of each SPP can be under the control of a different administrator.

Figure 34 shows how multiple SPPs are used to protect multiple subnets.

Figure 34: Multiple SPPs, multiple subnets

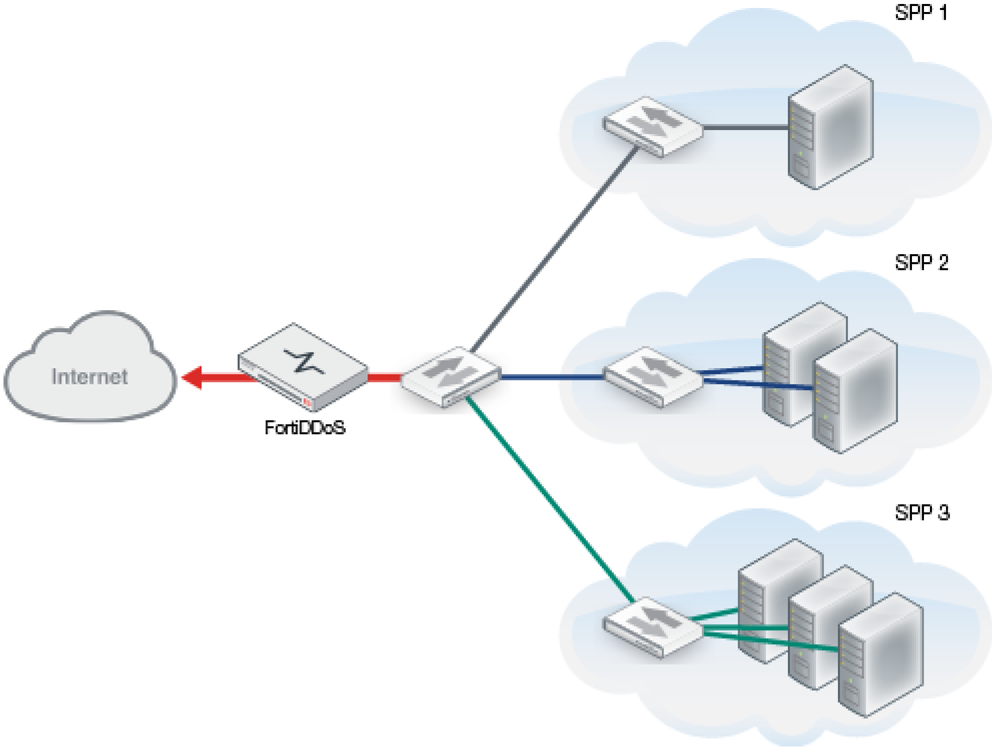

Working with the FortiDDoS Monitor graphs

The FortiDDoS system records data points for monitored thresholds every five minutes. The data point is the highest rate observed in any one second during the five minute window. Figure 35 illustrates this.

Figure 35: Maximum values recorded

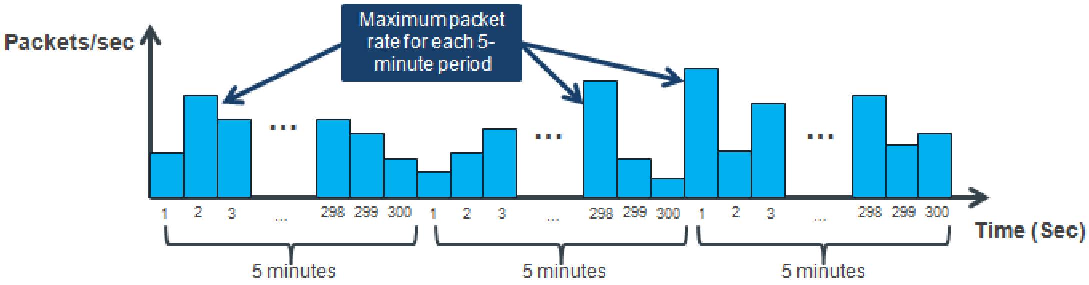

The FortiDDoS graphs are plots of data points. The reporting framework uses resolution periods to fit data points in time-based graphs. In a graph with a five minute resolution period, the graph is based on a plot of the rates or counts recorded for the regular five minute window. In a graph with a one hour resolution period, the graph is based on a plot of the rates or counts for the highest rate among the points recorded in a one hour window—in other words, the highest rate among the 12 five minute windows reported in the hour.

Table 13 lists resolution periods used for report periods.

Table 13: Data resolution periods

Graph Period	Resolution Period
1 hour	5 minutes
8 hours	5 minutes
1 day	5 minutes
1 week	1 hour
1 month	3 hours
1 year	45 hours

Note: The data displayed in a graph is current as of the time the last data point was written. For example, a 1-hour graph with a 5-minute resolution is current as of the time the last 5-minute resolution data point was stored. Traffic in the most recent 0-5 minutes might not have been registered yet. Similarly, for a 1-year graph with a data resolution of 45 hours, data for the last 0-45 hours might not have been registered yet.

Working with the FortiDDoS attack log

The monitoring and reporting framework is designed to maximize the processing resources that are available for preventing attacks, rather than forensics. In order to conserve resources to withstand multi-gigabyte attacks, the system records only data that it can use to improve security, not all possible Layer 3, Layer 4, and Layer 7 data. As a result, reporting tools such as the DDoS attack log do not always include detailed traffic parameter information. Outside of specific scenarios, the system does not report source and destination IP addresses and ports, protocols, and so on, for every dropped or blocked packet.

It is not uncommon for a FortiDDoS system that is monitoring a 1 Gbps traffic flow to be the target of a 700 Mbps SYN flood for 8 hours. If the system stored every source and destination IP address, port, and protocol, the logging demands (via hard disk, syslog, or SNMP trap) would soon overwhelm the disk or network.

By concentrating its resources on dropping attack traffic and maintaining service, FortiDDoS allows you to focus your attention elsewhere and still provide you with helpful and relevant information when an attack is underway.

A typical workflow for investigating FortiDDoS attack events

Whenever there is an attack, you should investigate until you fully understand why packets were dropped, and you know whether the attack event is a false positive.

A typical FortiDDoS attack investigation includes the following steps:

1. Identify the destination and source.
2. Identify the type of attack.
3. Identify the attack size.
4. Analyze Layer 3, Layer 4, and Layer 7 parameters to understand the attack method.

Step 1: Identifying the destination and source

Most of the statistics graphs identify the SPP and the direction of the attack, so, if there is only one subnet in the attacked SPP, you can easily determine the attack destination.

If the SPP contains more than one subnet, you can use the following reports to determine the attack destination:

- Execute Summary report
- Attack Graph Dashboard
- DDoS Attack Logs

The following reports can be used to determine the attack source:

- Executive Summary report
- Attack Graphs dashboard
- DDoS Attack Logs

Note: DDoS attacks are often spoofed attacks. Source information is not provided as it is irrelevant.

Step 2: Identifying the type of attack

The following reports can be used to determine the type of attack:

- Executive Summary report
- Attack Graphs dashboard
- DDoS Attack Logs

Table 14 describes DDoS attack types and identifies the FortiDDoS events to look for.

Table 14: Types of attacks

Attack	Description	Threshold to monitor/adjust	Events to watch
SYN attack	A spike in packets on a specific TCP port. In most cases, the source address is spoofed.	Layer 3 - TCP protocol (6) Layer 4 - TCP ports on which the server is listening and ports that are allowed by the firewall and ACL Layer 4 - SYN	Protocol 6 Flood SYN Flood Zombie Flood Port Flood
Source flood	A single source sends excessive number of IP packets.	Layer 3 – Most active source	Source Flood
Zombie attack	A spike in TCP packets from legitimate IP addresses.	Layer 3 – TCP protocol (6) Layer 4 – TCP ports on which the server is listening and ports that are allowed by the firewall and ACL Layer 4 – SYN Layer 4 – Established connections per destination (estab-per-dst) Layer 4 - SYN per source (syn-per-src)	Layer 3 Protocol 6 SYN Flood Zombie Flood Port Flood SYN Flood from Source
Fragment flood	An excessive number of fragmented packets.	Layer 3 – Fragmented packets	Fragment Flood
ICMP flood	An excessive number of ICMP packets.	Layer 3 – ICMP protocol (1) Layer 4 – ICMP type and code combinations that are allowed by the firewall and ACL	Protocol 1 Flood Layer 4 ICMP Flood of a specific type and code

Attack	Description	Threshold to monitor/adjust	Events to watch
Smurf attack	Traffic that appears to originate from the target server's own IP address or somewhere on its network. Targeted correctly, it can flood the network with pings and multiple responses.	Layer 3 – ICMP protocol (1) Layer 4 – ICMP type and codes combinations that are allowed by the firewall and ACL	Protocol 1 Flood ICMP Flood of Echo-Request/Response Type (Type= 0, Code = 0)
MyDoom attack	Excessive number of HTTP packets zombies.	Layer 3 – TCP protocol (6) Layer 4 – TCP port 80 Layer 4 – SYN Layer 4 – New connections Layer 4 – Established connections	Protocol 6 Flood SYN Flood Zombie Flood Port Flood
HTTP GET attack	Excessive number of HTTP packets from zombies.	Layer 3 – TCP protocol (6) Layer 4 – TCP ports on which the server is listening and ports that are allowed by the firewall and ACL Layer 4 – SYN Layer 4 – New connections Layer 4 – Concurrent connections per source Layer 7 – HTTP Methods Layer 7 – URL	Protocol 6 Flood SYN Flood Zombie Flood Port Flood TCP Connection Flood HTTP Method Flood URL Flood

Attack	Description	Threshold to monitor/adjust	Events to watch
Slow connection attack	Legitimate IP sources send legitimate TCP connections but do it slowly and remain idle, which fills up the server's connection table memory.	Layer 3 – TCP protocol (6) Layer 4 – TCP ports on which the server is listening and ports that are allowed by the firewall and ACL Layer 4 – SYN Layer 4 – New connections Layer 4 - Concurrent connections per source	Layer 3 Protocol 6 SYN Flood Zombie Flood Port Flood Concurrent Connections/ Source
UDP flood attack	An excessive number of UDP packets.	Layer 3 – UDP protocol (17) Layer 4 – UDP ports on which the server is listening and ports which are allowed by the firewall and ACL	Protocol 17 Flood Port Flood
Slammer attack	An excessive number of packets on UDP Port 1434.	Layer 3 – UDP protocol (17) Layer 4 – UDP port 1434	Protocol 17 UDP Flood Port Flood – 1434

Attack	Description	Threshold to monitor/adjust	Events to watch
Fraggle attack	Spoofed UDP packets to a list of broadcast addresses. Usually the packets are directed to port 7 on the target machines, which is the echo port. Other times, it is directed to the CHARGEN port. Sometimes a hacker is able to set up a loop between the echo and CHARGEN port.	Layer 3 – ICMP protocol (1) Layer 3 – UDP protocol (17) Layer 4 – UDP echo port (7) Layer 4 – Daytime Protocol port (13) Layer 4 – Quote of the Day (QOTD) port (17) Layer 4 – UDP Character Generator protocol (CHARGEN) (19) Layer 4 – ICMP Type/Codes specific to host/port not available	Protocol 1 Flood Protocol 17 Flood UDP Port 7 Flood UDP Port 13 Flood UDP Port 17 Flood UDP Port 19 Flood ICMP Flood of Port Not Available Type, Code (3,3) ICMP Flood of Host Not Available Type, Code (3,1)
DNS Port Flood	An excessive number of packets on UDP port 53.	Layer 3 - UDP protocol (17) Layer 4 - UDP port 53	Protocol 17 UDP Flood UDP Port 53 Flood ICMP Port/Host not available Flood
DNS Query Flood	A spike in DNS queries and occurrences of query data.	Layer 7 - DNS query-related thresholds	DNS Query Flood

Step 3: Identify the attack size

You can use the Monitor graphs to analyze the dimensions of the attack: increases in throughput and drops.

Step 4: Analyze attack parameters in each OSI layer

You can use the DDoS Attack log or the Monitor graphs to analyze aggregate throughput and drops due to Layer 3, Layer 4, and Layer 7 FortiDDoS rate thresholds or ACL rules.

1. Start using the following graphs to identify the layer at which the attack is happening:
   - Aggregate Flood Drops
   - Aggregate ACL Drops

- Anomaly Drops statistics

2. Drill down further by accessing statistics specific to each layer and attack type.

Chapter 2: Getting Started

This chapter provides the basic workflow for getting started with a new deployment.

Basic steps:

1. Install the appliance.
2. Configure the management interface.
3. Configure the following basic network settings:
   - Administrator password
   - System date and time
   - Network interfaces
   - DNS
4. Test connectivity.
5. Complete product registration, install your license, and update the firmware.
6. Deploy the system in Detection Mode for 2-7 days.
7. Generate traffic statistics, review them, and set SPP thresholds to the system recommended values.
8. Continue to monitor throughput rates and attacks, and adjust thresholds as needed.
9. Deploy the system in Prevention Mode.
10. Back up this basic configuration so that you have a restore point.

Tips:

- Configuration changes are applied to the running configuration as soon as you save them.
- Configuration objects are saved in a configuration management database. You cannot change the name of a configuration object after you have initially saved it.
- You cannot delete a configuration object that is referenced in another configuration object (for example, you cannot delete an address if it is used in a policy).

Step 1: Install the appliance

This Handbook assumes you have already installed the appliance into a hardware rack and used the appropriate cables to connect the traffic interfaces to your network.

The FortiDDoS system is deployed inline (between the Internet and your local network resources). Consecutively numbered ports belong to port pairs: Use an odd port numbers (1, 3, 5, and so on) for the LAN-side connection and an even port number (2, 4, 6, and so on) for the WAN-side connection. For example, port1 and port2 are a pair. The port1 interface is connected to a switch that connects servers in the local network; the port2 interface is connected to the network path that receives traffic from the Internet.

For information on hardware appliances, refer to the FortiDDoS hardware manuals.

Step 2: Configure the management interface

You use the management port for remote administrator access from the web user interface (web UI) or command line interface (CLI).

Figure 36 shows the web UI.

Figure 36: Web UI

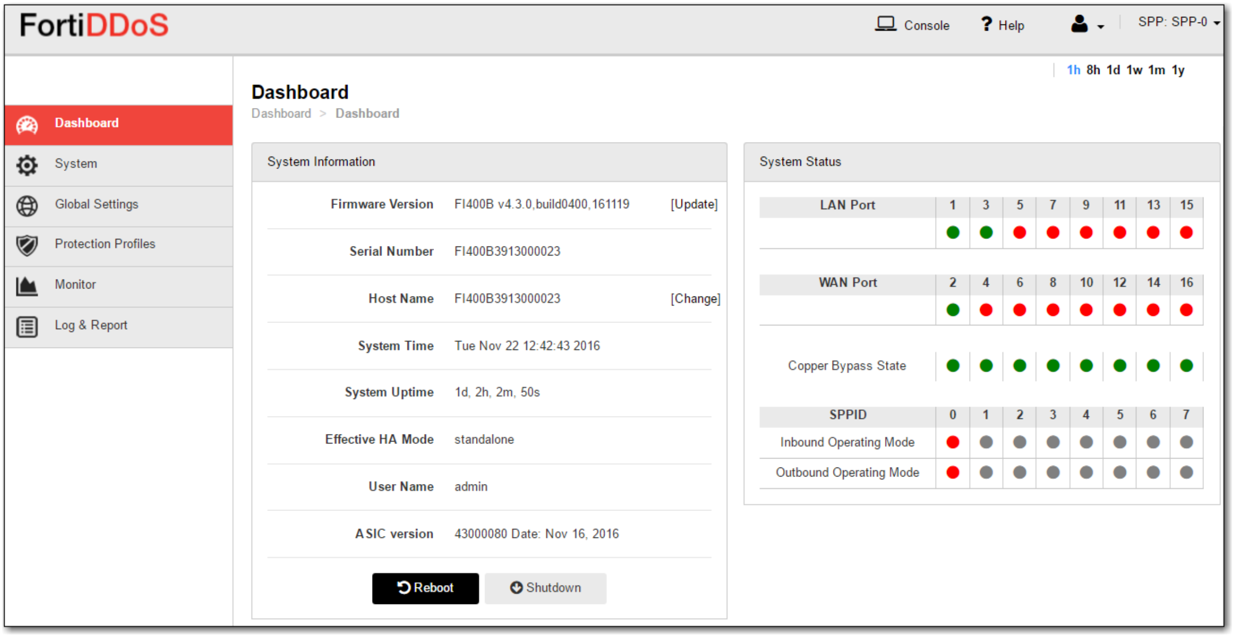

Dashboard

Dashboard > Dashboard

System Information

Firmware Version	FI400B v4.3.0,build0400,161119	[Update]
Serial Number	FI400B3913000023	
Host Name	FI400B3913000023	[Change]
System Time	Tue Nov 22 12:42:43 2016	
System Uptime	1d, 2h, 2m, 50s	
Effective HA Mode	standalone	
User Name	admin	
ASIC version	43000080 Date: Nov 16, 2016	

Reboot Shutdown

System Status

LAN Port	1	3	5	7	9	11	13	15
	●	●	●	●	●	●	●	●
WAN Port	**2**	**4**	**6**	**8**	**10**	**12**	**14**	**16**
	●	●	●	●	●	●	●	●
Copper Bypass State	●	●	●	●	●	●	●	●
SPPID	**0**	**1**	**2**	**3**	**4**	**5**	**6**	**7**
Inbound Operating Mode	●	●	●	●	●	●	●	●
Outbound Operating Mode	●	●	●	●	●	●	●	●

You configure the following basic settings to get started so that you can access the web UI from a remote location (like your desk):

- Static route—Specify the gateway router for the management subnet so you can access the web UI from a host on your subnet.
- IP address—Assign a static IP address for the management interface. The IP address is the host portion of the web UI URL. For example, the default IP address for the management interface is 192.168.1.99 and the default URL for the web UI is https://192.168.1.99.
- Access—Services for administrative access. We recommend HTTPS, SSH, SNMP, PING.

Before you begin the management interface configuration:

- You must know the IP address for the default gateway of the management subnet and the IP address you plan to assign the management interface.
- For your initial setup, you must have access to the machine room in which the physical appliance has been installed. You must connect a cable to the management port to get started.
- You need a laptop with an RJ-45 Ethernet network port, a crossover Ethernet cable, and a web browser (Microsoft Internet Explorer 8.0 or newer, or Mozilla Firefox 20 or newer). To minimize scrolling, the monitor resolution should be 1280 x 1024 or better.

- Configure the laptop Ethernet port with the static IP address 192.168.1.2 and a netmask of 255.255.255.0. These settings enable you to access the web UI as if from the same subnet as the FortiDDOS in its factory configuration state.
- Use the crossover cable to connect the laptop Ethernet port to the management port.

To connect to the web UI:

1. On your laptop, open the following URL in your web browser:
   https://192.168.1.99/

   The system presents a self-signed security certificate, which it presents to clients whenever they initiate an HTTPS connection to it.
2. Verify and accept the certificate, and acknowledge any warnings about self-signed certificates.
   The system displays the administrator login page.
3. Enter the username **admin** and no password.

The system displays the dashboard.

Note: It is not recommended to use Internet Explorer version 9 and 10. If you login to FortiDDoS GUI on Internet Explorer 11 from Windows 10 system, perform the following actions on IE 11 browser settings:

1. Go to Settings > Internet options.
2. Click Settings under Browsing history.
3. Select 'Every time I visit the webpage' option under 'Check for newer versions of stored pages:'.

To configure a static route:

- Go to System > Network > Static Route.

For details, refer to the online help or see Configuring static routes.

To configure the IP address and access services:

1. Go to System > Network > Interface.
2. Double-click the row for mgmt1 to display the configuration editor.
3. Use CIDR notation to specify the IP address/netmask, and enable services related to administrative access.
4. Save the configuration.

The system processes the update and disconnects your HTTP session because the interface has a new IP address and therefore the web UI has a new URL. At this point, you should be able to connect to the web UI from a host on the management subnet you just configured. You can go back to your desk to verify connectivity by attempting to open the web UI at the new address. You could see the status of configuration and link under Configured Status and Link Status column.

Figure 37: Network interface configuration page

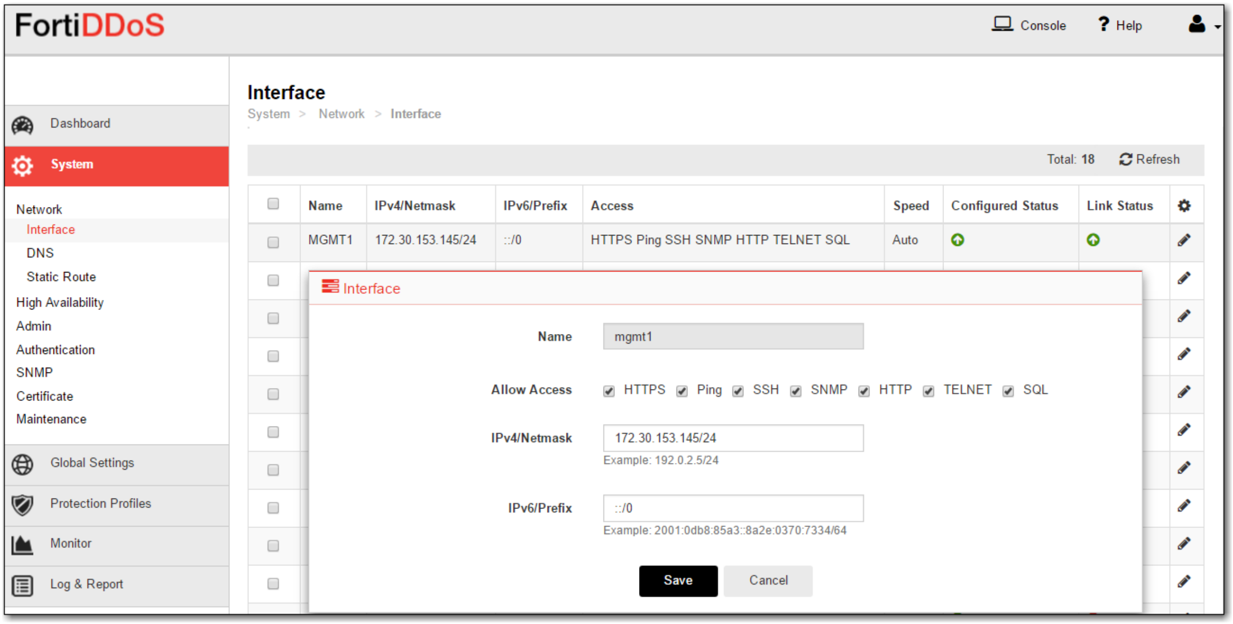

Name	IPv4/Netmask	IPv6/Prefix	Access	Speed	Configured Status	Link Status
MGMT1	172.30.153.145/24	::/0	HTTPS Ping SSH SNMP HTTP TELNET SQL	Auto	●	●

Interface

Name: mgmt1

Allow Access: ☑ HTTPS ☑ Ping ☑ SSH ☑ SNMP ☑ HTTP ☑ TELNET ☑ SQL

IPv4/Netmask: 172.30.153.145/24
Example: 192.0.2.5/24

IPv6/Prefix: ::/0
Example: 2001:0db8:85a3::8a2e:0370:7334/64

For details, refer to the online help or see Configuring network interfaces.

To complete the procedures in this section using the CLI:

1. Use an SSH client such as PuTTY to make an SSH connection to 192.168.1.99 (port 22).
2. Acknowledge any warnings and verify and accept the SSH key.

   The system displays the administrator login prompt.
3. Enter the username **admin** and no password.
4. Use the following command sequence to configure the static route:

```
config system default-gateway
  edit 1
    set gateway 172.30.153.254
end
```

5. Use the following command sequence to configure the management interface:

```
config system interface
   edit mgmt1
     set ip <address/mask>
     set allowaccess {https ping ssh snmp http
       telnet sql}
end
```

The system processes the update and disconnects your SSH session because the interface has a new IP address. At this point, you should be able to connect to the CLI from a host on the management subnet you just configured. You can go back to your desk to verify the configuration.

Step 3: Configure basic network settings

The system supports network settings for various environments. To get started, you configure the following basic settings:

- Administrator password—You must change the password for the **admin** account.
- Network interfaces—If necessary. The FortiDDoS appliance is deployed inline. In effect, it is a Layer 2 Bridge, so you do not configure IP addresses for its traffic interfaces. By default, the system is configured to autonegotiate speed/duplex. If desired, you can configure fixed speed/duplex settings.
- DNS—You must specify a primary and secondary server for system DNS lookups.
- System date and time—We recommend you use NTP to maintain the system time.

To change the admin password:

- Go to System > Admin > Administrator tab.

For details, refer to the online help or see Managing administrator users.

To configure network interfaces:

- Go to System > Network > Interface.

For details, refer to the online help or see Configuring network interfaces.

To configure DNS:

- Go to System > Network > DNS.

For details, refer to the online help or see Configuring DNS.

To configure system time:

- Go to System > Maintenance > Time Zone tab.

For details, refer to the online help or see Configuring system time.

To complete the procedures in this section using the CLI:

1. Use a command sequence similar to the following to change the administrator password:

   ```
   config system admin
      edit admin
         set password <new-password_str> ''
   end
   ```

2. Use a command sequence similar to the following to configure network interface speed/duplex:

   ```
   config system interface
      edit port1
         set speed
            {10full|10half|100full|100half|1000full|1000half|auto}
   end
   ```

3. Use a command sequence similar to the following to configure DNS:

   ```
   config system dns
      set primary <address_ipv4>
      set secondary <address_ipv4>
   end
   ```

4. Use a command sequence similar to the following to configure NTP:

   ```
             config system time ntp
                set ntpsync enable
                set ntpserver {<server_fqdn> | <server_ipv4>}
                set syncinterval <minutes_int>
             end
   ```

   Or use a command syntax similar to the following to set the system time manually:

   ```
             config system time manual
                set zone <timezone_index>
                set daylight-saving-time {enable | disable}
             end
                execute date <time_str> <date_str>
   ```

Step 4: Test connectivity to protected servers

Use ping and traceroute to test connectivity to protected servers.

To test connectivity from the FortiDDOS system to the protected server:

- Run the following commands from the CLI:

```
execute ping <destination_ip4>
execute traceroute <destination_ipv4>
```

To test connectivity from the protected server to the FortiDDOS system:

1. Enable ping on the network interface.
2. Use the ping and traceroute utilities available on the protected server to test connectivity to the FortiDDOS network interface IP address.

For troubleshooting tips, see Chapter 12: Troubleshooting.

Step 5: Complete product registration, licensing, and upgrades

Your new FortiDDoS appliance comes with a factory image of the operating system (firmware). However, if a new version has been released since factory imaging, you might want to install the newer firmware before continuing the system configuration.

Before you begin:

- Register—Registration is required to log into the Fortinet Technical Support site and download firmware upgrade files. For details, go to http://kb.fortinet.com/kb/documentLink.do?externalID=12071.
- Check the installed firmware version—Go to **Dashboard**.
- Check for upgrades—Major releases include new features, enhancements, and bug fixes. Patch releases can include enhancements and bug fixes. Download the release notes at http://docs.fortinet.com/fortiddos/. Download firmware upgrades at https://support.fortinet.com/.

To upgrade the firmware:

- Go to Dashboard > System Information.

For details, refer to the online help or see Updating firmware.

Step 6: Deploy the system in Detection Mode

You initially deploy the system in Detection Mode. In Detection Mode, the system operates with high (factory default) thresholds and does not drop any packets.

The system needs about 2 to 7 days of attack-free learning in Detection Mode to learn typical traffic patterns so it can set the initial thresholds. The length of the initial learning period depends upon the seasonality of traffic (its predictable or expected variations) and how representative of normal traffic conditions the learning period is.

Weekends alone are an insufficient learning period for businesses that have substantially different traffic during the week. Thus, it is better to start the learning period on a weekday. In most cases, 7 days is sufficient to capture the weekly seasonality in traffic.

Basic steps

1. Go to Global Settings > Service Protection Profiles > Config and configure SPP names and IDs.
2. Go to Global Settings > Service Protection Profiles > SPP Policy and configure SPP subnets.
3. Go to Protection Profiles > SPP Settings and ensure, for each SPP, that the system is deployed in Detection Mode (factory default).

For details, refer to the online help or see Configuring an SPP policy and Configuring SPP settings.

Step 7: Generate traffic statistics and set the configured minimum thresholds

At the end of the initial learning period, you can adopt system-recommended thresholds (usually lower than the factory default).

Basic steps

1. Go to Protection Profiles > Traffic Statistics > Generate and generate statistics for the selected SPP and time period.
2. Go to Protection Profiles > Traffic Statistics > Details and review the maximum packet rates generated for the SPP.

The values represent the maximum packet rate observed during the selected period. For example, during each 1-hour period, there are 12, 5-minute observation periods. FortiDDoS captures a maximum rate for each 5-minute interval. The generated threshold is the highest maximum rate that was captured among the 12 observation periods.

3. Go to Protection Profile > Thresholds > System Recommendation and set the configured minimum thresholds to the system recommended settings.

   For each OSI layer you specify two settings:

   - Percentage: The configured minimum threshold is the generated baseline rate multiplied by this percentage.
   - Low Traffic Threshold: The system uses this value instead for the configured minimum threshold if it is higher.

Tip: When you are getting started, we recommend that you accept the defaults for the adjustment percentages and low traffic thresholds.

For details, refer the online help or see the following sections:

- Generating baseline traffic statistics
- Displaying baseline traffic statistics
- Modifying thresholds

Step 8: Monitor the system and become familiar with logs and reports

For your initial deployment, continue to use Detection Mode for a day or two during which you review logs for potential false positives and false negatives.

Basic steps

1. Go to the Monitor menu and review throughput rates. Start with aggregate graphs and then use the more detailed graphs to drill in on patterns of interest or concern.
2. Go to Log & Report > Log Access > Logs > DDoS Attack Log and become familiar with the log table and how to use log filters.
3. Go Log & Report > Executive Summary and become familiar with the Executive Summary dashboard including DDoS Attack Graphs dashboard.

For details, refer to the online help or see the following sections:

- Chapter 5: Monitor Graphs
- Using the DDoS attack log table
- Using the DDoS Attack Log dashboard

Figure 38: Sample Attack Graph dashboard

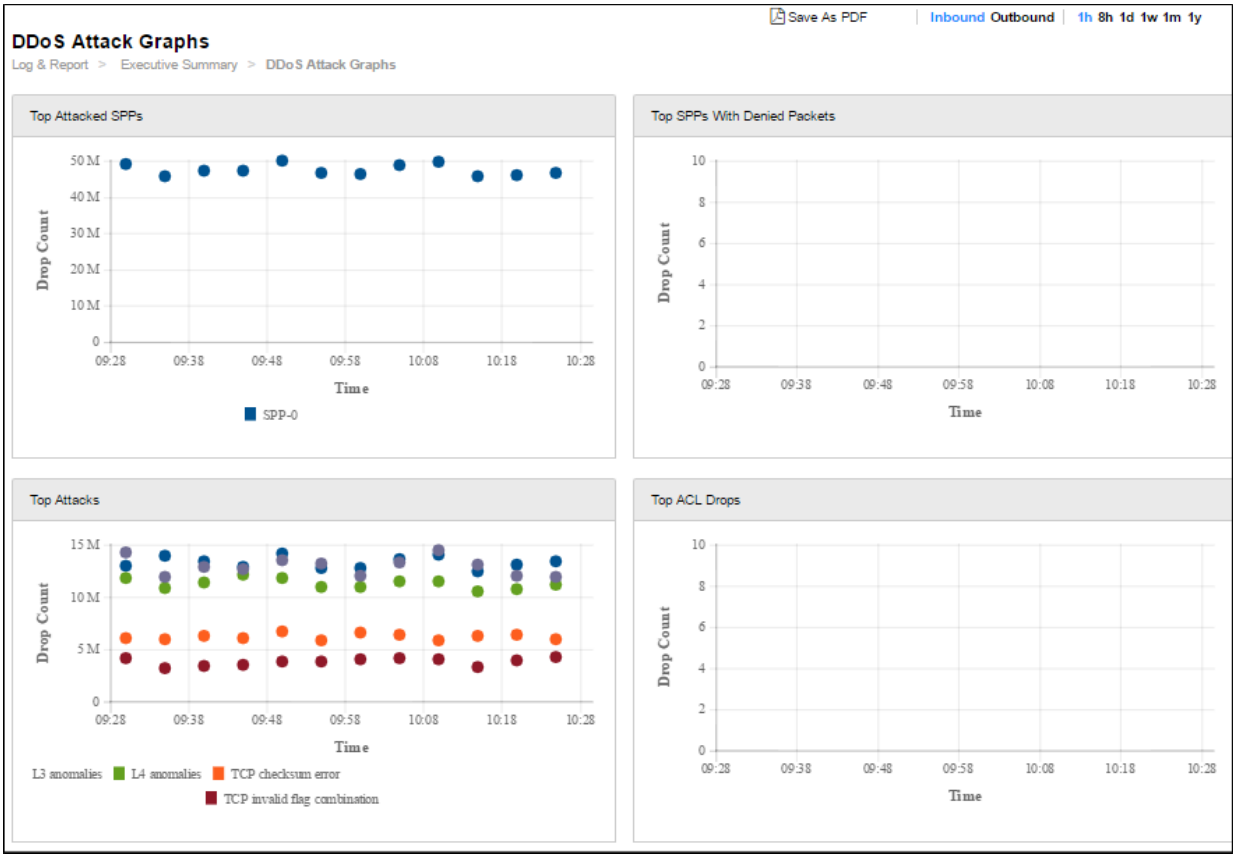

For details, refer to the online help or see Using the DDoS Attack Graph dashboard.

Step 9: Deploy the system in Prevention Mode

After you have set the statistical baseline and evaluated the configured minimum thresholds, you change to Prevention Mode. In Prevention Mode, the system uses the configured minimum threshold in its calculations that determine the estimated thresholds. The estimated thresholds are rate limits that are enforced by packet drops. The estimated thresholds are also the triggers for reporting flood attacks and entering SYN flood attack mitigation mode.

As needed, you repeat the tuning: monitor observed throughput, estimated thresholds, and drops; adjust the configured minimum thresholds; monitor; adjust.

Basic steps

1. Go to Protection Profiles > SPP Settings and change the configuration to Prevention Mode. Do this for each SPP.
2. On the Protection Profiles > SPP Settings > TCP tab, enable the recommended TCP session state anomalies options.
3. Continue to monitor traffic.
4. Tune the configuration if necessary. Go to Protection Profiles > Thresholds > Thresholds to set rates manually or Protection Profiles > Thresholds > System Recommendation to adjust percentages applied at OSI layers or to adjust the low traffic threshold.

For details, refer to the online help or see Configuring SPP settings and Modifying thresholds.

Step 10: Back up the configuration

Once you have tested your basic installation and verified that it functions correctly, create a backup. This “clean” backup is a reference point that has many benefits, including:

- Troubleshooting—You can use a tools such as a tool such as diff to compare a problematic configuration with this baseline configuration.
- Restarting—You can rapidly restore your system to a simple yet working point.
- Rapid deployment—You can use the configuration file as a template for other FortiDDOS systems to the extent it makes sense to do so. For example, you might use the same network infrastructure configuration (DNS, SNMP, log, syslog), the same general settings, and more or less the same ACL rules, but SPP settings and SPP thresholds are usually appropriate only to the subnet in which the system has been deployed. You can use any text editor to edit the plain text configuration file and import it into another FortiDDOS system. Be sure to change unique identifiers, such as the management IP address and sometimes other local network settings that differ from one deployment to another.

To backup the system configuration:

- Go to System > Maintenance > Backup & Restore tab.

For details, refer to the online help or see Backing up and restoring the configuration.

Chapter 3: Global Settings

This chapter includes the following topics:

Configuring SPP policy settings

This topic includes the following information:

- SPP basics
- SPP configuration overview
- Configuring SPP IDs
- Configuring the SPP switching policy
- Configuring an SPP policy

SPP basics

A *Service Protection Profile (SPP)* is a class for the counters and thresholds that protect a particular subnet. When the FortiDDoS system receives traffic, the *SPP policy* assigns the packets to an SPP based on source or destination IP address. The system monitors and maintains Layer 3, Layer 4, and Layer 7 data for each SPP.

You can configure 7 SPPs and 511 SPP policy rules.

When possible, consider creating two SPPs for each subnet you intend to protect: a primary SPP and an alternate SPP to be used when the packet rate to the primary SPP becomes high. You can enable an *SPP switching policy* to switch from the primary profile to the alternate profile when the packet rate reaches the maximum packet rate for the primary profile.

If desired, you can use one SPP for many rules. For example, you can create a protection profile named SPP-1 and apply it to two policy rules: Rule 1 protecting subnet 192.168.1.0/24 and Rule 2 protecting subnet 192.168.2.0/24.

SPP configuration overview

The SPP configuration objects are associated. Configure them in the order listed.

Basic steps

1. Configure multiple SPP IDs.
2. (Optional) Enable the SPP switching policy feature if you want to enable it in policy rules.
3. Configure the SPP policy that associates SPP IDs with the subnet address/mask.

The FortiDDoS system maintains traffic history for each SPP. It uses this data to generate recommended thresholds, dynamically adjust thresholds, and generate traffic statistics. If you change the SPP policy configuration or the resources it monitors, the data can become skewed. For example, if you remove a subnet from the profile, or change the servers that are deployed in the subnet, or change the services offered by those servers, the traffic history becomes less relevant.

Fortinet strongly recommends that you reset the traffic history for a profile before you make any significant changes to its configuration. Go to Protection Profiles > Factory Reset > Factory Reset. If you do not reset traffic statistics, changes to an SPP policy can result in counter-intuitive data accumulated in the longer reporting periods (year, month). For example, if a subnet belonged to the default SPP-0 before you assigned it to SPP-1, a report filtered by SPP-1 includes the SPP-0 traffic history for that subnet.

Best Practice: Provision a separate SPP for UDP
The only mitigation mechanisms for UDP attacks are source tracking and rate limiting. The source tracking feature detects attacks from a single source or a limited number of sources. However, if the attack is distributed, the rate limiting feature limits all UDP traffic in the SPP, including legitimate traffic. It cannot limit only the UDP traffic that is associated with the attack.

Unlike TCP traffic, the system does not track the state of UDP and DNS traffic. This means that for UDP traffic, it does not differentiate requests from responses. If you use thresholds for UDP ports that are too low, the system might block harmless UDP traffic. For example, in a DNS request, the destination UDP port is 53 and the source port is a randomly chosen UDP port. The response uses a source port of 53 and the destination port is the source port from the original request, which creates a lot of outbound traffic on many destination UDP ports. If you lower the traffic thresholds for ports other than 53, FortiDDoS might block legitimate responses.

To avoid these types of false positives, create an SPP that regulates UDP traffic only and set its outbound UDP port thresholds to reasonably high values. Alternatively, to avoid issues with DNS in FortiDDoS deployments, consider using DNS forwarders.

Configuring SPP IDs

The SPP name and ID are used in the FortiDDoS configuration and in report aggregations and filters.

We recommend you configure names that help you remember the purpose or other characteristics of the profile. For example, you can name one profile `web_servers`, another profile `DNS_servers`, and so on.

Before you begin:

- You must have Read-Write permission for Global Settings.

To configure SPP IDs:

1. Go to Global Settings > Service Protection Profiles > Config.
2. Click **Add** to display the configuration editor.
3. Complete the configuration as described in Table 15.
4. Save the configuration.

Table 15: SPP ID configuration

Settings	Guidelines
Name	Cannot contain spaces.
ID	Number between 1 and 7. 0 is reserved for SPP-0, the default profile.

To configure with the CLI, use a command sequence similar to the following:

```
config spp
   edit <spp_name>
      config ddos spp setting
         set spp-id <id>
      end
   next
end
```

Configuring the SPP switching policy

You can use the SPP switching policy option to enable the FortiDDoS system to switch to the alternate profile when the traffic rate exceeds a packet/second threshold that you specify in the SPP policy configuration.

For example, you can:

- Use the SPP switching policy to toggle automatically between a primary profile that handles low levels of traffic and a secondary profile that enforces stringent thresholds.
- Pair a primary profile that is deployed in Detection Mode with a secondary profile that is deployed in Prevention Mode.
- Pair an SPP to itself, if the only goal is to have an alarm or other signaling action as a result of the threshold breach.

When the system switches to the secondary profile, it monitors and regulates traffic for the subnet using the secondary profile as long as the packet/second rate remains above the switching policy threshold. After traffic has remained steadily below it for a timeout period that you specify, the system switches back to the primary profile.

Before you begin:

- You must have Read-Write permission for Global Settings.
- After you have enabled the switching policy feature, you can specify it in an SPP policy.

To configure the switching policy:

1. Go to Global Settings > Service Protection Profiles > Switching Policy.
2. Complete the configuration as described in Table 16.
3. Save the configuration.

Table 16: SPP switching policy configuration

Settings	Guidelines
Switching	• Enable • Disable
Timeout	When the system switches to the secondary profile, it monitors and regulates traffic for the subnet using the secondary profile as long as the packet/second rate remains above the switching policy threshold. After traffic has remained steadily below it for a timeout period that you specify, the system switches back to the primary profile. The default value is 255 seconds and it is recommended to set a higher value if you modify.

To configure with the CLI, use a command sequence similar to the following:

```
config ddos global spp-switching-policy
   set switching {enable | disable}
   set timeout <integer>
end
```

Configuring an SPP policy

An SPP policy rule matches the source or destination IP address in packets received at a FortiDDoS interface to an SPP. The policy makes a determination as to which SPP the packets belong to. The packets then are added to that SPP's counters, and the packets are subject to that SPP's security thresholds.

The system matches traffic to rules in the SPP policy table from top to bottom. The first rule that matches is applied, so be sure to order rules for specific servers before rules for the subnet that contains the address. If no rules match, the packets belong to SPP-0.

The system uses SPP-0 to monitor and regulate the following types of packets:

- Packets that do not match any policy rule.
- Packets that have a corrupt IP header.

SPP-0 is a catch-all profile and its traffic statistics are affected by the traffic that FortiDDoS assigns to it by default. Therefore, we recommend that you do not associate protected subnets with SPP-0. This practice ensures that all known traffic is included in non-default subnets and non-default SPPs.

Before you begin:

- You must have configured SPP IDs.
- You must have enabled the SPP switching policy feature if you want to configure it for the SPP policy.
- You must have Read-Write permission for Global Settings.

To configure an SPP policy:

1. Go to Global Settings > Service Protection Profiles > SPP Policy.
2. Click **Add** to display the configuration editor.
3. Complete the configuration as described in Table 17.
4. Save the configuration.

After you have saved the configuration, you can reorder the list. Click anywhere in the row to select it. Ctrl-Click to deselect it. Drag the table rows into the order you want them.

Table 17: SPP policy configuration

Settings	Guidelines
Name	Configuration name that describes the subnet.
Subnet ID	A value between 1 and 511 that identifies the subnet.
IP version	• IPv4 • IPv6
IP address/mask	IP address and CIDR-formatted subnet mask. For example, 192.0.2.0/24, `2001:DB8:12AB`
SPP profile	Select the profile. We recommend that you not associate subnets with the default SPP profile SPP-0. This practice ensures that all known traffic is included in non-default subnets and non-default SPPs. SPP-0 functions as a catch-all profile. Its traffic statistics include traffic that FortiDDoS assigns to it by default.
Comments	Add comments describing the purpose of the SPP policy so that other administrators are aware of its intended use.
SPP Switching	
Enable SPP Switching	• Enable • Disable
Alternate Service Protection Profile	Select the secondary SPP. If you simply want a notification that the traffic level has exceeded the SPP switching threshold without switching the SPP, select the primary SPP.
Threshold	Maximum packet rate (packets per second) for the primary profile. SPP Switching Threshold is the sum of both inbound and outbound pps rates to that SPP. When traffic exceeds this rate, the system switches to the secondary SPP. The default threshold value is high. The unit of the Measurement Threshold value would be assigned based on your selection of **SPP Switching Threshold Measurement Unit** option under Global setting > Settings > Settings > General tab. **Note**: We suggest you use the sum of all Protocols used, typically, 1(ICMP) 6(TCP), 17(UDP) and 50 (GRE), for example.

To configure with the CLI, use a command sequence similar to the following:

```
config ddos global spp-policy
   edit <rule_name>
      set subnet-id <entry_index>
      set ip-version {IPv4 | IPv6}
      set ip <address_ip/mask>
      set spp <spp_name>
      set enable-alt-spp {enable | disable}
      set alt-spp <spp_name>
      set switching-threshold <rate>
end
```

To change the order of rules:

```
config ddos global spp-policy
   move <entry_index> after <entry_index>
end
```

Configuring global settings

Global settings specify system behavior and detection settings that apply to all traffic, in contrast to SPP settings, which apply to traffic to and from the subnet matched in the SPP policy.

Before you begin:

- You must have Read-Write permission for Global Settings.

To configure global settings:

1. Go to Global Settings > Settings > Settings.
2. Click a tab to display its configuration page. Complete the configuration as described in the following sections:
   - General tab
   - Deployment tab
   - Blocking Period tab
   - Slow Connection tab
3. Save the configuration.

General tab

Use this tab to configure miscellaneous global settings.

Table 18: General tab settings

Settings	Guidelines
Source MAC Address Aggressive Aging	MAC address used to send TCP resets to the protected server when aggressive aging is triggered. By default, the system uses the MAC address of the management interface (mgmt1), but the MAC address displayed in the web UI is `00:00:00:00:00:00`. If you change this setting, the system uses the MAC address you specify.

Settings	Guidelines
HTTP Anomaly	Select one or more of the following HTTP anomaly responses: • known-method-anomaly—Drop HTTP traffic that uses a one of the eight methods: GET, HEAD, OPTIONS, PUT, POST, CONNECT, DELETE, or TRACE. By default all 8 are treated as valid and therefore no Monitor graphs are provisioned, even if there are drops. • unknown-method-anomaly—Drop HTTP traffic that uses a method other than one of the following: GET, HEAD, OPTIONS, PUT, POST, CONNECT, DELETE, or TRACE. For example, TEST or PROPFIND. Generates the attack log message 'Unknown Method Anomaly'. • invalid-HTTP-version-anomaly—Drop HTTP traffic with an HTTP version other than one of the following: 0.9, 1.0, or 1.1. Generates the attack log message 'Invalid HTTP Version Anomaly'. • do-not-parse-HTTP-0.9—Allow HTTP version 0.9 traffic. By default, it is dropped. The default setting is to treat HTTP 0.9 packets as HTTP packets and further parse. For more information about protocol anomalies, see Understanding FortiDDoS protocol anomaly protection.
Geo Location	The geolocation policy feature enables you to block traffic from the countries you specify, as well as anonymous proxies and satellite providers, whose geolocation is unknown. Select one of the following options to determine how geolocation rules in the Global ACL can be configured: • Allow all and deny some—You use the Global ACL rulebase to deny specified countries, anonymous proxies, and satellite providers. • Deny all and allow some—You use the Global ACL rulebase to allow specified countries, anonymous proxies, and satellite providers. **Note**: In this mode, countries are denied as Source or Destination. Be sure to add the country where the FortiDDoS appliance resides to the Allowed country list or all traffic will be blocked through FortiDDoS. Rules are based on the configured Geolocation address objects. See Configuring Geolocation addresses.

Settings	Guidelines
Local Address Antispoofing	These rules can be used to prevent attacks that spoof your internal addresses. Enable one or more antispoofing rules that consult the local address configuration: • Inbound source must not be local address—Blocks inbound packets that have a source address inside the network. The source address is definitely spoofed. • Inbound destination must be local address—Blocks inbound packets that do not have a destination in your network. The destination address is illegitimate. • Outbound source must be local address—Blocks outbound packets with a spoofed address. Reduces the risk of your network being used in spoof attacks. • Outbound destination must not be local-address—Blocks outbound packets with a destination inside your local network. Rules are based on the addresses you add to the Local address configuration. See Configuring Local addresses.
Drop HTTP Header Range	Enable to drop sessions when the HTTP request includes the HTTP Range header. The Range header can be abused by attackers to exhaust HTTP server resources. Disabled by default. Enable this feature if you know that your protected HTTP servers do not use the Range header, or when your protected network is being attacked with methods that exploit HTTP Range header behavior.
Invalid ICMP Anomaly	Packets dropped due to invalid ICMP type/code anomaly.
SPP Switching Threshold Measurement Unit	Select PPS or Mbps.
Persistent HTTP transactions	A simple HTTP transaction is one where the client makes a single request for HTTP content within a TCP session. Persistent connections allow the browser / HTTP client to utilize the same connection for different object requests to the same host name. If Persistent HTTP Transactions feature is enabled, FortiDDoS checks for application level conformity in every packet of a TCP connection. This functionality is similar to 4.2.x. If this feature is disabled (default in 4.3.0), checks are limited to the first transaction of a TCP connection. It is recommendation to use the disabled state to avoid HTTP anomalies, especially due to IP fragmentation and TCP segmentation.

To configure with the CLI, use a command sequence similar to the following:

```
config ddos global setting
   set ip-v6-prefix <ip_prefix>
   set source-mac-address-aggressive-aging <address>
   set http-anomaly {http-version-0-9 invalid-method-anomaly
       known-method-anomaly unknown-method-anomaly}
   set drop-http-header-range {enable|disable}
   set geolocation {deny-all-allow-some|allow-all-deny-some}
   set local-address-anti-spoofing {inbound-source-must-not-
       be-local-address inbound-destination-must-be-local-
       address outbound-source-must-be-local-address
       outbound-destination-must-not-be-local-address}
   set dns-response-size <bytes>
   set persistent-http-transactions {enable|disable}
end
```

Deployment tab

Use this tab to configure settings related to where in the network the FortiDDoS appliance is deployed.

Table 19: Deployment tab settings

Settings	Guidelines
Link Down Synchronization	• Wire—If one link in a peer port pair goes down, take down the other link, synchronizing the link state. When the down link becomes available, both links are brought up. • Hub—Do not synchronize the link state. **Note**: The system is restarted when you change this setting.
Power Fail Bypass Mode	• Fail Open—Default. Use to enable built-in bypass. The interfaces form a wire and pass traffic through without performing any monitoring or prevention tasks. • Fail Closed—Use with an external bypass unit or for the primary node in an HA active-passive deployment. Interfaces do not pass traffic. The external bypass system can detect the outage and forward traffic around the FortiDDoS. Applicable only for FortiDDoS models with copper ports or fixed LC connectors. For more information, see Built-in bypass.
Asymmetric Mode	Enable when deployed in a network segment where traffic can take asymmetric routes. This option is not enabled by default. Special considerations and configuration changes are required. See Understanding FortiDDoS Asymmetric Mode for TCP.

Settings	Guidelines
Allow Inbound SYN/ACK	Enable only when you enable Asymmetric Mode. When there is asymmetric traffic, the system might receive inbound SYN/ACK packets. When this option is enabled, these packets are treated as if there is a valid connection on which to accept data (if the connection does not already exist).
Tap Mode	Enable when deployed out-of-path in conjunction with a FortiBridge or FortiTap appliance. This option is not enabled by default. **Note**: The system is rebooted when you change this setting. Special considerations and configuration changes are required. See Tap Mode deployments.
Signaling Mode	• Customer Premises • Service Provider See Chapter 11: Service Provider Signaling Deployments.

To configure with the CLI, use a command sequence similar to the following:

```
config ddos global setting
   set link-down-synchronization {wire|hub}
   set power-fail-bypass-mode {fail-open|fail-closed}
   set asymmetric-mode {enable|disable}
   set asymmetric-mode-allow-inbound-synack {enable|disable}
   set tap-mode {enable|disable}
   set signaling-mode {customer-premises|service-provider}
end
```

Blocking Period tab

Use this tab to configure blocking period settings.

Table 20: Blocking Period tab settings

Settings	Guidelines
Blocking Period for All Attacks	How long to block traffic from any source when an attack threshold is triggered. The default is 15 seconds. The valid range is 1 to 15. When blocking period is over, the threshold is checked again. For example, if the value is 15, when the system detects a flood, it blocks the packets associated with that flood for 15 seconds. If the traffic still qualifies as an attack after the blocking period expires, it blocks the traffic for another 15 seconds, and so on. For more information about blocking periods, see Blocking and Reducing false positives.
Blocking Period for Identified Sources	How long to block all traffic from a source IP address associated with an attack event. The default is 60 seconds. The valid range is 1 to 65,535. When an attack threshold is triggered, the system multiplies the packet rate from the source of the blocked packets by the value of the source multiplier. If the calculated rate exceeds the value of the most-active-source threshold, the system identifies the IP address of the source as a source attacker.
Extended Blocking Period for Identified Sources	How long to block all traffic from a source IP associated with an attack resulting in the number of dropped packets specified in the next setting. The default is 60 seconds. The valid range is 1 to 65,535. Additional blocking period when the attack results in packet drops that exceed your specified threshold.
Drop Threshold to Extend Blocking Period for Identified Sources	Number of dropped packets that trigger the extended blocking period. The default is 5,000 dropped packets.

To configure with the CLI, use a command sequence similar to the following:

```
config ddos global setting
   set blocking-period <int>
   set source-blocking-period <int>
   set extended-blocking-period <int>
   set drop-threshold-within-blocking-period <int>
end
```

Slow Connection tab

Use this tab to specify slow connection detection settings (threshold and observation period). Select one of the following options:

- None: Do not monitor for slow connection attacks.
- Moderate: Uses predefined thresholds to detect slow connection attacks.
- Aggressive: Uses more aggressive (lower) thresholds to detect slow connection attacks.
- User-defined: Enables advanced users to specify custom thresholds to detect slow connection attacks.

Note: To reduce false positives, Fortinet recommends you initially set the option to moderate and switch to aggressive only if required.

When the User-defined option is selected, the defaults are the maximum values, so we recommend you use the predefined Moderate and Aggressive values as guidelines to help you specify your own settings.

The following table summarizes the predefined thresholds for slow connection detection.

Table 21: Slow connection detection thresholds

Setting	Moderate	Aggressive
Slow TCP connection byte threshold	512 bytes	2048 bytes
Slow TCP connection observation period	30 seconds	15 seconds

For more information about slow connection detection, see Slow connection detection and aggressive aging.

To configure with the CLI, use a command sequence similar to the following:

```
config ddos global setting
   set slow-connection-type {aggressive|moderate|none|user-
       defined}
end
```

Configuring HTTP service port settings

By default, the FortiDDoS system listens for HTTP traffic on service port 80. If the servers in your network use nonstandard ports for HTTP traffic, you can configure the system to listen for HTTP on nonstandard service ports. You can configure up to 8 HTTP service ports. When http service ports are configured, it is required to run system recommendation or manually update the threshold range.

Before you begin:

- You must have Read-Write permission for Global Settings.

To configure HTTP service port settings:

1. Go to Global Settings > Settings > HTTP Service Ports.
2. Click **Add** to display the configuration editor.
3. Complete the configuration as described in Table 22.
4. Save the configuration.

Table 22: HTTP service ports configuration

Settings	Guidelines
Name	Configuration name. Must not contain spaces.
Enable	Select to enable the service port configuration.
Port	Specify the port number.

To configure with the CLI, use a command sequence similar to the following:

```
config ddos global http-service-ports
   edit 1
      set enable-port {enable | disable}
      set port-number <port>
   end
```

The system recommended threshold procedure excludes HTTP service ports from the port configuration blocks that it generates. When user-configured HTTP service ports are enabled, the packet rate thresholds for the user-configured ports are set to a high rate. If an HTTP service port configuration is subsequently disabled or deleted, the threshold remains at the high rate until you change it manually or perform the system recommended threshold procedure.

To manually configure detection thresholds for the nonstandard service ports:

1. Go to Protection Profiles > Thresholds > Thresholds.
2. Select **TCP Ports** from the Type drop-down list.
3. Configure the threshold and save the configuration.

Configuring service provider signaling

The Service Provider Signaling feature enables small/medium businesses and enterprises to work with participating service providers to route traffic through a "scrubbing station" in the service provider network (SPN) before it is forwarded through the WAN link to the customer premises network (CPN).

For details on deployments with signaling between FortiDDoS devices, see Chapter 11: Service Provider Signaling Deployments.

For information on deployments with signaling to Verisign OpenHybrid, see the FortiDDoS Deployment Guide for Cloud Signaling with Verisign OpenHybrid.

Note: You must use mgmt1 port for signaling.

Before you begin:

- You must have Read-Write permission for Global Settings.

To configure service provider signaling:

1. Go to Global Settings > Settings > Signaling.
2. Click **Add** to display the configuration editor.
3. Complete the configuration as described in Table 23.
4. Save the configuration.

Table 23: Signaling configuration

Settings	Guidelines
Customer Premises FDD	
Service Provider Device Type	• FortiDDoS—If the service provider uses FortiDDoS, select this option and complete the fields described next. • Third Party—If the service provider does not use FortiDDoS, select this option and specify the account ID, shared secret, and URL expected by the third party.

Settings	Guidelines
Name	Configuration name. Must not contain spaces.
Service Provider Serial Number	Serial number of the FortiDDoS in the service provider network. The serial number configuration is case sensitive. FI200B0914000128 is not same as fI200B0914000128. Be careful to enter the serial number exactly as it is provided to you.
Shared Secret/API Key	Must match the string configured on the SPN FortiDDoS. **Note**: Once entered, the Shared Secret/API Key is not displayed on GUI nor in CLI and cannot be recovered. If forgotten, a new matching key must be entered for the paired devices.
Service Provider Address Type	• IPv4 • IPv6
Service Provider IP Address	IP address of the SPN FortiDDoS management interface.
Service Provider	
Name	Configuration name. Must not contain spaces.
Customer Premises FDD S/N	Serial number of the FortiDDoS in the customer premises network. The serial number configuration is case sensitive. FI200B0914000128 is not same as fI200B0914000128. Be careful to enter the serial number exactly as it is provided to you.
Shared Secret	Must match the string configured on the CPN FortiDDoS.
Customer Premises FDD IP Version	• IPv4 • IPv6
Customer Premises FDD IP address	IP address of the CPN FortiDDoS management interface.

CLI commands:

```
CP-FDD # config ddos global service-provider-devices
CP-FDD (service-provid~r) # edit SP-FDD
CP-FDD (SP-FDD) # set enable-sp-device enable
CP-FDD (SP-FDD) # set spp-device-type FDD
CP-FDD (SP-FDD) # set serial-number FI800B3913800024
CP-FDD (SP-FDD) # set shared-secret/Authorization-key test1
CP-FDD (SP-FDD) # set ipv4-address 172.30.153.125
CP-FDD (SP-FDD) # end
```

Configuring IP reputation settings

The FortiGuard IP Reputation Service is a licensed subscription service that maintains data on IP addresses and network IP ranges that pose a threat to your network. After you purchase IP Reputation, you register the FortiDDoS appliance serial number. Then, you can download the IP reputation list or schedule updates.

After you have enabled the feature, the FortiDDoS system downloads the most recent definitions file and then maintains updates for it according to the schedule you configure. To use over-the-wire updates, the management port must be able to access the Internet. Alternatively, you can obtain the IP reputation definitions file and upload it using the dashboard License Information.

The License Information portlet displays the status of the most recent update. If the download is successful and new definitions are available, the lists are replaced; otherwise, the previous list remains in use.

You can configure how the FortiDDoS system receives scheduled updates.

Information about packets denied by Local Address Anti-spoofing rules is reported in the following graphs and reports:

- Graphs (Monitor > ACL Drops > Layer 3, Monitor > Layer 3 > Address Denied)
- Executive Summary dashboard (Log & Report > Report Browse > Executive Summary)
- Reports (Log & Report > Report Configuration)

Before you begin:

- You must have Read-Write permission for Global Settings.

To configure IP reputation settings:

1. Go to Global Settings > IP reputation.
2. Complete the configuration as described in Table 24.
3. Save the configuration.

Table 24: IP reputation configuration

Settings	Guidelines
Status	• Enable—Enable scheduled updates. • Disable—Disable scheduled updates.
Override server IP	• Enable—Enable to specify the override server IP address. • Disable—To not use an override server address.
Schedule type	• Every—Schedule periodic updates. Specify the time to perform the update. • Daily—Schedule daily updates. Specify the time of day to perform the update. • Weekly—Schedule weekly updates. Specify the day and time to perform the update.

Settings	Guidelines
Category	Select an IP reputation subscription category. If you use IP reputation, we recommend you select **DDoS** reputation data. You can select from the following choices: • DDoS • Anonymous Proxies
Tunneling	Enable to use a web proxy server IP address.
Tunneling IP address	Web proxy server IP address.
Port	Port for the web proxy server.
User Name	Administrator user name for the web proxy server.
Password	Password for the web proxy server.

To configure with the CLI, use a command sequence similar to the following:

```
config ddos global ip-reputation
   set ip-reputation-status {enable | disable}
   set override-server-ip {enable | disable}
   set ip-reputation-ip-address <override_server_address>
   set ip-reputation-schedule-type {hourly | daily | weekly}
   [set schedule-hour <hour_int>]
   [set schedule-weekdays {sunday | monday | tuesday
       ...|saturday}]
   set ip-reputation-category {ddos anonymous-proxies}
   set tunneling-status {enable | disable}
   set tunneling-address <tunneling_address>
   set tunneling-port <tunneling_port_int>
   set tunneling-username <tunneling_user_str>
   set tunneling-password <tunneling_pwd>
end
```

Configuring proxy IP settings

FortiDDoS can take account of the possibility that a source IP address might be a proxy IP address, and adjust the threshold triggers accordingly. If a source IP address is determined to be a proxy IP address, the system adjusts thresholds for Most Active Source, SYN per source, and Concurrent Connections per Source by a multiplier that you specify.

You can configure either or both of the following methods to determine whether source IP address is a proxy IP address:

- Concurrent connection count—Used when there are many users behind a web proxy or NAT device like an enterprise firewall.
- HTTP headers—Used when there are many users behind a Content Delivery Network (CDN), such as Akamai.

Before you begin:

- You must have Read-Write permission for Global Settings.

To configure proxy IP settings:

1. Go to Global Settings > Proxy IP.
2. Complete the configuration as described in Table 25.
3. Save the configuration.

Table 25: Proxy IP configuration

Settings	Guidelines
Proxy IP threshold factor	Specify a multiplier when the source IP address is identified as a proxy IP address. For example, if you specify 32, and the Most Active Source threshold is 1000, then the Most Active Source threshold applied to proxy IP addresses is 32 * 1000 or 32,000. The default is 128. The maximum is 32,768.
Proxy IP list status	Displays the date and time when the list was last updated.
Detect proxy IP by number of connections	
Concurrent connections per source	Every 5 minutes, the system records the IP addresses of sources with more than this number of concurrent connections to test whether those sources might be using a proxy IP address. The default is 100 concurrent connections.
Percent present	Threshold that determines whether the source IP address is regarded as a proxy IP address. For example, the default is 30. After the observation period, the IPs whose numbers of concurrent connections have been 30% of the time above 100 are identified as proxy IPs.

Settings	Guidelines
Observation period	• Past Week—Uses data from the past week to determine whether a source IP address is a proxy IP address. • Past Month—Uses data from the past month.
Generate proxy IP list	Select to generate the list of detected proxy IP addresses. This list is useful for identifying IP addresses that the system has treated as a proxy but are actually attackers. You can add these kinds of IP addresses to an ACL to block their traffic.
Detect proxy IP using headers	
Proxy HTTP header type	Select HTTP headers that indicate a proxy address might be in use: • true-client-IP • x-forwarded-for (selecting this option also enables parsing of x-true-client-ip and x-real-ip headers) **Tip**: Shift-click to select multiple items.

To configure with the CLI, use a command sequence similar to the following:

```
config ddos global proxy-ip-setting
   set auto-proxy-ip-status {enable | disable}
   set proxy-ip-percent-present <integer>
   set proxy-ip-observation-period {past-week | past-month}
   set header-proxy-ip-status {enable | disable}
   set header-proxy-type {true-client-ip X-Forwarded-For}
   set proxy-ip-threshold-factor <integer>
end
```

Configuring address objects for global ACLs

This section includes the following information:

- Configuring Local addresses
- Configuring IPv4 addresses
- Configuring IPv6 addresses
- Configuring Geolocation addresses

Configuring Local addresses

Unlike the IP/IPv6 address configuration, the local address configuration is not used in the Global ACL policy. Instead, it is used by the Local Address Anti-spoofing rules configured on the Global Settings > Settings page. The anti-spoofing ACL leverages your knowledge of the local address space to prevent spoof attacks to and from local addresses.

You can enable one or more rules that consult the local address configuration:

- Inbound source must not be local address—Blocks inbound packets that have a source address inside the network. The source address is definitely spoofed.
- Inbound destination must be local address—Blocks inbound packets that do not have a destination in your network. The destination address is illegitimate.
- Outbound source must be local address—Blocks outbound packets with a spoofed address. Reduces the risk of your network being used in spoof attacks.
- Outbound destination must not be local address—Blocks outbound packets with a destination inside your local network.

Information about packets denied by Local Address Anti-spoofing rules is reported in the following graphs and reports:

- Graphs (Monitor > ACL Drops > Layer 3, Monitor > Layer 3 > Address Denied)
- Executive Summary dashboard (Log & Report > Executive Summary)
- Reports (Log & Report > Report Configuration)

Basic steps

1. Configure local addresses.
2. Enable Local Address Anti-spoofing rules.

Before you begin:

- You must have Read-Write permission for Global Settings.

To configure local addresses:

1. Go to Global Settings > Address > [Local Address Config | Local Address Config IPv6].
2. Click **Add** to display the configuration editor.
3. Complete the configuration as described in Table 26.
4. Save the configuration.

Table 26: Local address configuration

Settings	Guidelines
Name	Configuration name. Must not contain spaces.
IPv4	
IP-Netmask	Specify an address/mask pattern using CIDR notation.
IPv6	
IPv6-Prefix	Specify an IPv6 prefix to define a local address block.

To configure with the CLI, use a command sequence similar to the following:

```
config ddos global local-address
   edit <address_name>
      set ip-netmask <address_ipv4mask>
   end
```

Configuring IPv4 addresses

You create address objects to identify IPv4 addresses and subnets that you want to match in the following policy rulebases:

- Global ACL
- Do Not Track

Before you begin:

- You must have Read-Write permission for Global Settings.

To configure IPv4 addresses:

1. Go to Global Settings > Address > Address Config.
2. Click **Add** to display the configuration editor.
3. Complete the configuration as described in Table 27.
4. Save the configuration.

Table 27: IPv4 address configuration

Settings	Guidelines
Name	Configuration name. Must not contain spaces.

Settings	Guidelines
Type	• IP address—Create an entry for an individual IP address. • IP netmask—Create an entry for a subnet using an IP address/mask notation. **Note**: In the Global ACL for IPv4 addresses, you can add “deny rules” based on specified IP addresses or IP netmask configuration objects; you can add “allow rules” based on IP address configuration objects only.
Address	Specify an IP address or an address/mask pattern using CIDR notation.

To configure with the CLI, use a command sequence similar to the following:

```
config ddos global address
   edit <address_name>
      set type {ip-netmask | ip-address}
      set ip-netmask <address_ipv4mask>
      set ip-address <address_ipv4>
   end
```

Configuring IPv6 addresses

You create address objects to identify IPv6 addresses and subnets that you want to match in the following policy rulebases:

- Global ACL
- Do Not Track

Before you begin:

- You must have Read-Write permission for Global Settings.

To configure IPv6 addresses:

1. Go to Global Settings > Address > Address Config IPv6.
2. Click **Add** to display the configuration editor.
3. Complete the configuration as described in Table 28.
4. Save the configuration.

Table 28: IPv6 address configuration

Settings	Guidelines
Name	Configuration name. Must not contain spaces.

Settings	Guidelines
Type	• IPv6 Address—Create an entry for an individual IP address. • IPv6 Prefix—Create an entry for a subnet using an IPv6 address/prefix notation. **Note**: The restriction noted for the Global ACL for IPv4 addresses does not apply. In the Global ACL for IPv6 addresses, you can add “deny rules” or “allow rules” based on either IPv6 address IPv6 Prefix objects.
Address	Specify an IPv6 address. The address must fall within an address space specified by the IPv6 prefix set in global settings.
Prefix	Specify an IPv6 prefix using an IP address/prefix notation. The prefix must be consistent with the IPv6 prefix set in global settings.

To configure with the CLI, use a command sequence similar to the following:

```
config ddos global address-v6
   edit <address_name>
      set type {ipv6-network | ipv6-address}
      set ipv6-network <address_ipv6mask>
      set ipv6-address <address_ipv6>
end
```

Configuring Geolocation addresses

You create Geolocation configuration objects for the locations that you want to match in the Global ACL rulebase. Geolocation addresses include countries as well as anonymous proxies and satellite providers.

Information about packets denied by Global ACL Geolocation rules is reported in the following graphs and reports:

- Graphs (Monitor > ACL Drops > Layer 3, Monitor > Layer 3 > Address Denied)
- Executive Summary dashboard (Log & Report > Report Browse > Executive Summary)
- Reports (Log & Report > Report Configuration > Report Configuration)

Before you begin:

- You must have Read-Write permission for Global Settings.
- You should know the Geolocation Policy setting on the Global Settings > Settings configuration page. In the Global ACL, the action for Geolocation source addresses can only be one of Deny or Accept, depending on the Global Settings > Settings option. Knowing the setting for your deployment informs the geolocation addresses you create.

To configure Geolocation addresses:

1. Go to Global Settings > Address > Address Config.
2. Click **Add** to display the configuration editor.

3. Complete the configuration as described in Table 29.
4. Save the configuration.

Table 29: Geolocation address configuration

Settings	Guidelines
Name	Configuration name. Must not contain spaces.
Type	Select **Geolocation** to create an entry for a location, anonymous proxy, or satellite provider.
Geolocation	Select a location, anonymous proxy, or satellite provider.

To configure with the CLI, use a command sequence similar to the following:

```
config ddos global address
   edit <address_name>
      set type geo-location
      set geo-location <country_code>
   end
```

Note: The “country_code” for Anonymous Proxy is `A1`; the code for Satellite Provider is `A2`.

Configuring a Do Not Track policy

You can specify IP addresses that FortiDDoS does not restrict or track. Packets matching the Do Not Track policy are forwarded without inspection.

Before you begin:

- You must have configured address objects that you want to match in policy rules. See Configuring address objects for global ACLs.
- You must have Read-Write permission for Global Settings.

To configure a Do Not Track policy:

1. Go to Global Settings > Do Not Track Policy > [Do Not Track Policy | Do Not Track Policy IPv6].
2. Click **Add** to display the configuration editor.
3. Complete the configuration as described in Table 30.
4. Save the configuration.

Table 30: Do Not Track policy configuration

Settings	Guidelines
Name	Configuration name. Must not contain spaces.
IP address	Select an address object.
Do not track action	• Do not track—Never drop or block packets to/from these IP addresses; do not include them in the statistics for continuous learning and threshold estimation. • Track and Allow—Never drop or block packets to/from these IP addresses; include them in the statistics for continuous learning and threshold estimation.

To configure with the CLI, use a command sequence similar to the following:

```
config ddos global {do-not-track-policy | do-not-track-
    policy-v6}
   edit <do_not_track_name>
      set do-not-track-IP-address <address_object>
      set do-not-track-action {track-and-allow | do-not-
          track}
end
```

Configuring a global ACL policy

This section describes usage and configuration steps for the global access control list (ACL) policy. It includes the following information:

- Using the global ACL to block dark and bogon addresses
- Using a whitelist to reduce false positives
- Configuring a global ACL policy
- Configuring a global distress ACL for protocol traffic

Using the global ACL to block dark and bogon addresses

A bogon is an informal name for an IP packet on the public Internet that claims to be from an area of the IP address space that is reserved but not yet allocated or delegated by the Internet Assigned Numbers Authority (IANA) or a delegated Internet registry. The areas of unallocated address space are called “bogon space” or “dark address space”.

The term “bogon” stems from hacker jargon, where it is defined as the quantum of “bogosity”, or the property of being bogus. A bogon packet is frequently bogus both in the conventional sense of being forged for illegitimate purposes, and in the hackish sense of being incorrect, absurd, and useless.

In a private network, this could mean undefined private addresses should not be expected as source or destination. For example, if an enterprise uses only the 192.168.3.x range within its private domain, and then any other private addresses such as 192.168.1.x, 192.168.2.x and 192.168.4.x through 92.168.254.x are illegal. Use of these addresses usually means stealth activity that is mostly performed by worms.

In a public network, this would mean all bogon-prefixes should not appear as source or destination. A bogon prefix is a route that should never appear in the Internet routing table. A packet routed over the public Internet (not including over VPN or other tunnels) should never have a source address in a bogon range. These are commonly found as the source addresses of DDoS attacks.

Bogon prevention is a component of anti-spoofing. The following site is informative:

http://www.cymru.com/Documents/bogon-dd.html

You can configure FortiDDoS to block these types of addresses by adding them to its global ACL policy or to the SPP ACL policy.

Examples:

- To deny spoofing packets from the Internet with the source address 192.168.x.x, create an address object and global ACL rule to block 192.168.0.0/16.
- To deny outbound spoofing packets (that is, to deny addresses that are not in your inside LAN) with the source address as private addresses 172.16.x.x, create an address object and global ACL rule to block 172.16.0.0/16.
- To deny the address range 10.x.x.x altogether because it is “dark” both inside and outside your network, create an address object and global ACL rule to block 10.0.0.0/8.

Using a whitelist to reduce false positives

You can create a whitelist that includes IP addresses that are known to be acceptable, even if they exceed set thresholds. For example, devices that perform backups have a high traffic profile because they need to establish many connections or send a large number of packets to perform their tasks.

FortiDDoS does not track connections for items that are allowed by the global ACL. However, it does track the source and associated traffic for items that are configured as **Track & Allow** in the SPP ACL.

Configuring a global ACL policy

The global ACL policy establishes allow and deny rules for traffic based on source IP address.

Packets from IP addresses that are denied or allowed by ACLs do not affect the statistics for continuous learning for source addresses. However, other characteristics of the packets, such as protocols and ports, are included in the corresponding statistics.

Information about packets denied by the global ACL policy is reported in the following graphs and reports:

- Graphs (Monitor > ACL Drops > Layer 3, Monitor > Layer 3 > Address Denied)
- Executive Summary dashboard (Log & Report > Executive Summary)
- Reports (Log & Report > Report Configuration)

Before you begin:

- You must have configured address objects that you want to match in policy rules. See Configuring address objects for global ACLs.
- You must have Read-Write permission for Global Settings.

To configure a global ACL policy:

1. Go to Global Settings > Access Control List > [Access Control List | Access Control List IPv6].
2. Click **Add** to display the configuration editor.
3. Complete the configuration as described in Table 31.
4. Save the configuration.

Table 31: Access control list configuration

Settings	Guidelines
Name	Configuration name. Must not contain spaces.
Source address	Select an address object.

Settings	Guidelines
Action	Deny - Block traffic from traffic matching the address object from Address > Address Config, which may be an IP Address, subnet or Geolocation. **Note**: The action for Geolocation source addresses depends on the Geolocation Policy option on the Global Settings > Settings page. The user interface restricts the action you specify here to the logical action in the Global Settings setting. **Note**: You cannot set policies to Deny a subnet (for example,1.2.3.4/24) and Allow or Track and Allow a smaller subnet or IP address (for example, 1.2.3.5,). You need to split the subnets.

To configure with the CLI, use a command sequence similar to the following:

```
config ddos global {acl | acl6}
   edit <entry_index>
      set source-address <address_name>
      set action {accept | deny}
end
```

Configuring a global distress ACL for protocol traffic

You can use the SPP ACL to block traffic associated with a specific protocol. However, in some cases, it is more efficient for the system to block specific protocols before processing the SPP configuration. For example, to prevent large-scale, brute force attacks using UDP or ICMP, you can block these protocols on a subnet that is monitored by more than one SPP.

Drops based on the global ACL for protocol traffic are not included in traffic graphs or reports.

Before you begin:

- You must have Read-Write permission for Global Settings.

Unlike other ACLs, the global distress ACL blocks traffic even in detection mode.

To configure a global distress ACL:

1. Go to Global Settings > Access Control List > Advanced Settings > Distress ACL.
2. Click **Add** to display the configuration editor.
3. Complete the configuration as described in Table 32.
4. Save the configuration.

Table 32: Distress ACL configuration

Settings	Guidelines
Name	Configuration name. Must not contain spaces.
IP-Netmask	Specify the IP address and CIDR-formatted subnet mask, separated by a forward slash (/), such as 192.0.2.5/24. Dotted quad formatted subnet masks are not accepted.
Protocol	Specify a protocol number.

CLI:

```
config ddos global distress-acl
   edit <No.>
      set ip-netmask <address_ipv4mask>
      set protocol <protocol_int>
   next
end
```

Configuring a bypass MAC address list

In a deployment with a bypass switch such as FortiBridge, the bypass switch passes heartbeat packets to test the health of the FortiDDoS traffic interfaces. If the heartbeats packets are not passed, bypass mode is triggered.

You must configure an address list that allows heartbeat packets from the bypass switches to be passed through the FortiDDoS interfaces. The heartbeats are Layer 2 packets, so the system allows traffic based on the MAC addresses you configure.

Every FortiDDoS link pair can be connected via a FortiBridge link pair. For example, you can use a FortiBridge link to bridge the Port 1/Port 2 link pair and another FortiBridge link to bridge the Port 3/Port 4 link pair. Each link pair is associated with a pair of MAC addresses. Therefore, if you are using two links, you configure four MAC addresses.

You can add up to 16 MAC addresses to the bypass list. (Only 8 for FortiDDoS 200B, which has only 4 pairs of network interfaces.)

Before you begin:

- You must know the MAC addresses for the bypass switch.
- You must have Read-Write permission for Global Settings.

To configure a bypass MAC address list:

1. Go to Global Settings > Bypass MAC.
2. Click **Add** to display the configuration editor.
3. Complete the configuration as described in Table 33.
4. Save the configuration.

Table 33: Bypass MAC address list configuration

Settings	Guidelines
Name	Configuration name. Must not contain spaces.
MAC address	Specify the MAC address. **Note**: You can view MAC addresses for FortiBridge on its status page. If the bypass switches are from the same vendor, the most significant 24-bits of their MAC addresses are the same.

To configure with the CLI, use a command sequence similar to the following:

```
config ddos global bypass-mac
edit <entry_index>
set mac-address <address>
end
```

Configuring blacklisted domains

You can use Blacklisted Domains option to deny a large set of blacklisted domains.

To perform various functions under Blacklisted Domains:

1. Go to Global Settings > Blacklisted Domains.
2. Select the option based on the requirement:
   - **Upload**: Choose and upload the file with the list of blacklisted domains.
   - **Download**: Save the file with the list of blacklisted domains to your system.
   - **Add Domain**: Add a new blacklisted domain and click **Add** to include in the existing list.
   - **Delete Domain**: Enter the specific domain address to remove from the existing list and click **Delete.**

Figure 39: Blacklisted Domains

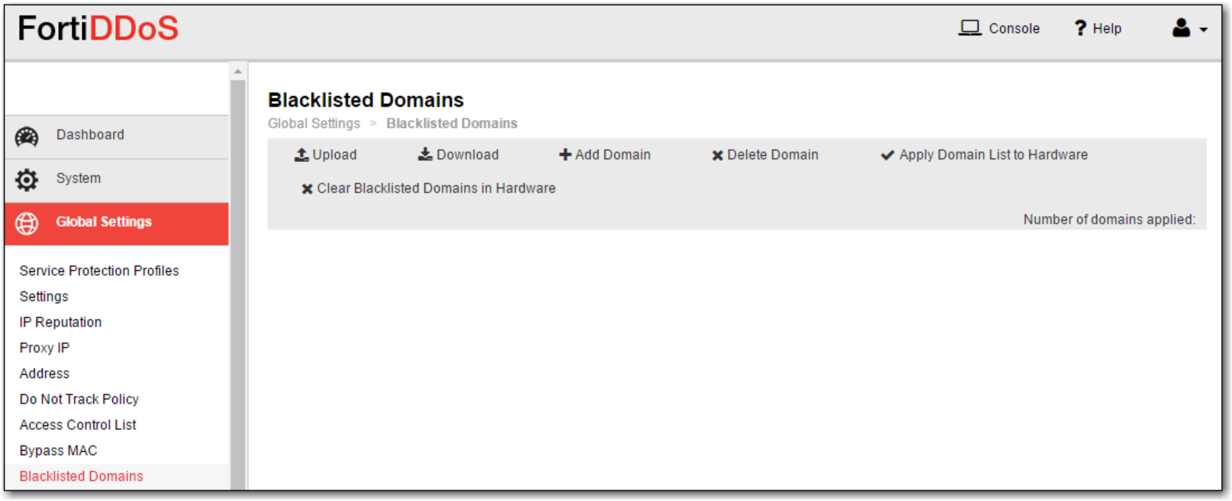

3. Click **Apply Domain list to Hardware** to apply the above changes or click **Clear Blacklisted Domains in Hardware** to remove the entire list of blacklisted domains.
   Number of domains applied shows the total count of blacklisted domains in the hardware.

To apply the Blacklisted Domain changes using the CLI, execute the corresponding commands:

Upload: `# execute dns-blacklist upload tftp =<filename> <serverip>`

Download: `# execute dns-blacklist download tftp <filename> <serverip>`

Add Address:`# execute dns-blacklist append domain<domain-name>`

Delete Address: `# execute dns-blacklist delete domain <domain-name>`

Apply Address list to Hardware: `# execute dns-blacklist domain apply`

Configuring blacklisted IPv4 addresses

You can use Blacklisted IPv4 Address option to deny a large set of blacklisted IPv4 addresses.

To perform various functions under Blacklisted IPv4 Addresses:

1. Go to Global Settings > Blacklisted IPv4 Addresses.
2. Select the option based on the requirement:
   - **Upload**: Choose and upload the file with the list of blacklisted addresses.
   - **Download**: Save the blacklisted address list to your system.
   - **Add Address**: Add a new address and click **Add** to include in the existing list.
   - **Delete Address**: Enter the specific address to remove from the existing list and click **Delete**.

Figure 40: Blacklisted IPv4 addresses

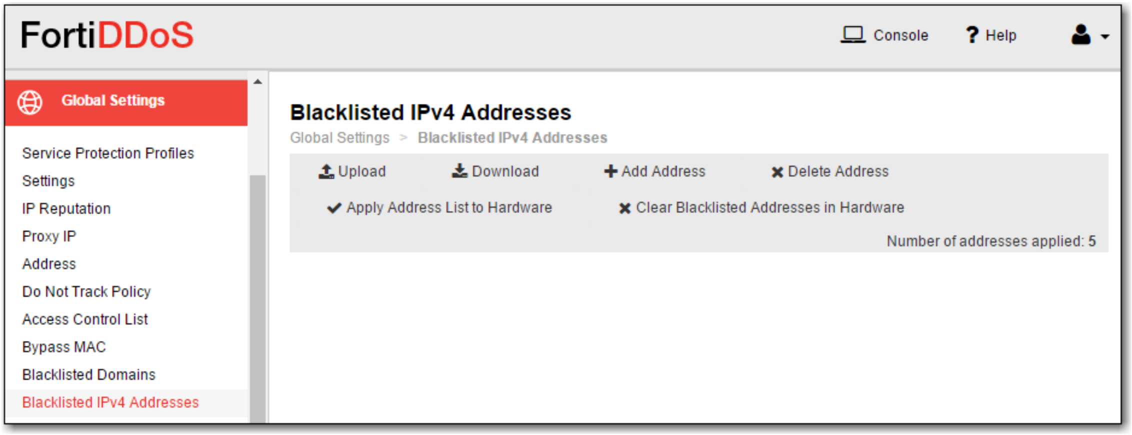
FortiDDoS
Console
Help
Global Settings
Service Protection Profiles
Settings
IP Reputation
Proxy IP
Address
Do Not Track Policy
Access Control List
Bypass MAC
Blacklisted Domains
Blacklisted IPv4 Addresses
Blacklisted IPv4 Addresses
Global Settings > Blacklisted IPv4 Addresses
Upload
Download
Add Address
Delete Address
Apply Address List to Hardware
Clear Blacklisted Addresses in Hardware

Number of addresses applied: 5

3. Click **Apply Address list to Hardware** to apply the above changes or click **Clear Blacklisted Addresses in Hardware** to remove the entire list of blacklisted addresses.

To apply the Blacklisted IPv4 changes using the CLI, execute the corresponding commands:

Upload: `# execute ipv4-blacklist upload tftp`

Download: `# execute ipv4-blacklist download tftp <filename> <serverip>`

Add Address:`# execute ipv4-blacklist append address<address>`

Delete Address: `# execute ipv4-blacklist delete address <address>`

Apply Address list to Hardware: `# execute ipv4-blacklist domain apply`

Using the preset anomaly detection setting

By default, packets with the following anomalies are dropped:

- IP first fragments for IPv4 and IPv6 packets that do not pass a DoS attack check.
- TCP packets with control flags = 0 and sequence number = 0.
- TCP SYN packets with source port between 0-1023.
- Packets with MAC source address = destination address.
- IPv4 and IPv6 packets where source address = destination address.
- IPv6 fragments smaller than the minimum size.
- Fragmented ICMP packets.
- TCP fragments with offset value of 1.
- UDP and TCP packets with source port = destination port.
- TCP packets with SYN and FIN bits set.
- TCP packets with FIN, URG and PSH bits set and Seq number = 0.

These drops are not included in logs or reports.

Packets with these anomalies are well understood to be either harmful or useless, so we recommend that you maintain the default setting. You can use the CLI to disable/enable it if needed for testing or debugging.

To enable/disable detection for this set of protocol anomalies:

Use the following command:

```
execute dos-control {enable | disable}
```

Chapter 4: Service Protection Profiles (SPP)

This chapter includes the following topics:

Configuring SPP settings

The SPP Settings configuration includes key feature settings that might vary among SPPs.

Before you begin:

- You must have a good understanding of the features you want to enable. Refer to Chapter 1: Key Concepts.
- You must have Read-Write permission for Protection Profile settings.

To configure SPP settings:

1. Go to Protection Profiles > SPP Settings.
2. Select the SPP you want to configure from the drop-down list.
3. Click a tab to display its configuration page. Complete the configuration as described in the following sections:
   - General tab
   - Source Tracking tab
   - TCP tab
   - DNS Anomalies tab
   - DNS Feature Controls tab
4. Save the configuration.

Note: FortiDDoS 600B and 900B do not support DNS ACLs, DNS anomaly detection, or DNS flood mitigation.

General tab

Use this tab to configure:

- Detection or Prevention Mode
- SYN flood mitigation mode
- Adaptive mode and adaptive limit

The following table provides guidelines.

Table 34: SPP configuration - General tab

Settings	Guidelines
Operating Mode	

Settings	Guidelines
Inbound operating mode	Set the mode for traffic received from WAN-side interfaces: • Detection—Logs events and builds traffic statistics for the profile but does not limit or block traffic. • Prevention—Limits and blocks traffic that exceeds thresholds.
Outbound operating mode	Set the mode for traffic received from LAN-side interfaces: • Detection—Logs events and builds traffic statistics for the profile but does not limit or block traffic. • Prevention—Limits and blocks traffic that exceeds thresholds.
SYN Flood Mitigation Mode	
SYN flood mitigation direction	Enable/disable the feature for one or both traffic directions. **Note**: If you do not enable SYN flood mitigation, and the TCP session feature control SYN Validation option is enabled, then during a flood, packets from sources not in the legitimate IP address table are not given the opportunity to complete the antispoofing challenge. The packets will be dropped.
SYN with payload direction	Enable/disable the feature in one or both traffic directions. A SYN packet with payload is an anomaly that indicates an attack known as a Tsunami SYN Flood. We recommend you enable this feature for inbound traffic. The only reason to disable is if you are running tests with tools that generate SYN packets with payload. Drops due to this anomaly are logged as L4 Anomaly events and included in the Layer 4 Anomaly graphs.

Settings	Guidelines
SYN flood mitigation mode	• ACK cookie—Sends the client two ACK packets: one with a correct ACK number and another with a wrong number. The system determines whether the source is spoofed based on the client's response. If the client's response indicates that the source is not spoofed, FortiDDoS allows the connection and adds the source to the legitimate IP address table. Fortinet recommends this option if you have enough bandwidth in the reverse direction of the attack. • SYN cookie—Sends a SYN/ACK with a cookie value in the TCP sequence field. If it receives an ACK back with the right cookie, an RST/ACK packet is sent and the IP address is added to the legitimate IP address table. If the client then retries, it succeeds in making a TCP connection. Fortinet recommends this option if you cannot use ACK Cookie and you anticipate high volume attacks. Fortinet recommends this option if you cannot use ACK Cookie and you anticipate high volume attacks. • SYN retransmission—Drops the initial SYNs to force the client to send a SYN again. If the expected number of retransmitted SYNs arrive within the predetermined time period, the system considers the source to be legitimate. FortiDDoS then allows the connection to go through and adds the source to the legitimate IP address table. Fortinet recommends this option if you cannot use ACK Cookie and you anticipate low volume attacks.
Adaptive Mode	
Adaptive mode	• Fixed—Does not use the adaptive limit. The configured minimum thresholds are the maximum limits. • Adaptive—Uses the adaptive limit. The configured minimum thresholds multiplied by the adaptive limit are the maximum limits.
Adaptive limit	A percentage of the configured minimum threshold that establishes the upper limit of the estimated threshold. The adaptive limit is an upper rate limit beyond which the system blocks all traffic. The valid range is 100% to 300%. For example, the default is 150%. The system uses the dynamic threshold estimation algorithm to raise the calculated threshold up to 150% of the value of the configured minimum threshold. Thus, if the inbound threshold for Protocol 17 (UDP) is 10,000, the threshold never falls below 10,000 and never exceeds 15,000. When the adaptive limit is 100, the system does not use dynamic threshold estimation to adjust thresholds.

To configure with the CLI, use a command sequence similar to the following:

```
config spp
   edit <spp_name>
      config ddos spp setting
         set inbound-operating-mode {detection|prevention}
         set outbound-operating-mode {detection|prevention}
         set syn-flood-mitigation-direction {inbound|outbound}
         set syn-flood-mitigation-mode {syn-cookie | ack-cookie |
             syn-retransmission}
         set syn-with-payload-direction {inbound|outbound}
         set adaptive-mode adaptive
         set adaptive-limit <percent_int>
      end
end
```

Source Tracking tab

Use this tab to configure packet count multipliers for identified source attackers and Layer 7 HTTP and DNS attacks.

The following table provides guidelines.

Table 35: SPP configuration - Source Tracking tab

Settings	Guidelines
Source multiplier inbound / outbound	Applies the specified multiplier to the packet count for traffic with a source IP address that the system has identified as the source of a flood. In effect, the multiplier makes traffic from the source violate thresholds sooner. The default is 2. For example, if the most active source threshold is 100 packets per second, and the source multiplier is 4, an identified source attacker will violate the threshold if it sends 26 packets per second. Because incoming traffic is more likely to be the source of a threat, you can configure different multipliers for incoming and outgoing traffic.
Layer 7 multiplier inbound / outbound	Applies the specified multiplier to the packet count for traffic that the system has detected is related to a Layer 7 HTTP flood. The system tracks HTTP headers (URL or Host, Referer, Cookie or User-Agent header) and associates traffic with matching headers with the attack. The default is 2. **Note**: When both Source flood and Layer 7 flood conditions are met, the packet count multipliers are compounded. For example, when there is a User Agent flood attack, a source is sending a User-Agent that is overloaded. If the Source multiplier is 4 and the Layer 7 multiplier is 64, the total multiplier that is applied to such traffic is 4 x 64 = 256. In effect, each time the source sends a Layer 7 packet with that particular User-Agent header, FortiDDoS considers each packet the equivalent of 256 packets.

To configure with the CLI, use a command sequence similar to the following:

```
config spp
   edit <spp_name>
      config ddos spp setting
         set source-multiplier-inbound <integer>
         set source-multiplier-outbound <integer>
         set layer-7-multiplier-inbound <integer>
         set layer-7-multiplier-outbound <integer>
      end
end
```

TCP tab

Use this tab to configure:

- TCP state anomaly detection
- Aggressive aging

Table 34 provides summary guidelines.

Special guidelines apply to TCP session feature control when the system is deployed in Detection Mode or Asymmetric Mode. Make sure you understand the recommendations in Understanding FortiDDoS Detection Mode or Understanding FortiDDoS Asymmetric Mode for TCP.

Table 36: SPP configuration - TCP tab

Settings	Guidelines
TCP session feature control	Select one or more of the following options to detect TCP state anomalies: • Sequence validation—The FortiDDoS TCP state machine ensures that TCP sequence numbers for the packets within a session are valid. • SYN validation—Required to support SYN Flood Mitigation. If SYN Validation is not enabled, packets are not dropped during a SYN flood. If SYN Validation is enabled, during a SYN flood, the TCP state machine allows only TCP SYNs from IP addresses in the legitimate IP address (LIP) table (sources that have done a three-way handshake in the recent past). SYNs from source IP addresses that do not have an entry in the LIP table must pass a SYN Flood Mitigation challenge to be added to the LIP table. • State transition anomalies validation—The TCP state machine ensures that TCP state transitions follow the rules. For example, if an ACK packet is received when FortiDDoS has not observed a SYN/ACK packet, it is a state transition anomaly. • Foreign packet validation—The TCP state machine drops TCP packets without an existing TCP connection and reports them as a foreign packet. In most cases, the foreign packet validation is useful for filtering out junk, but enabling it is not important. The number of foreign packets can be high, so the system does not store the source and destination of each packet. Therefore, you might not be able to determine the origin of a foreign packet. Foreign packet drops are logged in the DDoS Attack Log (State Anomalies event). **Important**: See TCP session state anomalies for recommended settings for Detection Mode, Prevention Mode, and Asymmetric Mode.

Settings	Guidelines
	The configuration also enables you to allow some TCP state sequences that the system would otherwise detect as a TCP state anomaly when the anomaly detection options are enabled. Select the following options to enable them: • Allow tuple reuse—Allow this exception to Sequence validation (if enabled). Otherwise, these are logged as a “State Anomalies: Outside window” event. When the “allow” option is enabled, the FortiDDoS TCP state machine updates the TCP entry when a tuple is reused. This update occurs only during the closed or close-wait, fin-wait, time-wait states, when the connection is just about to retire. Useful in testing environments where test equipment reuses tuples in rapid succession. Enabled by default. Cannot be disabled. • Allow duplicate SYN-in-SYN-SENT—Allow this exception to Sequence validation (if enabled). Otherwise, these are logged as a “State Anomalies: Outside window” event. When the “allow” option is enabled, the TCP state machine allows duplicate TCP SYN packets during the SYN-SENT state. It allows this type of packet even if the sequence numbers are different. Disabled by default. We suggest you enable this in Detection Mode. • Allow duplicate SYN-in-SYN-RECV—Allow this exception to Sequence validation (if enabled). Otherwise, these are logged as a “State Anomalies: Outside window” event. When the “allow” option is enabled, the TCP state machine allows duplicate TCP SYN packets during the SYN-RECV state. It allows this type of packet even if the sequence numbers are different. Disabled by default. Normally these violations are not expected in real-world traffic but might be seen in test environments. • Allow SYN anomaly, Allow SYN-ACK anomaly, Allow ACK anomaly, Allow RST anomaly, Allow FIN anomaly—Allow these exceptions to State transition anomalies (if enabled). Otherwise, these are logged as a “State Anomalies: State transition error” event. When the “allow” options are enabled, the TCP state machine allows duplicate TCP packets during any other state even if the sequence numbers are different from the existing connection entry. This is equivalent to allowing the packet without updating an existing connection entry with the new information from the allowed packet. Disabled by default. Normally these violations are not expected in real-world traffic but might be seen in test environments. In most cases, these options should remain disabled to enforce TCP compliance.

Settings	Guidelines
Aggressive aging TCP connections feature control	Select to enable aggressive aging options: • Layer 7 flood—Sends a TCP RST to the destination server to reset idle connections when a Layer 7 flood is detected. • High concurrent connections per source—Sends a TCP RST to the destination server to reset idle connections from the identified source when the maximum is reached for the concurrent-connection-per-source threshold. • Track slow TCP connections—Sends a TCP RST to the destination server to reset idle connections from the identified source when the slow connection attack thresholds set on the Global Settings > Settings page are reached. Slow connection detection events are logged in the DDoS Attack Log (Slow Connection: Agrressive Aging). For more information, see Aggressive aging.
Source blocking for slow connections	Enable/disable applying the source blocking when slow connection attacks are detected.

To configure with the CLI, use a command sequence similar to the following:

```
config spp
   edit <spp_name>
      config ddos spp setting
        set tcp-session-feature-control {sequence-validation
            syn-validation state-transition-anomalies-
            validation foreign-packet-validation allow-
            tuple-reuse allow-duplicate-syn-in-syn-sent
            allow-duplicate-syn-in-syn-recv allow-syn-
            anomaly allow-syn-ack-anomaly allow-ack-anomaly
            allow-rst-anomaly allow-fin-anomaly}
        set aggressive-aging-feature-control {layer7-flood
            high-concurrent-connection-per-source track-
            slow-tcp-connections}
        set source-blocking-for-slow-connections
            {enable|disable}
      end
end
```

DNS Anomalies tab

Use this tab to configure DNS traffic anomaly detection. We recommend you enable all DNS anomaly detection options. The configuration is "open" for testing and troubleshooting purposes.

The following table describes the settings.

Table 37: SPP configuration - DNS Anomalies tab

Settings	Guidelines
DNS header anomaly	• Invalid op-code—Invalid value in the OpCode field. • SP, DP both 53—Normally, all DNS queries are sent from a high-numbered source port (49152 or above) to destination port 53, and responses are sent from source port 53 to a high-numbered destination port. If the header has port 53 for both, it is probably a crafted packet. • Illegal flag combination—Invalid combination in the flags field.
DNS query anomaly	• Query bit set—DNS query with the query reply (QR) bit set to 1. In a legitimate query, QR=0. • RA bit set—DNS query with the recursion allowed (RA) bit set. The RA bit is set in responses, not queries. • Null query—DNS query in which the question, answer, additional, and name server counts are 0. • QDCNT not 1 in query—Number of entries in the question section of the DNS packet is normally 1. Otherwise, it might be an exploit attempt.
DNS response anomaly	• QCLASS in reply—DNS response with a resource specifying a CLASS ID reserved for queries only (QCLASS). • Query bit not set—DNS response with the query reply (QR) bit set to 0. In a legitimate response, QR=1. • QTYPE in reply—DNS response with a resource specifying a TYPE ID reserved for queries only (QTYPE). • QDCNT not 1 in response—Number of entries in the question section of the DNS packet is normally 1. Otherwise, it might be an exploit attempt.
DNS buffer overflow anomaly	• TCP Message too long—TCP query or response message that exceeds the maximum length specified in the message header. • Label length too large—Query or response with a label that exceeds the maximum length (63) specified in the RFC. • UDP message too long—UDP query or response message that exceeds the maximum length specified in the message header. • Name too long—DNS name that exceeds 255 characters. This can cause problems for some DNS servers.

Settings	Guidelines
DNS exploit anomaly	• Pointer loop—DNS message with a pointer that points beyond the end of data (RFC sec4.1.4). This is an exploit attempt. • Class is not IN—A query/response in which the question/resource address class is not IN (Internet Address). Although allowed by the RFC, this is rare and might indicate an exploit attempt. • Message ends prematurely—A message that ends prematurely might indicate an exploit attempt. • Zone transfer—An asynchronous Transfer Full Range (AXFR) request (QTYPE=252) from untrusted networks is likely an exploit attempt. • Empty UDP message—An empty message might indicate an exploit attempt. • TCP Buffer underflow—A query/response with less than two bytes of data specified in the two-byte prefix field.
DNS info anomaly	• Type ALL used—Detects a DNS request with request type set to ALL (QTYPE=255).
DNS data anomaly	• Invalid type, class—A query/response with TYPE or CLASS reserved values. • TTL too long—TTL value is greater than 7 days (or 604800 seconds). • Extraneous data—A query/response with excess data in the packet after valid DNS data. • Name length too short—A query/response with a null DNS name.

To configure with the CLI, use a command sequence similar to the following:

```
config spp
   edit <spp_name>
      config ddos spp setting
        set dns-anomaly-feature-control {premature-end-of-
            packet extraneous-data long-ttl qclass-in-reply
            qtype-in-reply class-not-in invalid-class-type
            type-all zone-transfer pointer-loop long-name-
            length long-label-length short-name-length qr-
            bit-set qr-bit-not-set null-query qd-count-not-
            one-in-response qd-count-not-one-in-query
            illegal-flag-combination ra-bit-set invalid-op-
            code empty-udp tcp-message-long udp-message-long
            sp-equals-dp tcp-buffer-underflow }
      end
end
```

DNS Feature Controls tab

⚠ Most DNS Feature controls must be disabled if the FortiDDoS appliance is operating in Asymmetric traffic mode and cannot see both directions of the DNS traffic. See Understanding FortiDDoS Asymmetric Mode for more information.

Use this tab to configure DNS feature controls. We recommend you enable all mitigation features. The configuration is "open" for testing and troubleshooting purposes. For an overview of DNS features, see Understanding FortiDDoS DNS attack mitigation.

The following table provides guidelines.

Table 38: SPP configuration - DNS Feature Control tab

Settings	Guidelines
Match responses with queries (DQRM)	Enable/disable the DNS query response match (DQRM) table.
Allow only valid queries under flood (LQ)	Enable/disable the legitimate query (LQ) table.
Validate TTL for queries from the same IP under flood	Enable/disable the time-to-live (TTL) table.
DNS UDP Anti-spoofing Method inbound/outbound	Enable/disable anti-spoofing checks for inbound and outbound UDP DNS traffic.
DNS flood mitigation mode inbound/outbound	Specify the antispoofing method if the source IP address is not already in the legitimate IP table (LIP): • Force TCP (TC=1)—Return a DNS response to the client that has the DNS Truncate bit set and no response record data. A properly implemented DNS client will respond to the spoofed response by retrying the original DNS query using TCP port 53. • DNS Query Retransmission—Drop packets and test for valid retransmission. A valid client is expected to retransmit the queries within preset time windows.
Generate response from cache under flood	Enable/disable DNS caching.

Settings	Guidelines
Force TCP or forward to server when no cache response available	If DNS caching is enabled, one of the following behaviors must be configured: • Force TCP (TC=1)—Return a DNS response to the client that has the DNS Truncate bit set and no response record data. A properly implemented DNS client will respond to the spoofed response by retrying the original DNS query using TCP port 53. • Forward to Server—Forward the DNS query to the DNS server.
Duplicate query check before response	Enable/disable checks for repeated queries from the same source. If enabled, under non-flood conditions, the system checks and drops repeated UDP/TCP queries from the same source if it sends them at a rate greater than 3 per second. Under flood conditions, the duplicate query check is done for TCP and not for UDP.
Block identified sources	Enable/disable source IP address blocking periods for violators of any DNS-protection feature (ACL, anomalies, or flood meters). Disabled by default. DNS floods are often spoofed, so we do not recommend blocking an identified source to avoid punishing legitimate clients. Instead, we recommend you rely on the other DNS protection methods. They make packet-by-packet determinations, and are not prone to false positives. The configuration is open, allowing you to enable source blocking if you want to experiment with it in your network. If enabled, when a threshold is reached, packets are dropped and the identified sources are subject to the Blocking Period for Identified Sources configured on the Global Settings > Settings page. When disabled, DNS per-source thresholds (dns-query-per-source and dns-packet-track-per-source) are not tracked and the following graohs will not be updated: • Layer 7 > DNS > Query Per Source > DNS Query Per Source Egress Max Packet Rate/Sec • Layer 7 > DNS > Suspicious Sources > Packet Track Per Source Egress Max Packet Rate/Sec

Settings	Guidelines
Restrict DNS Queries to Specific Subnets	Enable/disable restriction to DNS queries from unwanted sources from the Internet. This feature allows service providers to protect their recursive or open DNS resolvers. In a typical deployment, the service provider will keep their open resolvers on the LAN or the protected side and their own customers will be on the Internet side. On the Internet side, there will also be the authoritative servers, and rogue DNS clients who send unwanted DNS query floods. To ensure that the DNS queries are only allowed from the service provider's customers, the service provider must add those subnets into the Restricted Subnets list. By restricting the DNS queries to specific subnets, the service provider can avoid responding to unwanted queries and thus protecting DNS infrastructure from getting overloaded.

To configure with the CLI, use a command sequence similar to the following:

```
config spp
   edit <spp_name>
      config ddos spp setting
         set dns-authentication-direction {inboud|outbound}
         set dns-flood-mitigation-mode-inbound {TC-equal-one|dns-
             retransmission}
         set dns-flood-mitigation-mode-outbound {TC-equal-one|dns-
             retransmission}
         set match-response-with-queries {enable|disable}
         set validate-ttl-for-queries-from-the-same-ip
             {enable|disable}
         set generate-response-from-cache-under-flood
             {enable|disable}
         set allow-only-valid-queries-under-flood {enable|disable}
         set source-blocking-in-dns {enable|disable}
         set duplicate-query-check {enable|disable}
         set force-tcp-or-forward-to-server {force-tcp|forward-to-
             server}
         set restrict-dns-queries-to-specific-subnets
             {enable/disable}
      end
end
```

Managing baseline traffic statistics

This section includes the following information:

- Baseline traffic statistics overview
- Generating baseline traffic statistics
- Displaying baseline traffic statistics

Baseline traffic statistics overview

The baseline traffic statistics are the maximum value (rate or count) measured by the counter during the observation period. The system saves data points every five minutes. During a 1-hour period, for example, there are 12, 5-minute observation periods. FortiDDoS saves a data point for each 5-minute interval. If you choose a 1-hour period, the system generates the maximum value across these 12 periods of 5-minute intervals.

The baseline statistics are used to establish the configured minimum threshold and ultimately the absolute maximum rate limit. Figure 41 illustrates the relationship between the baseline statistics, threshold settings, and monitor graphs.

In Figure 41:

1. The generated baseline statistic for the most-active-source threshold is 9774 packets/second.
2. The generated baseline statistic is multiplied by the Layer 3 percentage adjustment on the System Recommendation page. The default is 300%.
3. The product of the baseline and the percentage adjustment determines the configured minimum threshold. 9774x 300% = 29322 packets/second.
4. The configured minimum threshold is displayed on its monitor graph.
5. On the monitor graph, the estimated threshold is the top line. The estimated threshold can go no higher than the product of the configured minimum threshold and the adaptive limit. 29322 * 150% = 43983 packets/second.

Figure 41: Relationship baseline traffic statistics-thresholds

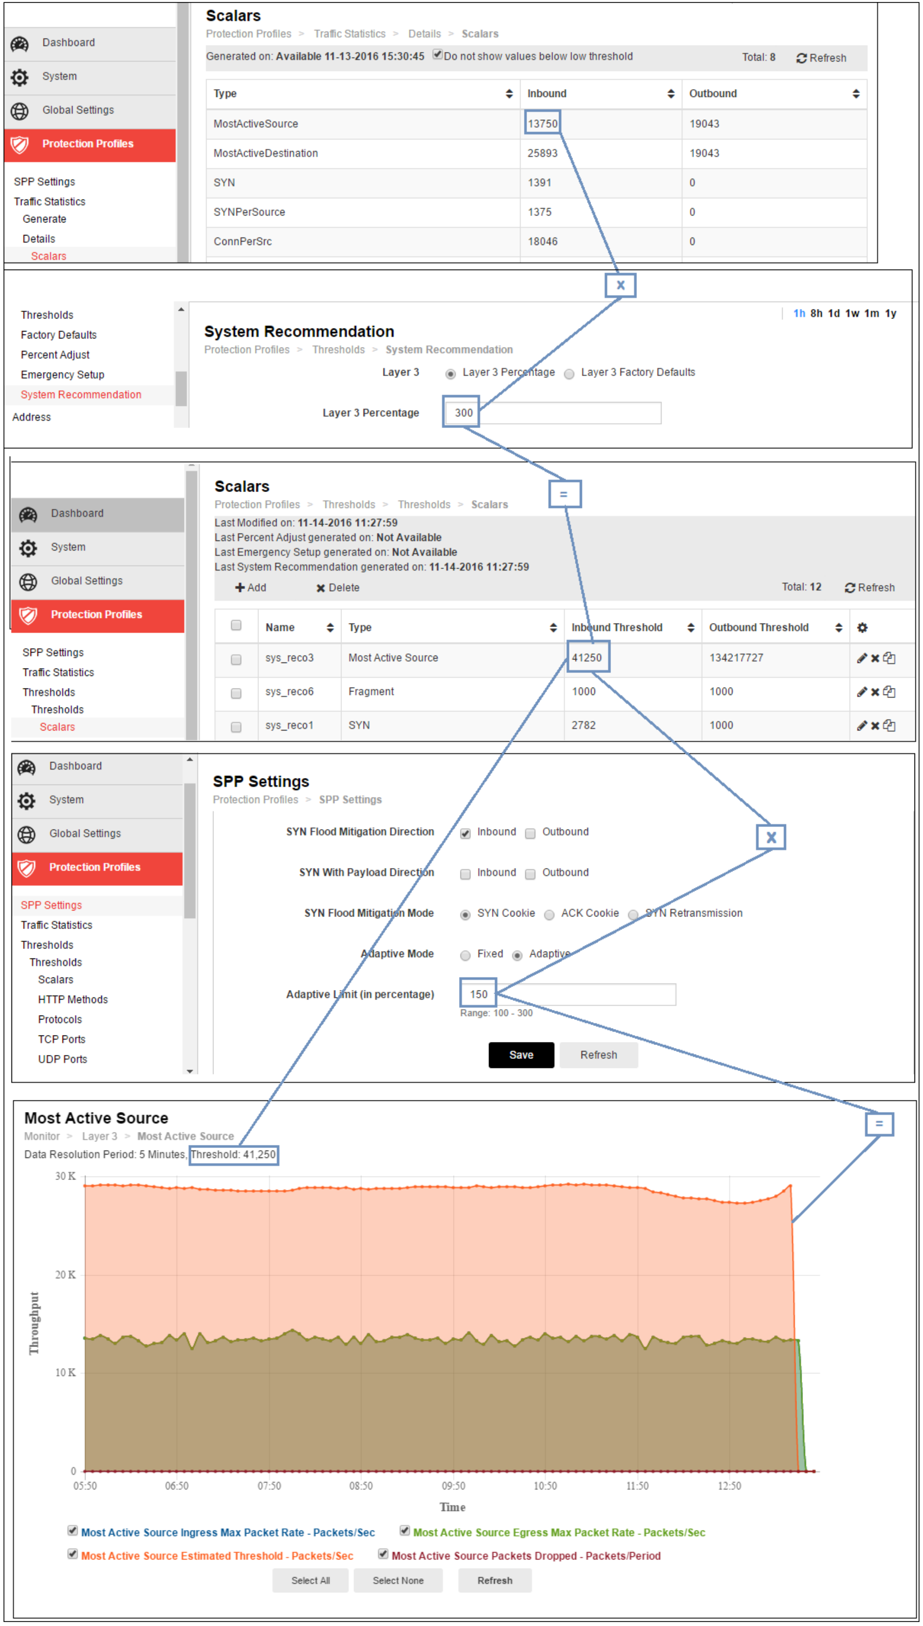

Scalars

Protection Profiles > Traffic Statistics > Details > Scalars

Generated on: **Available 11-13-2016 15:30:45** ☑Do not show values below low threshold Total: 8 Refresh

Type	Inbound	Outbound
MostActiveSource	13750	19043
MostActiveDestination	25893	19043
SYN	1391	0
SYNPerSource	1375	0
ConnPerSrc	18046	0

System Recommendation

Protection Profiles > Thresholds > System Recommendation

Layer 3 ● Layer 3 Percentage ○ Layer 3 Factory Defaults

Layer 3 Percentage 300

Scalars

Protection Profiles > Thresholds > Thresholds > Scalars

Last Modified on: **11-14-2016 11:27:59**
Last Percent Adjust generated on: **Not Available**
Last Emergency Setup generated on: **Not Available**
Last System Recommendation generated on: **11-14-2016 11:27:59**

Add Delete Total: **12** Refresh

□	Name	Type	Inbound Threshold	Outbound Threshold	
□	sys_reco3	Most Active Source	41250	134217727	
□	sys_reco6	Fragment	1000	1000	
□	sys_reco1	SYN	2782	1000	

SPP Settings

Protection Profiles > SPP Settings

SYN Flood Mitigation Direction ☑ Inbound □ Outbound

SYN With Payload Direction □ Inbound □ Outbound

SYN Flood Mitigation Mode ● SYN Cookie ○ ACK Cookie ○ SYN Retransmission

Adaptive Mode ○ Fixed ● Adaptive

Adaptive Limit (in percentage) 150
Range: 100 - 300

Save Refresh

Most Active Source

Monitor > Layer 3 > Most Active Source

Data Resolution Period: 5 Minutes, Threshold: 41,250

☑ Most Active Source Ingress Max Packet Rate - Packets/Sec ☑ Most Active Source Egress Max Packet Rate - Packets/Sec
☑ Most Active Source Estimated Threshold - Packets/Sec ☑ Most Active Source Packets Dropped - Packets/Period

Select All Select None Refresh

Generating baseline traffic statistics

You can generate baseline traffic statistics based on the following observation periods:

- Past 1 hour
- Past 8 hours
- Past 1 day
- Past 1 week
- Past 1 month
- Past 1 year

Use a time period that is representative of typical traffic volume and has had no attacks.

Before you begin:

- You must have Read-Write permission for Protection Profile settings.
- Note that the FortiDDoS hardware is accessed when you generate traffic statistics or set system recommended thresholds. Do not perform multiple operations simultaneously.

To generate baseline traffic statistics:

1. Go to Protection Profiles > Traffic Statistics> Generate.
2. Select the SPP you want to configure from the drop-down list.
3. Select the time period from the drop-down list.
4. Select **Generate**.
5. Save the configuration.
6. It takes about ten minutes for the process to complete. Click **Refresh** to track the status. The process is complete when the status shows "Available" and a timestamp.

To configure with the CLI, use a command sequence similar to the following:

```
config spp
   edit <spp_name>
      config ddos spp threshold-report
         set generate {enable | disable}
         set report-period {last-hour | last-8-hours | last-
             24-hours | last-week | last-month | last-year}
      end
```

Displaying baseline traffic statistics

You can review the statistics that are the basis of the system recommended thresholds.

Before you begin:

- You must have generated the report. See Generating baseline traffic statistics.
- You must have Read-Write permission for Protection Profile settings.

To display baseline traffic statistics

1. Go to Protection Profiles > Traffic Statistics > Details.
2. Select the SPP of interest from the drop-down list.
3. Select the type of statistics from the drop-down list.
4. Select the time period from the drop-down list.

Note: By default, the system does not display parameters with counts lower than the following.

Layer	Low threshold
3	100
4	500
7	200

Clear the **Do not show values below low threshold option** if you want to see these low counts.

Managing thresholds

This section includes the following information:

- Using system recommended thresholds
- Modifying threshold settings
- Adjusting minimum thresholds by percentage
- Configuring an emergency setup
- Restoring factory default threshold settings

Using system recommended thresholds

We recommend you use the system recommendation feature to set thresholds for most types of traffic. The system recommendation procedure sets the configured minimum threshold to a percentage of the generated baseline rates.

You use the Protection Profiles > Thresholds > System Recommendation page to set the multiplier for each OSI layer. The resulting configured minimum thresholds are populated on the Protection Profiles > Thresholds > Thresholds page. As you become a FortiDDoS expert, you can tune the thresholds on the Protection Profiles > Thresholds > Threshold page.

Table 40 explains how the system recommendation feature sets thresholds.

Table 39: How the system recommendation feature sets thresholds

Threshold Group	Notes
Scalar thresholds	• Thresholds are set to either the observed maximum multiplied by the Layer 3 or Layer 4 percentage, or to the low traffic threshold, whichever is higher. • The system recommended procedure sets the following L7 scalar meters to the system maximum rate (not traffic history times Layer 7 adjustment percentage): DNS Query per Source, DNS Packet Track per Source (Suspicious Sources on Monitor Graphs), DNS Question Count. • The system recommended procedure sets the following L3 and L4 scalar meters to the system maximum rate (not traffic history times Layer 3 or 4 adjustment percentage): Most Active Destination, New Connections.
Protocol thresholds	• The system recommendation procedure does not set the threshold for TCP protocol (6) and UDP protocol (17).

Threshold Group	Notes
TCP/UDP Port, ICMP Type/Code	• Packet rates vary across ports, SPPs, and traffic direction. • All contiguous TCP/UDP ports or ICMP type/codes that have the same inbound and outbound traffic rates are grouped into ranges. • We limit the number of ranges to 512 to optimize the internal configuration database. • The system recommendation procedure uses an algorithm to generate a set of ranges and packet rate thresholds for them. The algorithm is based on the following factors: • The recorded baseline traffic for ports or type/code from 0 to 64K. • If the traffic is below the low traffic value, the low traffic value is considered the baseline. • Otherwise, the recorded baseline rates are multiplied by the Layer 4 adjustment percentage. • The resulting rates are divided by 512 to determine a round-up factor. • Rates are rounded up to next multiple of round-up factor. • If the number of ranges is below 512, the thresholds are set. • Otherwise, the rates are rounded to the next multiple of round-up factor, and so on, until the number of ranges is below 512. Then, the thresholds are set. • The system recommendation procedure does not set the threshold for widely used TCP service ports 20-23, 25, 53, 80, 110, 139, 443 and 590; or TCP/UDP SIP ports 5060 and 5061. It does not set the threshold for user-configured HTTP service ports. The thresholds for these are set to high values. • For FortiDDoS models that support DNS features (all models except 600B and 900B), the system recommendation procedure does not set a threshold for UDP port 53 because there are more granular DNS counters to detect floods. For 600B and 900B, the procedure does set a threshold for port 53.
HTTP Method	• Thresholds are set to either the observed maximum multiplied by the Layer 7 percentage, or to the low traffic threshold, whichever is higher.
URL, Host, Cookie, Referer, User-Agent	• The rate meters for URLs and HTTP headers are based on indexes. • Packet rates vary across these indexes, SPPs, and traffic direction, depending on the time the baseline is taken. • The “observed maximum” used by the system recommendation procedure is the packet rate for the 95th percentile of observed rates for all indexes (excluding indexes with zero traffic), unless the number of indexes is unusually low. If low, the highest rate for all indexes is used. • Thresholds are set to either the observed maximum multiplied by the Layer 7 percentage, or to the low traffic threshold, whichever is higher.

Before you begin:

- You must have generated traffic statistics for a learning period. Ensure that the traffic statistics report that you generate for use with System Recommendation is for a period that is free of attacks and that it is long enough to be a representative period of activity. If necessary, reset statistics for the SPP before initiating the learning period.

- You must have Read-Write permission for Protection Profile settings.
- Note that the FortiDDoS hardware is accessed when you generate traffic statistics or set system recommended thresholds. Do not perform multiple operations simultaneously.

To generate the system recommended thresholds:

1. Go to Protection Profiles > Thresholds > System Recommendation.
2. Select the SPP you want to configure from the drop-down list.
3. Complete the configuration as described in Table 41.
4. Click **Save** to generate the system recommended thresholds.
5. Go to Protection Profiles > Thresholds > Thresholds and review the thresholds.

Table 40: Adjusting the system recommended thresholds

Settings	Guidelines
Layer <N> adjustment	• Percentage—Multiply the generated rates by the specified percentage to compute the recommended thresholds. • Factory default— Use factory default values instead of the recommended values. The factory default values are high so that the appliance can be placed inline and not immediately drop traffic.
Layer <N> percentage	Multiply the generated maximum rates by the specified percentage to compute the recommended thresholds. For example, if the value is 100%, the threshold is equal to the generated maximum rate. If it is 300%, the threshold is three times the generated maximum rate. The default adjustment for Layer 3 is 300. The default for Layer 4 is 200. The default for Layer 7 is 200. The valid range is 100 to 500.
Layer <N> low traffic threshold	Specify a minimum threshold to use instead of the recommended rate when the recommended rate is lower than this value. This setting is helpful when you think that the generated maximum rates are too low to be useful. The default is 1000. For example, assume the generated maximum packet rate for inbound Layer 4 TCP packets is 2,000 and the outgoing rate is 3,000. The value of **Layer 4 percentage** is 300 (percent) and the value of **Layer 4 low traffic threshold** is 8,000. In this example, the recommended threshold for inbound packets is 8,000 (2,000 * 300% = 6,000). However, because 6,000 is less than the low traffic threshold of 8,000, the system sets the threshold to 8,000.In this example, the recommended threshold for outbound packets is 9,000 (3,000 * 300% = 9,000). Because 9,000 is greater than the low traffic threshold of 8,000, the system sets the threshold to 9,000.

To configure with the CLI, use a command sequence similar to the following:

```
config spp
  edit <spp_name>
    config ddos spp threshold-adjust
      set threshold-adjustment-type system-recommendation
      set threshold-system-recommended-report-period
      {1-hour | 8-hours | 1-day | 1-week | 1-month | 1-
          year}
      set threshold-system-recommended-layer-3 {layer3-
          percentage | layer3-factory-defaults}
      set threshold-system-recommended-layer-3-percentage
          <percent>
      set threshold-system-recommended-layer-3-low-traffic
          <integer>
      set threshold-system-recommended-layer-4 {layer4-
          percentage | layer4-factory-defaults}
      set threshold-system-recommended-layer-4-percentage
          <percent>
      set threshold-system-recommended-layer-4-low-traffic
          <integer>
      set threshold-system-recommended-layer-7 {layer7-
          percentage | layer7-factory-defaults}
      set threshold-system-recommended-layer-7-percentage
          <percent>
      set threshold-system-recommended-layer-7-low-traffic
          <integer>
    end
```

Note the following:

To avoid generating too many high (>1023) TCP and UDP port ranges, the approach for System Recommendation has been changed in 4.3.0.

The System Recommendations now creates a single port range from 1024 through 65535 and assigns inbound and outbound thresholds which are calculated as the maximum packet rate of all these ports for that SPP.

Note that traffic originating from high ports and terminating on “service” ports (<1024) is always as associated with the service port in either direction. That ensures that the graphs and reports are showing HTTP (80) or SMTP (25) application traffic rather than randomly-selected high ports used by the application or firewall.

Other changes include the following:

- Outbound Most Active Source and Outbound UDP ports (<1024) are set to factory high thresholds. This was done to remove “noise” from outbound MAS and port traffic that is usually in Detection mode but is included in some aggregate graphs and reports which can be confusing to users.
- 5060 and 5061 TCP/UDP ports are now rate limited, based on the System Recommendations above

Prior to setting System Recommendations, Traffic Statistics should be checked for high traffic on ports above 1024. This could include:

- Alternate HTTP/s service ports like TCP 8080, 8081, 8000, 4443, 10443, etc.
- UDP port 4500, used by IPSEC NAT Traversal
- TCP/UDP ports 5060/5061 for SIP
- Blocks of high ports used for specific custom application purposes (like gaming)

After System Recommendations have been created, it may be necessary to manually configure high port ranges if:

- Specific high ports have very high traffic and you want to ensure the remaining ports use lower thresholds
- Specific high ports are in use but have significantly different data rates for the highest rate to lowest rate and compared to unused ports.

Modifying threshold settings

You use the Protection Profiles > Thresholds > Thresholds page to review system recommended thresholds and to make manual adjustments as you fine tune the configuration.

One of the key features of the FortiDDoS solution is the availability of system recommended thresholds that are adapted automatically according to statistical trends and tested heuristics. We recommend that in most cases, you should rely on the system intelligence. In some cases, such as demonstration, test, and troubleshooting situations, you might want to specify user-defined values for one or more thresholds. The threshold configuration is open, and can be updated manually.

Before you begin:

- You must have an expert understanding of packet rates and other Layer 3, Layer 4, and Layer 7 parameters that you want to set manually. Refer to Understanding FortiDDoS rate limiting thresholds .
- You must have Read-Write permission for Protection Profile settings.

To configure threshold settings:

1. Go to Protection Profiles > Thresholds > Thresholds.
2. Select the SPP you want to configure from the drop-down list.
3. Select the type of statistics from the drop-down list.
4. Double-click the row for the threshold you want to edit or click **Add** to create a new entry.
5. Set thresholds for inbound and outbound traffic for the settings described in Table 42.
6. Save the configuration.

Table 41: Threshold settings configuration

Settings	Guidelines	Graphs
Scalars		
syn	Packet/second rate of SYN packets received. Threshold for a SYN Flood event. When total SYNs to the SPP exceeds the threshold, the SYN flood mitigation mode tests are applied to all new connection requests from IP addresses that are not already in the legitimate IP address table.	Layer 4

Settings	Guidelines	Graphs
new-connections	Connection/second rate of new connections. Threshold for zombie floods (when attackers hijack legitimate IP addresses to launch DDoS attacks). When it detects a zombie flood, FortiDDoS blocks all new connection requests for the configured blocking period. In order to be effective, the **new-connections** threshold should always be higher than the **syn** threshold. We recommend that you use the FortiDDoS generated threshold unless you have a specific reason to change it.	Layer 4
syn-per-src	Packet/second rate of SYN packets from any one source. No single source in an SPP is allowed to exceed this threshold. Threshold for a SYN Flood From Source event.The system applies the blocking period for identified sources.	Layer 4
most-active-source	Packet/second rate for the most active source. A source that sends packets at a rate that surpasses this threshold is considered a threat. Threshold for a source flood. No single source in an SPP is allowed to exceed this threshold, and the system applies the blocking period for identified sources.	Layer 3
concurrent-connections-per-source	Count of TCP connections from a single source. The TCP connection counter is incremented when a connection moves to the established state and decremented when a sessions is timed out or closes. This threshold is used to identify suspicious source IP behavior. An inordinate number of connections is a symptom of both slow and fast TCP connection attacks. The system applies the blocking period for identified sources. If the aggressive aging **high-concurrent-connection-per-source** option is enabled, the system also sends a TCP RST to the server to reset the connection.	Layer 4
syn-per-dst	Packet/second rate for SYN packets to a single destination. When the per-destination limits are exceeded for a particular destination, the SYN flood mitigation mode tests are applied to all new connection requests to that particular destination. Traffic to other destinations is not subject to the tests.The system applies the blocking period for identified sources.	Layer 4
method-per-source	Packet/second rate for Method packets (GET, HEAD, OPTION, POST, etc) from a single Source. When the per-source limits are exceeded for a particular source, the system applies the blocking period for identified sources. The connection to the server may also be RST if Protection Profiles > SPP Settings > TCP Tab: Aggressive Aging TCP Connections Feature Control: Layer 7 Flood is enabled.	Layer 7
most-active-destination	Packet/second rate for the most active destination. A destination that is sent packets at this rate is considered under attack. Threshold for a destination flood.	Layer 3
fragment	Packet/second rate of fragmented packets received. Although the IP specification allows IP fragmentation, excessive fragmented packets can cause some systems to hang or crash.	Layer 3

Settings	Guidelines	Graphs
dns-query	Queries/second. Threshold for a DNS Query Flood event.	Layer 7
dns-question-count	Question count/second. Threshold for a DNS Question Flood event.	Layer 7
dns-mx-count	Packet/second rate of DNS queries for MX records (QTYPE=15). Threshold for a DNS MX Flood event.	Layer 7
dns-all	Packet/second rate of DNS queries for all DNS records (QTYPE=255). Threshold for a DNS ALL Flood event.	Layer 7
dns-zone-xfer	Packet/second rate of DNS zone transfer (AXFR) queries (QTYPE=252). Threshold for a DNS Zone Transfer Flood event.	Layer 7
dns-fragment	Packet/second rate of fragmented packets received. Threshold for a DNS Fragment Flood event.	Layer 7
dns-query-per-src	Packet/second rate of normal DNS queries from any one source. No single source in an SPP is allowed to exceed this threshold. Threshold for a DNS Query Per Source flood event. The system applies the blocking period for identified sources.	Layer 7
dns-packet-track-per-src	Packet/second rate of a source that demonstrates suspicious activity (a score based on heuristics that count fragmented packets, response not found in DQRM, or queries that generate responses with RCODE other than 0). Threshold for a DNS Suspicious Sources flood event. The system applies the blocking period for identified sources.	Layer 7
HTTP Methods		
HTTP/1.1 uses the following set of common methods: • GET • HEAD • OPTIONS • TRACE • POST • PUT • DELETE • CONNECT	Packet/second rate for the specified HTTP method. Threshold for an HTTP method flood attack. When the maximum rate is reached, the system drops packets matching the parameter. If the aggressive aging **layer7-flood** option is enabled, the system also sends a TCP RST to the server to reset the connection.	Layer 7
Protocols		

Settings	Guidelines	Graphs
Protocol Start / End	Packet/second rate for the specified protocol. Threshold for a Protocol Flood event. When you specify a threshold for protocols, enter a range, even if you are specifying a threshold for a single protocol. For example, to set a threshold for protocol 6, enter 6 for both Protocol Start and Protocol End.	Layer 3
TCP Ports		
Port Start / End	Packet/second rate for the specified TCP port. Threshold for a Port Flood event. Monitoring the packet rate for ports is helpful to prevent floods against a specific application such as HTML, FTP, SMTP or SQL. TCP accommodates 64K (65,536) ports, most of which may never be used by a particular server. Conversely, a server might see most or all of its traffic on a small group of TCP ports. For this reason, globally assigning a single threshold to all ports generally does not provide useful protection. However, you can globally set a (usually low) TCP Port Threshold for all TCP ports and then manually configure a higher threshold for the ports your protected network is using. When you specify a threshold for ports, you enter a range, even if you are specifying a threshold for a single port. For example, to set a threshold for port 8080, enter 8080 for both Port Start and Port End.	Layer 4
UDP Ports		
Port Start / End	Packet/second rate for the specified UDP port.Threshold for a Port Flood event. When you specify a threshold for ports, you enter a range, even if you are specifying a threshold for a single port. For example, to set a threshold for port 53, enter 53 for both Port Start and Port End.	Layer 4
ICMP Types/Codes		
ICMP Type/Code Start/End	Packet/second rate for the specified ICMP type/code range. The ICMP header includes an 8-bit type field, followed by an 8-bit code field. Threshold for an ICMP Type/Code Flood event. A popular use for ICMP is the “Echo groping” message (type 8) and its corresponding reply (type 0), which are often useful tools to test connectivity and response time. In some cases, this message and reply can also be used as an attack weapon to effectively disable a target system’s network software. Take care when you set the ICMP type 0 and type 8 thresholds to ensure the desired functionality is preserved.	Layer 4
HTTP		

Settings	Guidelines	Graphs
URL	Packet/second rate for packets with the specified URL match. When the maximum rate is reached, the system drops packets matching the parameter. If the aggressive aging **layer7-flood** option is enabled, the system also sends a TCP RST to the server to reset the connection. Specify the URL for a specific website. Botnets make it easy to launch attacks on specific URLs. When such an attack happens, FortiDDoS can isolate the URL and limit just the traffic that is associated with it, while all other traffic is unaffected. The URL is found in the website's HTTP GET or POST operations. For example, the URL for http://www.website.com/index.html is **/index.html**. When you specify a threshold for a URL, the system generates a corresponding hash index value. FortiDDoS displays the hash index value in the list of URL thresholds. Make note of it. You use the hash value to select this URL elsewhere in the web UI. To view statistics associated with the threshold, go to Monitor > Specific Graphs > URLs, and then, for **Please enter URL/Hash index**, enter either the original URL you specified or the hash index value. The valid range of hash index values for URLs is 0-32k per SPP.	Layer 7
	You can use the special prefix `sys_reco_v` to create hash index ranges that aggregate URLs that you are interested in only as an aggregate. For example, assume your team wants to pay close attention to a five websites, and all others can be treated essentially the same. With the first five, your configuration is specific, so you know the website URL and the corresponding hash index, and you can use FortiDDoS to track it specifically. The system does not track the others with specificity, but you can track, as an aggregate, whether those sites experience rising and falling rates, including attacks. 1. Create entries for the five priority websites and note their hash index numbers. Let's assume the hash index numbers are 1, 20, 21, 39, 40. 2. Create ranges to aggregate the gaps: a. The first gap is from 2-19, so you create a configuration named `sys_reco_v2_19`. This includes hash numbers 2 through 19. b. The second gap is from 22-38, so you create a configuration named `sys_reco_v22_38`. c. The next gap is from 41 to the end of the range, so you create a configuration named `sys_reco_v41_8192`. **Note**: You cannot carve out a small block out of a large block. If you want to use hash index values that are already in use, you must delete the existing range and then create two ranges.	

Settings	Guidelines	Graphs
Host, Referer, Cookie, User-Agent headers	Packet/second rate for packets with the specified header matches. When the maximum rate is reached, the system drops packets matching the parameter. If the aggressive aging **layer7-flood** option is enabled, the system also sends a TCP RST to the server to reset idle connections. A connection is deemed idle if it has not sent traffic in the last 2 minutes. Specify HTTP header values. With the advent of botnets, it is easy to launch attacks using scripts. Most of the scripts use the same code. The chances that they all use the same Host, Referer, Cookie, or User-Agent header fields is very high. When such an attack happens, FortiDDoS can easily isolate the four headers among many and limit traffic associated with that specific header, while all other traffic is unaffected. As with URL hash indexes, you can use the `sys_reco_v` prefix to define hash index ranges that aggregate header values you are not specifically interested in. The valid range of hash index values is 0-511 for each setting for each SPP: Host, Referer, Cookie, User-Agent.	Layer 7

To configure with the CLI, use a command sequence similar to the following:

```
config spp
   edit <spp_name>
      config ddos spp scalar-threshold
         edit <threshold_name>
            set type {syn |syn-per-src | most-active-source |
                concurrent-connections-per-source | most-
                active-destination | method-per-source |
                fragment | new-connections | syn-per-dst |
                dns-query | dns-question-count | dns-mx-count
                | dns-all | dns-zone-xfer | dns-fragment |
                dns-query-per-src | dns-packet-track-per-src}
            set inbound-threshold <integer>
            set outbound-threshold <integer>
      end
      config ddos spp protocol-threshold
         edit <threshold_name>
            set protocol-start <protocol_int>
            set protocol-end <protocol_int>
            set inbound-threshold <integer>
            set outbound-threshold <integer>
      end
```

Adding TCP or UDP Port Ranges

In 4.3.0, after the System Recommendations are created, there will only be one range for TCP and UDP “high” (>1023) ports labeled as “sys_reco_v1024_65535”.

If you use specific and/or want to exclude specific high ports, you must enter these manually. You cannot have overlapping port ranges. To add a port or range, first delete the existing range.

For example, if you want to allow Port 4500 for high traffic and leave all others as default:

1. Delete the port range “sys_reco_v1024_65535”.
2. Add port '4500':
   - Name: IPSEC
   - Port Start: 4500
   - Port End: 4500
   - Inbound Threshold: as required to system max of 16,777,215
   - Outbound Threshold: as required to system max of 16,777,215
3. Replace deleted range with two ranges:
   - Add Range
   - Name: Default1024_4499
   - Port Start: 1024
   - Port End: 4499
   - Inbound Threshold: 500
   - Outbound Threshold: 500
   - Add Range
   - Name: DefaultAbove4500
   - Port Start: 4501
   - Port End: 65535
   - Inbound Threshold: 500
   - Outbound Threshold: 500

Note the following:

- Name labels can be alphanumeric plus “-“ and “_” only, 35 characters maximum.
- It is not necessary to follow the system label syntax of “sys_reco_vXXX_YYYYY” for ports or protocols. You must follow this for all other thresholds.
- It is not necessary to place the ranges in numerical order. You can use the Port Start column to sort as required.

Adjusting minimum thresholds by percentage

You can arbitrarily adjust SPP thresholds by percentage. This is useful when you expect a spike in legitimate traffic (for example, because of a news story or an advertising campaign). You can adjust the thresholds by as much as 300%.

Before you begin:

- Go to Protection Profiles > Thresholds > Thresholds and note the settings so that you can later verify the adjustment procedure or subsequently reset the thresholds to the values before the adjustment procedure.
- You must have Read-Write permission for Protection Profile settings.

To adjust minimum thresholds by percentage:

1. Go to Protection Profiles > Thresholds > Percent Adjust.
2. Select the SPP of interest from the drop-down list.
3. Specify a percentage in the text box. For example, to increase the threshold by 20 percent, enter `20`. To decrease it by 20 percent, enter `-20`.
4. Save the configuration.
5. Go to Protection Profiles > Thresholds > Thresholds and verify that the adjustment has been applied.

To configure with the CLI, use a command sequence similar to the following:

```
config spp
   edit <spp_name>
        config ddos spp threshold-adjust
           set threshold-adjustment-type percent-adjust
           set threshold-percent-adjust <percent_int>
        end
```

Configuring an emergency setup

You can use the emergency setup option to adjust only certain key thresholds based on empirical knowledge. You can expect these adjustments to protect against common attacks. For example, if you are already under attack, you can use emergency setup to deploy the unit without an initial learning period.

Before you begin:

- You must have Read-Write permission for Protection Profile settings.

To configure an emergency setup:

1. Go to Protection Profiles > Thresholds > Emergency Setup.
2. Select the SPP you want to configure from the drop-down list.
3. Adjust the defaults listed in Table 43 according to your empirical knowledge.
4. Save the configuration.

Table 42: Emergency setup configuration

Settings	Default
Inbound SYN Threshold	1000
Outbound SYN Threshold	1000
Inbound SYN/Source Threshold	1000
Outbound SYN/Source Threshold	1000
Inbound Most Active Source Threshold	10,000
Outbound Most Active Source Threshold	134217727
Inbound Concurrent Connections per Source Threshold	1000
Outbound Concurrent Connections per Source Threshold	1000

To configure with the CLI, use a command sequence similar to the following:

```
edit <spp_name>
   config ddos spp threshold-adjust
      set threshold-adjustment-type easy-setup
      set threshold-easy-setup-inbound-syn-threshold
          <integer>
      set threshold-easy-setup-outbound-syn-threshold
          <integer>
      set threshold-easy-setup-inbound-syn-per-source-
          threshold <integer>
      set threshold-easy-setup-outbound-syn-per-source-
          threshold <integer>
      set threshold-easy-setup-inbound-most-active-source-
          threshold <integer>
      set threshold-easy-setup-outbound-most-active-source-
          threshold <integer>
      set threshold-easy-setup-inbound- concurrent-
          connections-per-source-threshold <integer>
      set threshold-easy-setup-outbound- concurrent-
          connections-per-source-threshold <integer>
      set threshold-easy-setup-inbound- concurrent-invite-
          per-source-threshold <integer>
      set threshold-easy-setup-outbound- concurrent-invite-
          per-source-threshold <integer>
   end
```

Restoring factory default threshold settings

In some situations, you might want to reset thresholds for an SPP. For example:

- You want to ensure that the application does not drop any packets due to rate thresholds. (The factory default values are high so that the appliance can be placed inline and not immediately drop traffic.)
- You are conducting a demonstration or test, or you are troubleshooting an issue.

Table 44 summarizes “factory reset” options.

Table 43: “Factory reset” options

Task	Menu
Reset the threshold configuration for an SPP.	See below.
Reset the threshold configuration and clear traffic history for an SPP.	Protection Profiles > Factory Reset > Factory Reset
Reset the system to its factory state. All SPPs, statistics, and logs will be deleted.	See Resetting the system.

Before you begin:

- You must have Read-Write permission for Protection Profile settings.

To reset SPP threshold settings:

1. Go to Protection Profiles > Thresholds > Factory Defaults.
2. Select the SPP you want to configure from the drop-down list.
3. Select **Set to Factory Defaults**.
4. Save the configuration.

To configure with the CLI, use a command sequence similar to the following:

```
config spp
   edit <spp_name>
      config ddos spp threshold-adjust
         set threshold-adjustment-type factory-defaults
         set threshold-factory-defaults {enable | disable}
      end
```

Using SPP ACL policies

This section includes the following information:

- Configuring SPP ACL address objects
- Configuring SPP ACL service objects
- Configuring an SPP ACL policy

Configuring SPP ACL address objects

You create SPP ACL address configuration objects to identify IP addresses and subnets that you want to match in SPP ACL policies.

Before you begin:

- You must have Read-Write permission for Protection Profile settings.

To configure the addresses:

1. Go to Protection Profiles > Address > [Address Config | Address Config IPv6].
2. Select the SPP you want to configure from the drop-down list.
3. Click **Add** to display the configuration editor.
4. Complete the configuration as described in Table 44.
5. Save the configuration.

Table 44: Address configuration

Settings	Guidelines
Name	Configuration name. Must not contain spaces.
Address	Specify an IP address.

To configure with the CLI, use a command sequence similar to the following:

```
config spp
   edit <spp_name>
      config ddos spp {address | address6} <address_name>
          spp-source-ip-address {<address_ipv4> | <address_
          ipv6}
         edit <name>
            ...
      end
```

Configuring SPP ACL service objects

You configure service objects identify the services that you want to match in SPP ACL policies.

Before you begin:

- You must have Read-Write permission for Protection Profile settings.

To configure service objects:

1. Go to Protection Profiles > Service > Service Config.
2. Select the SPP you want to configure from the drop-down list.
3. Click **Add** to display the configuration editor.
4. Complete the configuration as described in Table 45.
5. Save the configuration.

Table 45: Service object configuration

Settings	Guidelines
Fragment	
Fragment	No parameters. If you configure an ACL rule to match the Fragment service object, you are creating a rule to deny fragmented packets. Some Internet technologies, such as multimedia streaming, rely on fragmentation. Ensure that you understand your network and its packet behavior before you use the ACLs for fragmented packets.
Protocol	
Protocol Start / End	When you configure a service object for protocols, you enter a range, even if you are specifying a single protocol. For example, to configure a service object for protocol 6, enter 6 for both Protocol Start and Protocol End. Networks use of some of the protocols, such as 1 (ICMP), TCP (6), and UDP (17), ubiquitously. Ensure that you understand your network and its packet behavior before you use the ACLs for protocols.
TCP Port	
Port Start / End	When you configure a service object for ports, you enter a range, even if you are specifying a single port. For example, to configure a service object for port 8080, enter 8080 for both Port Start and Port End.
UDP Port	
Port Start / End	When you configure a service object for ports, you enter a range, even if you are specifying a single port. For example, to configure a service object for port 53, enter 53 for both Port Start and Port End.
ICMP Types/Code	

Settings	Guidelines
ICMP Type/Code Start / End	The header of Internet Control Message Protocol packets include an 8-bit type field, followed by an 8-bit code field. The value of this field can be read as a hexadecimal number.
URL, Host, Referer, Cookie, User Agent	
HTTP-Param	A matching value for the selected URL or HTTP header. When you create a service that specifies a URL to deny, enter the text that follows the protocol and the web address. For example, if you enter `http://www.website.com/index.html` in a browser to access a specific URL, enter `/index.html`. Because the number of possible URLs is infinite, FortiDDoS stores these values in a hash table. Up to 32,767 such hash indexes are allowed. If there are duplicate hash-indexes, the most recent URL that corresponds to a hash index overwrites any previous URLs in the URL field. However, all the URLs affect the threshold and maximum packet rate calculations and all URLs that hash to the same index are denied if the hash index is blocked. Similarly, if there is an attack that corresponds to a hash index, all URLs that hash to the same location are dropped. You can deny traffic by specifying the following HTTP header field types: Host, Referer, Cookie, and User-Agent. This is useful when a specific hash-index is under attack. FortiDDoS allows the source to establish the TCP connection with the server. However, when FortiDDoS detects the specified hash-index, it denies the packet and sends an RST packet to the server to aggressively age the connection. The appliance treats all subsequent packets from the source on that TCP connection as foreign packets and blocks the source for the specified blocking period.
DNS	
DNS-All	No parameters. If you configure an ACL rule to match the DNS-All service object, you are creating a rule to deny DNS queries for all DNS records (QTYPE=255). This DNS QTYPE is a query for all resource records. Some references, such as Wireshark, call this an "any" query. Most users should not be permitted to do all/any queries.
DNS-Fragment	No parameters. If you configure an ACL rule to match the DNS-Fragment service object, you are creating a rule to deny fragmented packets.
DNS-MX	No parameters. If you configure an ACL rule to match the DNS-MX service object, you are creating a rule to deny DNS queries for MX records.

Settings	Guidelines
DNS-Zone-Transfer	No parameters. If you configure an ACL rule to match the DNS-Zone-Transfer service object, you are creating a rule to deny DNS zone transfer (AXFR) queries (QTYPE=252). This DNS QTYPE is a query that initiates transfer of an entire zone file from the master name server to secondary name servers. Most users should not be permitted to do zone transfers.

To configure with the CLI, use a command sequence similar to the following:

```
config spp
   edit <spp_name>
      config ddos spp service
        edit <service_name>
           set type {cookie | dns-all | dns-fragment | dns-
               mx | dns-zone-xfer | fragment | host | icmp-
               type-code | protocol | referer | tcp-port |
               udp-port | url | user-agent}
           [set dns-rcode-start <int_start>]
           [set dns-rcode-end <int_end>]
           [set protocol-start <int_start>]
           [set protocol-end <int_end>]
           [set tcp-port-start <int_start>]
           [set tcp-port-end <int_end>]
           [set udp-port-start <int_start>]
           [set udp-port-end <int_end>]
           [set icmp-type <integer>[
           [set icmp-code <integer>]
           [set http-param <http_para_str>]
      end
```

Configuring an SPP ACL policy

An SPP ACL policy establishes allow/deny rules for traffic that matches the following data:

- IP Address
- Fragment
- Protocol
- TCP Port
- UDP Port
- ICMP Type/Code
- URL
- HTTP header field: Host, Referer, Cookie, User Agent
- DNS-All
- DNS-Fragment

- DNS-MX
- DNS-Zone-Transfer

ACL rules match a single data point, not multiple conditions. Rules are evaluated from the top of the table to the bottom. If a rule matches, it is applied and subsequent rules are not consulted. In most cases, you should order deny rules before allow rules.

Information about packets denied by an SPP ACL policy is reported in the following graphs and reports:

- Graphs (Monitor > ACL Drops)
- Executive Summary dashboard (Log & Report > Executive Summary)
- Reports (Log & Report > Report Configuration)

Before you begin:

- You must have configured address objects and service objects that you want to match in policy rules. See Configuring SPP ACL address objects and Configuring SPP ACL service objects.
- You must have Read-Write permission for Protection Profile settings.

To configure an ACL policy:

1. Go to Protection Profiles > Access Control List > Access Control List.
2. Select the SPP you want to configure from the drop-down list.
3. Click **Add** to display the configuration editor.
4. Complete the configuration as described in Table 46.
5. Save the configuration.

Table 46: Access control list configuration

Settings	Guidelines
Name	Configuration name. Must not contain spaces.
Type	• Address • Address IPv6 • Service
Address / Address IPv6	
Source Address	Select an address configuration object.
Address Action	• Deny—Drop traffic that matches the address object. • Track and Allow—Allow the traffic and include it in the statistics for continuous learning and threshold estimation. • Restrict DNS Queries to Specific Subnets—Restricts DNS queries from unwanted sources from the Internet. By restricting the DNS queries to specific subnets, the ISP can avoid responding to unwanted queries and thus protecting its DNS infrastructure from getting overloaded.
Service	

Settings	Guidelines
Direction	• inbound • outbound **Tip**: Shift-click to select multiple items.
Service	Select a service configuration object.
Service Action	• Deny—Drop traffic that matches the service object.

To configure with the CLI, use a command sequence similar to the following:

```
config spp
   edit <spp_name>
      config ddos spp acl
        edit <No.>
           set type {v4address | v6address | service}
           set direction {outbound inbound}
           set source-address <address_name>
           set v6address <address_name>
           set address-action {deny | track-and-allow |
               restrict DNS to specific subnets}
           set service <service_name>
           set service-action deny
      end
```

Performing a factory reset of SPP settings

In some situations, you might want to both reset thresholds for a particular SPP and clear its traffic history. For example:

- You have changed the SPP policy network address configuration.
- You have deployed a new server or service in the SPP network. Therefore, you want the appliance to recalculate the baseline.
- You are conducting a demonstration or test, or you are troubleshooting an issue.

Table 47 summarizes “factory reset” options.

Table 47: “Factory reset” options

Task	Menu
Reset the threshold configuration for an SPP.	Protection Profiles > Thresholds > Factory Defaults
Reset the threshold configuration and clear traffic history for an SPP.	See below.
Reset the system to its factory state. All SPPs, statistics, and logs will be deleted.	See Resetting the system.

Before you begin:

- You must have Read-Write permission for Protection Profile settings.

To perform a factory reset of SPP threshold settings and traffic history:

1. Go to Protection Profiles > Factory Reset > Factory Reset.
2. Select the SPP you want to reset from the drop-down list.
3. Select the reset check box.
4. Save the configuration.

To configure with the CLI, use a command sequence similar to the following:

```
config spp
   edit <entry_index>
      config ddos spp factoryreset
         set reset enable
      end
```

FAQ: SPP Settings

Service Ports

This section discusses some of the questions that users often have about service port ACLs and port rate limit thresholds.

How does FortiDDoS identify UDP services?

Ports 0-1023 are assigned by IANA to well known services. For example, UDP port 53 is assigned to DNS. When you configure ACL or threshold rules for well known UDP services in the 0-1023 range, configure rules for the IANA-assigned port. You do not configure rules for the associated, unassigned ports used by the client (these are numbered above 1023). For example, for DNS, configure an inbound rule for port 53 and outbound rule for port 53.

The user interface label for the ACL service setting shows "Destination Port." This is misleading. Beginning with release 4.1.6, the UDP service is identified when either the source or destination port is the well known port. The inbound and outbound traffic shown in Figure 42, for example, is identified as port 53 (DNS) traffic.

For an ACL deny rule, UDP service identification means the packets are denied if either the source port or the destination port matches the well known port. If you use an ACL policy to deny port 53, you are denying all DNS service traffic in the direction specified in your rule. If you want to deny inbound DNS service to an SPP, but the SPP has internal clients making outbound DNS queries to resolve addresses, we recommend that you not use the ACL (which would result in inbound DNS response traffic being dropped).

Instead, we recommend that you use rate limiting thresholds that allow inbound responses to outbound queries but at low rates to prevent DNS floods. You can set a low inbound threshold for DNS (UDP port 53) rather than deny inbound DNS service. The system recommended thresholds will set limits consistent with your baseline traffic. If you set user-defined thresholds for UDP ports, keep in mind these guidelines on how FortiDDoS tracks UDP service traffic. The inbound and outbound packet counters are incremented when traffic for the service is identified (by either source or destination port). Think of it as a service rate limit rather than a port rate limit.

The most active source outbound threshold and UDP port 0 -1023 outbound threshold is set to a higher value as per system recommendation.

Figure 42: UDP service ports for DNS

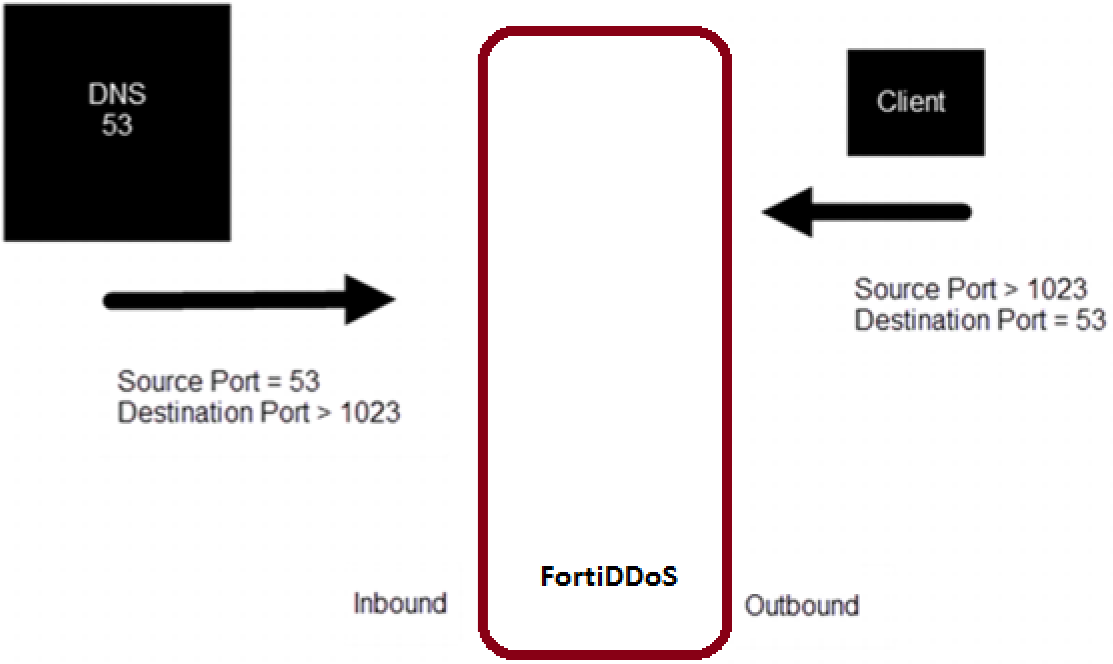

Chapter 5: Monitor Graphs

This chapter includes the following topics:

Monitor graphs overview

You use the Monitor graphs to track trends in throughput rates, source and destination traffic, connections, and drops related to FortiDDoS detection and prevention settings.

The Monitor graphs menu includes the following categories:

- Port statistics
- SPP statistics
- Aggregate drops
- Flood drops
- ACL drops
- Anomaly drops
- Hash attack drops
- Out of Memory drops
- Layer 3 threshold rate meters and counters
- Layer 4 threshold rate meters and counters
- Layer 7 threshold rate meters and counters

Each category includes one or more graphs, and each graph plots multiple queries. In addition, graphs can be queried by SPP, time, and traffic direction (when relevant).

The multiple views and granular filters are useful for comparing and contrasting trends broadly, and for drilling into details. For example, you can use the Aggregate drops graph to get an overall picture on security events and see whether to review ACL graphs, flood graphs, or anomalies graphs next.

Figure 43 is an example of a monitor graph. It shows the following information for the selected SPP, parameter, period, and direction:

- Threshold—The configured minimum threshold (matches the setting on the Protection Profiles > Thresholds > Thresholds page).
- Throughput—A graph of the throughput rate for the selected protocol during the time period.
- Packets dropped—A graph of packets dropped because the threshold was exceeded.
- Packets blocked—A graph of packets blocked due to the application of blocking periods.
- Data resolution—Whether data points for the graph are rolled up in 5 minute, 1 hour, 3 hour, or 45 hour windows.

Figure 43: Graph of inbound TCP traffic

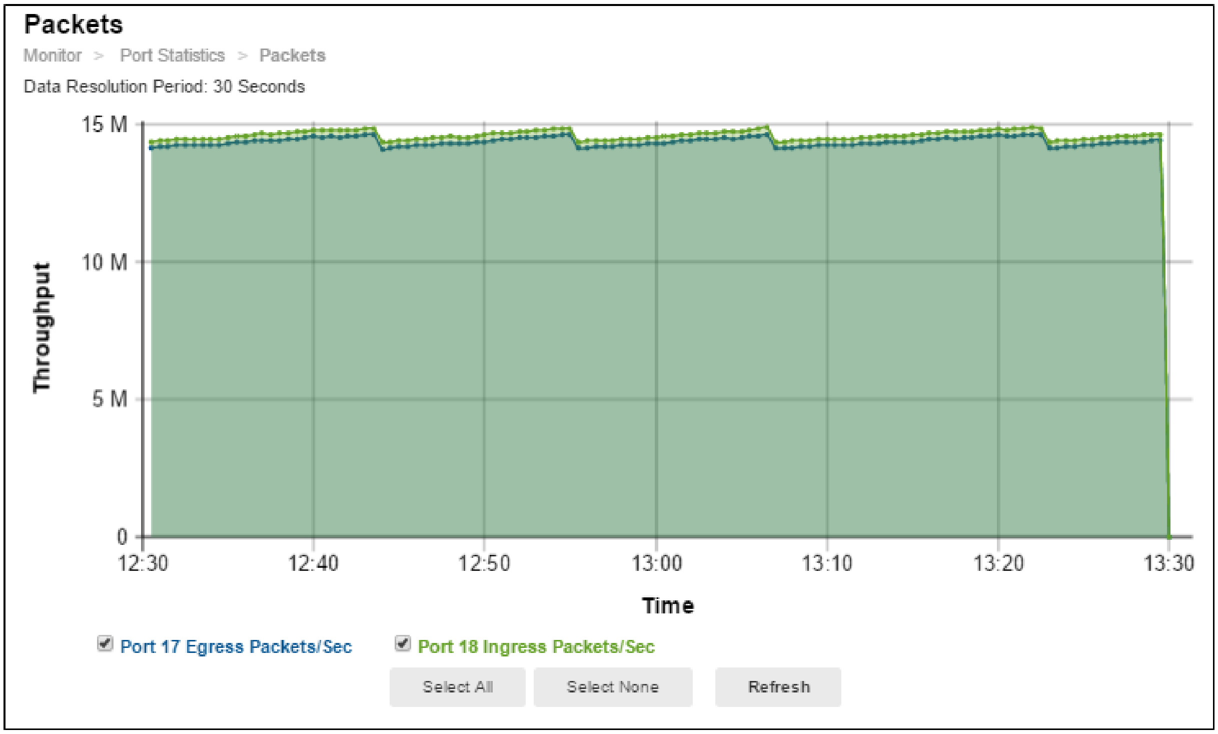

Packets

Monitor > Port Statistics > Packets

Data Resolution Period: 30 Seconds

Drilling-down

In aggregate graphs, the legend includes a graph icon 📈 that is a link to a graph of the item of interest. For example, shows the Aggregate Drops graph. The graph indicates that there are Flood drops, so you can click the graph icon next to Flood Drops in the legend to display the Flood Drops graphs.

Figure 44: Aggregate Drops Graph

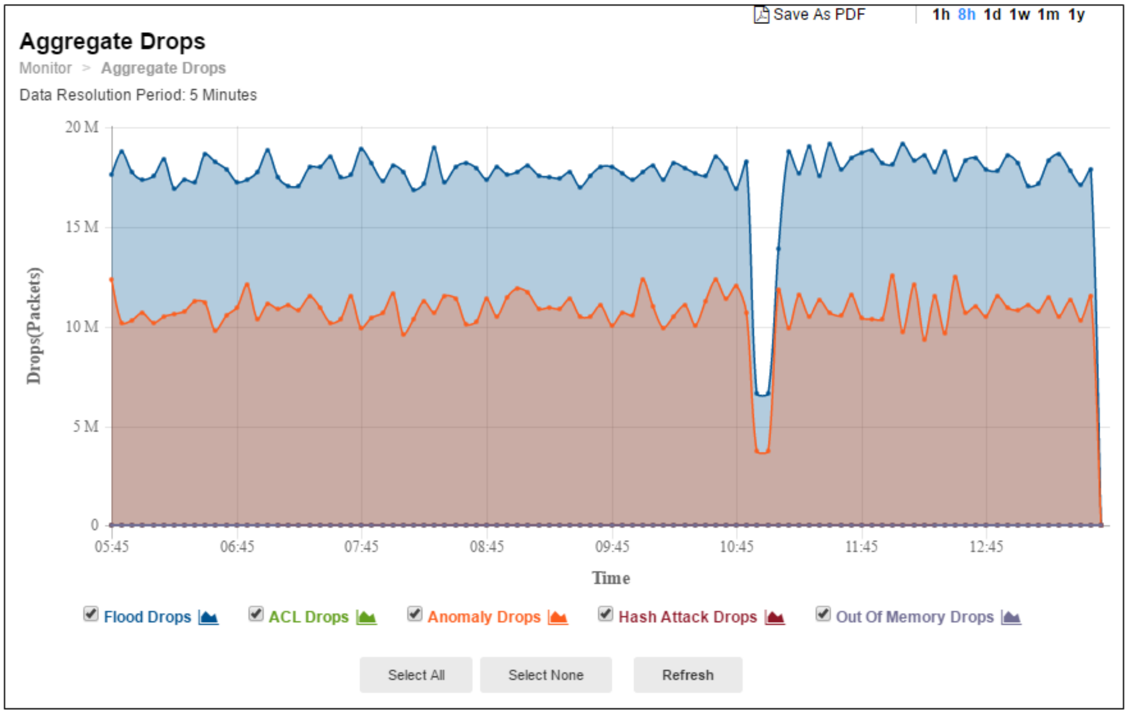

Data point details

Figure 45 shows an infotip that is displayed when the mouse pointer hovers over a point in the graph. The infotip has details about that data point.

Figure 45: Tooltip information for point on graph line

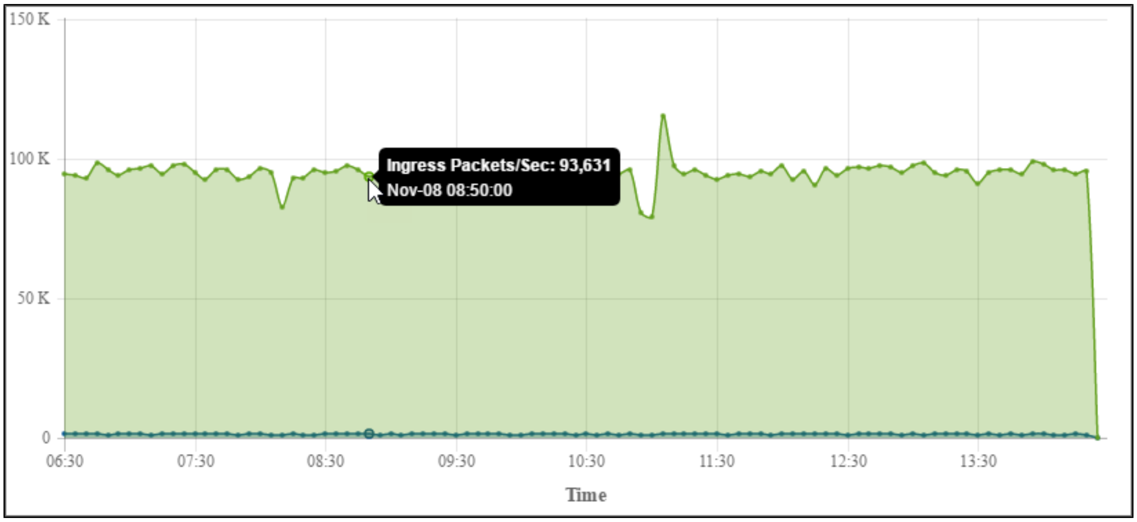

Using the Port Statistics graphs

You use the Port Statistics graphs to monitor network interface throughput. FortiDDoS ports are configured as network interface pairs. You configure odd-numbered ports for the LAN-side connection and even-numbered ports for the WAN-side connection.

The Port Statistics set includes the following graphs:

- Packets—Throughput in packets per second.
- Bits—Throughput in bits per second.

Figure 46 shows the Packets graph.

You can customize the following query terms: interface pair, period, and direction. The Port Statistics graphs displays the 30-second average data rate of Ingress and Egress packets. As the time period lengthens, the maximum value of the averages is used for the next period. All other graphs under Monitor displays the peak value during the time period (varying from 5 minutes - 45 hours). Therefore, these can show higher rates than the Port Statistics.

Note: Port statistics are not maintained per SPP. You cannot query by SPP, and port statistics are not reset when you reset SPP statistics. Port statistics are global data that gets reset only when you perform a complete factory reset and reformat the log disk.

Before you begin:

- You must have Read-Write permission for Log & Report settings.

To display the graphs:

- Go to Monitor > Port Statistics > [Selection].

Figure 46: Port Statistics: Packets graph

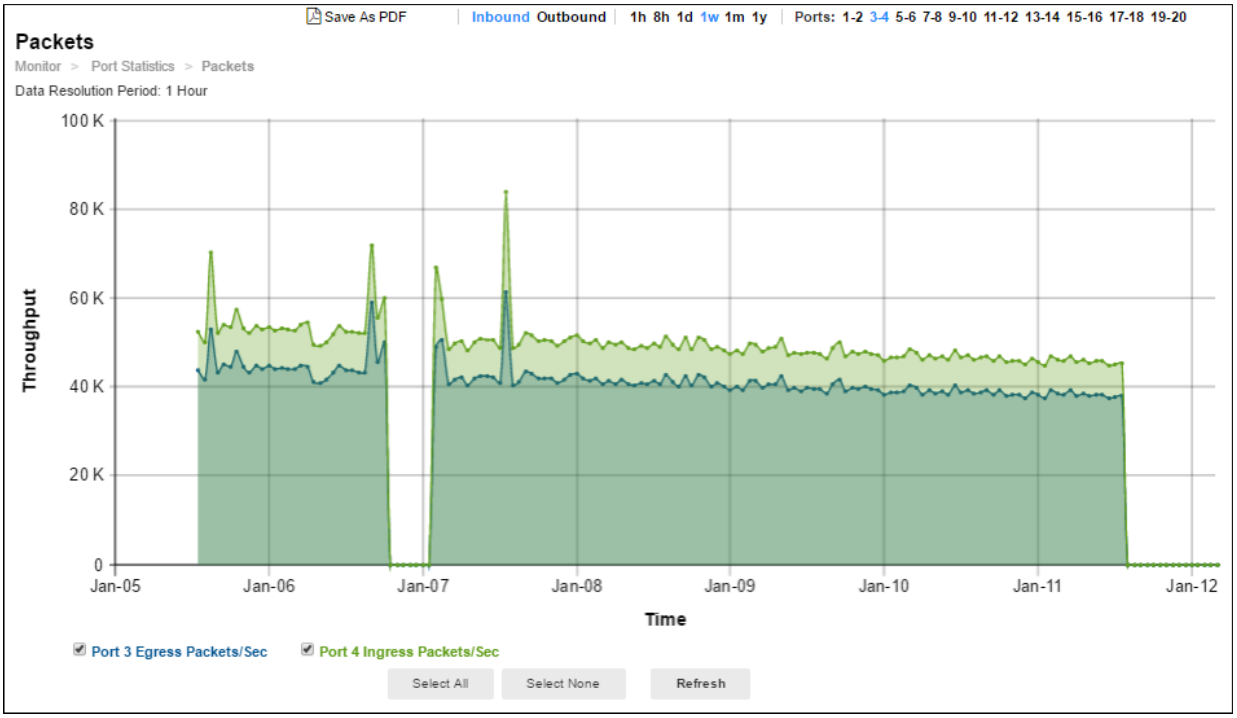

Data Resolution Period: 1 Hour

Note: For Port Statistics graphs, the data resolution for the 1 hour period is 30 seconds.

Using the SPP Statistics graphs

You use the SPP Statistics graphs to monitor overall throughput for a selected SPP.

The SPP Statistics set includes the following graphs:

- Packets—Throughput in packets per period.
- Bits—Throughput in bits per period.

Figure 47 shows the Packets graph.

You can customize the following query terms: SPP, period, and direction.

Before you begin:

- You must have Read-Write permission for Log & Report settings.

To display the graphs:

- Go to Monitor > SPP Statistics > [Selection].

Sample Graph

Figure 47: SPP Statistics: Packets graph

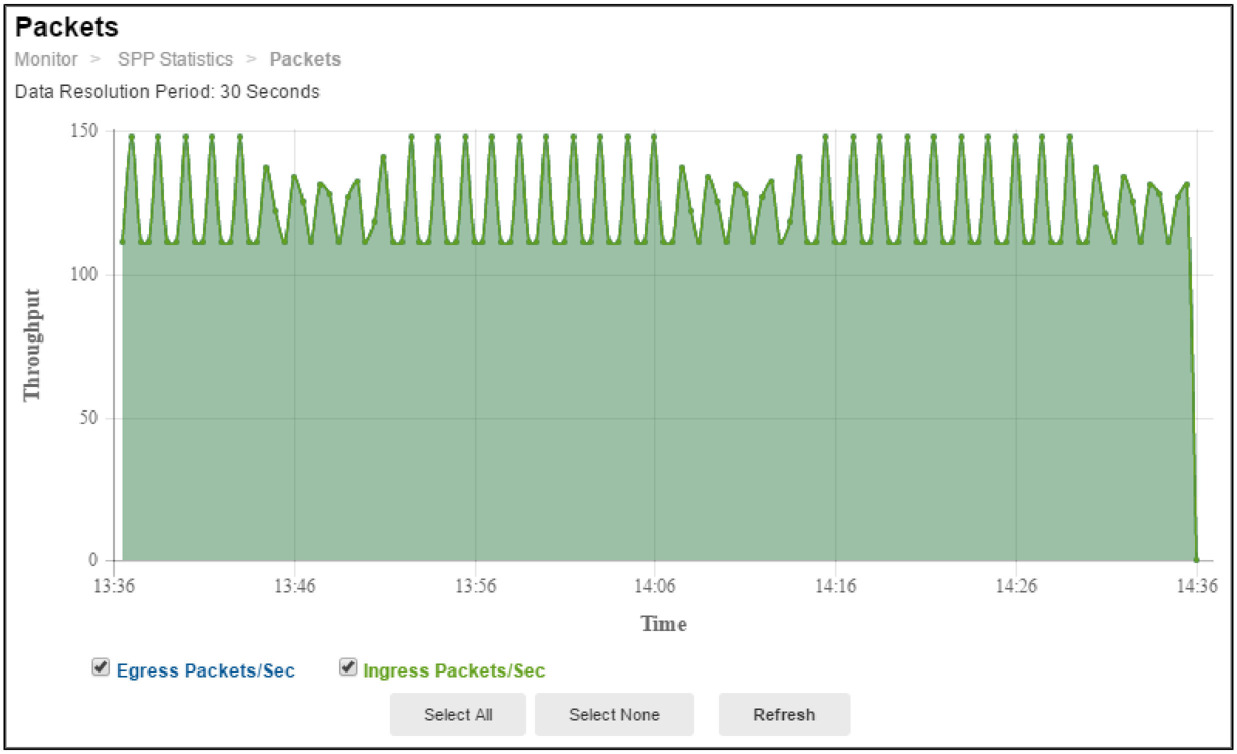

Packets

Monitor > SPP Statistics > Packets

Data Resolution Period: 30 Seconds

Using the Aggregate Drops graph

You use the Aggregate Drops graph to monitor trends in drops over time.

The Aggregate Drops graph plots the following data:

- Flood Drops—Aggregate of drops due to packet rate thresholds.
- ACL Drops—Aggregate of drops due to ACL rules.
- Anomaly Drops—Aggregate of drops due to anomaly detection methods.
- Hash Attack Drops—Aggregate of drops due to built-in rules that detect hash attacks on the FortiDDoS system itself.
- Out of Memory Drops—Aggregate of drops due to built-in rules that detect memory attacks on the FortiDDoS system itself.

Figure 48 shows the Aggregate Drops graph.

You can customize the following query terms: SPP, period.

Before you begin:

- You must have Read-Write permission for Log & Report settings.

To display the graph:

- Go to Monitor > Aggregate Drops.

Figure 48: Aggregate Drops graph

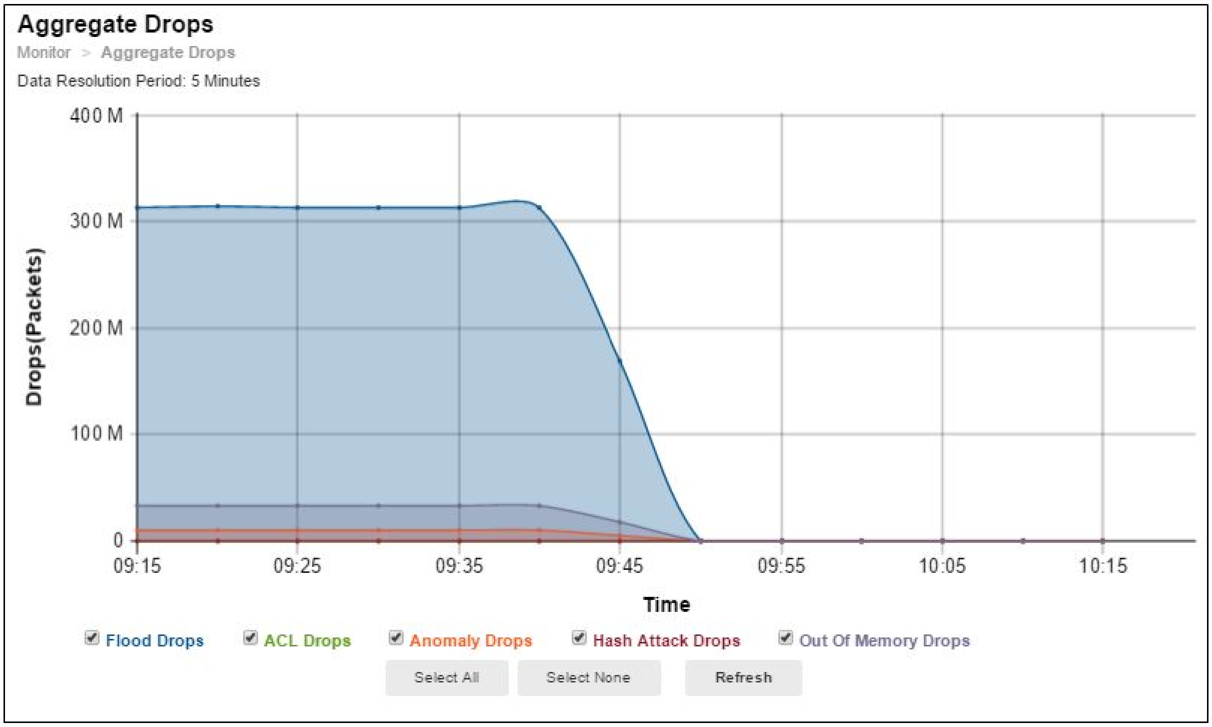

Aggregate Drops

Monitor > Aggregate Drops

Data Resolution Period: 5 Minutes

Using the Flood Drops graphs

You use the Flood Drops graphs to monitor drops due to SPP packet rate thresholds that detect flood attacks. Table 48 summarizes the statistics displayed in the graphs.

You can customize the following query terms: SPP and period.

Before you begin:

- You must have Read-Write permission for Log & Report settings.

To display the graph:

- Go to Monitor > Flood Drops > [Selection].

Click the graph icon 📈 in the legend to drill down from aggregate statistics to more specific queries.

Table 48: Flood Drops graphs

Statistic	Description
Aggregate	
Layer 3	Aggregation of drops due to SPP Layer 3 thresholds.
Layer 4	Aggregation of drops due to SPP Layer 4 thresholds.
Layer 7	Aggregation of drops due to SPP Layer 7 thresholds.
Layer 3	
Protocols	Aggregation of drops due to protocols thresholds. These counters track the packet rate for each protocol.
Fragmented Packets	Drops due to the **fragment** threshold. This counter tracks packet fragments received.
Source Flood	Drops due to the **most-active-source** threshold. This counter tracks packets from source IP addresses.
Destination Flood	Drops due to the **most-active-destination** threshold. This counter tracks packets to protected IP addresses.

Statistic	Description
Layer 4	
SYN	Drops due to **syn** threshold. This counter tracks the SYN packet rate for all traffic belonging to the SPP.
TCP Ports	Aggregation of drops due to the thresholds for TCP ports.
UDP Ports	Aggregation of drops due to the thresholds for UDP ports.
ICMP Types/Codes	Aggregation of drops due to the SPP thresholds for ICMP types/codes.
Zombie Flood	Drops due to the **new-connections** threshold, which sets a limit for legitimate IPs. FortiDDoS assumes a zombie flood is underway when the number of allowed legitimate IP addresses during a SYN flood exceeds a set threshold. These packets indicate that non-spoofed IP addresses are creating a DDoS attack by generating a large number SYN packets.
SYN Per Source Flood	Drops due to the **syn-per-source** threshold. This counter tracks SYN packets for each source.
Connections Per Source	Drops due to the **concurrent-connections-per-source** threshold.
SYN Per Destination	Drops due to the **syn-per-dst** threshold.
Slow Connection	Drops due to slow connection detection and blocking of identified sources of slow connection attacks.
Layer 7 > Aggregate	
DNS	Aggregation of drops due to DNS thresholds.
HTTP	Aggregation of drops due to HTTP thresholds.
Layer 7 > DNS	
Unsolicited DNS Response Drop	Drops when a DNS response is received but there is no DNS query entry in the DNS query response matching table.
LQ Drops	Drops during a Query Flood when the query is not in the legitimate query (LQ) table.
TTL Drop	Drops during a Query Flood when a source IP address sends a repeated DNS query for the same destination before the TTL has expired. It is expected that the query should not be repeated until the TTL expires.
Cache Drop	Drops during a Query Flood when a response was served from the cache or because a response was not found in the cache and the system is configured to drop such queries.

Statistic	Description
Spoofed IP Drop	Drops due antispoofing checks.
Unexpected Query Drop	Drops due to duplicate query checks.
Query Per Source Drop	Drops due to the **dns-query-per-src** threshold. This counter tracks DNS queries from source IP addresses.
Suspicious Sources Drop	Drops due to the **dns-packet-track-per-src** threshold. This counter tracks sources that demonstrate suspicious activity (a score based on heuristics that count fragmented packets, response not found in DQRM, or queries that generate responses with RCODE other than 0).
Fragment Drop	Drops due to **dns-fragment** threshold for TCP traffic.
TCP Query Drop	Drops due to the **dns-query threshold** for TCP traffic.
TCP Question Drop	Drops due to the **dns-question-count** threshold for TCP traffic.
TCP MX Drop	Drops due to the **dns-mx-count** threshold for TCP traffic.
TCP All Drop	Drops due to the **dns-all** threshold for TCP traffic.
TCP Zone Transfer Drop	Drops due top the **dns-zone-xfer** threshold for TCP traffic.
Layer 7 > HTTP	
Method Flood	Aggregation of drops due to the thresholds for HTTP Methods.
URL Flood	Aggregation of drops due to thresholds for URLs.
Host Flood	Aggregation of drops due to thresholds for Host headers.
User Agent Flood	Aggregation of drops due to thresholds for User Agent headers.
Cookie Flood	Aggregation of drops due to thresholds for Cookie headers.
Referer Flood	Aggregation of drops due to thresholds for Referer headers.
Methods per Source Flood	Packet/second rate for Method packets (GET, HEAD, OPTION, POST, etc) from a single Source. When the per-source limits are exceeded for a particular source, the system applies the blocking period for identified sources. The connection to the server may also be RST if Protection Profiles > SPP Settings > TCP Tab: Aggressive Aging TCP Connections Feature Control: Layer 7 Flood is enabled.

Sample Graphs

Figure 49: Flood Drops Aggregate graph

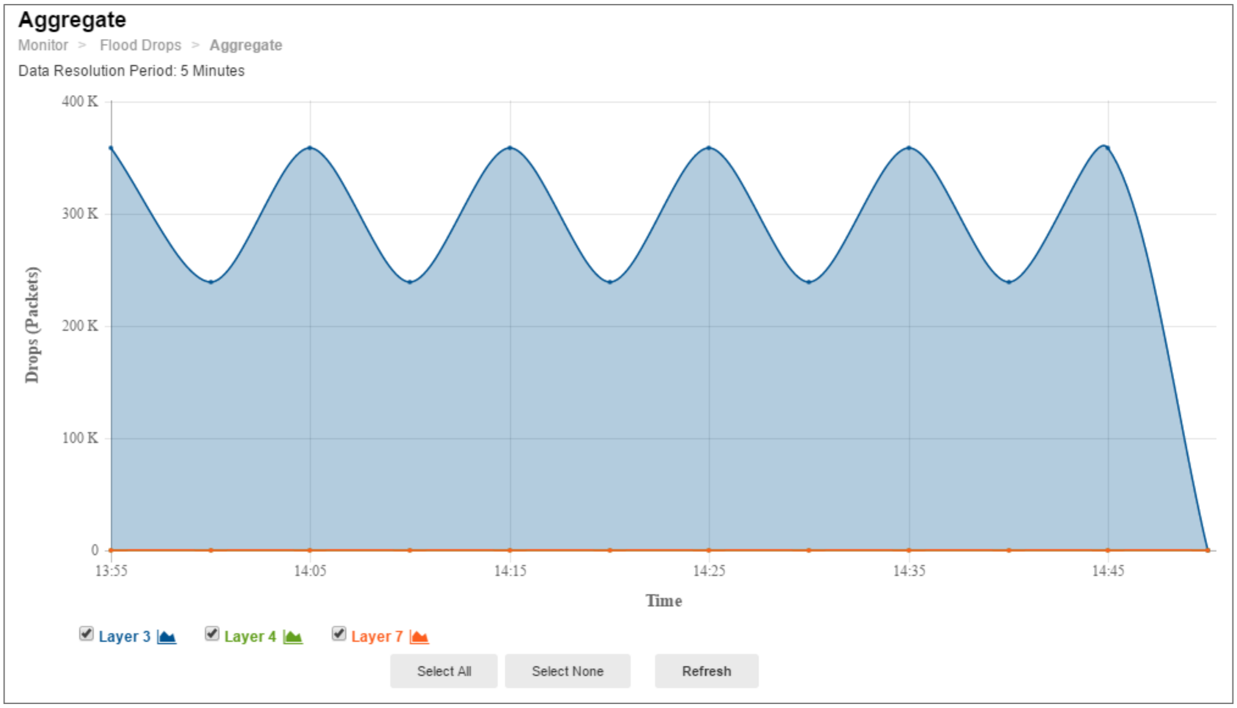

Aggregate

Monitor > Flood Drops > Aggregate

Data Resolution Period: 5 Minutes

Figure 50: Flood Drops Layer 3 graph

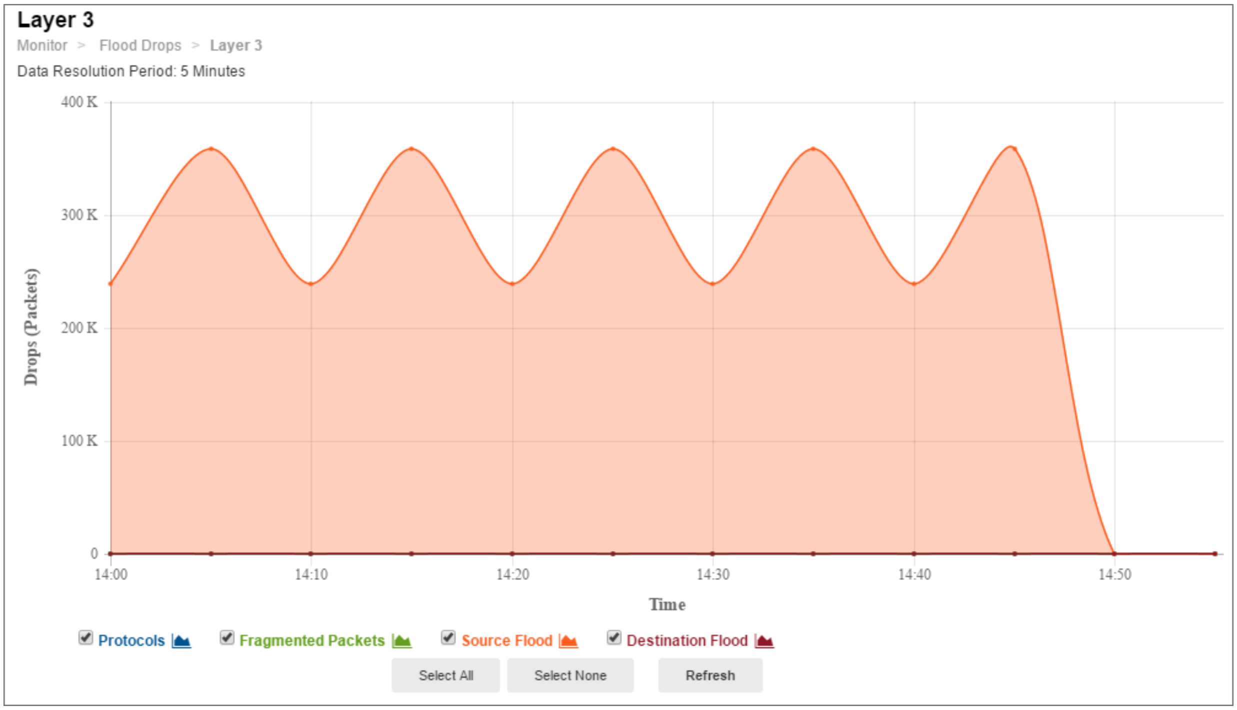

Layer 3

Monitor > Flood Drops > Layer 3

Data Resolution Period: 5 Minutes

Figure 51: Flood Drops Layer 4 graph

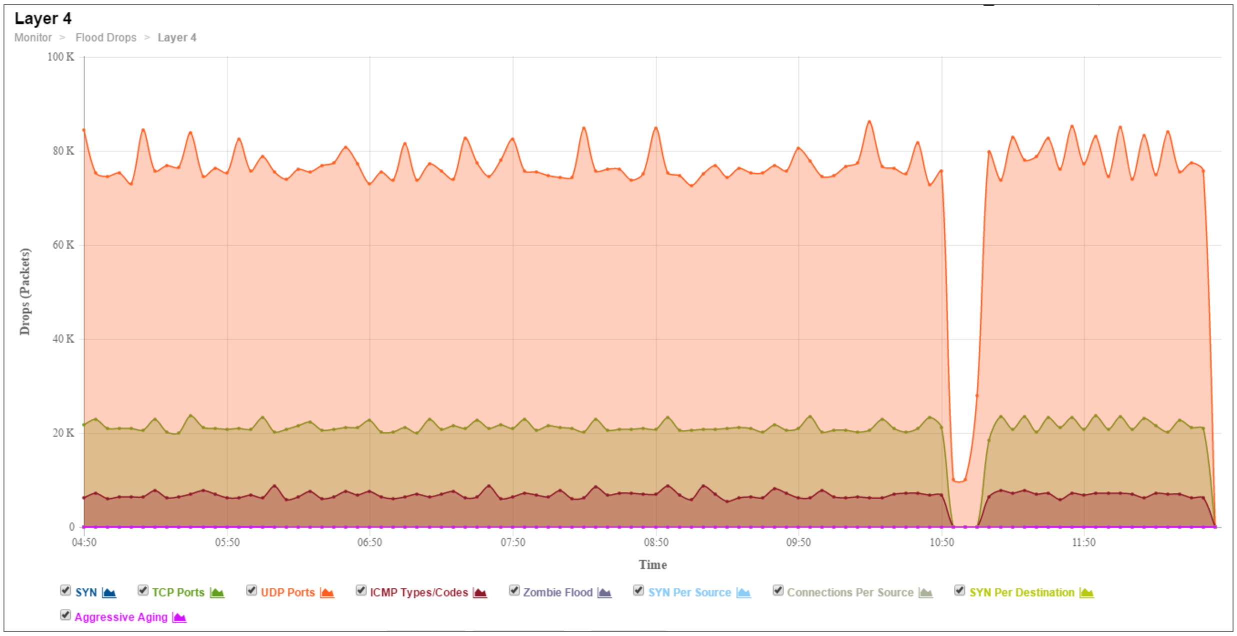

Figure 52: Flood Drops Layer 7 Aggregate

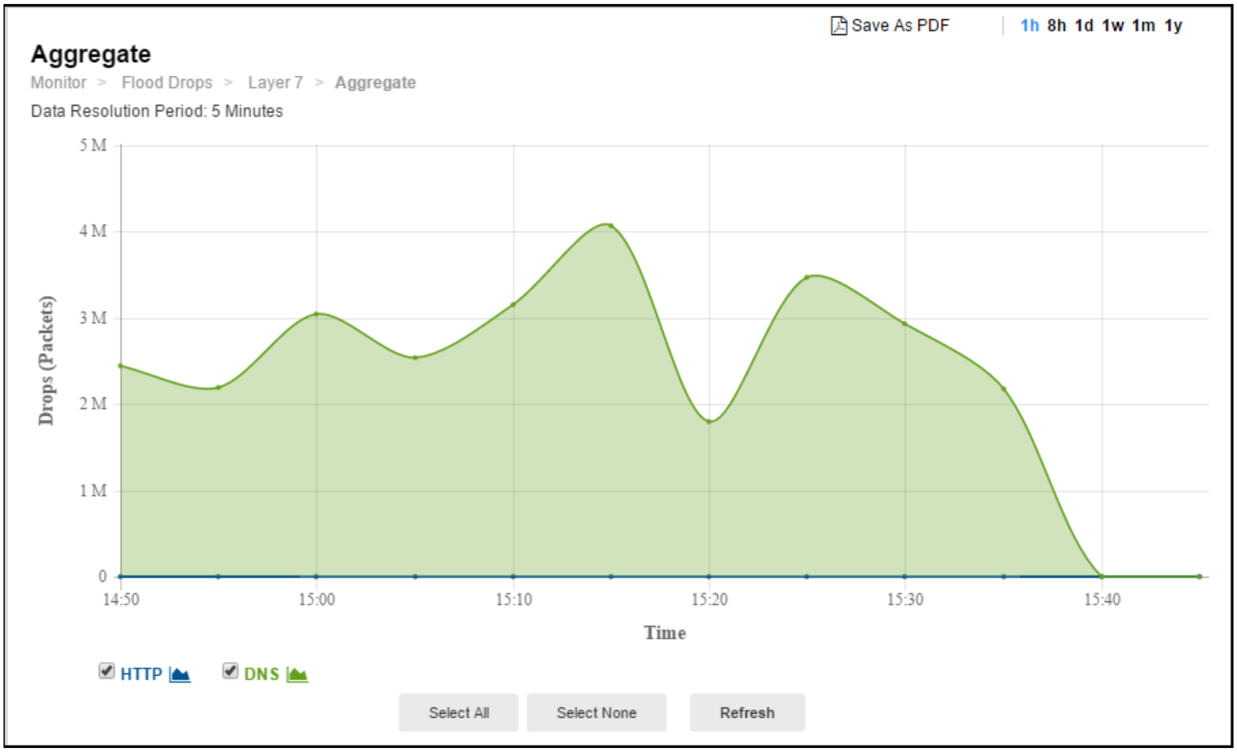

Data Resolution Period: 5 Minutes

Figure 53: Flood Drops Layer 7 HTTP

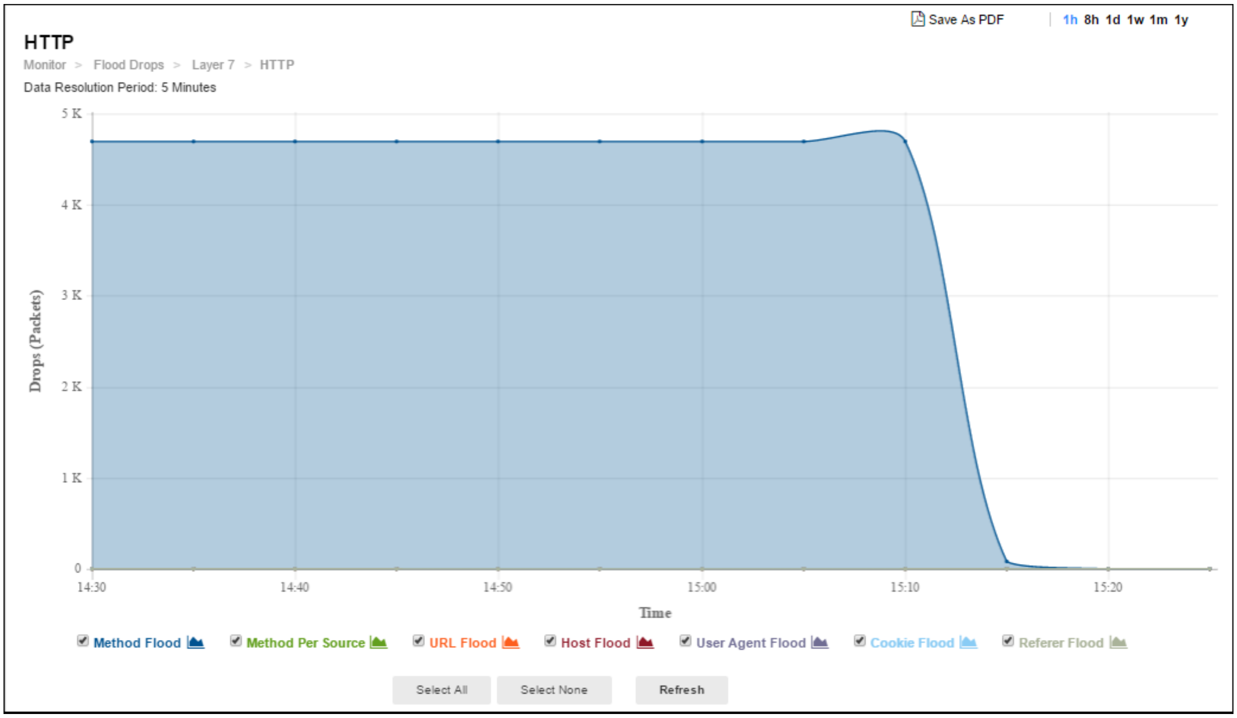

Data Resolution Period: 5 Minutes

Figure 54: Flood Drops Layer 7 DNS

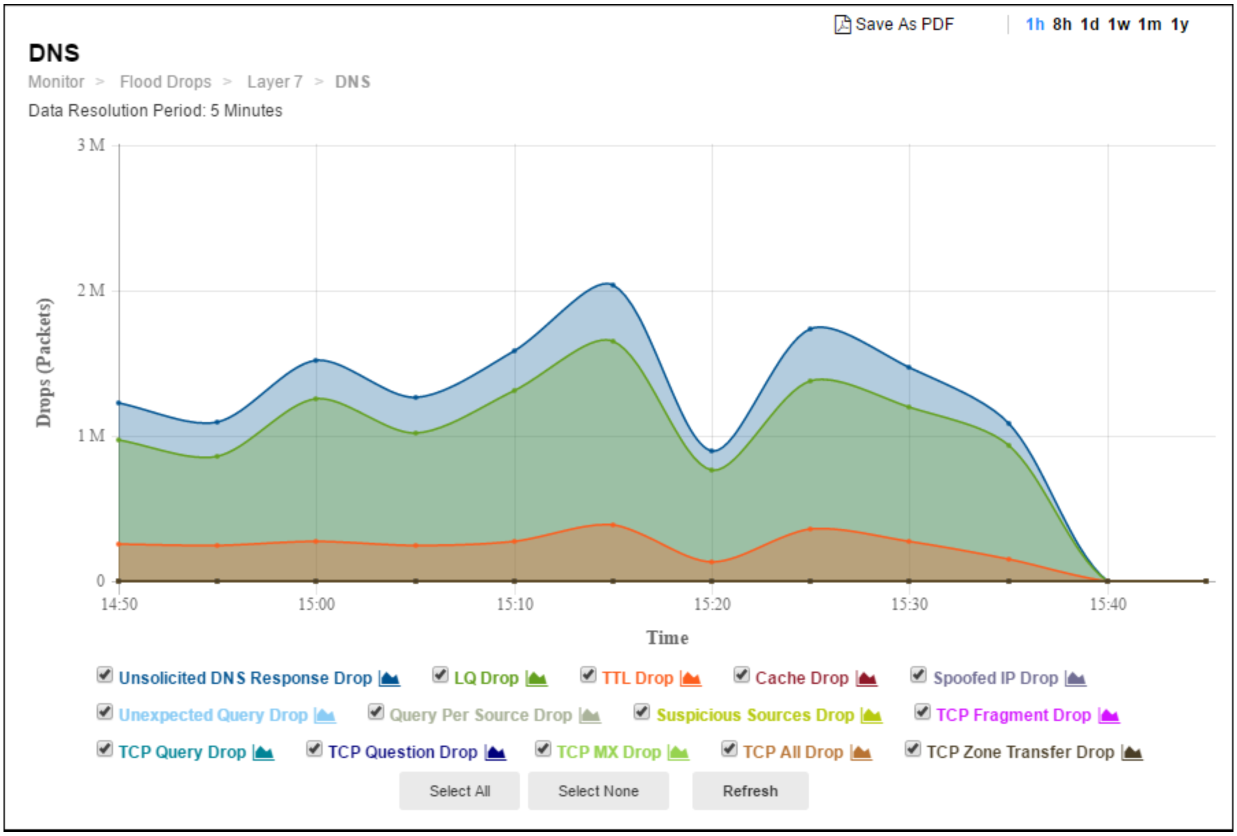

Using the ACL Drops graphs

You use the ACL Drops graphs to monitor drops due to Global ACL and SPP ACL rules.

Note: If ACL is set and unset within 5 minutes, any drops associated with this ACL will be shown under Flood Drops instead of ACL drops.

Table 49 summarizes the statistics displayed in the graphs. You can customize the following query terms: SPP and period.

Before you begin:

- You must have Read-Write permission for Log & Report settings.

To display the graph:

- Go to Monitor > ACL Drops > [Selection].

Click the graph icon 📈 in the legend to drill down from aggregate statistics to more specific queries.

Table 49: ACL Drops graphs

Statistic	Description
Aggregate	
Layer 3	An aggregation of drops due to ACL rules based on Layer 3 parameters.
Layer 4	An aggregation of drops due to ACL rules based on Layer 4 parameters.
Layer 7	An aggregation of drops due to ACL rules based on Layer 7 parameters.
Layer 3	
Protocol Denied Drops	Drops due to ACL rules based on service protocol.
Fragmented Packet Denied Drops	Drops due to ACL rules based on service object **Fragment**.
Address Denied	Drops due to ACL rules based on IP address, geolocation, IP reputation, or local address anti-spoofing rules.
Layer 4	
TCP Denied Drops	Drops due to ACL rules based on service object **TCP-Port**.
UDP Denied Drops	Drops due to ACL rules based on service object **UDP-Port**.
ICMP Type/Code Denied Drops	Drops due to ACL rules based on service object **ICMP-Type-Code**.
Layer 7 > Aggregate	
DNS	Drops due to DNS ACL rules.
HTTP	Drops due to HTTP ACL rules.
Layer 7 > DNS	
Frag Drops	Drops due to ACL rules based on service object **DNS-Fragment**.
MX Drops	Drops due to ACL rules based on service object **DNS-MX**.
Qtype All Drops	Drops due to ACL rules based on service object **DNS-ALL**.
Zone Transfer Drops	Drops due to ACL rules based on service object **DNS-Zone-Transfer**.
Query Restricted t Specific Subnet Drops	Drops due to an ACL.
Query ACL Drops	Drops due to an ACL.

Statistic	Description
Layer 7 > HTTP	
URL Denied Drops	Drops due to ACL rules based on service object **URL**.
Host Denied Drops	Drops due to ACL rules based on service object **Host**.
Cookie Denied Drops	Drops due to ACL rules based on service object **Cookie**.
Referer Denied Drops	Drops due to ACL rules based on service object **Referer**.
User Agent Denied Drops	Drops due to ACL rules based on service object **User-Agent**.

Figure 55: ACL Drops Aggregate graph

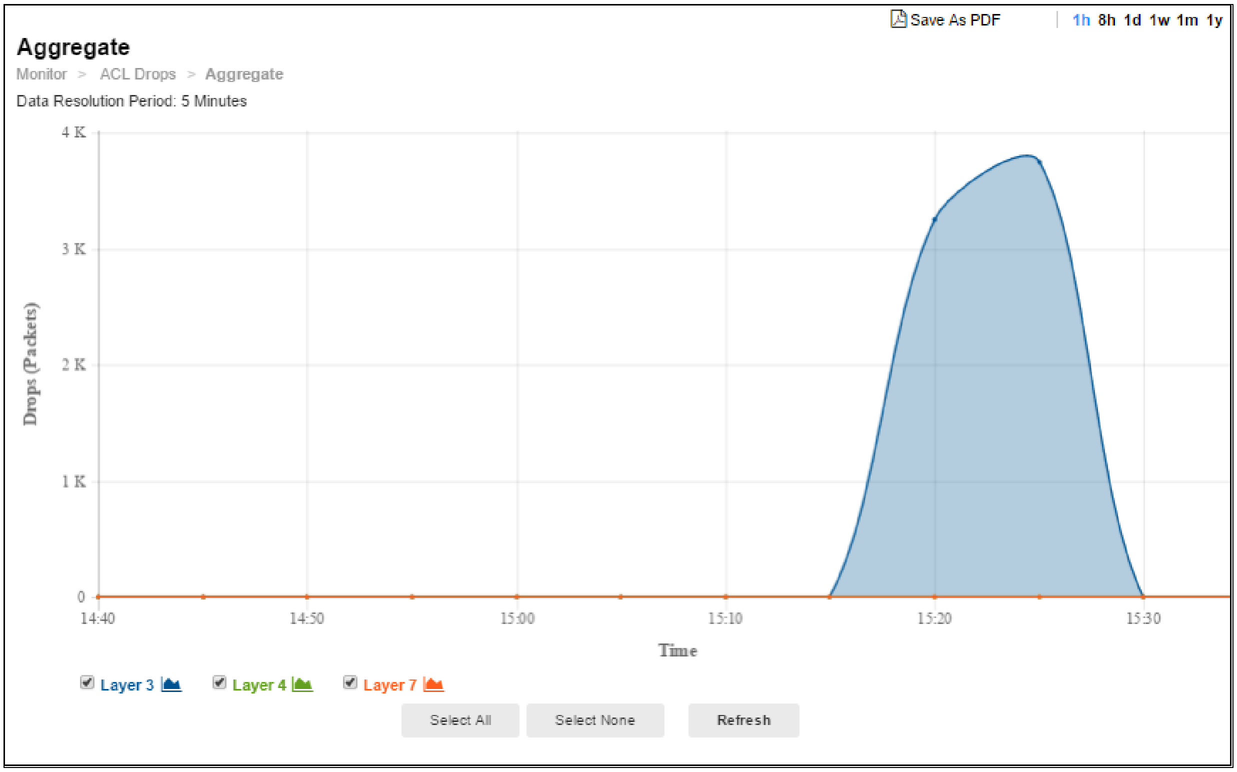

Figure 56: Layer 3 ACL Drops graph

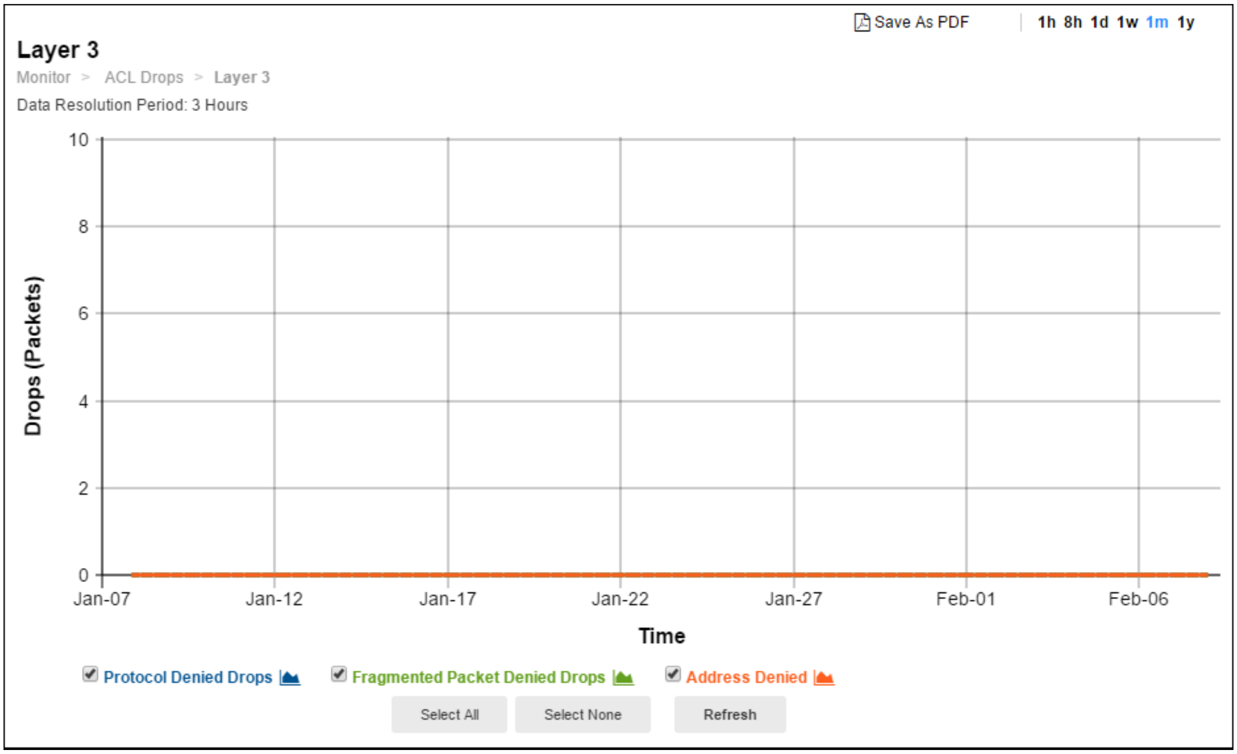

Figure 57: Layer 4 ACL Drops graph

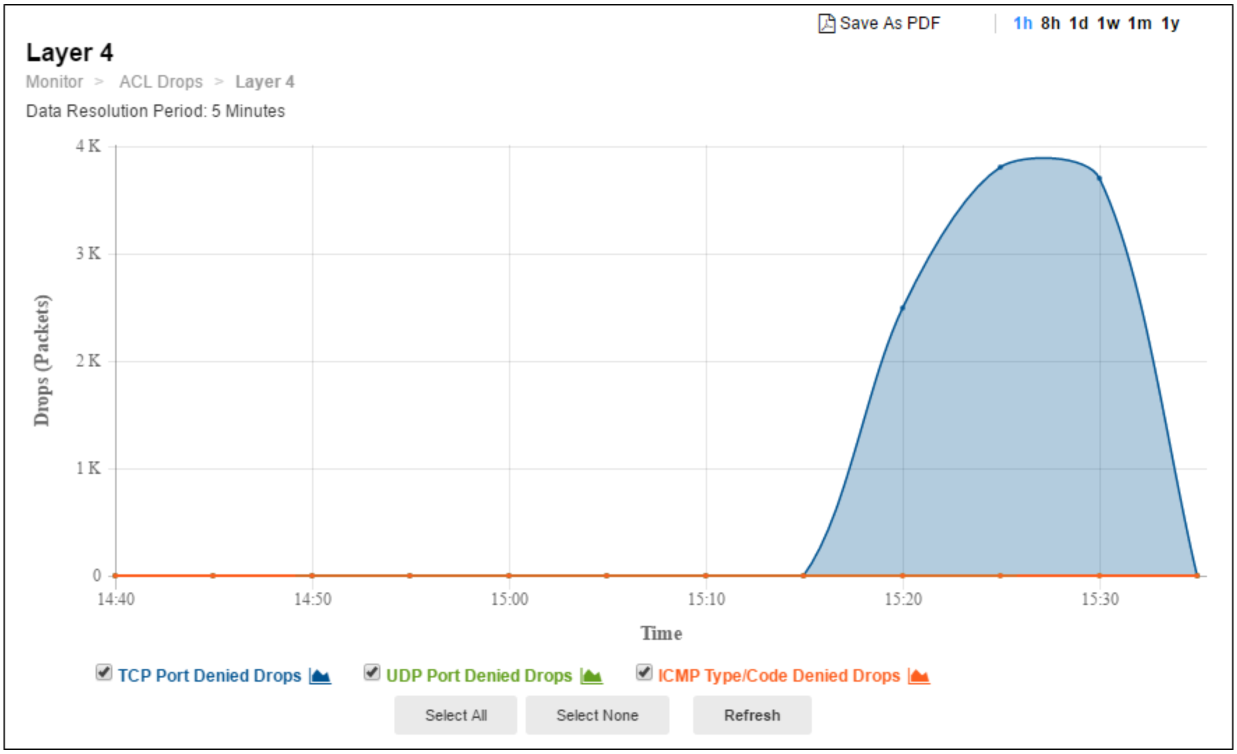

Layer 4

Monitor > ACL Drops > Layer 4

Data Resolution Period: 5 Minutes

Figure 58: Layer 7 ACL Drops Aggregate graph

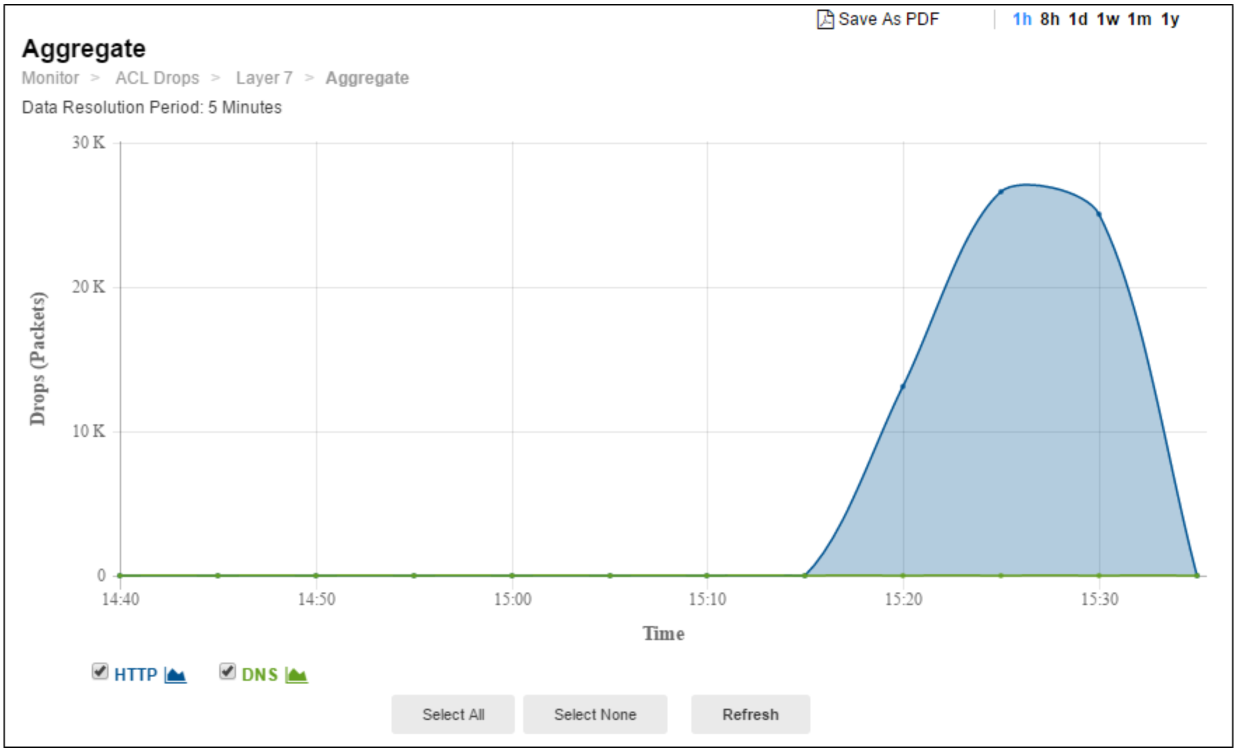

Aggregate

Monitor > ACL Drops > Layer 7 > Aggregate

Data Resolution Period: 5 Minutes

Figure 59: Layer 7 HTTP ACL Drops graph

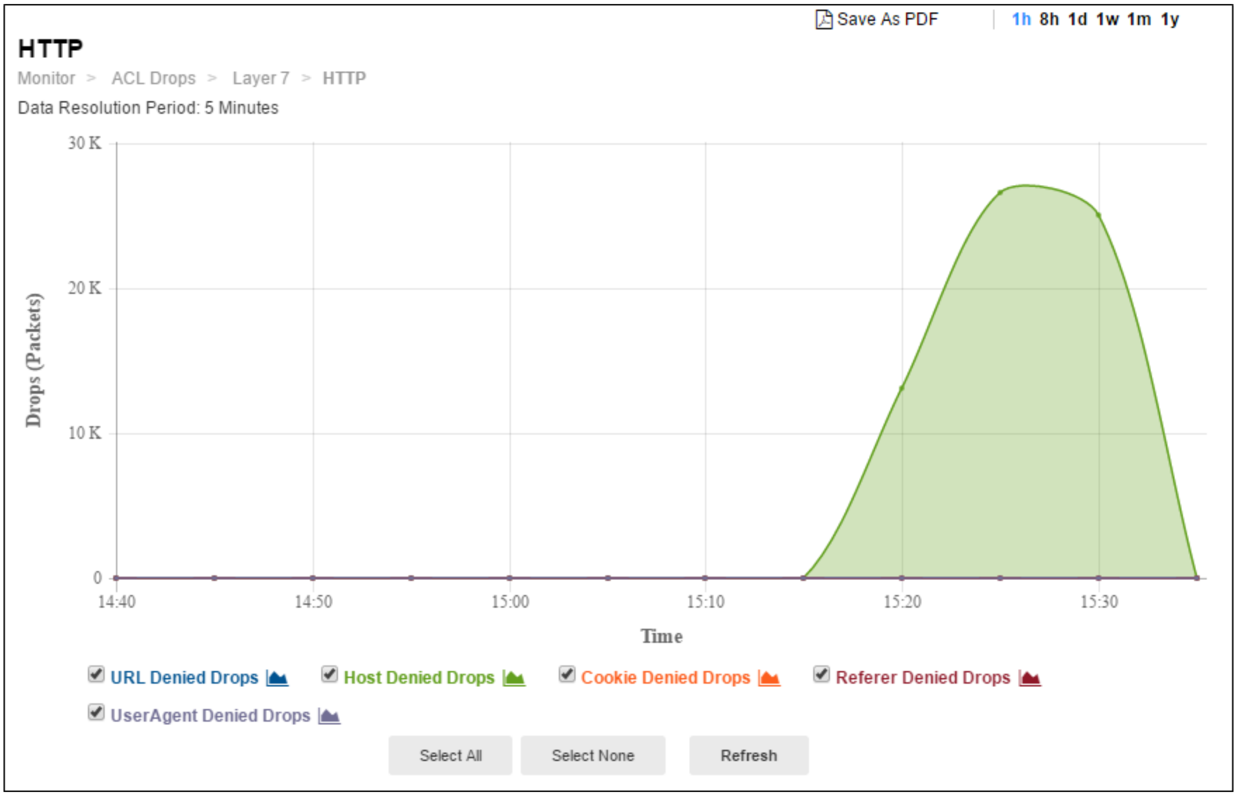

Figure 60: Layer 7 DNS ACL Drops graph

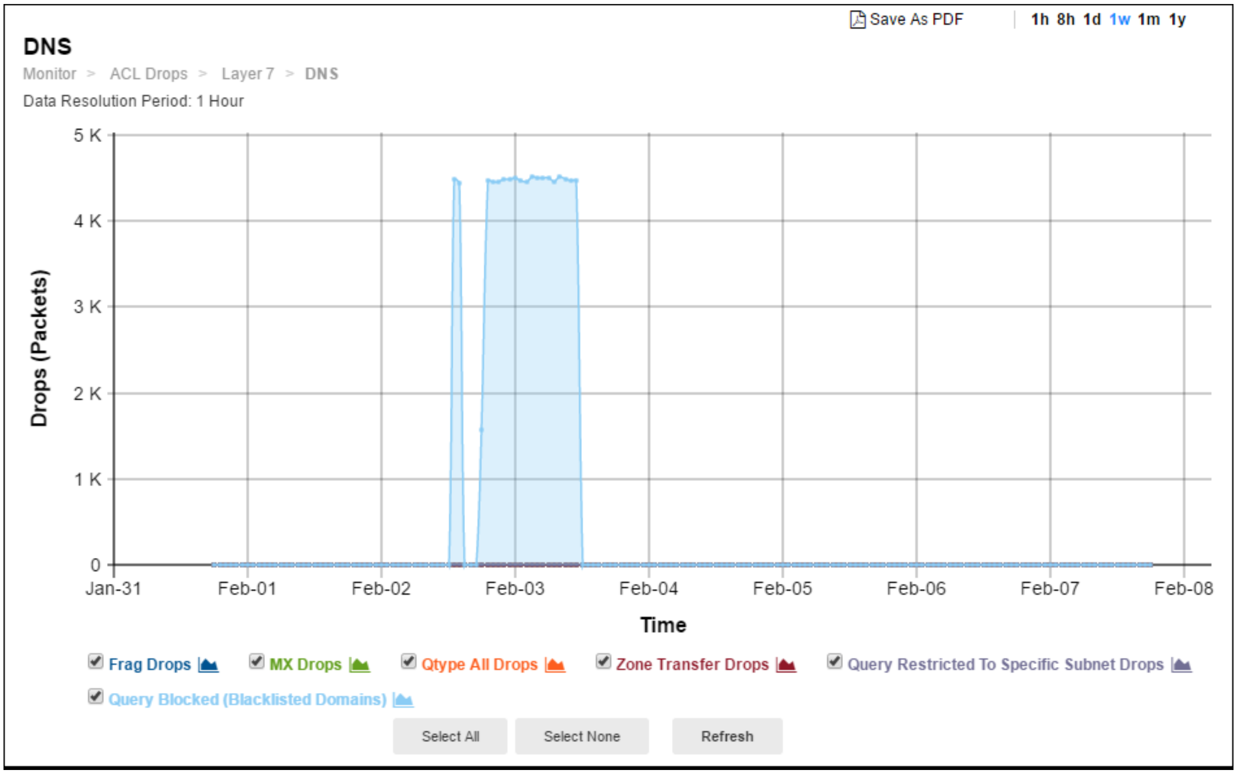

Using the Anomaly Drops graphs

You use the Anomaly Drops graphs to monitor drops due to Layer 3, Layer 4, and Layer 7 protocol anomalies. Table 51 summarizes the statistics displayed in the graph.

You can customize the following query terms: SPP, period, direction.

Before you begin:

- You must have Read-Write permission for Log & Report settings.

To display the graph:

- Go to Monitor > Anomaly Drops > [Selection].

Click the graph icon 📈 in the legend to drill down from aggregate statistics to more specific queries.

Table 50: Anomaly Drops graphs

Statistic	Description
Aggregate	
Layer 3	Aggregation of Layer 3 anomaly drops.
Layer 4	Aggregation of Layer 4 anomaly drops.
Layer 7	Aggregation of Layer 7 anomaly drops.
Layer 3	
IP Header Checksum Error	Drops due to checksum errors.

Statistic	Description
Layer 3	Drops due to the Layer 3 anomalies, including: • IP version other than 4 or 6 • Header length less than 5 words • End of packet (EOP) before 20 bytes of IPV4 Data • Total length less than 20 bytes • EOP comes before the length specified by Total length • End of Header before the data offset (while parsing options) • Length field in LSRR/SSRR option is other than (3+(n*4)) where n takes value greater than or equal to 1 • Pointer in LSRR/SSRR is other than (n*4) where n takes value greater than or equal to 1 • For IP Options length less than 3 • Reserved flag set • More fragments
Source and Destination Address Match	Drops due to an anomaly that is detected when source and destination addresses are the same (LAND attack).
Source/Destination as Localhost	Drops due to an anomaly that is detected when source or destination address is the same as the localhost (loopback address spoofing).
Layer 4 > Aggregate	
Header	Drops due to Layer 4 header anomalies.
State	Drops due to TCP state anomalies.
Layer 4 > Header	
TCP Checksum Error	Drops due to checksum errors.
UDP Checksum Error	Drops due to checksum errors.
ICMP Checksum Error	Drops due to checksum errors.
TCP Invalid Flag Combination	Drops due to invalid flag combinations, such as SYN/RST.
Invalid ICMP Type/Code	Trend in packets dropped due to invalid ICMP type/code anomaly.

Statistic	Description
Anomaly Detected	Drops due to Layer 4 anomalies, including: • Other header anomalies, such as incomplete packet • Urgent flag is set then the urgent pointer must be non-zero • SYN or FIN or RST is set for fragmented packets • Data offset is less than 5 for a TCP packet • End of packet is detected before the 20 bytes of TCP header • EOP before the data offset indicated data offset • Length field in Window scale option other than 3 in a TCP packet • Missing UDP payload • Missing ICMP payload • SYN with payload
Layer 4 > State	
Forward Transmission Not Within Window	Drops due to packets outside the receiver's TCP or UDP windows (when the Protection Profiles > SPP Settings > TCP session feature control **seq-validation** option is enabled).
Reverse Transmission Not Within Window	Drops due to packets outside the receiver's TCP or UDP windows (**seq-validation**).
TCP State Transition	Drops due to packets that violate the TCP Protocol state transition rules or sequence numbers (**state-transition-anomalies-validation**).
Foreign Packets	Drops due to packets that do not belong to a known TCP connection (**foreign-packet-validation**). For example, when the system receives a packet for a connection that has not been established with a SYN exchange.
Aggressive ageing due to Slow Connection	Drops due to slow connection.
Layer 7 > Aggregate	
HTTP	Aggregate of drops due to HTTP anomalies.
DNS	Aggregate of drops due to DNS anomalies.
Layer 7 > HTTP Header	
Known Method	Drops the packet if its method matches with any of the eight known opcodes but the same is not allowed method.
Unknown Method	Drops due to packets with an invalid HTTP method.

Statistic	Description
Invalid HTTP Version	Drops due to packets with an invalid HTTP version.
Range Present	Drops due to packets with a header range request. Present when Global Settings > Settings > Settings > General Tab > Drop HTTP Range Header is ENABLED.
Layer 7 > DNS > Aggregate	
Header	Aggregate of drops due to header anomalies.
Query	Aggregate of drops due to query message anomalies.
Response	Aggregate of drops due to response message anomalies.
Buffer Overflow	Aggregate of drops due to buffer overflow anomalies.
Exploit	Aggregate of drops due to DNS exploit anomalies.
Info	Aggregate of drops due to DNS info anomalies.
Data	Aggregate of drops due to DNS data anomalies.
Layer 7 > DNS > Header	
Invalid Opcode	Drops due to an invalid value in the OpCode field.
Illegal Flag	Drops due to an invalid combination in the flags field.
Source/Destination Both port 53	Drops due to source and destination port both being port 53.
Layer 7 > DNS > Query	
Query Bit Set	Drops due to a DNS query with the query reply (QR) bit set to 1. In a legitimate query, QR=0.
RA Bit Set	Drops due to a DNS query with the recursion allowed (RA) bit set. The RA bit is set in responses, not queries.
Null Query	Drops due to a DNS query in which the question, answer, additional, and name server counts are 0.
QD Count not One	Drops due to the number of entries in the question section of the DNS packet not equal to 1. Normally, a DNS query includes one question.
Layer 7 > DNS > Response	

Statistic	Description
QClass in Response	Drops due to a DNS response with a resource specifying a CLASS ID reserved for queries only (QCLASS).
Qtype in Response	Drops due to a DNS response with a resource specifying a TYPE ID reserved for queries only (QTYPE).
Query bit not set	Drops due to a DNS response with the query reply (QR) bit set to 0. In a legitimate response, QR=1.
QD count not 1	Drops due to the number of entries in the question section of the DNS packet not equal to 1. Normally, a DNS response includes one question.
Layer 7 > DNS > Buffer Overflow	
Message too long	Drops due to a TCP or UDP query or response message that exceeds the maximum length specified in the message header.
Name too long	Drops due to a DNS name that exceeds 255 characters.
Label length too large	Drops due to a query or response with a label that exceeds the maximum length (63) specified in the RFC.
Layer 7 > DNS > Exploit	
Pointer loop	Drops due to a DNS message with a set of DNS pointers that form a loop.
Zone Transfer	Drops due to an asynchronous Transfer Full Range (AXFR) request (QTYPE-E=252) from untrusted network.
Class not in	Drops due to a query or response in which the question/resource address class is not IN (Internet Address).
Empty message	Drops due to an empty UDP message.
Message ends prematurely	Drops due to a message that ends prematurely.
Layer 7 > DNS > Info	
Type all used	Drops due to a query for ALL resource records.
Layer 7 > DNS > Data	
Invalid type class	Drops due to invalid type or class data.
Extra data	Drops due to data in fields where no data is expected.
TTL too long	Drops due to a TTL value is greater than 7 days (or 604800 seconds).

Statistic	Description
Name length short	Drops due to a message with a null DNS name.

Sample Graphs

Figure 61: Anomaly Drops Aggregate graph

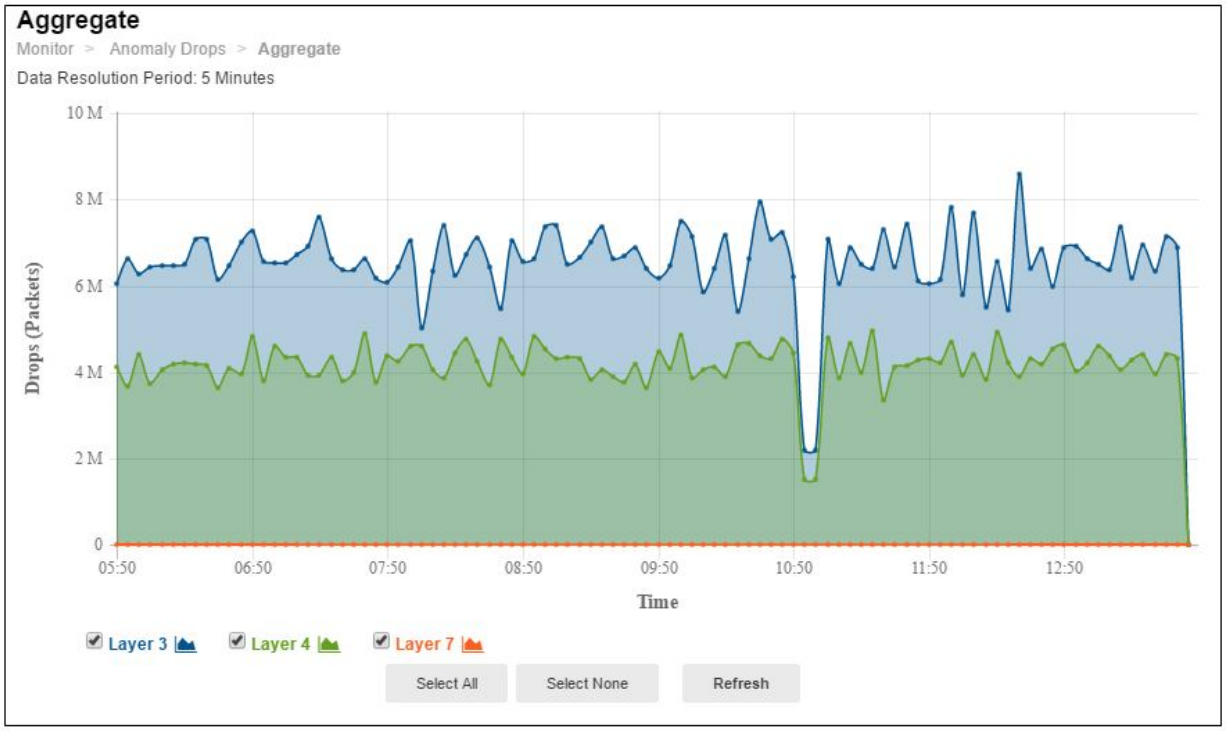
Aggregate

Monitor > Anomaly Drops > Aggregate

Data Resolution Period: 5 Minutes

Figure 62: Layer 3 Anomaly Drops graph

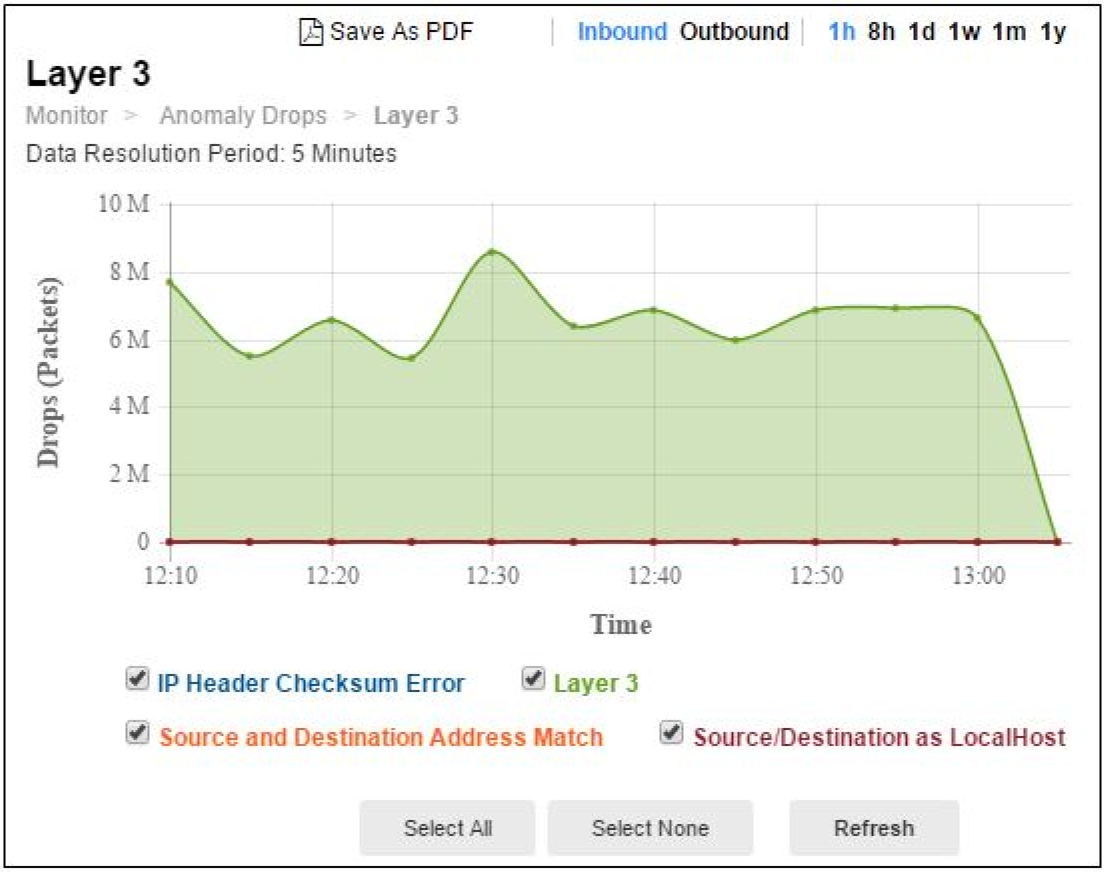

Layer 3

Monitor > Anomaly Drops > Layer 3

Data Resolution Period: 5 Minutes

Figure 63: Layer 4 Anomaly Drops Aggregate graph

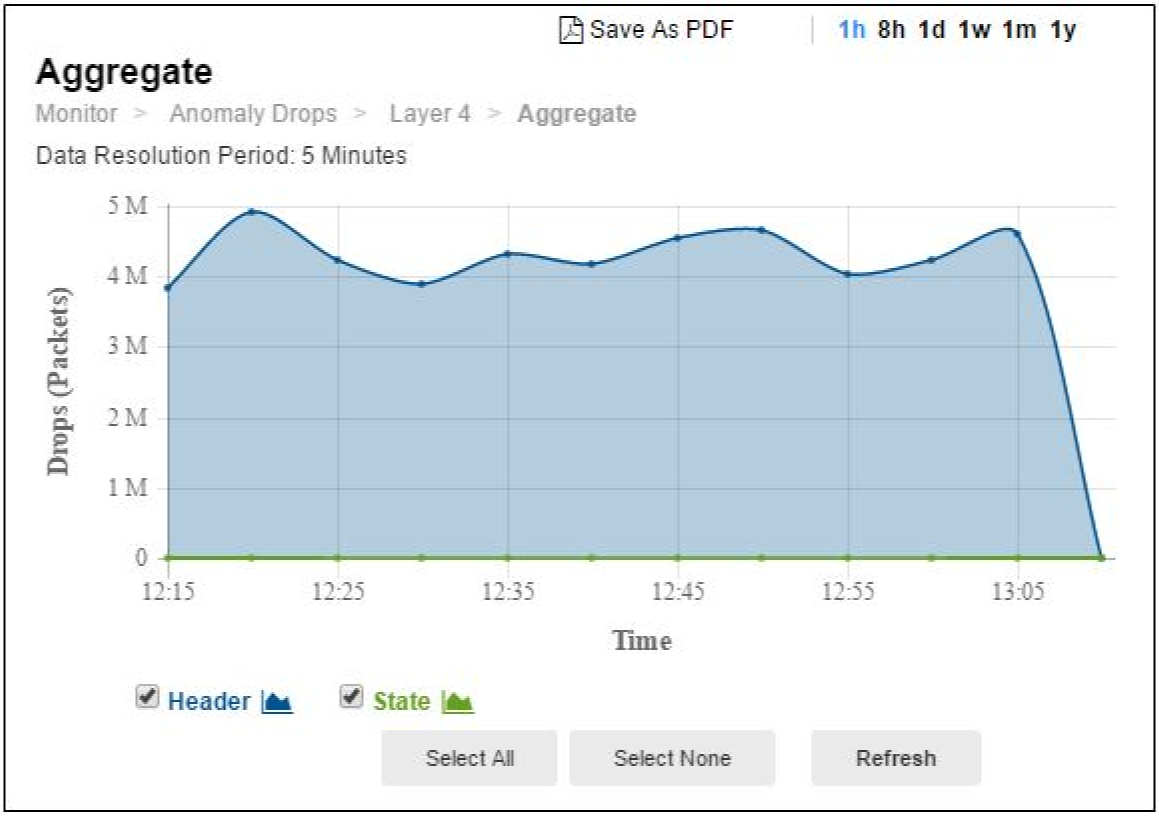

Aggregate

Monitor > Anomaly Drops > Layer 4 > Aggregate

Data Resolution Period: 5 Minutes

Figure 64: Layer 4 Header Anomaly Drops graph

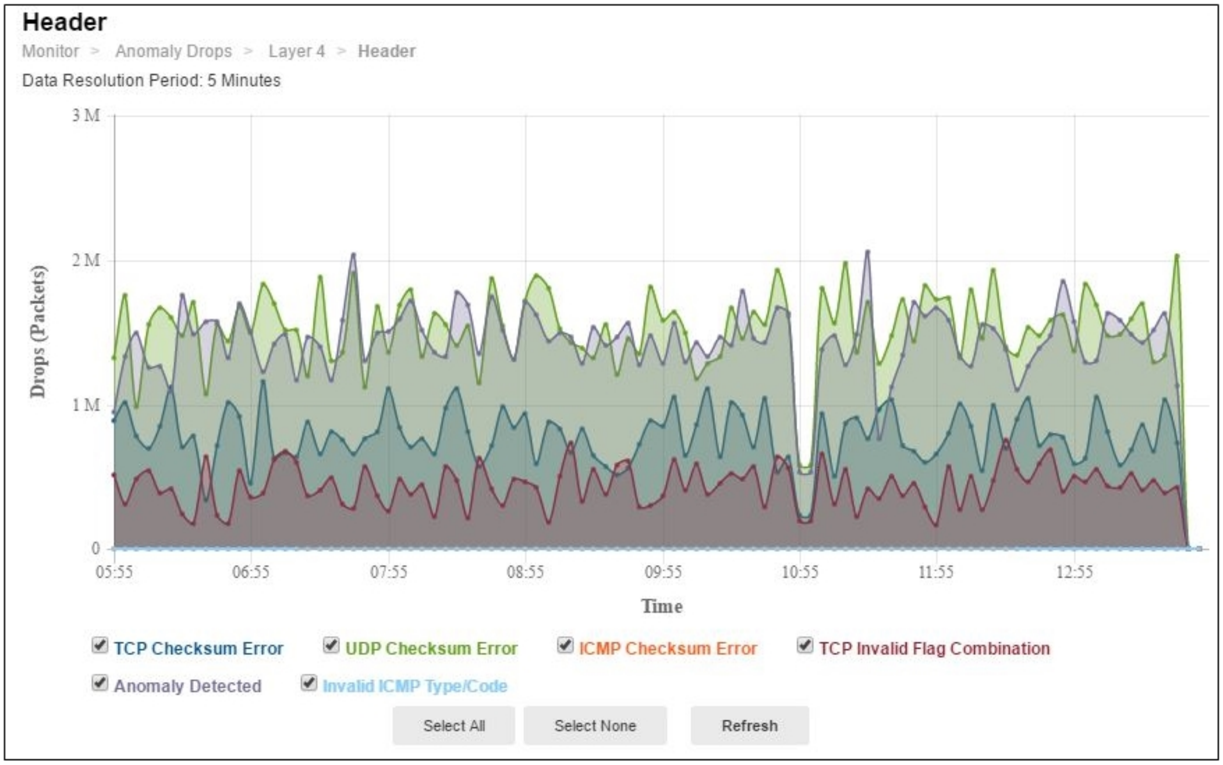

Header

Monitor > Anomaly Drops > Layer 4 > Header

Data Resolution Period: 5 Minutes

☑ TCP Checksum Error ☑ UDP Checksum Error ☑ ICMP Checksum Error ☑ TCP Invalid Flag Combination

☑ Anomaly Detected ☑ Invalid ICMP Type/Code

Figure 65: Layer 4 State Anomaly Drops graph

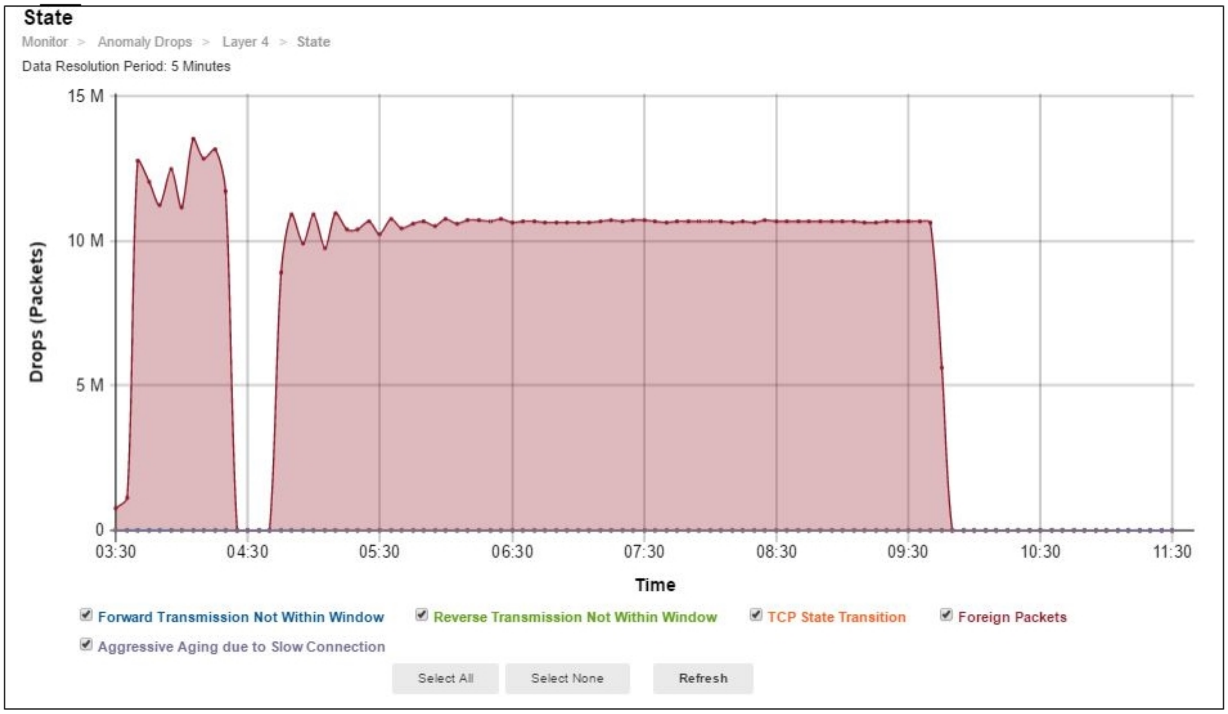

State

Monitor > Anomaly Drops > Layer 4 > State

Data Resolution Period: 5 Minutes

☑ Forward Transmission Not Within Window ☑ Reverse Transmission Not Within Window ☑ TCP State Transition ☑ Foreign Packets

☑ Aggressive Aging due to Slow Connection

Figure 66: Layer 7 Anomaly Drops Aggregate graph

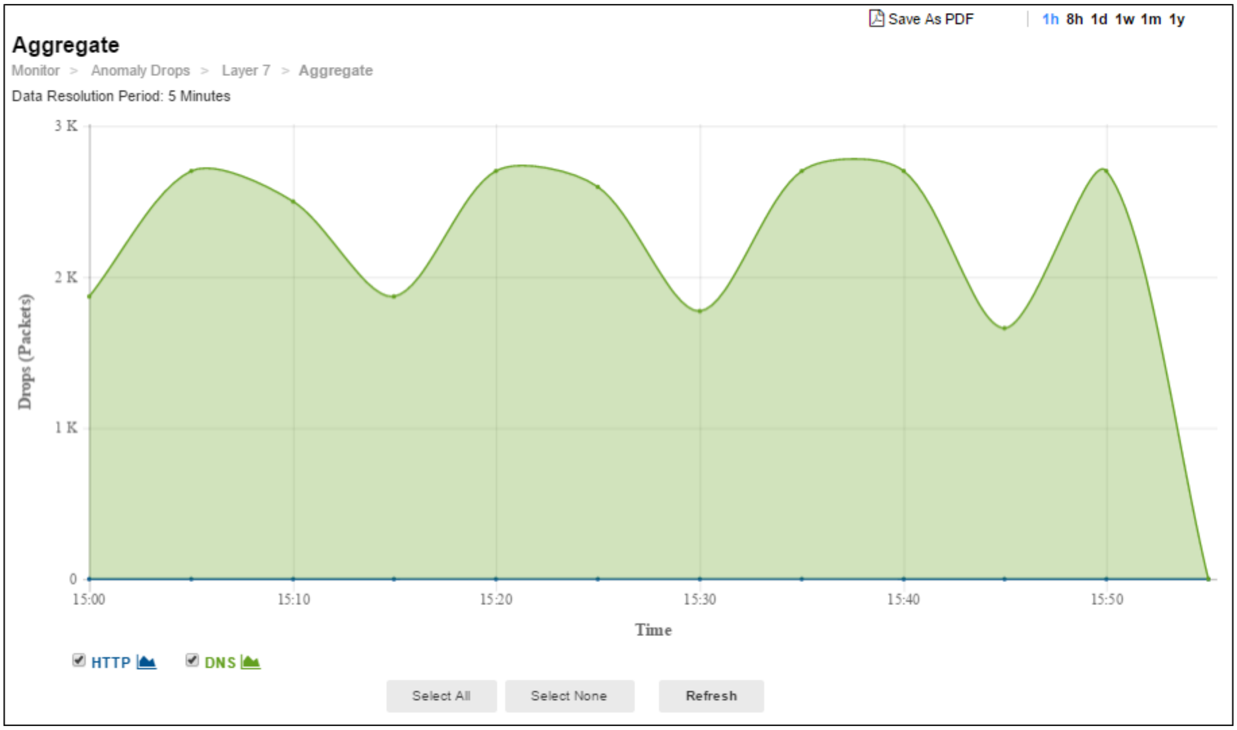

Figure 67: DNS Header Anomaly Drops graph

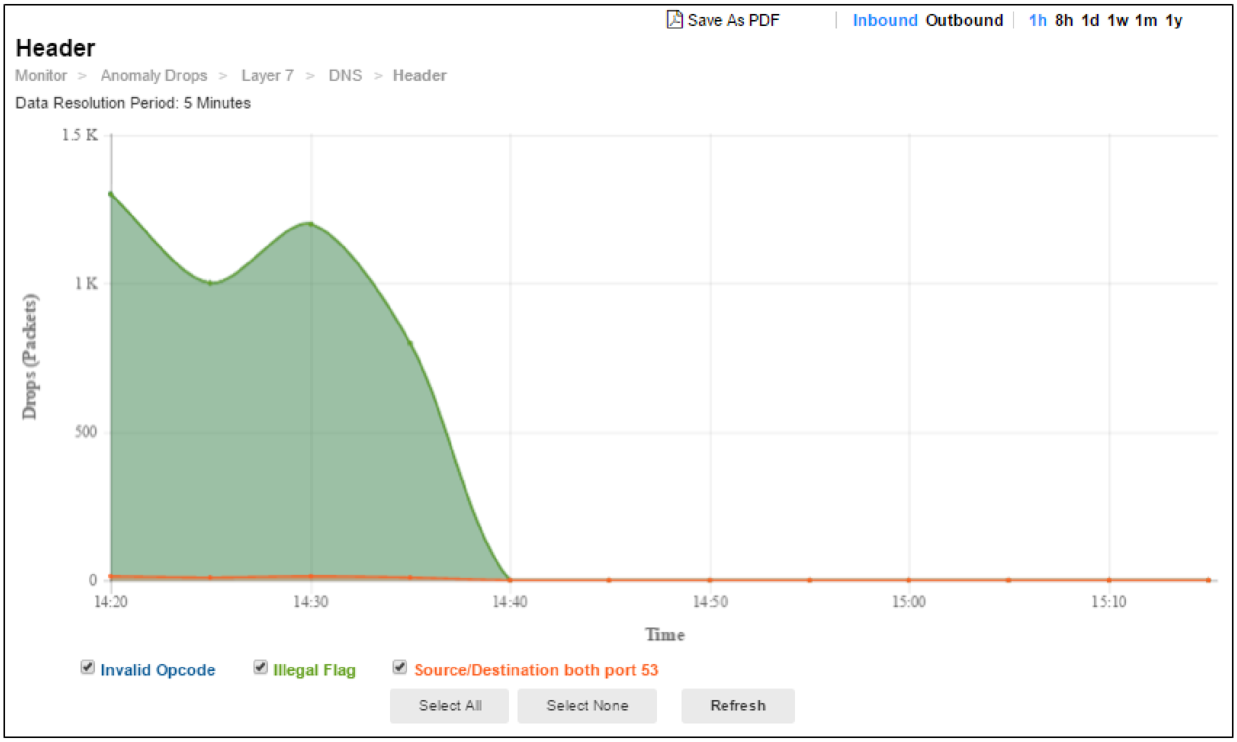

Figure 68: DNS Query Anomaly Drops graph

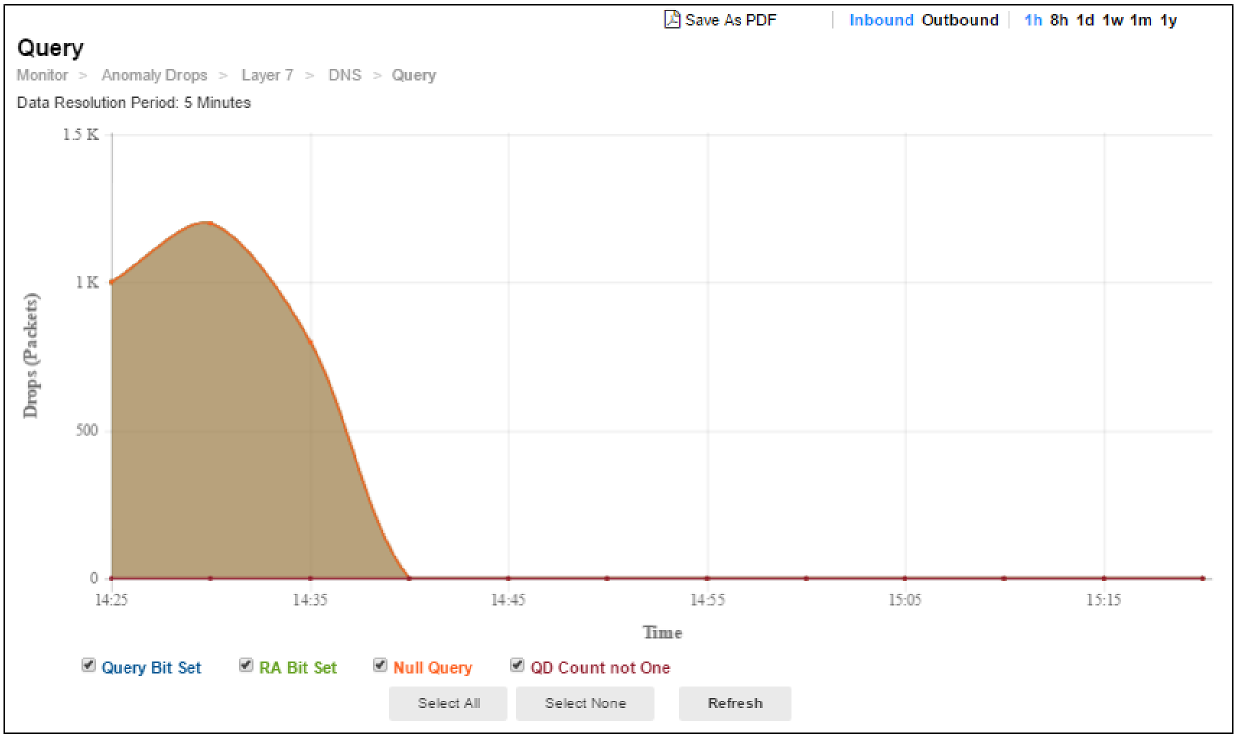

Figure 69: DNS Response Anomaly Drops graph

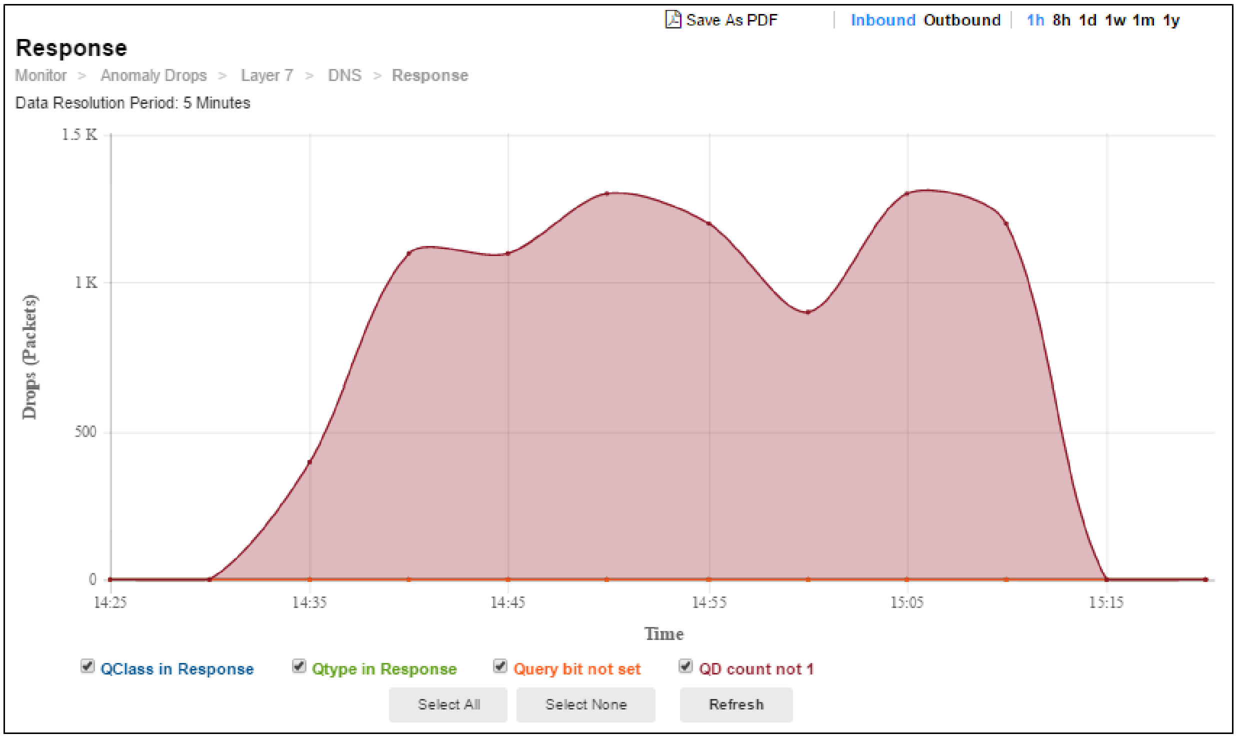

Figure 70: DNS Exploit Anomaly Drops graph

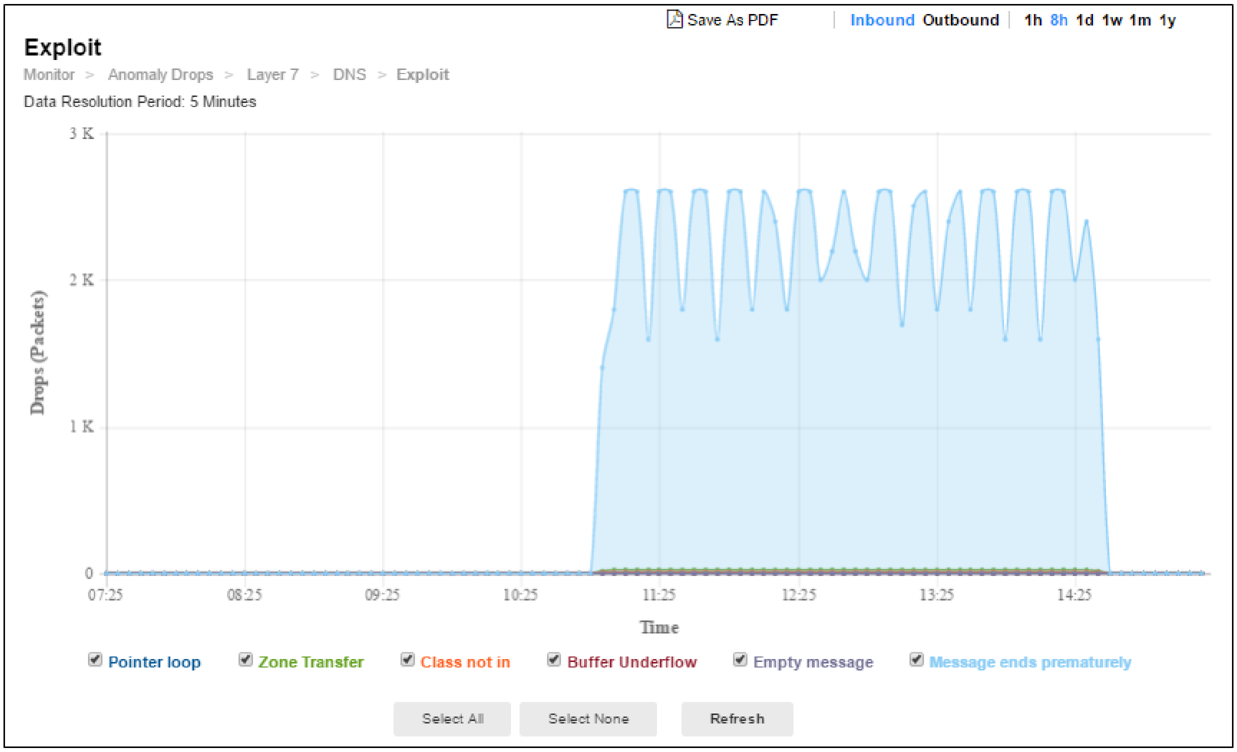

Exploit

Monitor > Anomaly Drops > Layer 7 > DNS > Exploit

Data Resolution Period: 5 Minutes

Figure 71: DNS Info Anomaly Drops graph

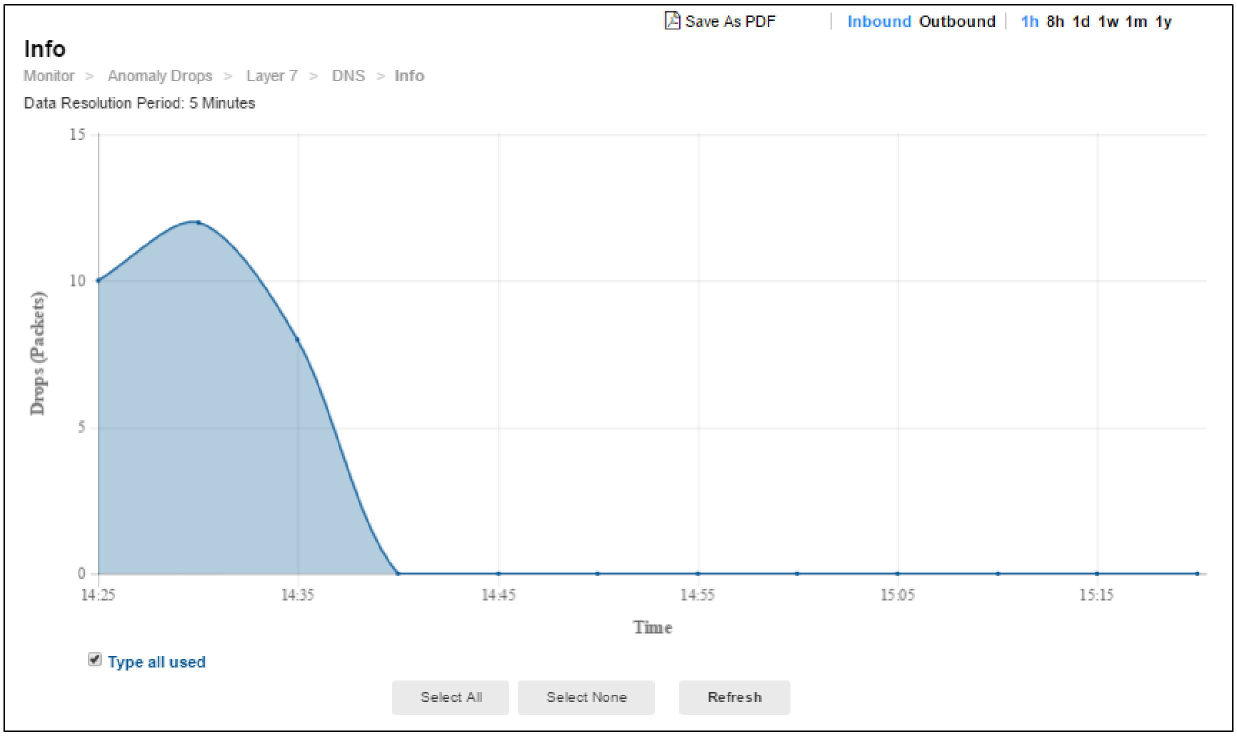

Info

Monitor > Anomaly Drops > Layer 7 > DNS > Info

Data Resolution Period: 5 Minutes

Figure 72: DNS Data Anomaly Drops graph

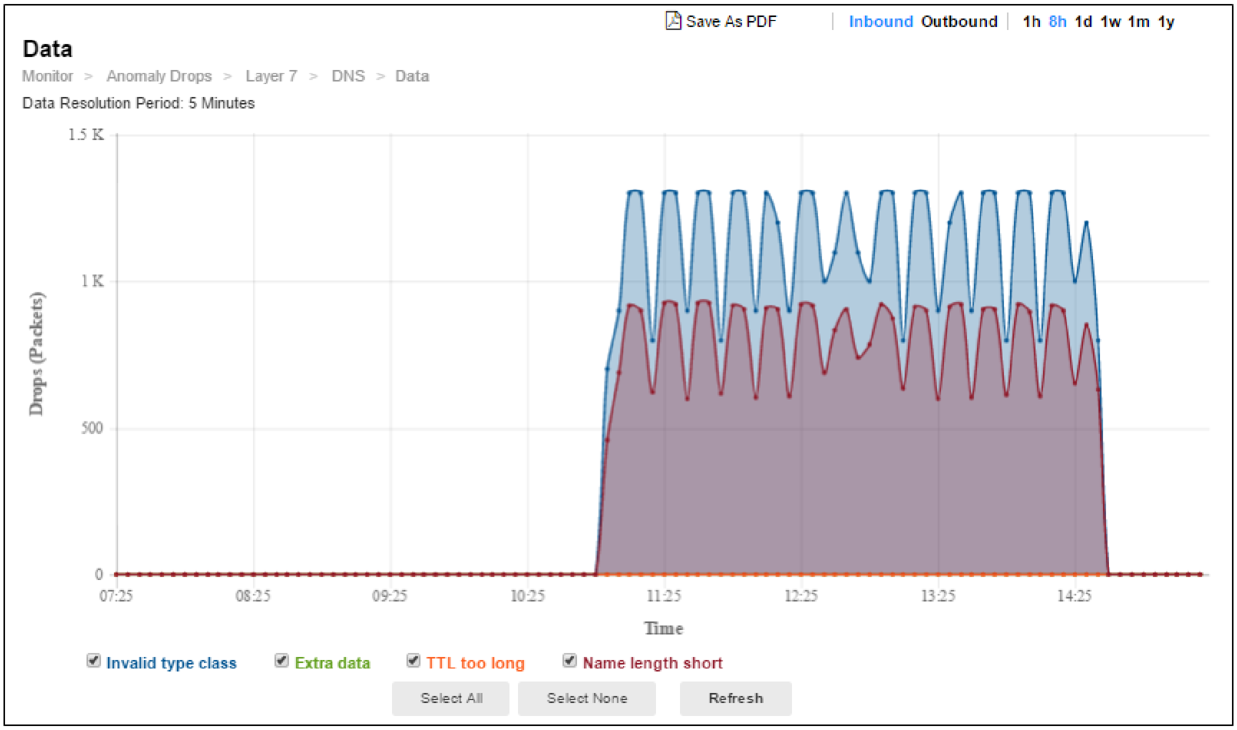

Data

Monitor > Anomaly Drops > Layer 7 > DNS > Data

Data Resolution Period: 5 Minutes

☑ Invalid type class ☑ Extra data ☑ TTL too long ☑ Name length short

Using the Hash Attack Drops graphs

You use the Hash Attacks Drops graphs to monitor drops due to hash attacks on the FortiDDoS system. If these graphs report any dropped traffic, contact Fortinet for assistance.

The Hash Attacks Drops set includes the following graphs:

- Aggregate
- Source Table
- Destination Table
- Connection Table
- DNS Query Response Table

Figure 72 shows a Hash Attack Drops graph.

You can customize the following query terms: SPP, period, direction.

Before you begin:

- You must have Read-Write permission for Log & Report settings.

To display the graphs:

- Go to Monitor > Hash Attack Drops > [Selection].

Click the graph icon 📈 in the legend to drill down from aggregate statistics to more specific queries.

Figure 73: Hash Attack Sample Graph

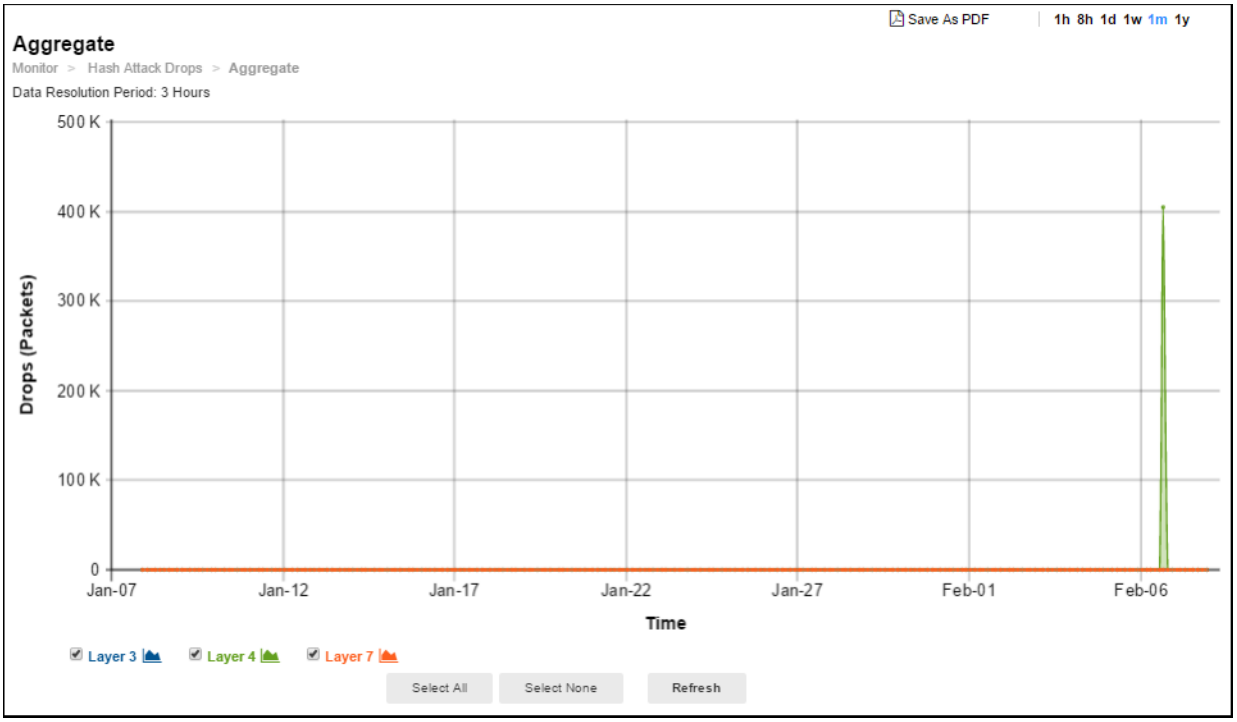

Using the Out of Memory Drops graphs

You use the Out of Memory Drops graphs to monitor drops due to memory attacks on the FortiDDoS system. If these graphs report any dropped traffic, contact Fortinet for assistance.

The Out of Memory Drops set includes the following graphs:

- Aggregate
- Source Table
- Destination Table
- Connection Table
- DNS Query Response Table

Figure 73 shows an Out of Memory Drops graph.

You can customize the following query terms: SPP, period, direction.

Before you begin:

- You must have Read-Write permission for Log & Report settings.

To display the graphs:

- Go to Monitor > Out of Memory Drops > [Selection].

Click the graph icon 📈 in the legend to drill down from aggregate statistics to more specific queries.

Figure 74: Out of Memory Drops - Sample graph

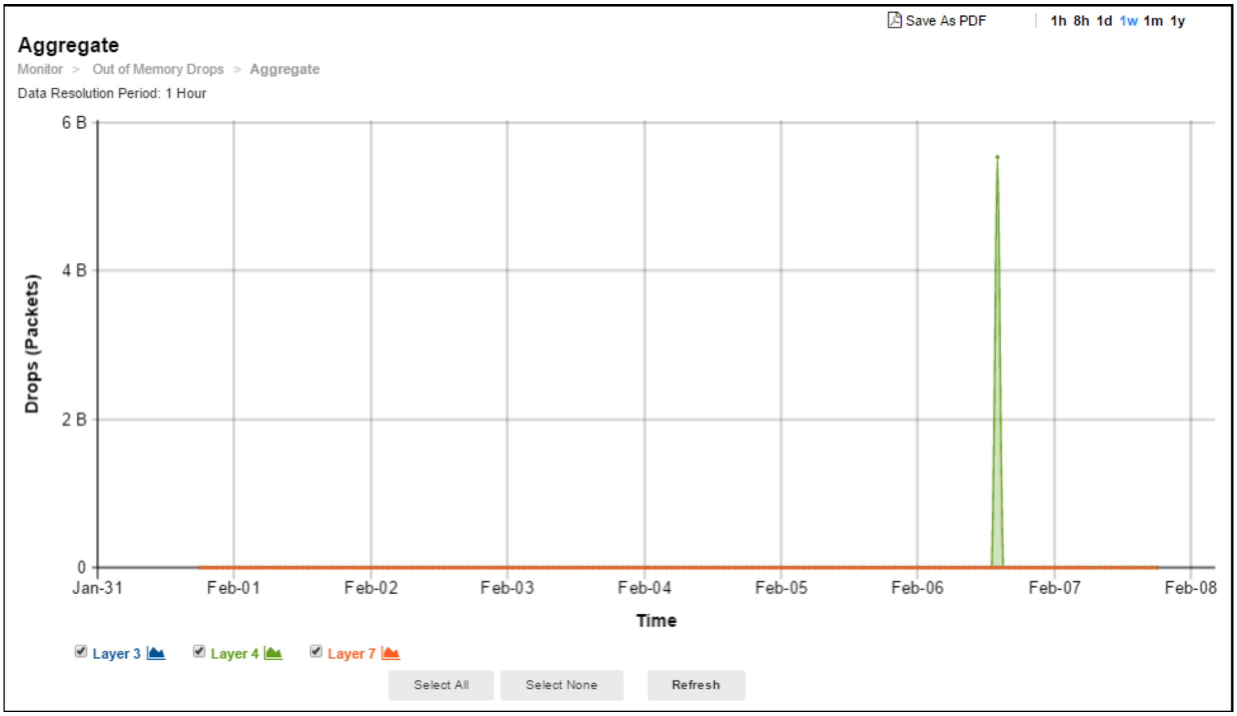

Aggregate

Monitor > Out of Memory Drops > Aggregate

Data Resolution Period: 1 Hour

Using the Layer 3 graphs

You use the Layer 3 graphs to monitor trends in Layer 3 counters. Table 51 summarizes the statistics displayed in each graph.

You can customize the following query terms: SPP, period, direction.

Before you begin:

- You must have Read-Write permission for Log & Report settings.

To display the graphs:

- Go to Monitor > Layer 3 > [Selection].

Table 51: Layer 3 graphs

Statistic	Description
Most Active Source	
Most Active Source Ingress Max Packet Rate - Packets/sec	Trend in observed ingress packet rate of the most active source address. Note that this is not necessarily a graph of the same source over time, but rather a trend of the rate for the most active source during each sampling period.
Most Active Source Egress Max Packet Rate - Packets/sec	Trend in observed egress packet rate of the most active source address. Note that this is not necessarily a graph of the same source over time, but rather a trend of the rate for the most active source during each sampling period.
Most Active Source Estimated Threshold - Packets/sec	Trend in the estimated threshold. In contrast to the configured minimum threshold, which is based on a snapshot of previously recorded data, the estimated threshold adjusts as more traffic is observed. It is almost always higher than the configured minimum threshold. It is based on algorithms designed to distinguish attack traffic from traffic increases that are the result of legitimate users accessing protected resources. Factors include historical data, trend, and seasonality.
Most Active Source Packets Dropped - Packets/sec	Trend in drops due to the effective rate limit for the **most-active-source** threshold.
Most Active Destination	
Most Active Destination Ingress Max Packet Rate - Packets/sec	Trend in observed ingress packet rate of the most active destination address. Note that this is not necessarily a graph of the same destination over time, but rather a trend of the rate for the most active destinations during each sampling period.

Statistic	Description
Most Active Destination Egress Max Packet Rate - Packets/sec	Trend in observed egress packet rate of the most active destination address. Note that this is not necessarily a graph of the same destination over time, but rather a trend of the rate for the most active destinations during each sampling period.
Most Active Destination Estimated Threshold - Packets/sec	Trend in the estimated threshold.
Most Active Destination Packets Dropped - Packets/sec	Trend in drops due to the effective rate limit for the **most-active-destination** threshold.
Count of Unique Sources	
Unique Sources	Trend in the count of unique source IP addresses in the session table. A spike in this graph indicates a possible DDoS attack.
Fragmented Packets	
Fragmented Ingress Max Packet Rate - Packets/sec	Trend in observed ingress packet rate of fragmented packets.
Fragmented Egress Max Packet Rate - Packets/sec	Trend in observed egress packet rate of fragmented packets.
Fragmented Packets Estimated Threshold - Packets/sec	Trend in the estimated threshold.
Fragmented Packets Dropped - Packets/Period	Trend in drops due to the effective rate limit for the **fragment** threshold.
Fragmented Packets Blocked - Packets/Period	Trend in drops due to blocking periods that were triggered when the system detected an attack based on the fragment counter.
Address Denied	
Geo Location Drops - Packets/Period	Trend in drops due to global ACL rules that deny traffic from specified Geolocations.
Denied Address Drops - Packets/Period	Trend in drops due to global ACL rules or SPP ACL rules that deny traffic from specified IP addresses.
IP Reputation Drops - Packets/Period	Trend in drops due to IP Reputation rules.
Local Address Anti-Spoof Denied Drops - Packets/Period	Trend in drops due to global anti-spoofing rules.

Statistic	Description
Protocols	
Protocol <Number> Ingress Max Packet Rate - Packets/sec	Trend in observed ingress packet rate for the specified protocol. A spike in this graph shows a possible protocol flood.
Protocol <Number> Egress Max Packet Rate - Packets/sec	Trend in observed egress packet rate for the specified protocol. A spike in this graph shows a possible protocol flood.
Protocol <Number> Packets Dropped - Packets/Period	Trend in packets dropped due to the effective rate limit.
Protocol <Number> Packets Blocked - Packets/Period	Trend in packets blocked due to related blocking periods.

Sample Graphs

Figure 75: Most Active Source

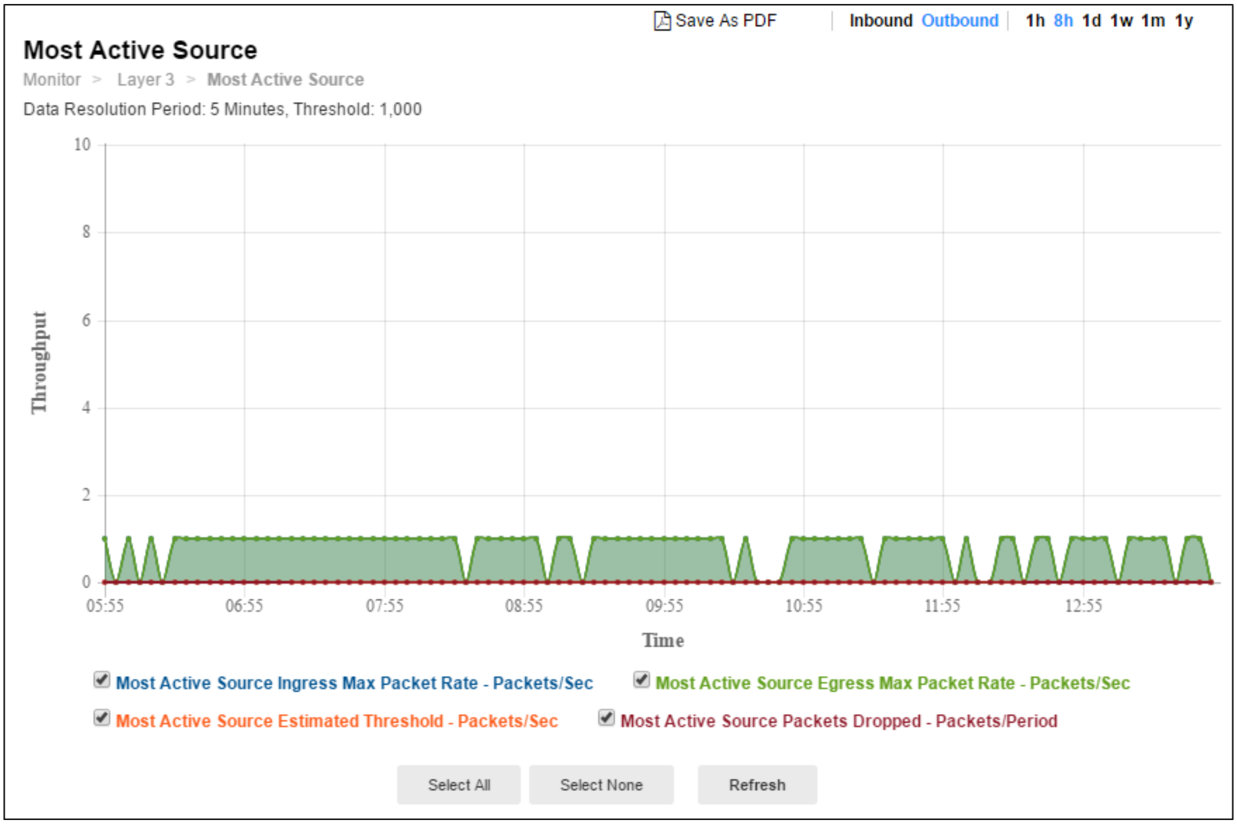

Most Active Source

Monitor > Layer 3 > Most Active Source

Data Resolution Period: 5 Minutes, Threshold: 1,000

☑ Most Active Source Ingress Max Packet Rate - Packets/Sec
☑ Most Active Source Egress Max Packet Rate - Packets/Sec
☑ Most Active Source Estimated Threshold - Packets/Sec
☑ Most Active Source Packets Dropped - Packets/Period

Figure 76: Most Active Destination

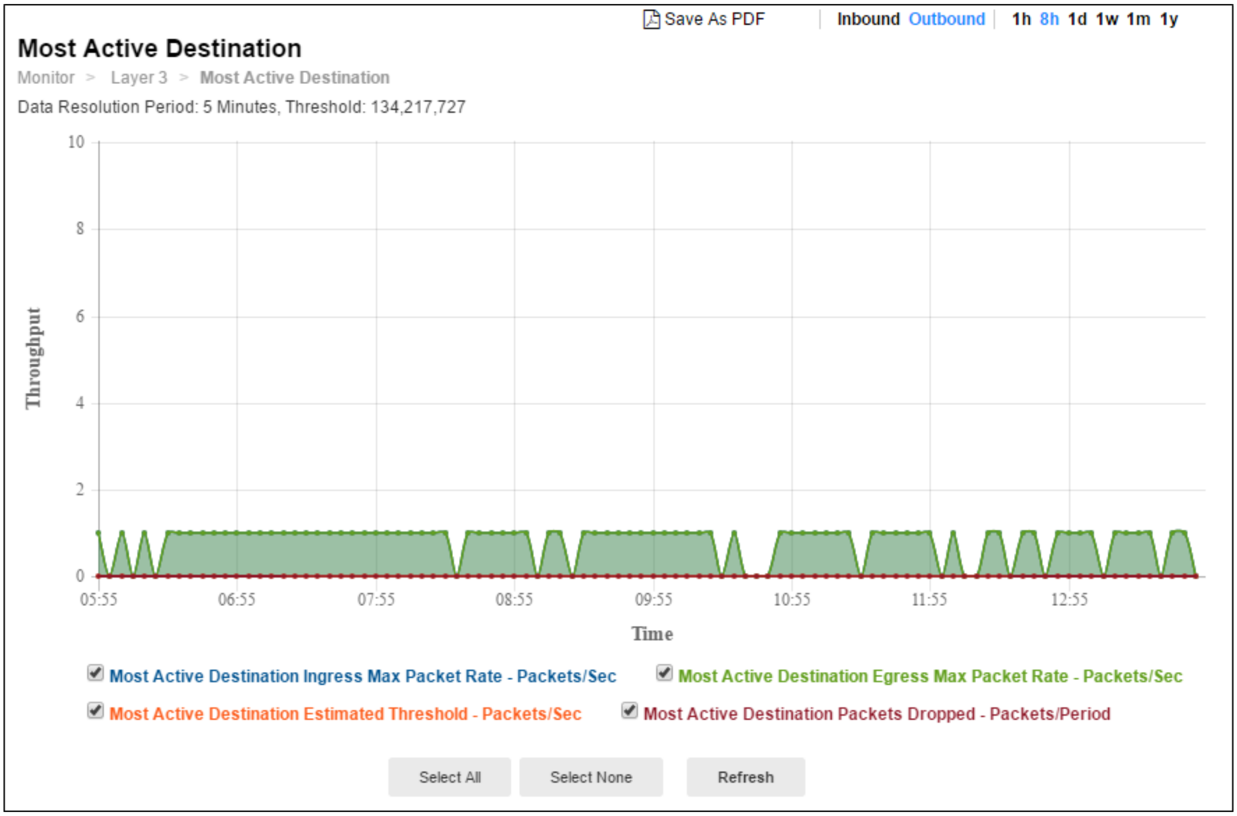

Most Active Destination

Monitor > Layer 3 > Most Active Destination

Data Resolution Period: 5 Minutes, Threshold: 134,217,727

☑ Most Active Destination Ingress Max Packet Rate - Packets/Sec ☑ Most Active Destination Egress Max Packet Rate - Packets/Sec

☑ Most Active Destination Estimated Threshold - Packets/Sec ☑ Most Active Destination Packets Dropped - Packets/Period

Figure 77: Count of unique sources

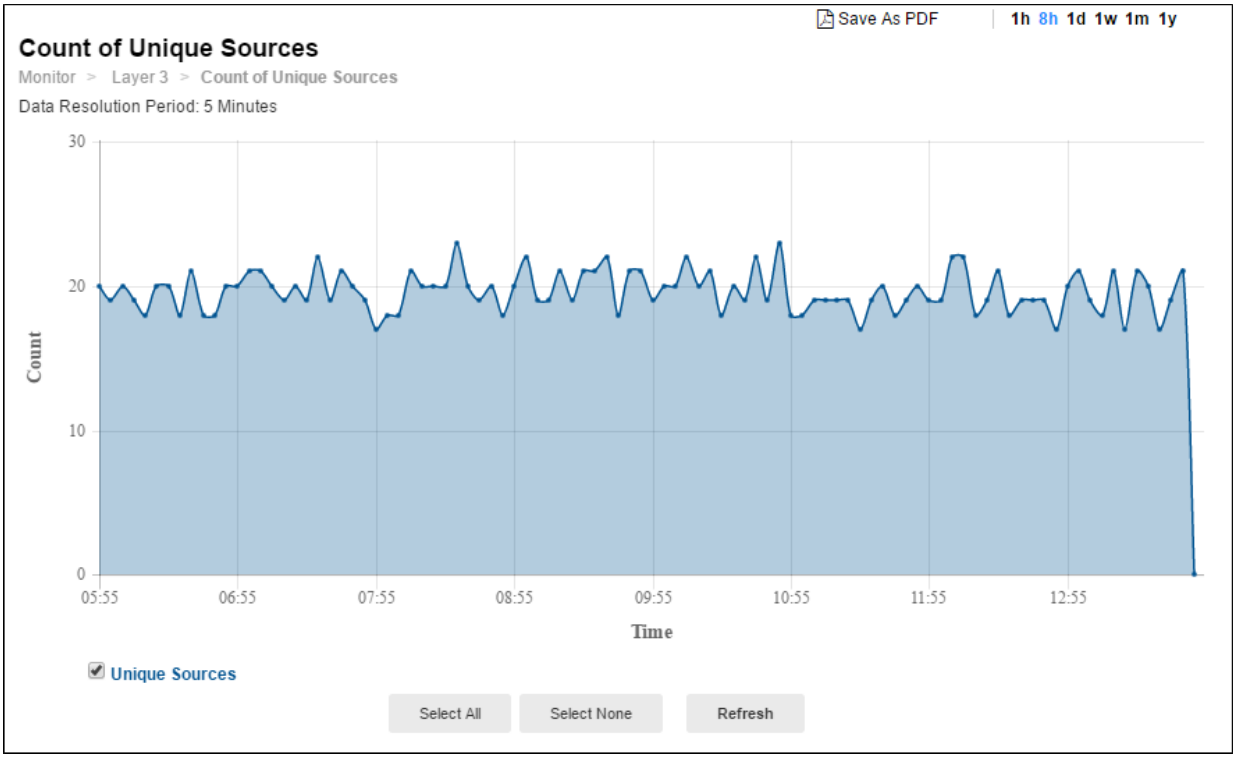

Count of Unique Sources

Monitor > Layer 3 > Count of Unique Sources

Data Resolution Period: 5 Minutes

☑ Unique Sources

Figure 78: Fragmented Packets

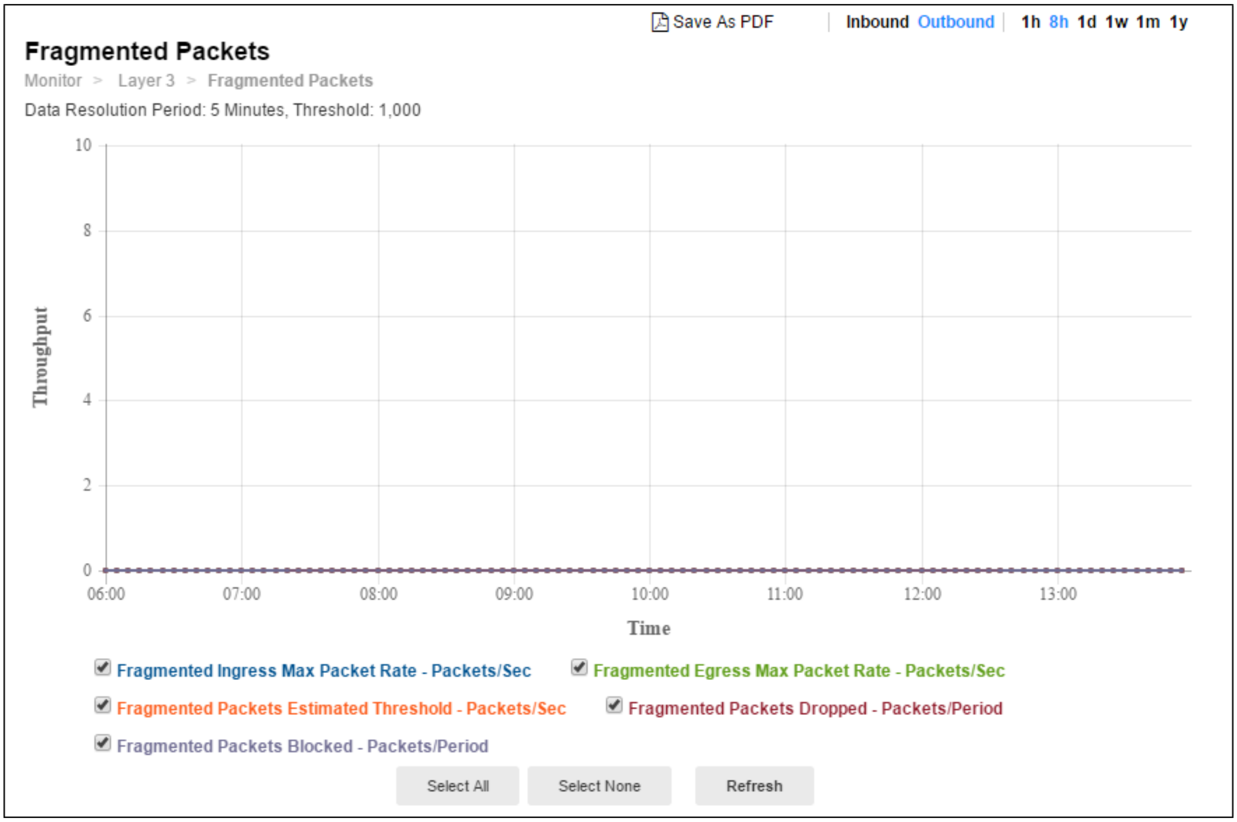

Fragmented Packets

Monitor > Layer 3 > Fragmented Packets

Data Resolution Period: 5 Minutes, Threshold: 1,000

☑ Fragmented Ingress Max Packet Rate - Packets/Sec
☑ Fragmented Egress Max Packet Rate - Packets/Sec
☑ Fragmented Packets Estimated Threshold - Packets/Sec
☑ Fragmented Packets Dropped - Packets/Period
☑ Fragmented Packets Blocked - Packets/Period

Figure 79: Address Denied

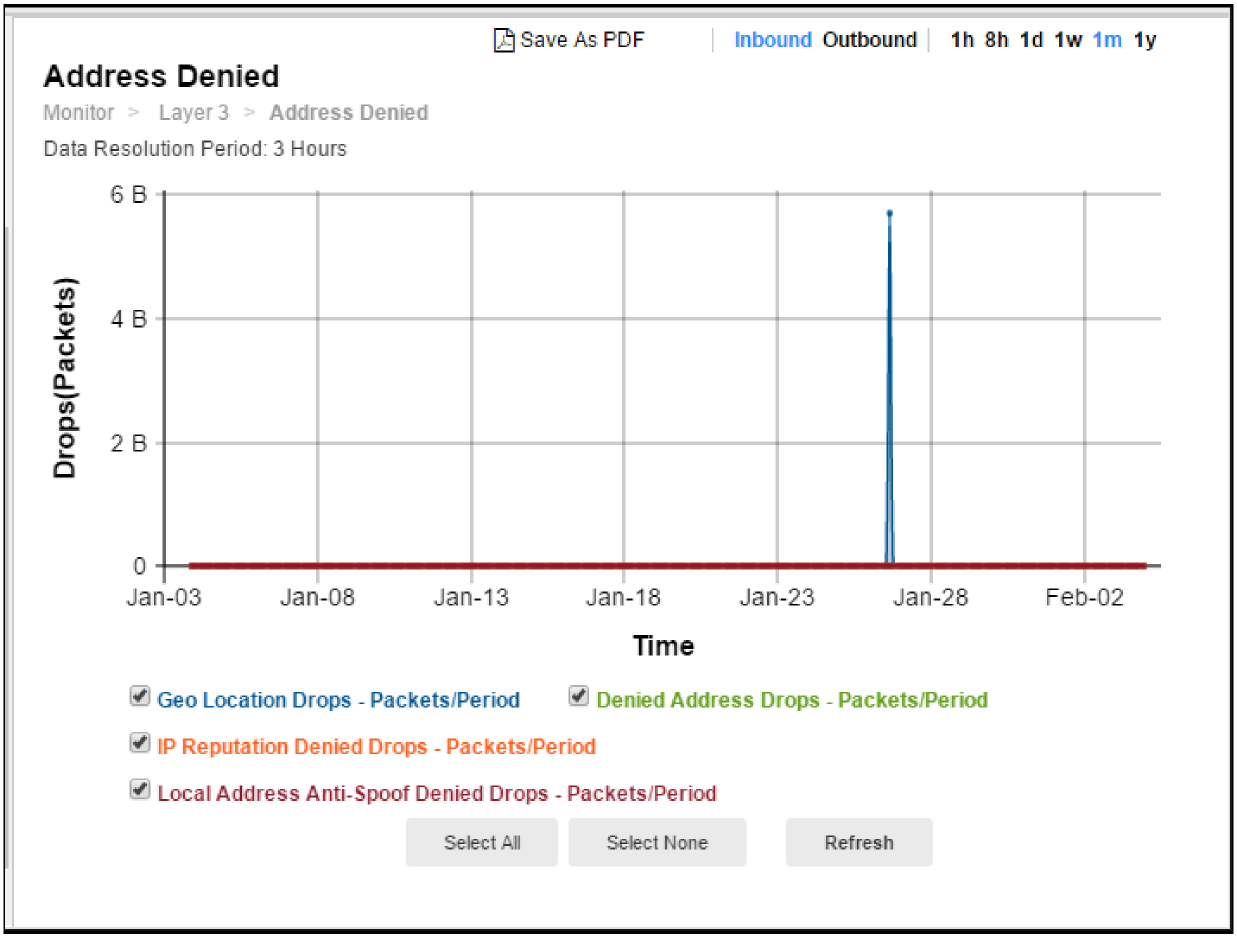

Address Denied

Monitor > Layer 3 > Address Denied

Data Resolution Period: 3 Hours

☑ Geo Location Drops - Packets/Period ☑ Denied Address Drops - Packets/Period

☑ IP Reputation Denied Drops - Packets/Period

☑ Local Address Anti-Spoof Denied Drops - Packets/Period

Figure 80: Protocols

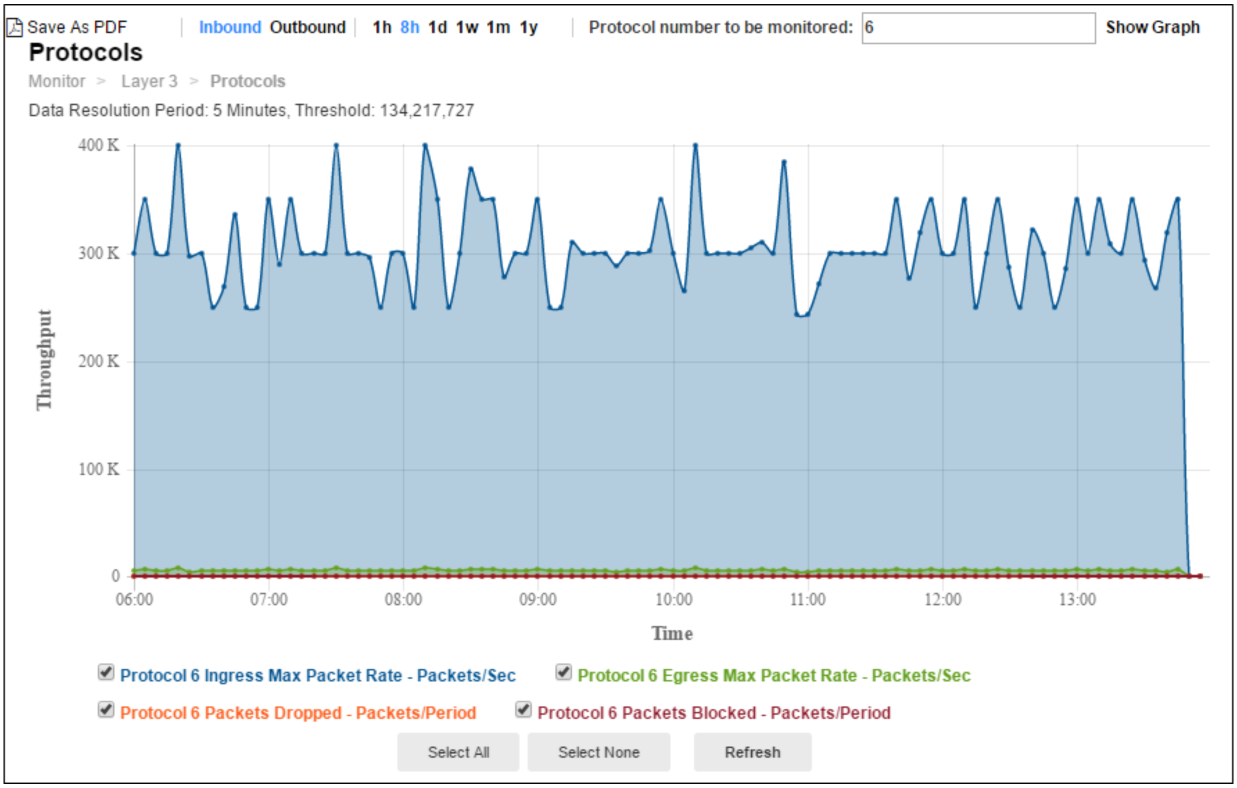

Save As PDF | Inbound Outbound | 1h 8h 1d 1w 1m 1y | Protocol number to be monitored: 6 Show Graph

Protocols

Monitor > Layer 3 > Protocols

Data Resolution Period: 5 Minutes, Threshold: 134,217,727

☑ Protocol 6 Ingress Max Packet Rate - Packets/Sec ☑ Protocol 6 Egress Max Packet Rate - Packets/Sec

☑ Protocol 6 Packets Dropped - Packets/Period ☑ Protocol 6 Packets Blocked - Packets/Period

Using the Layer 4 graphs

You use the Layer 4 graphs to monitor trends in Layer 4 counters. Table 52 summarizes the statistics displayed in each graph.

You can customize the following query terms: SPP, period, direction.

Before you begin:

- You must have Read-Write permission for Log & Report settings.

To display the graphs:

- Go to Monitor > Layer 4 > [Selection].

Table 52: Layer 4 graphs

Statistic	Description
SYN Packets	
SYN Ingress Max Packet Rate - Packets/sec	Trend in observed ingress maximum packet rate (SYN packets/second).
SYN Egress Max Packet Rate - Packets/sec	Trend in observed egress maximum packet rate (SYN packets/second).
SYN Packets Estimated Threshold - Packets/sec	Trend in the estimated threshold. In contrast to the configured minimum threshold, which is based on a snapshot of previously recorded data, the estimated threshold adjusts as more traffic is observed. It is almost always higher than the configured minimum threshold. It is based on algorithms designed to distinguish attack traffic from traffic increases that are the result of legitimate users accessing protected resources. Factors include historical data, trend, and seasonality.
SYN Packets Dropped - Packets/Period	Trend in drops due to the effective rate limit for the **syn** threshold.
SYN Per Source	
SYN Per Source Ingress Max Packet Rate - Packets/sec	Trend in observed ingress maximum rate of SYN packets from a single source. A spike in this graph shows a possible SYN attack from a single source or a few sources.
SYN Per Source Egress Max Packet Rate - Packets/sec	Trend in observed egress maximum rate of SYN packets from a single source. A spike in this graph shows a possible SYN attack from a single source or a few sources.

Statistic	Description
SYN Per Source Estimated Threshold - Packets/sec	Trend in the estimated threshold.
SYN Per Source Packets Dropped - Packets/Period	Trend in drops due to the effective rate limit for the **syn-per-source** threshold.
SYN Per Destination	
SYN Per Destination IngressMax Packet Rate - Packets/sec	Trend in observed ingress maximum rate of SYN packets to a single destination. A spike in this graph shows a possible SYN attack on a single destination or a few destinations.
SYN Per Destination Egress Max Packet Rate - Packets/sec	Trend in observed egress maximum rate of SYN packets to a single destination. A spike in this graph shows a possible SYN attack on a single destination or a few destinations.
SYN Per Destination Estimated Threshold - Packets/sec	Trend in the estimated threshold.
SYN Per Destination Drops - Packets/Period	Trend in drops due to the effective rate limit for the **syn-per-dst** threshold.
Connections Per Source	
Max Concurrent Connections Per Source	Trend in observed count of concurrent connections for the busiest source. A spike in this graph shows that a single source may be trying to establish too many connections.
Estimated Threshold for Concurrent Connections Per Source	Trend in the estimated threshold.
Concurrent Connections Dropped Per Source	Trend in drops due to the effective rate limit for the **concurrent-connections-per-source** threshold.
New Connections	
Max New Connections Establishment - Connections/sec	Trend in observed packet rate of new connections. A spike in this graph shows a possible concerted DoS or DDoS attack.
Estimated Threshold for Connections Establishment - Connections/sec	Trend in the estimated threshold.
New Connections Dropped - Drops/Period	Trend in drops due to the effective rate limit for the **new-connections** threshold.
Non-Spoofed IPs	

Statistic	Description
Entries in LIP Address Table	Trend in count of entries in the legitimate IP address table. **Note**: The legitimate IP address table is maintained and reported as a global count. (Please disregard the SPP selection menu on this page.) Therefore, the Non-Spoofed IPs graph is not reset when you reset SPP statistics.
Established Connections	
Established Connections	Trend in count of entries in the TCP state table that are in the established state (completed three-way handshake).
Number of Entries in TCP State Table	Trend in count of entries in the TCP state table, including half-open connections. If the values for the number of entries in the TCP state table are significantly higher than those for established connections, it shows a possible SYN flood attack.
Slow Connections	
Connections Dropped/Period	Connections dropped by slow connection detection.
TCP Ports	
TCP <Port> Ingress Max Packet Rate - Packets/sec	Trend in observed ingress maximum packet rate to the specified port. A spike in this graph shows a possible port flood.
TCP <Port> Egress Max Packet Rate - Packets/sec	Trend in observed egress maximum packet rate to the specified port. A spike in this graph shows a possible port flood.
TCP <Port> Packets Dropped - Packets/sec	Trend in packets dropped due to the effective rate limit.
TCP <Port> Packets Blocked - Packets/Period	Trend in packets blocked due to related blocking periods.
UDP Ports	
UDP <Port> Ingress Max Packet Rate - Packets/sec	Trend in observed ingress maximum packet rate to the specified port. A spike in this graph shows a possible port flood.
UDP <Port> Egress Max Packet Rate - Packets/sec	Trend in observed egress maximum packet rate to the specified port. A spike in this graph shows a possible port flood.
UDP <Port> Packets Dropped - Packets/Period	Trend in packets dropped due to the effective rate limit.
UDP <Port> Packets Blocked-Packets/Period	Trend in packets blocked due to related blocking periods.

Statistic	Description
ICMP Type/Code	
ICMP <Index> Egress Max Packet Rate - Packets/sec	Trend in observed egress maximum packet rate of packets with the specified ICMP type/code. A spike in this graph shows a possible ICMP flood.
ICMP <Index> Packets Dropped - Packets/Period	Trend in packets dropped due to the effective rate limit.
ICMP <Index> Packets Blocked - Packets/Period	Trend in packets blocked due to related blocking periods.

Sample Graphs

Figure 81: SYN Packets

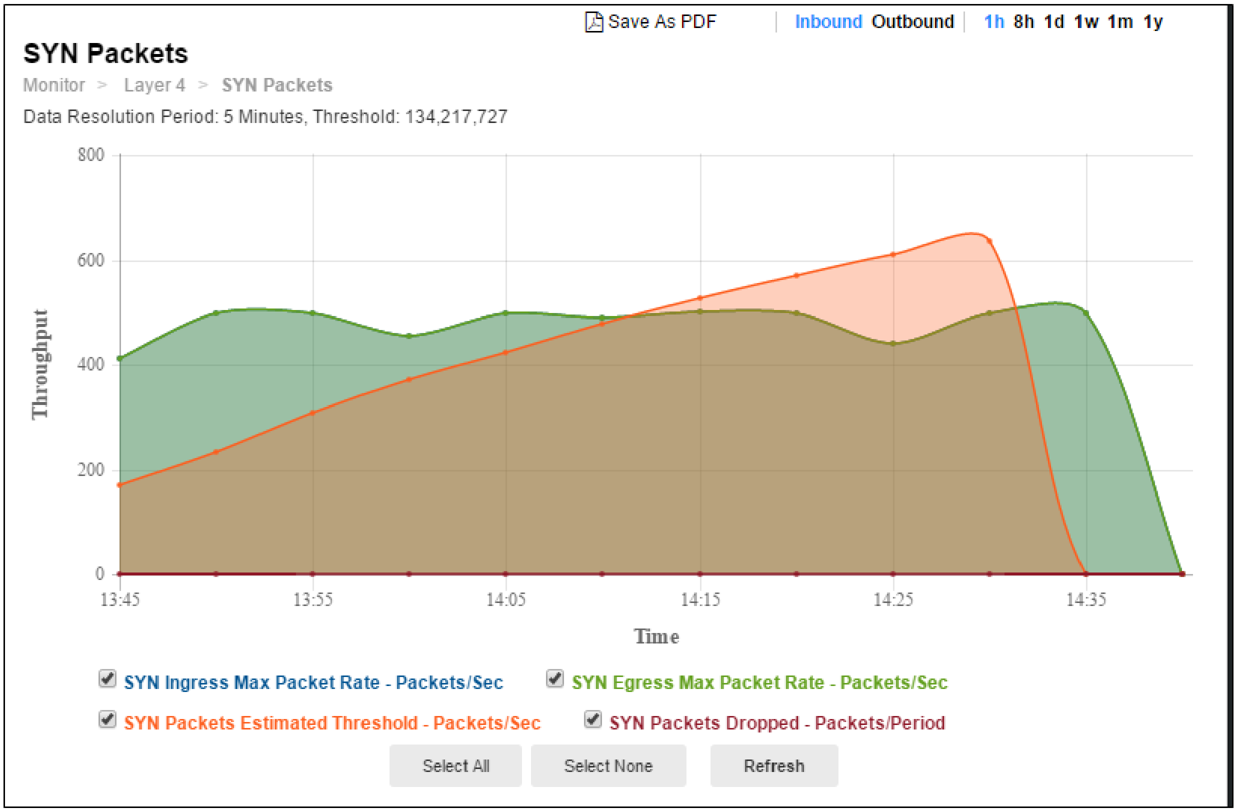

SYN Packets

Monitor > Layer 4 > SYN Packets

Data Resolution Period: 5 Minutes, Threshold: 134,217,727

Save As PDF
Inbound
Outbound
1h
8h
1d
1w
1m
1y
800
600
400
200
0
Throughput
13:45
13:55
14:05
14:15
14:25
14:35
Time
☑ SYN Ingress Max Packet Rate - Packets/Sec
☑ SYN Egress Max Packet Rate - Packets/Sec
☑ SYN Packets Estimated Threshold - Packets/Sec
☑ SYN Packets Dropped - Packets/Period
Select All
Select None
Refresh

Figure 82: SYN Per Source

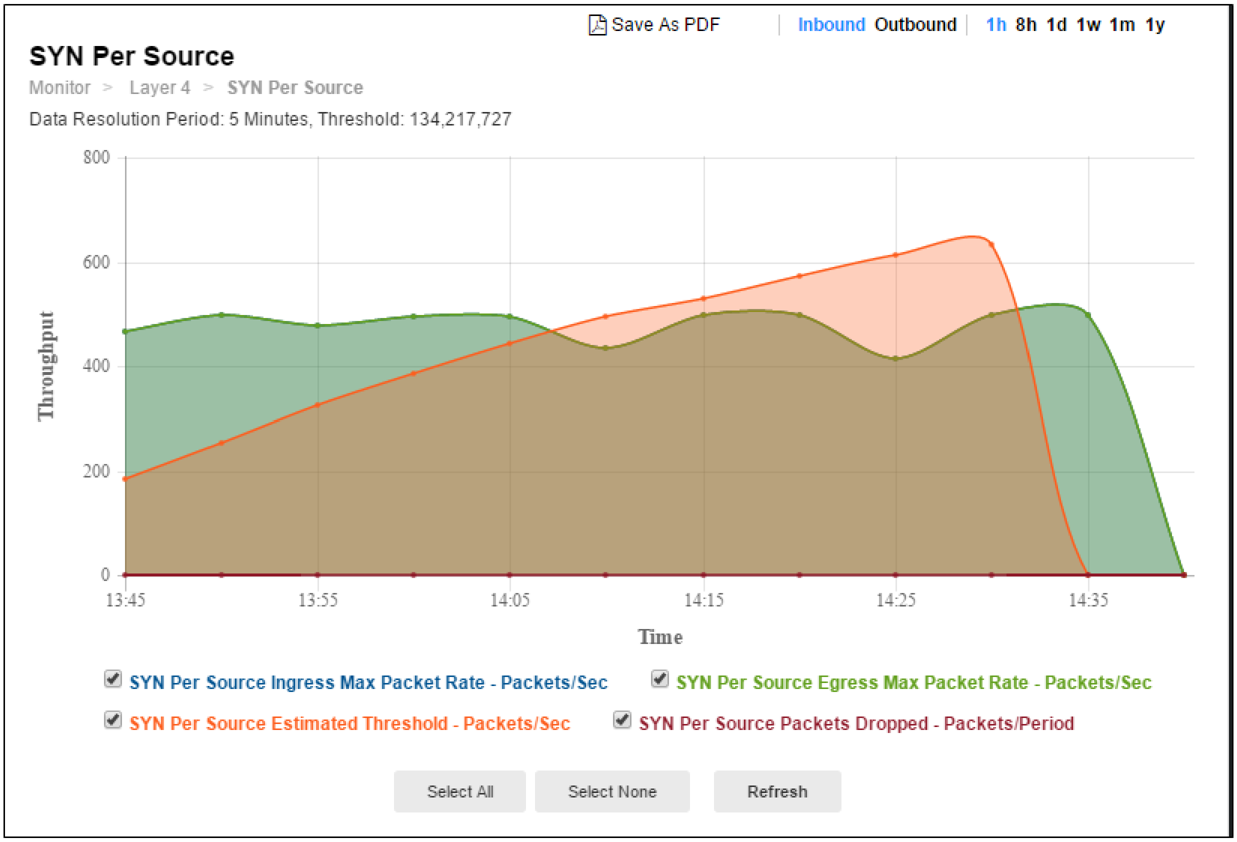

Save As PDF | Inbound Outbound | 1h 8h 1d 1w 1m 1y

SYN Per Source

Monitor > Layer 4 > SYN Per Source

Data Resolution Period: 5 Minutes, Threshold: 134,217,727

☑ SYN Per Source Ingress Max Packet Rate - Packets/Sec ☑ SYN Per Source Egress Max Packet Rate - Packets/Sec

☑ SYN Per Source Estimated Threshold - Packets/Sec ☑ SYN Per Source Packets Dropped - Packets/Period

Figure 83: SYN Per Destination

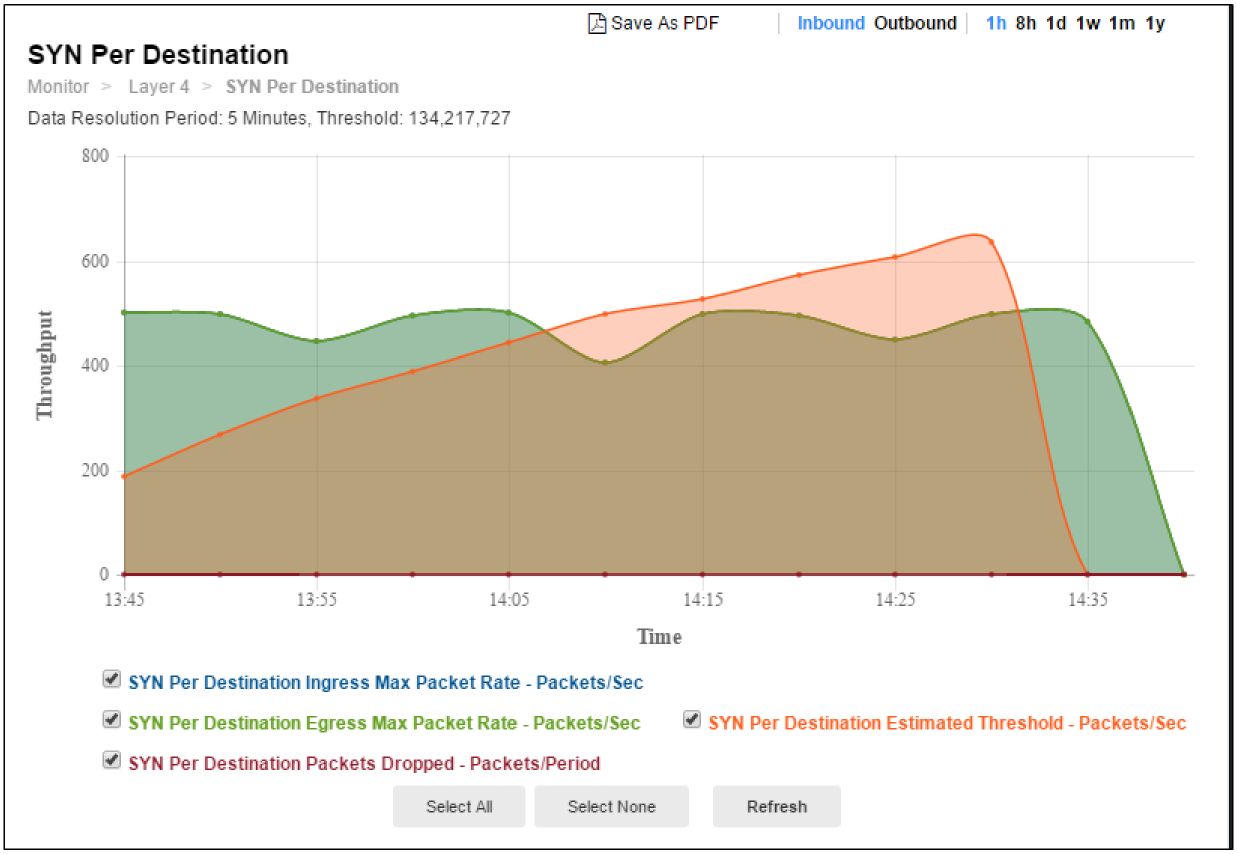

SYN Per Destination

Monitor > Layer 4 > SYN Per Destination

Data Resolution Period: 5 Minutes, Threshold: 134,217,727

☑ SYN Per Destination Ingress Max Packet Rate - Packets/Sec

☑ SYN Per Destination Egress Max Packet Rate - Packets/Sec

☑ SYN Per Destination Estimated Threshold - Packets/Sec

☑ SYN Per Destination Packets Dropped - Packets/Period

Figure 84: Connection Per Source

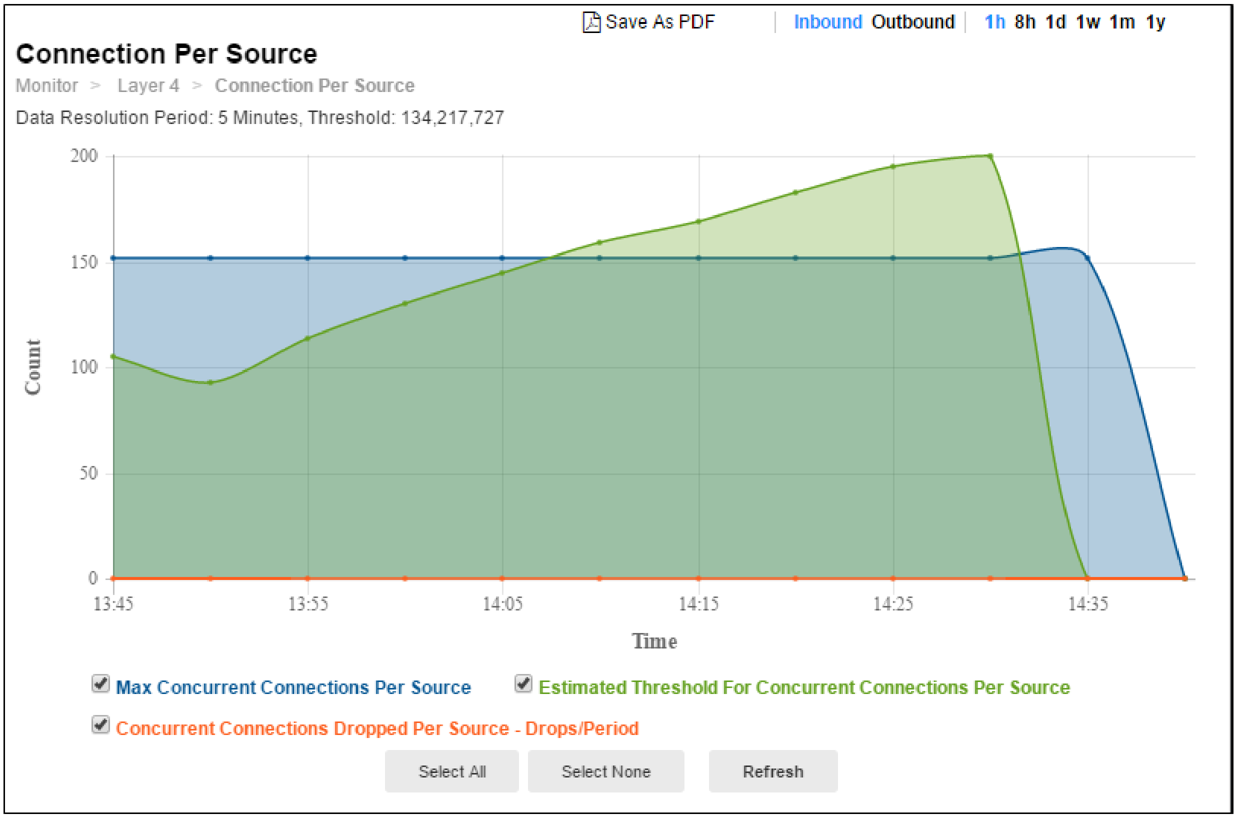

Save As PDF | Inbound Outbound | 1h 8h 1d 1w 1m 1y

Connection Per Source

Monitor > Layer 4 > Connection Per Source

Data Resolution Period: 5 Minutes, Threshold: 134,217,727

☑ Max Concurrent Connections Per Source ☑ Estimated Threshold For Concurrent Connections Per Source

☑ Concurrent Connections Dropped Per Source - Drops/Period

Figure 85: New Connections

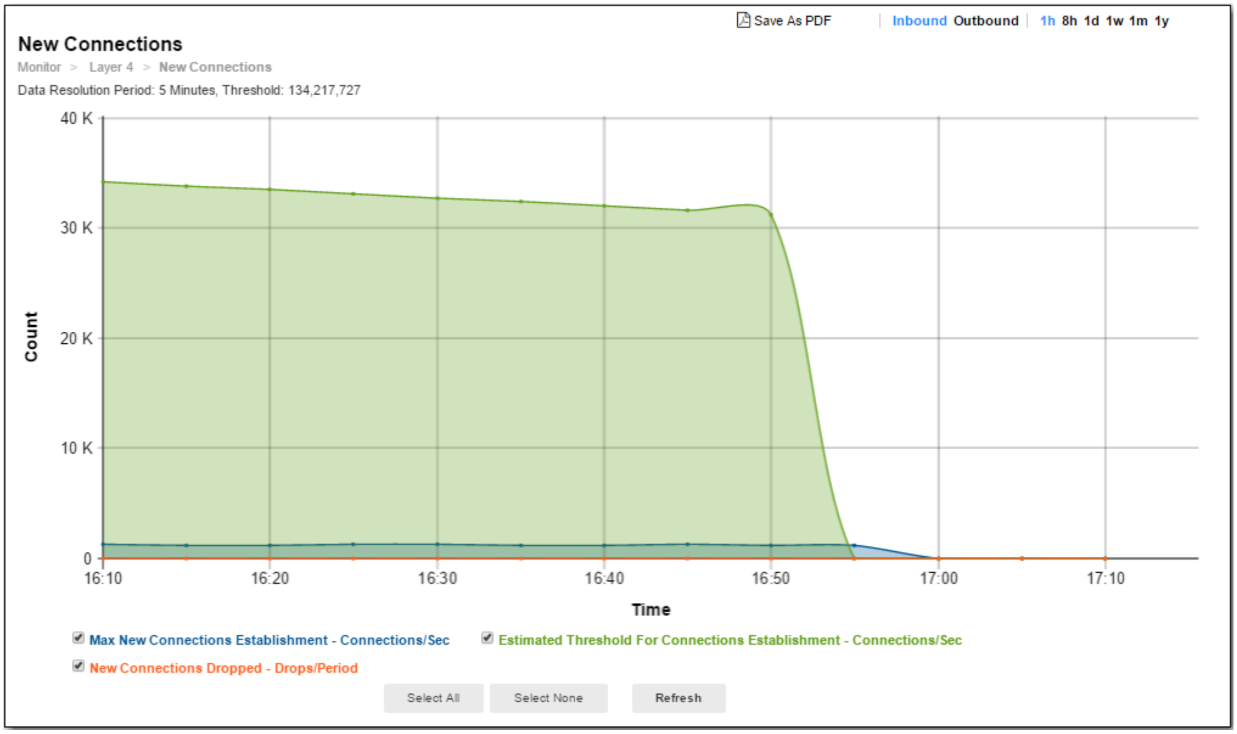

New Connections

Monitor > Layer 4 > New Connections

Data Resolution Period: 5 Minutes, Threshold: 134,217,727

Save As PDF
Inbound Outbound
1h 8h 1d 1w 1m 1y
40 K
30 K
20 K
10 K
0
Count
16:10
16:20
16:30
16:40
16:50
17:00
17:10
Time
☑ Max New Connections Establishment - Connections/Sec
☑ Estimated Threshold For Connections Establishment - Connections/Sec
☑ New Connections Dropped - Drops/Period
Select All
Select None
Refresh

Figure 86: Non-Spoofed IPs

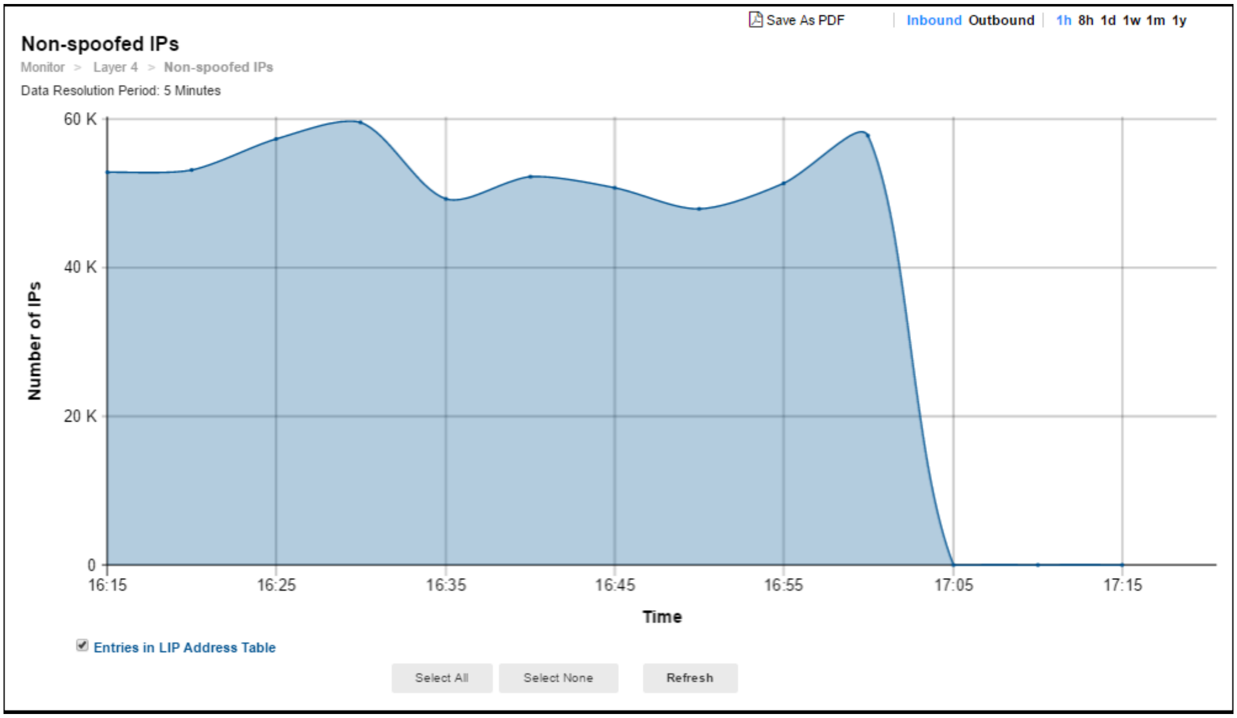

Non-spoofed IPs

Monitor > Layer 4 > Non-spoofed IPs

Data Resolution Period: 5 Minutes

Figure 87: Established Connections

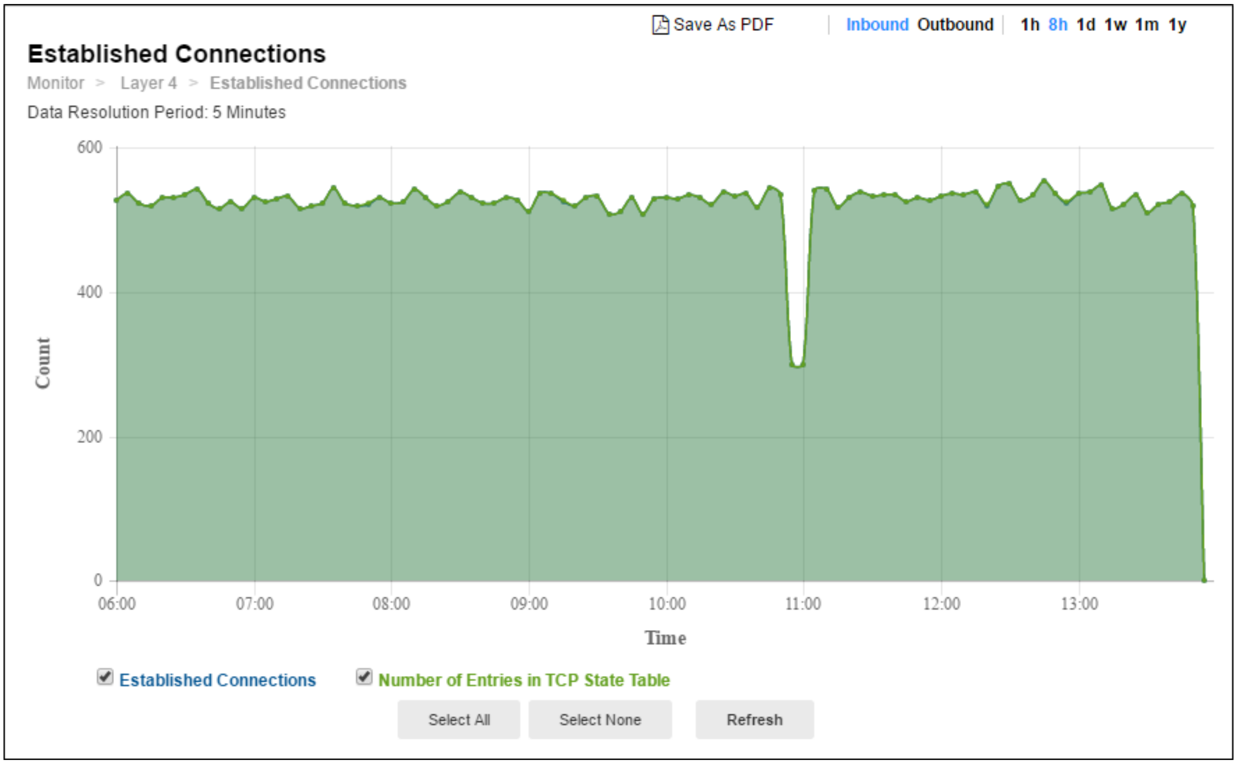

Established Connections

Monitor > Layer 4 > Established Connections

Data Resolution Period: 5 Minutes

Figure 88: Slow Connections

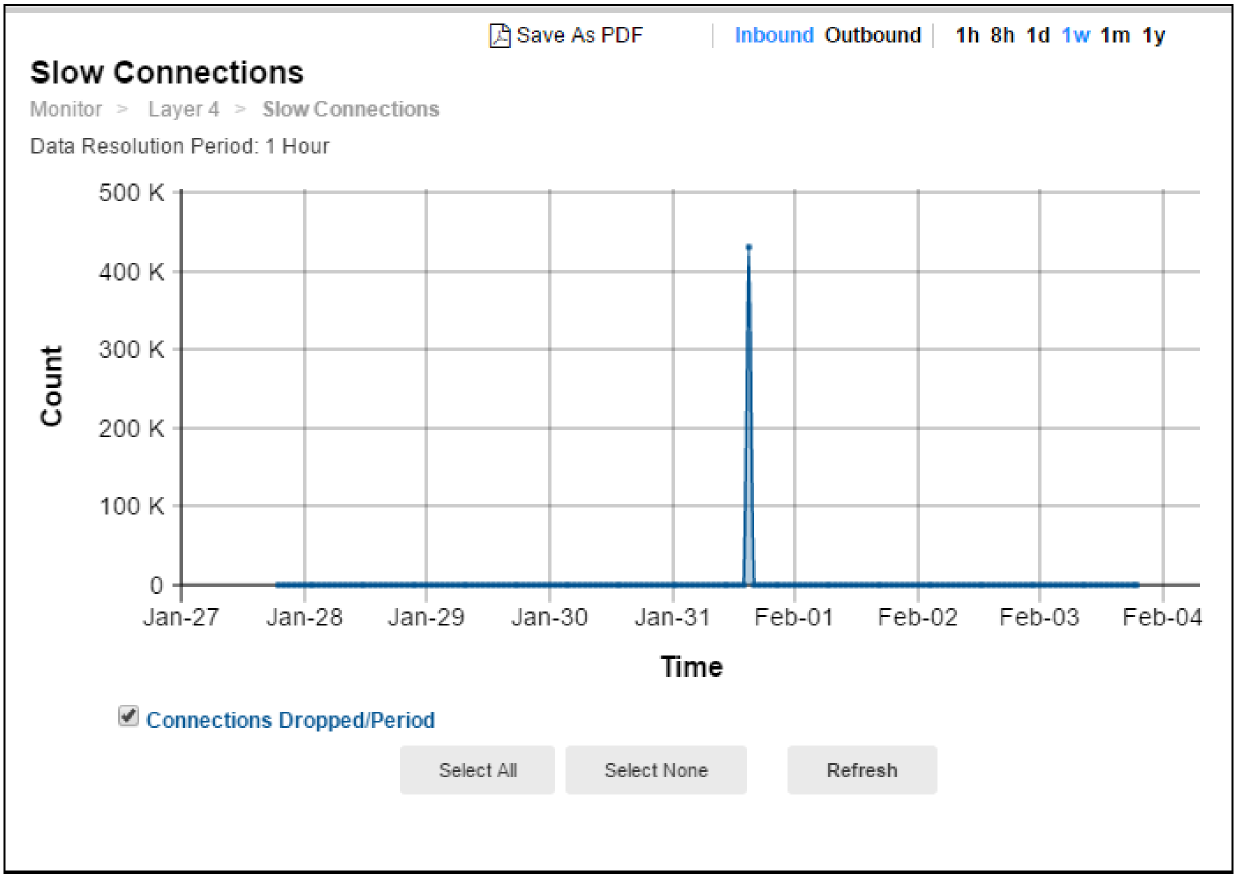

Figure 89: TCP Port

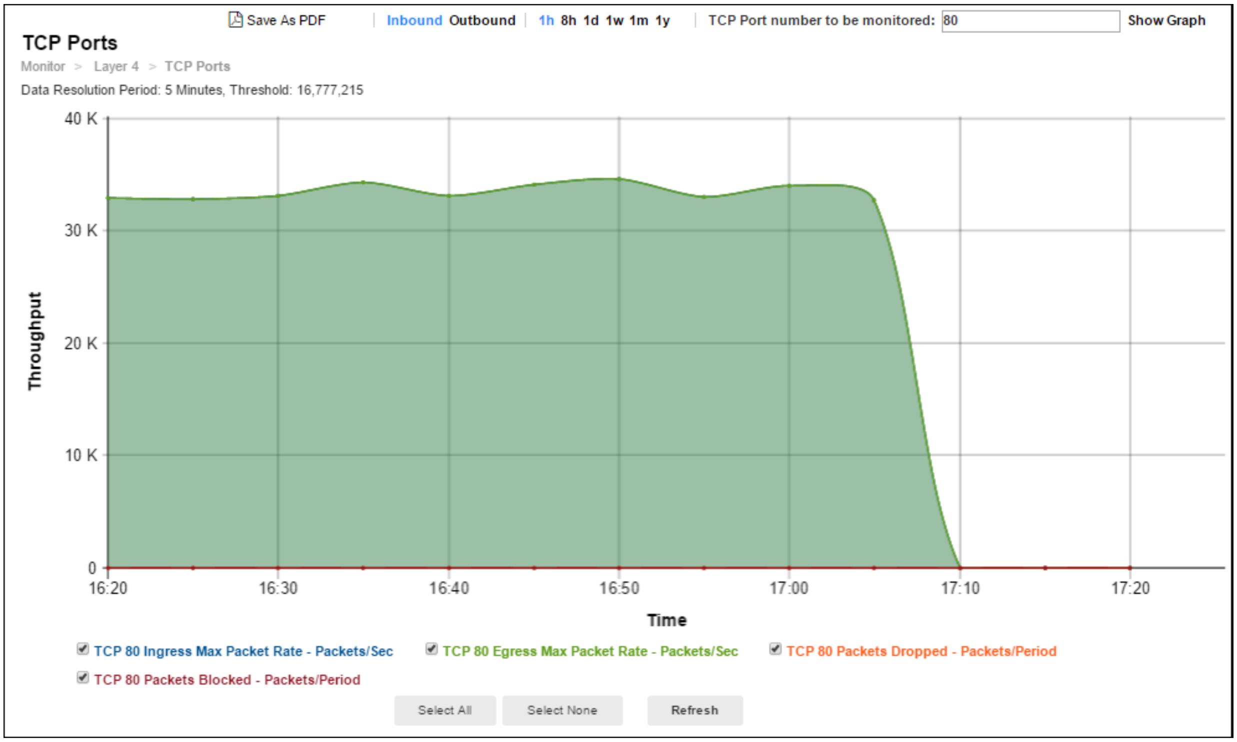

TCP Ports

Monitor > Layer 4 > TCP Ports

Data Resolution Period: 5 Minutes, Threshold: 16,777,215

TCP Port number to be monitored: 80

☑ TCP 80 Ingress Max Packet Rate - Packets/Sec

☑ TCP 80 Egress Max Packet Rate - Packets/Sec

☑ TCP 80 Packets Dropped - Packets/Period

☑ TCP 80 Packets Blocked - Packets/Period

Figure 90: UDP Ports

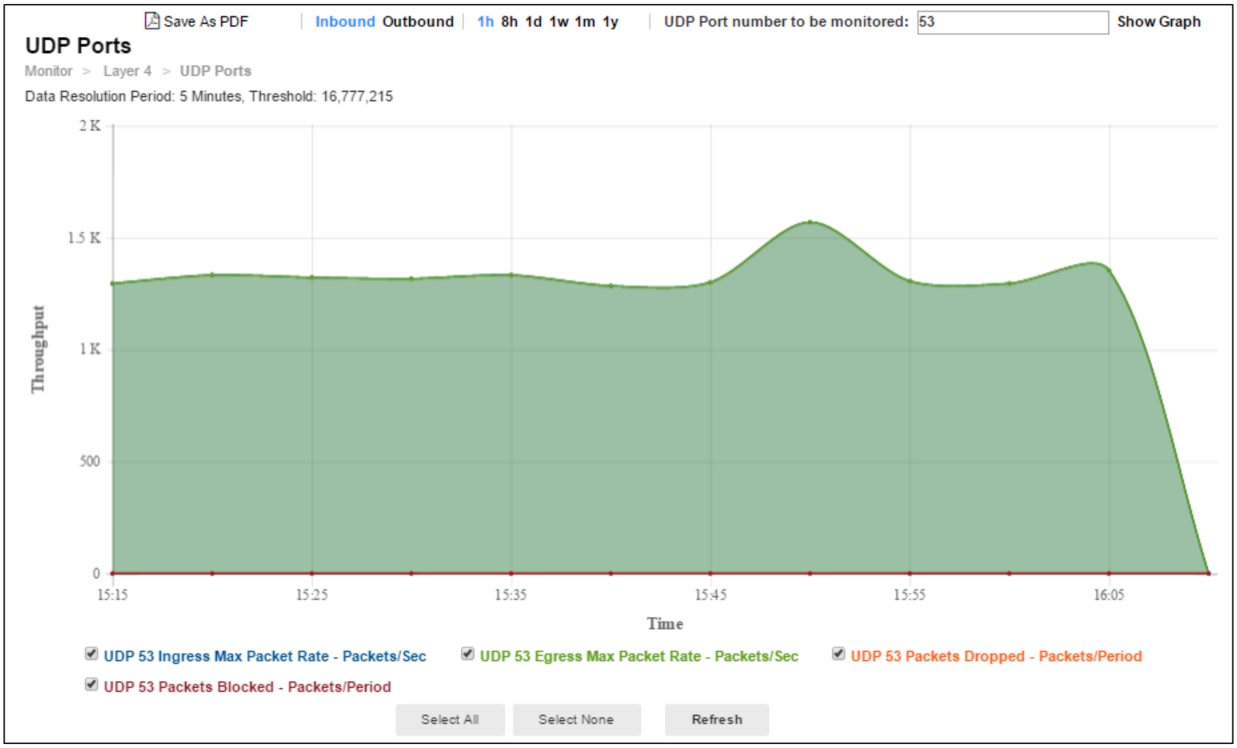

UDP Ports

Monitor > Layer 4 > UDP Ports

Data Resolution Period: 5 Minutes, Threshold: 16,777,215

UDP Port number to be monitored: 53

Save As PDF
Inbound Outbound
1h 8h 1d 1w 1m 1y
Show Graph
2 K
1.5 K
1 K
500
0
Throughput
15:15
15:25
15:35
15:45
15:55
16:05
Time
☑ UDP 53 Ingress Max Packet Rate - Packets/Sec
☑ UDP 53 Egress Max Packet Rate - Packets/Sec
☑ UDP 53 Packets Dropped - Packets/Period
☑ UDP 53 Packets Blocked - Packets/Period
Select All
Select None
Refresh

Figure 91: ICMP Types/Codes

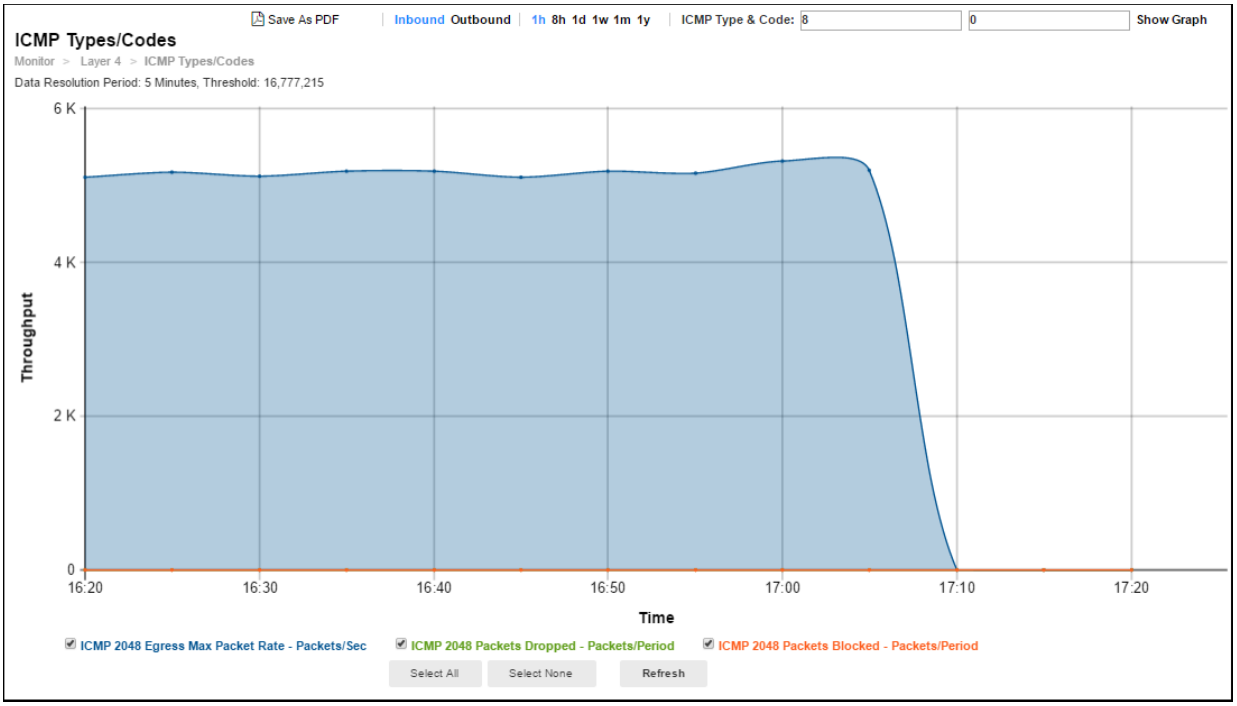

ICMP Types/Codes

Monitor > Layer 4 > ICMP Types/Codes

Data Resolution Period: 5 Minutes, Threshold: 16,777,215

ICMP Type & Code: 8 0

Using the Layer 7 graphs

You use the Layer 7 graphs to monitor trends in Layer 7 counters. Table 53 summarizes the statistics displayed in the graphs.

You can customize the following query terms: SPP, period, direction.

Before you begin:

- You must have Read-Write permission for Log & Report settings.

To display the graphs:

- Go to Monitor > Layer 7 > [Selection].

Table 53: Layer 7 graphs

Statistic	Description
HTTP Methods	
HTTP Method Ingress Max Packet Rate - Packets/sec	Trend in observed ingress maximum rate for the selected HTTP method. The following methods are available: • GET • HEAD • OPTIONS • TRACE • POST • PUT • DELETE • CONNECT
HTTP Method Egress Max Packet Rate - Packets/sec	Trend in observed egress maximum rate for the selected HTTP method. The following methods are available: • GET • HEAD • OPTIONS • TRACE • POST • PUT • DELETE • CONNECT

Statistic	Description
HTTP Method Estimated Threshold - Packets/sec	Trend in the estimated threshold. In contrast to the configured minimum threshold, which is based on a snapshot of previously recorded data, the estimated threshold adjusts as more traffic is observed. It is almost always higher than the configured minimum threshold. It is based on algorithms designed to distinguish attack traffic from traffic increases that are the result of legitimate users accessing protected resources. Factors include historical data, trend, and seasonality.
HTTP Method Packets Dropped	Trend in packets dropped due to the effective rate limit.
HTTP Methods per Source	
HTTP Methods per Source Ingress Max Packet Rate - Packets/Sec	Trend in observed ingress maximum rate for the HTTP method from a single source.
HTTP Methods per Source Egress Max Packet Rate - Packets/Sec	Trend in observed egress maximum rate for the HTTP method from a single source.
HTTP Methods per Source Estimated Threshold - Packets/Sec	Trend in the estimated threshold.
HTTP Methods per Source Packets Dropped - Packets/Period	Drops due to HTTP method per source threshold.
HTTP URLs	
URL <Index> Egress Max Packet Rate- Packets/Sec	Trend in observed egress maximum packet rate of packets with the specified request URL. A spike in this graph shows a possible URL flood.
URL <Index> Packets Dropped - Packets/Period	Trend in packets dropped due to the effective rate limit.
URL <Index>Packets Blocked- Packets/Period	Trend in packets blocked due to related blocking periods.
HTTP Hosts	

Statistic	Description
Host <Index> Ingress Max Traffic - Packets/Sec	Trend in observed ingress maximum packet rate of packets with the specified Host header. A spike in this graph shows a possible Host flood.
Host <Index> Egress Max Traffic - Packets/Sec	Trend in observed egress maximum packet rate of packets with the specified Host header. A spike in this graph shows a possible Host flood.
Host <Inde - Packets/Secx> Packets Dropped - Packets/Period	Trend in packets dropped due to the effective rate limit.
Host <Index> Packets Blocked - Packets/Period	Trend in packets blocked due to related blocking periods.
HTTP Referers	
Referer <Index> Ingress Max Traffic - Packets/Sec	Trend in observed ingress maximum packet rate of packets with the specified Referer header. A spike in this graph shows a possible Referer flood.
Referer <Index> Egress Max Traffic - Packets/Sec	Trend in observed egress maximum packet rate of packets with the specified Referer header. A spike in this graph shows a possible Referer flood.
Referer <Index> Packets Dropped - Packets/Period	Trend in packets dropped due to the effective rate limit.
Referer <Index> Packets Blocked - Packets/Period	Trend in packets blocked due to related blocking periods.
HTTP Cookies	
Cookie <Index> Ingress Max Traffic - Packets/Sec	Trend in observed ingress maximum packet rate of packets with the specified cookie header. A spike in this graph shows a possible Cookie flood.
Cookie <Index> Egress Max Traffic - Packets/Sec	Trend in observed egress maximum packet rate of packets with the specified cookie header. A spike in this graph shows a possible Cookie flood.
Cookie <Index> Packets Dropped - Packets/Period	Trend in packets dropped due to the effective rate limit.

Statistic	Description
Cookie <Index> Packets Blocked - Packets/Period	Trend in packets blocked due to related blocking periods.
HTTP User Agents	
User Agents <Index> Ingress Max Traffic - Packets/Sec	Trend in observed ingress maximum packet rate of packets with the specified user agent header. A spike in this graph shows a possible User Agent flood.
User Agents <Index> Egress Max Traffic - Packets/Sec	Trend in observed egress maximum packet rate of packets with the specified user agent header. A spike in this graph shows a possible User Agent flood.
User Agents <Index> Packets Dropped - Packets/Period	Trend in packets dropped due to the effective rate limit.
User Agents <Index> Packets Blocked - Packets/Period	Trend in packets blocked due to related blocking periods.
DNS Query	
Query Ingress Max Packet Rate/Sec	Trend in observed ingress maximum rate for DNS queries.
Query Egress Max Packet Rate/Sec	Trend in observed egress maximum rate for DNS queries.
Query Estimated Threshold/Sec	Trend in the estimated threshold.
TCP Query Dropped - Packets/Period	Drops due to the dns-query threshold.
Query Blocked (Restricted Subnets) - Packets/Period	Drops due to an ACL.
Query Blocked (ACLs) - Packets/Period	Drops due to an ACL.
DNS Query Per Source	
Query Per Source Ingress Max Packet Rate/Sec	Trend in observed ingress maximum rate for DNS queries from a single source.

Statistic	Description
Query Per Source Egress Max Packet Rate/Sec	Trend in observed egress maximum rate for DNS queries from a single source. **Note**: If **Block Identified Sources** setting is disabled under Protection Profiles > SPP Settings > DNS Feature Controls tab, DNS per-source thresholds (dns-query-per-source and dns-packet-track-per-source) are not tracked and this graph will not be updated.
Query Per Source Estimated Threshold/Sec	Trend in the estimated threshold.
Query Per Source Packets Dropped/Period	Drops due to the dns-query-per-src threshold.
DNS Suspicious Sources	
Packet-Track Per Source Ingress Max Packet Rate/Sec	Trend in observed ingress maximum rate for a source that demonstrates suspicious activity (a score based on heuristics that count fragmented packets, response not found in DQRM, or queries that generate responses with RCODE other than 0).
Packet-Track Per Source Egress Max Packet Rate/Sec	Trend in observed egress maximum rate for a source that demonstrates suspicious activity (a score based on heuristics that count fragmented packets, response not found in DQRM, or queries that generate responses with RCODE other than 0). **Note**: If **Block Identified Sources** setting is disabled under Protection Profiles > SPP Settings > DNS Feature Controls tab, DNS per-source thresholds (dns-query-per-source and dns-packet-track-per-source) are not tracked and this graph will not be updated.
Packet-Track Per Source Estimated Threshold/Sec	Trend in the estimated threshold.
Packet-Track Per Source Packets Dropped/Period	Drops due to the dns-packet-track-per-src threshold.
DNS Question Count	
Question Ingress Max Packet Rate/Sec	Trend in observed ingress maximum rate for DNS queries.
Question Egress Max Packet Rate/Sec	Trend in observed egress maximum rate for DNS queries.
Question Estimated Threshold/Sec	Trend in the estimated threshold.
TCP Question Dropped/Period	Drops due to the dns-question-count threshold.

Statistic	Description
DNS QType MX Count	
MX Ingress Max Packet Rate/Sec	Trend in observed ingress maximum rate for DNS queries.
MX Egress Max Packet Rate/Sec	Trend in observed egress maximum rate for DNS queries.
MX Estimated Threshold/Sec	Trend in the estimated threshold.
TCP MX Dropped - Packets/Period	Drops due to the dns-mx-count threshold.
MX Blocked - Packets/Period	Drops due to an ACL rule blocking MX.
DNS QType ALL	
ALL Ingress Max Packet Rate/Sec	Trend in observed ingress maximum rate for DNS queries.
ALL Egress Max Packet Rate/Sec	Trend in observed egress maximum rate for DNS queries.
ALL Estimated Threshold/Sec	Trend in the estimated threshold.
TCP All Dropped - Packets/Period	Drops due to the dns-all-count threshold.
All Blocked - Packets/Period	Drops due to an ACL rule blocking ALL.
DNS QType Zone Transfer	
Zone Transfer Ingress Max Packet Rate/Sec	Trend in observed ingress maximum rate for DNS queries.
Zone Transfer Egress Max Packet Rate/Sec	Trend in observed egress maximum rate for DNS queries.
Zone Transfer Estimated Threshold/Sec	Trend in the estimated threshold.
TCP Zone Transfer Dropped - Packets/Period	Drops due to the dns-zone-xfer threshold

Statistic	Description
Zone Transfer Blocked - Packets/Period	Drops due to DNS fragment ACL.
DNS Fragment	
Fragment Ingress Max Packet Rate/Sec	Trend in observed ingress maximum rate for DNS queries.
Fragment Egress Max Packet Rate/Sec	Trend in observed egress maximum rate for DNS queries.
Fragment Max Packet Rate/Sec	Trend in the estimated threshold.
Fragment Dropped - Packets/Period	Drops due to the dns-fragment threshold.
Fragment Blocked - Packets/Period	Drops due to DNS fragment ACL
DNS Unsolicited Response	
Unsolicited UDP Response Drop	Drops when an unsolicited DNS response is received (UDP).
Unsolicited TCP Response Drop	Drops when an unsolicited DNS response is received (TCP).
Unsolicited UDP Response same port Drop	Drops when an unsolicited DNS response is received at port 53 (UDP).
Unsolicited TCP Response same port Drop	Drops when an unsolicited DNS response is received at port 53 (TCP).
DNS Unexpected Query	
UDP Duplicate Query before Response Drop	Drop due to the duplicate query check (UDP).
TCP Duplicate Query before Response Drop	Drop due to the duplicate query check (TCP).
DNS LQ Drop	
DNS Query Flood Drop	Drops because a query was not found in the LQ table during a flood. The **dns-query** threshold triggered the mitigation.

Statistic	Description
DNS Question Flood Drop	Drops because a query was not found in the LQ table during a flood. The **dns-question** threshold triggered the mitigation.
DNS QType All Flood Drop	Drops because a query was not found in the LQ table during a flood. The **dns-all** threshold triggered the mitigation.
DNS QType Zone Transfer Flood Drop	Drops because a query was not found in the LQ table during a flood. The **dns-zone-xfer** threshold triggered the mitigation.
DNS QType MX Flood Drop	Drops because a query was not found in the LQ table during a flood. The **dns-mx** threshold triggered the mitigation.
DNS Query Flood due to Negative Response	DNS Query Flood due to negative response flood. (*Deprecated*)
DNS TTL	
DNS Query Flood Drop	Drops because a query was found in the TTL table during a flood. The **dns-query** threshold triggered the mitigation.
DNS Question Flood Drop	Drops because a query was found in the TTL table during a flood. The **dns-question** threshold triggered the mitigation.
DNS QType All Flood Drop	Drops because a query was found in the TTL table during a flood. The **dns-all** threshold triggered the mitigation.
DNS QType Zone Transfer Flood Drop	Drops because a query was found in the TTL table during a flood. The **dns-zone-xfer** threshold triggered the mitigation.
DNS QType MX Flood Drop	Drops because a query was found in the TTL table during a flood. The **dns-mx** threshold triggered the mitigation.
DNS Cache	
UDP Query Flood Drop Due to Response from Cache	Drops because a response was served from the cache during a flood. The **dns-query** threshold triggered the mitigation.
UDP Question Flood Drop Due to Response from Cache	Drops because a response was served from the cache during a flood. The **dns-question** threshold triggered the mitigation.
UDP QType All Flood Drop Due to Response from Cache	Drops because a response was served from the cache during a flood. The **dns-all** threshold triggered the mitigation.

Statistic	Description
UDP QType Zone Transfer Flood Drop Due to Response from Cache	Drops because a response was served from the cache during a flood. The **dns-zone-xfer** threshold triggered the mitigation.
UDP QType MX Flood Drop Due to Response from Cache	Drops because a query was served from the cache during a flood. The **dns-mx** threshold triggered the mitigation.
UDP Query Flood Drop Due to No Response from Cache	Drops because a response was not found in the cache during a flood. The **dns-query** threshold triggered the mitigation.
UDP Question Flood Drop Due to No Response from Cache	Drops because a response was not found in the cache during a flood. The **dns-question** threshold triggered the mitigation.
UDP QType All Flood Drop Due to No Response from Cache	Drops because a response was served from the cache during a flood. The **dns-all** threshold triggered the mitigation.
UDP QType Zone Transfer Flood Drop Due to No Response from Cache	Drops because a response was not found in the cache during a flood. The **dns-zone-xfer** threshold triggered the mitigation.
UDP QType MX Flood Drop Due to No Response from Cache	Drops because a response was not found in the cache during a flood. The **dns-mx** threshold triggered the mitigation.
DNS Spoofed IP	
UDP Query Flood Drop During TC=1 check	Drops due to anti-spoofing checks during a flood. The source IP address was not found in the legitimate IP address (LIP) table. The query was dropped and a response was sent with the TC bit set to force the client to retry the query over TCP. The **dns-query** threshold triggered the mitigation.
UDP Question Flood Drop During TC=1 check	Drops due to anti-spoofing checks during a flood. The **dns-question** threshold triggered the mitigation.
UDP QType All Flood Drop During TC=1 check	Drops due to anti-spoofing checks during a flood. The **dns-all** threshold triggered the mitigation.
UDP QType Zone Transfer Flood Drop During TC=1 check	Drops due to anti-spoofing checks during a flood. The **dns-zone-xfer** threshold triggered the mitigation.
UDP QType MX Flood Drop During TC=1 check	Drops because a response was not found in the cache during a flood. The **dns-mx** threshold triggered the mitigation.

Statistic	Description
UDP Query Flood Drop During Retransmission Check	Drops due to anti-spoofing checks during a flood. The source IP address was not found in the legitimate IP address (LIP) table. The query was dropped and the retransmission check was performed. The **dns-query** threshold triggered the mitigation.
UDP Question Flood Drop During Retransmission Check	Drops due to anti-spoofing checks during a flood. The **dns-question** threshold triggered the mitigation.
UDP QType All Flood Drop During Retransmission Check	Drops due to anti-spoofing checks during a flood. The **dns-all** threshold triggered the mitigation.
UDP QType Zone Transfer Flood Drop During Retransmission Check	Drops due to anti-spoofing checks during a flood. The **dns-zone-xfer** threshold triggered the mitigation.
UDP QType MX Flood Drop During Retransmission Check	Drops because a response was not found in the cache during a flood. The **dns-mx** threshold triggered the mitigation.
DNS Rcode	
DNS Rcode Ingress Max Packet Rate	Ingress maximum packet rate for the selected RCODE. An RCODE 0 indicates a successful query. All other RCODEs indicate errors.
DNS Rcode Egress Max Packet Rate	Egress maximum packet rate for the selected RCODE. An RCODE 0 indicates a successful query. All other RCODEs indicate errors.

Sample Graphs

Figure 92: HTTP Methods

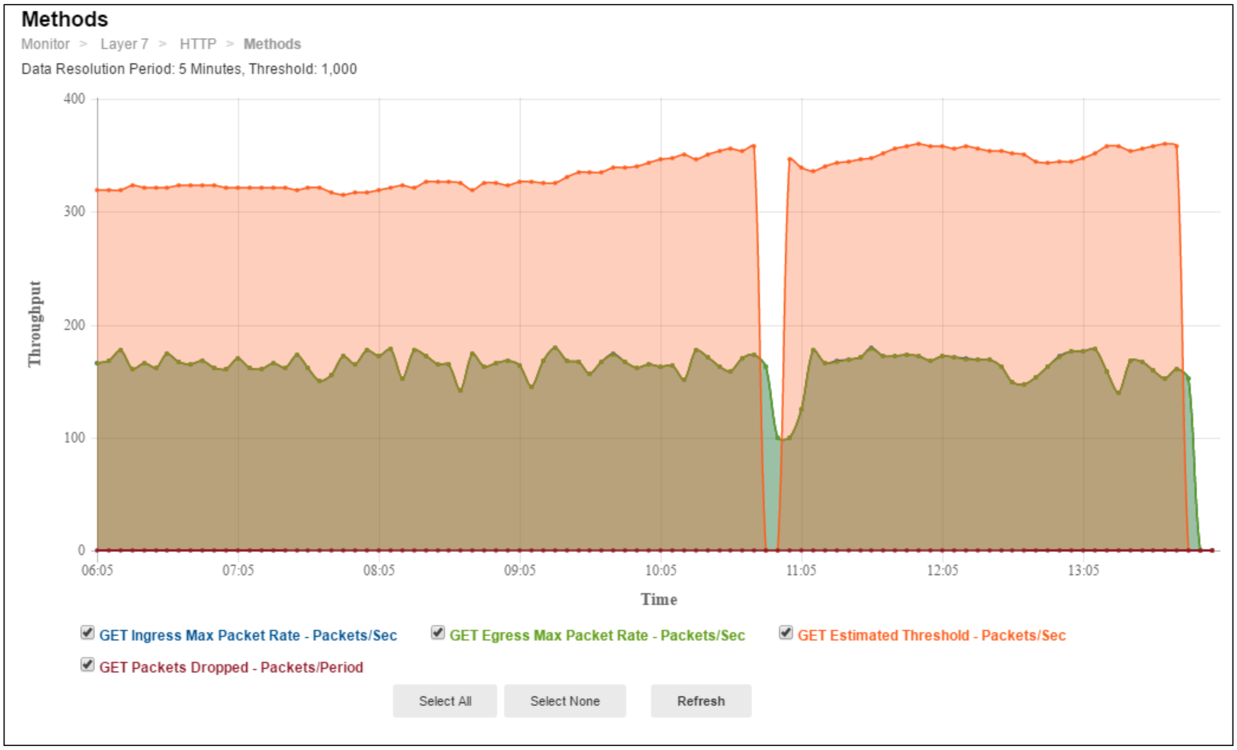

Methods

Monitor > Layer 7 > HTTP > Methods

Data Resolution Period: 5 Minutes, Threshold: 1,000

400
300
200
100
0
Throughput
06:05
07:05
08:05
09:05
10:05
11:05
12:05
13:05
Time
☑ GET Ingress Max Packet Rate - Packets/Sec
☑ GET Egress Max Packet Rate - Packets/Sec
☑ GET Estimated Threshold - Packets/Sec
☑ GET Packets Dropped - Packets/Period
Select All
Select None
Refresh

Figure 93: Methods Per Source

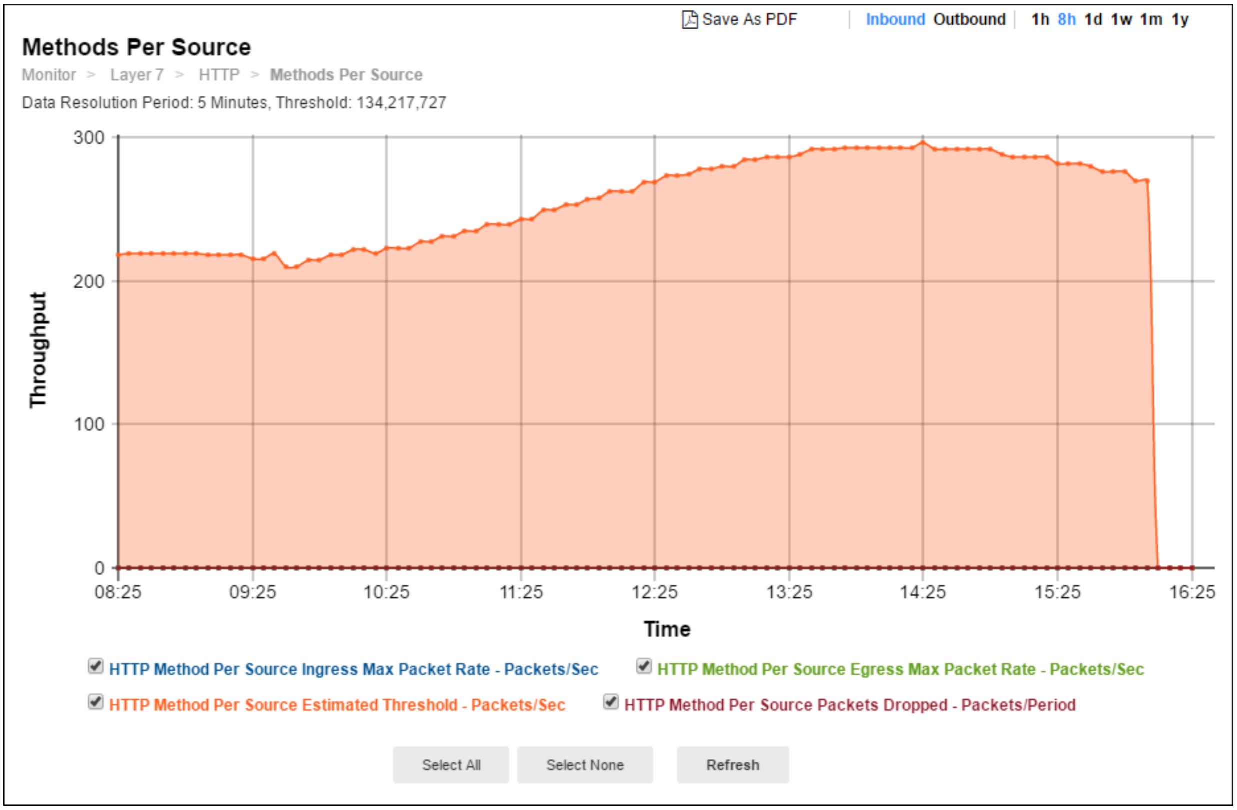

Save As PDF
Inbound Outbound
1h 8h 1d 1w 1m 1y
Methods Per Source
Monitor > Layer 7 > HTTP > Methods Per Source

Data Resolution Period: 5 Minutes, Threshold: 134,217,727

300
200
100
0
Throughput
08:25
09:25
10:25
11:25
12:25
13:25
14:25
15:25
16:25
Time
☑ HTTP Method Per Source Ingress Max Packet Rate - Packets/Sec
☑ HTTP Method Per Source Egress Max Packet Rate - Packets/Sec
☑ HTTP Method Per Source Estimated Threshold - Packets/Sec
☑ HTTP Method Per Source Packets Dropped - Packets/Period
Select All
Select None
Refresh

Figure 94: URLs

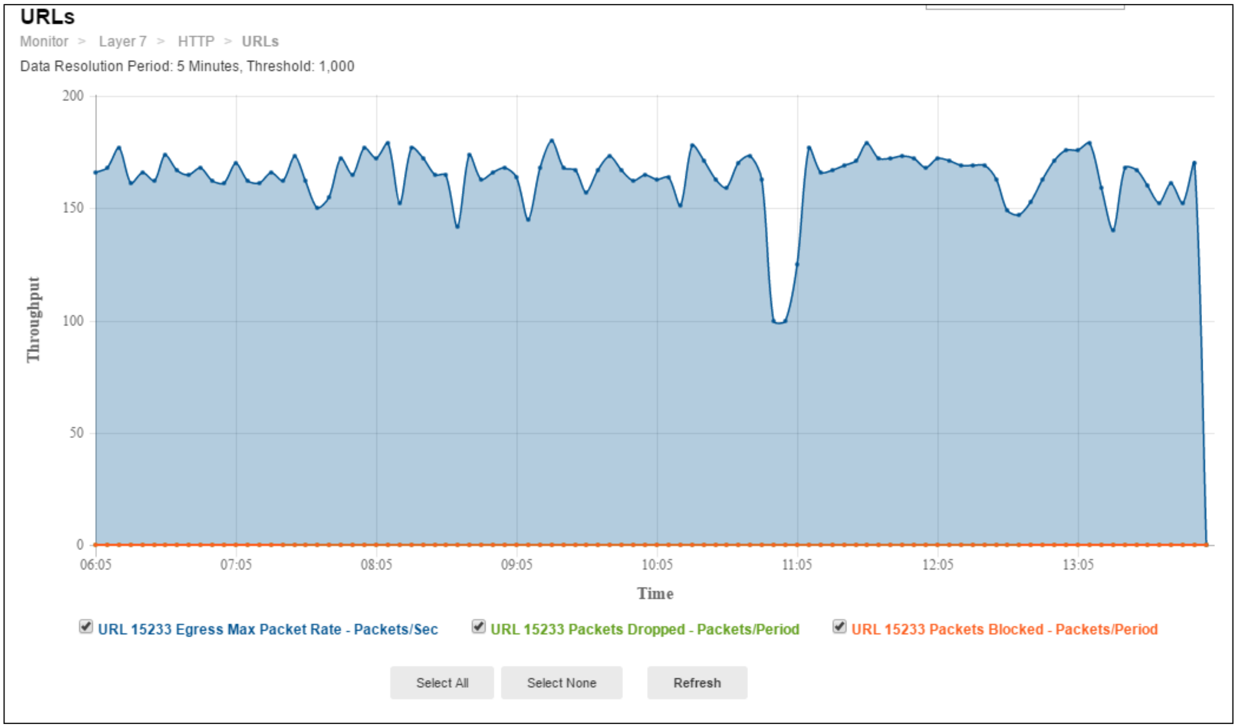

URLs
Monitor > Layer 7 > HTTP > URLs

Data Resolution Period: 5 Minutes, Threshold: 1,000

200
150
100
50
0
Throughput
06:05
07:05
08:05
09:05
10:05
11:05
12:05
13:05
Time
☑ URL 15233 Egress Max Packet Rate - Packets/Sec
☑ URL 15233 Packets Dropped - Packets/Period
☑ URL 15233 Packets Blocked - Packets/Period
Select All
Select None
Refresh

Figure 95: HTTP Hosts

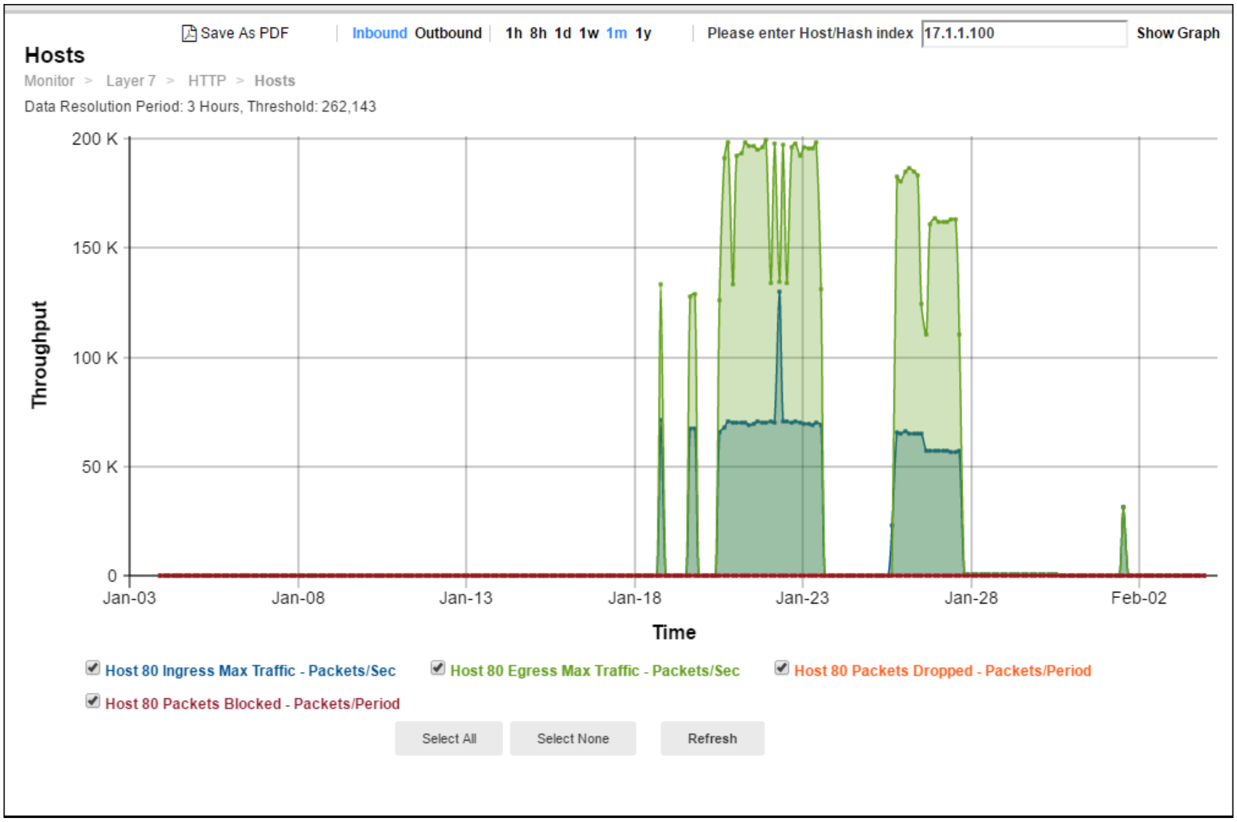

Hosts

Monitor > Layer 7 > HTTP > Hosts

Data Resolution Period: 3 Hours, Threshold: 262,143

Please enter Host/Hash index 17.1.1.100

Save As PDF
Inbound
Outbound
1h
8h
1d
1w
1m
1y
Show Graph
200 K
150 K
100 K
50 K
0
Throughput
Jan-03
Jan-08
Jan-13
Jan-18
Jan-23
Jan-28
Feb-02
Time
☑ Host 80 Ingress Max Traffic - Packets/Sec
☑ Host 80 Egress Max Traffic - Packets/Sec
☑ Host 80 Packets Dropped - Packets/Period
☑ Host 80 Packets Blocked - Packets/Period
Select All
Select None
Refresh

Figure 96: HTTP Referers

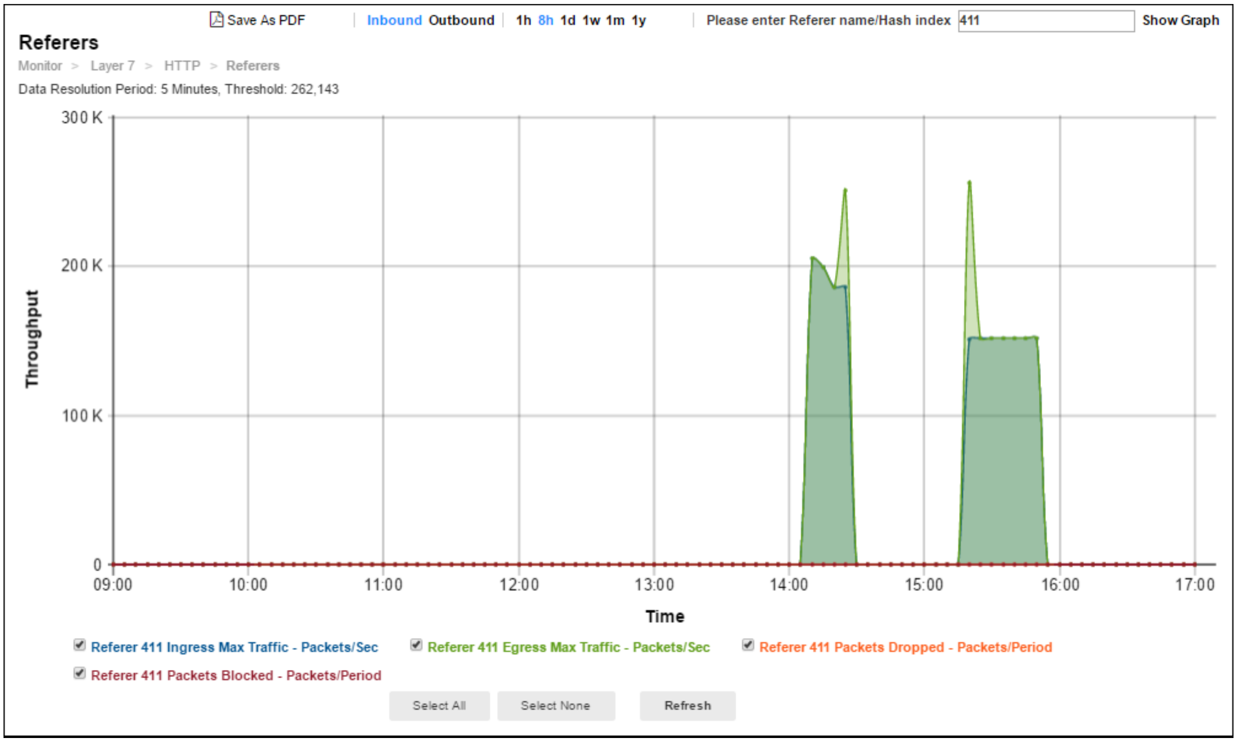

Referers

Monitor > Layer 7 > HTTP > Referers

Data Resolution Period: 5 Minutes, Threshold: 262,143

Please enter Referer name/Hash index 411

Save As PDF
Inbound
Outbound
1h
8h
1d
1w
1m
1y
Show Graph
300 K
200 K
100 K
0
Throughput
09:00
10:00
11:00
12:00
13:00
14:00
15:00
16:00
17:00
Time
☑ Referer 411 Ingress Max Traffic - Packets/Sec
☑ Referer 411 Egress Max Traffic - Packets/Sec
☑ Referer 411 Packets Dropped - Packets/Period
☑ Referer 411 Packets Blocked - Packets/Period
Select All
Select None
Refresh

Figure 97: HTTP Cookies

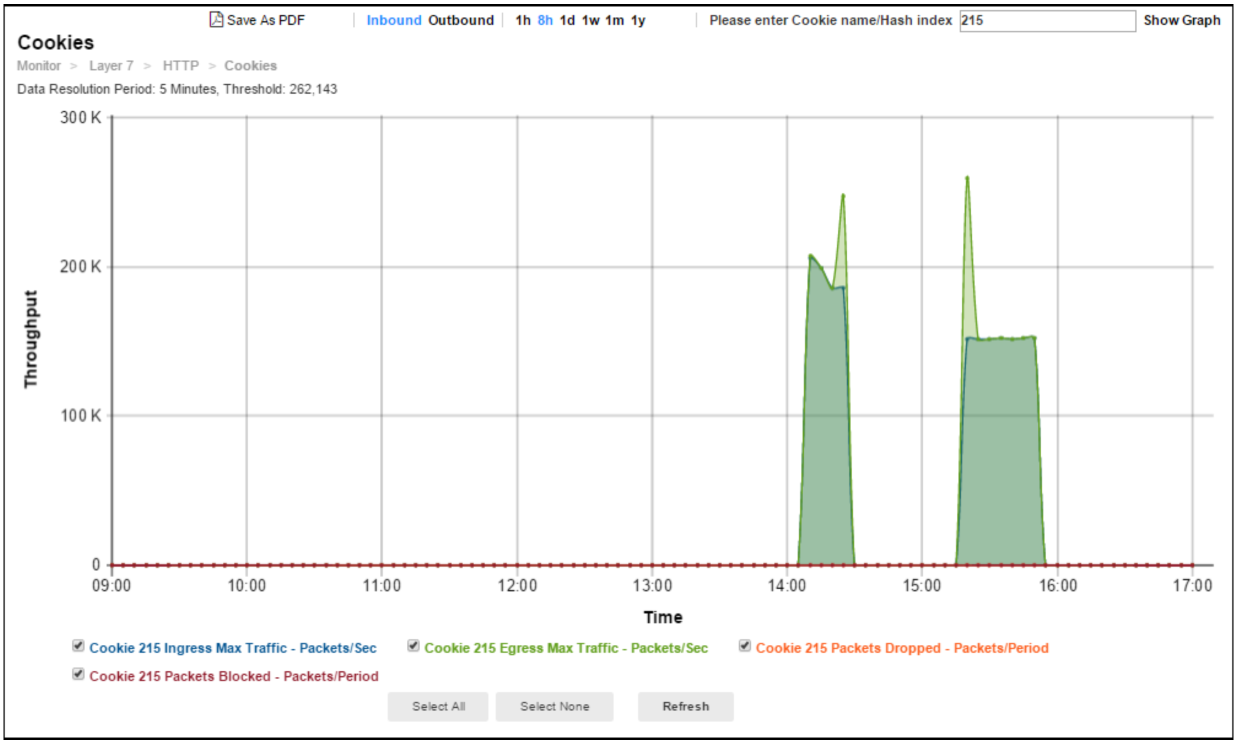

Cookies

Monitor > Layer 7 > HTTP > Cookies

Data Resolution Period: 5 Minutes, Threshold: 262,143

Please enter Cookie name/Hash index 215

Save As PDF
Inbound
Outbound
1h
8h
1d
1w
1m
1y
Show Graph
300 K
200 K
100 K
0
Throughput
09:00
10:00
11:00
12:00
13:00
14:00
15:00
16:00
17:00
Time
☑ Cookie 215 Ingress Max Traffic - Packets/Sec
☑ Cookie 215 Egress Max Traffic - Packets/Sec
☑ Cookie 215 Packets Dropped - Packets/Period
☑ Cookie 215 Packets Blocked - Packets/Period
Select All
Select None
Refresh

Figure 98: HTTP User Agents

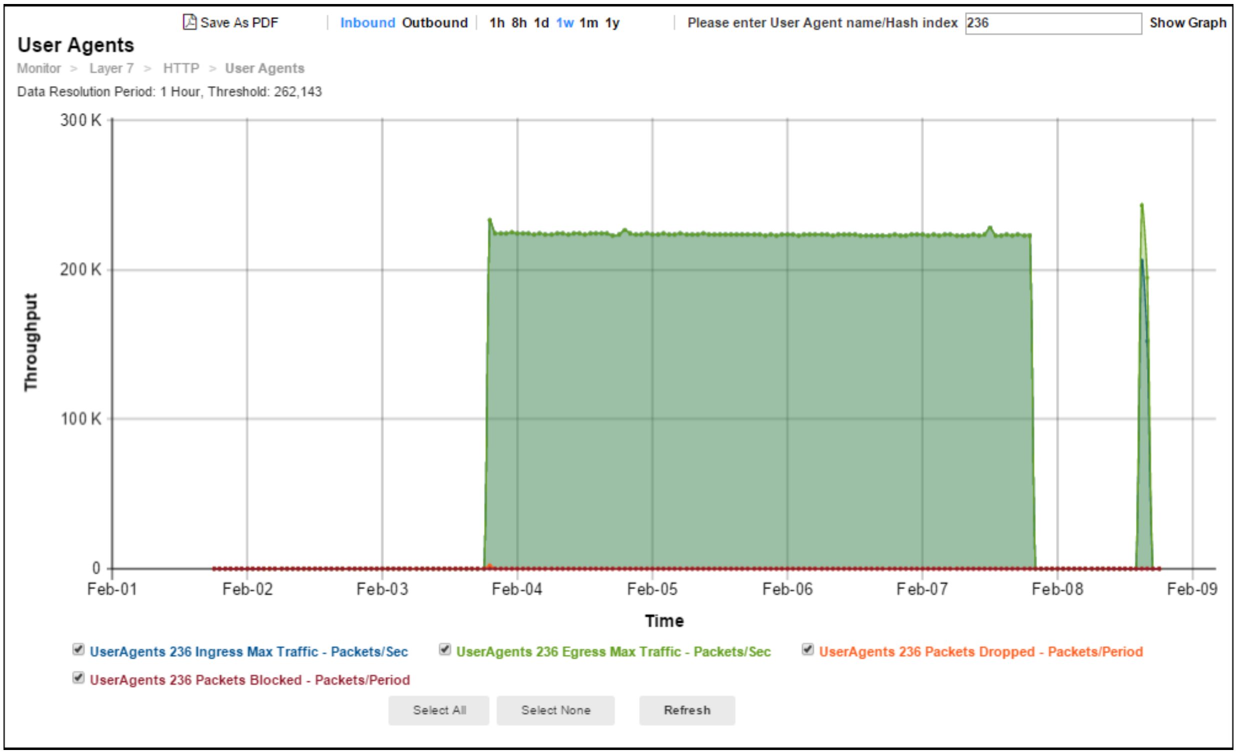

User Agents

Monitor > Layer 7 > HTTP > User Agents

Data Resolution Period: 1 Hour, Threshold: 262,143

Please enter User Agent name/Hash index 236

Save As PDF
Inbound
Outbound
1h
8h
1d
1w
1m
1y
Show Graph
300 K
200 K
100 K
0
Throughput
Feb-01
Feb-02
Feb-03
Feb-04
Feb-05
Feb-06
Feb-07
Feb-08
Feb-09
Time
☑ UserAgents 236 Ingress Max Traffic - Packets/Sec
☑ UserAgents 236 Egress Max Traffic - Packets/Sec
☑ UserAgents 236 Packets Dropped - Packets/Period
☑ UserAgents 236 Packets Blocked - Packets/Period
Select All
Select None
Refresh

Figure 99: DNS Query

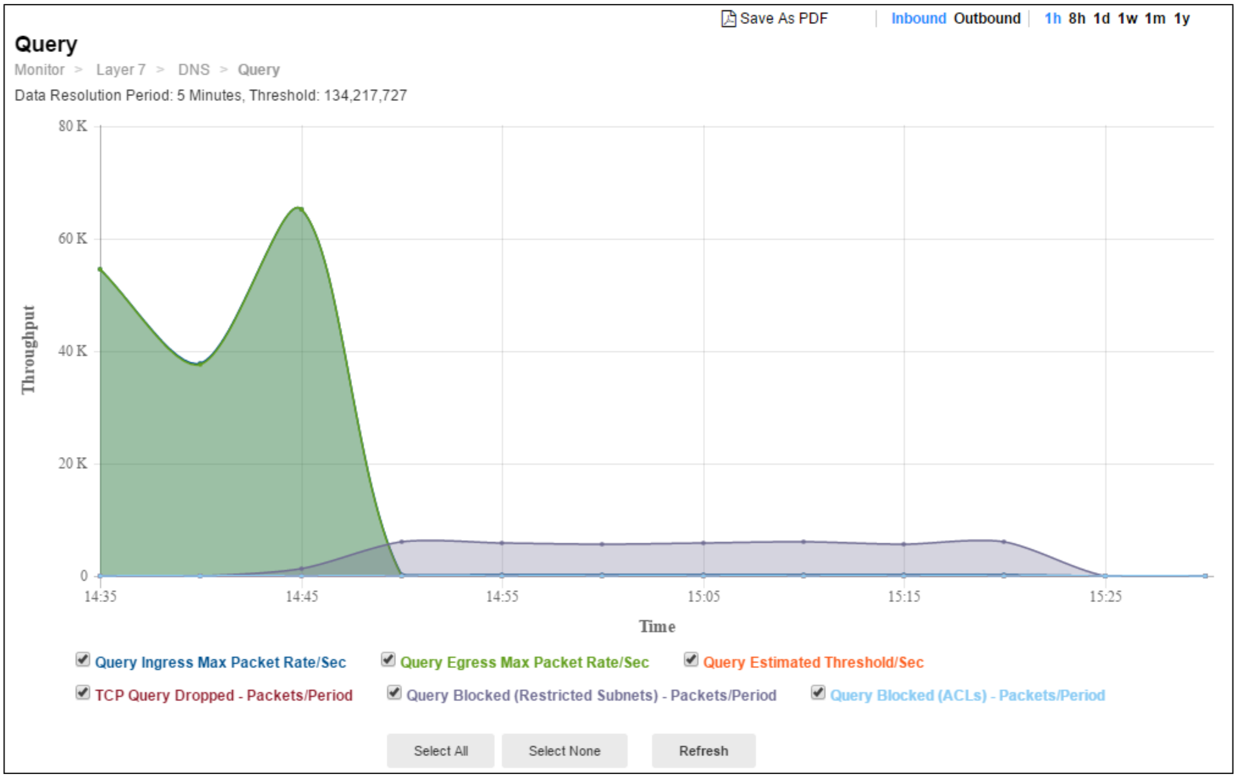
Save As PDF
Inbound Outbound
1h 8h 1d 1w 1m 1y
Query
Monitor > Layer 7 > DNS > Query
80 K
60 K
40 K
20 K
0
Throughput
14:35
14:45
14:55
15:05
15:15
15:25
Time
☑ Query Ingress Max Packet Rate/Sec
☑ Query Egress Max Packet Rate/Sec
☑ Query Estimated Threshold/Sec
☑ TCP Query Dropped - Packets/Period
☑ Query Blocked (Restricted Subnets) - Packets/Period
☑ Query Blocked (ACLs) - Packets/Period
Select All
Select None
Refresh

Data Resolution Period: 5 Minutes, Threshold: 134,217,727

Figure 100: DNS Query Per Source

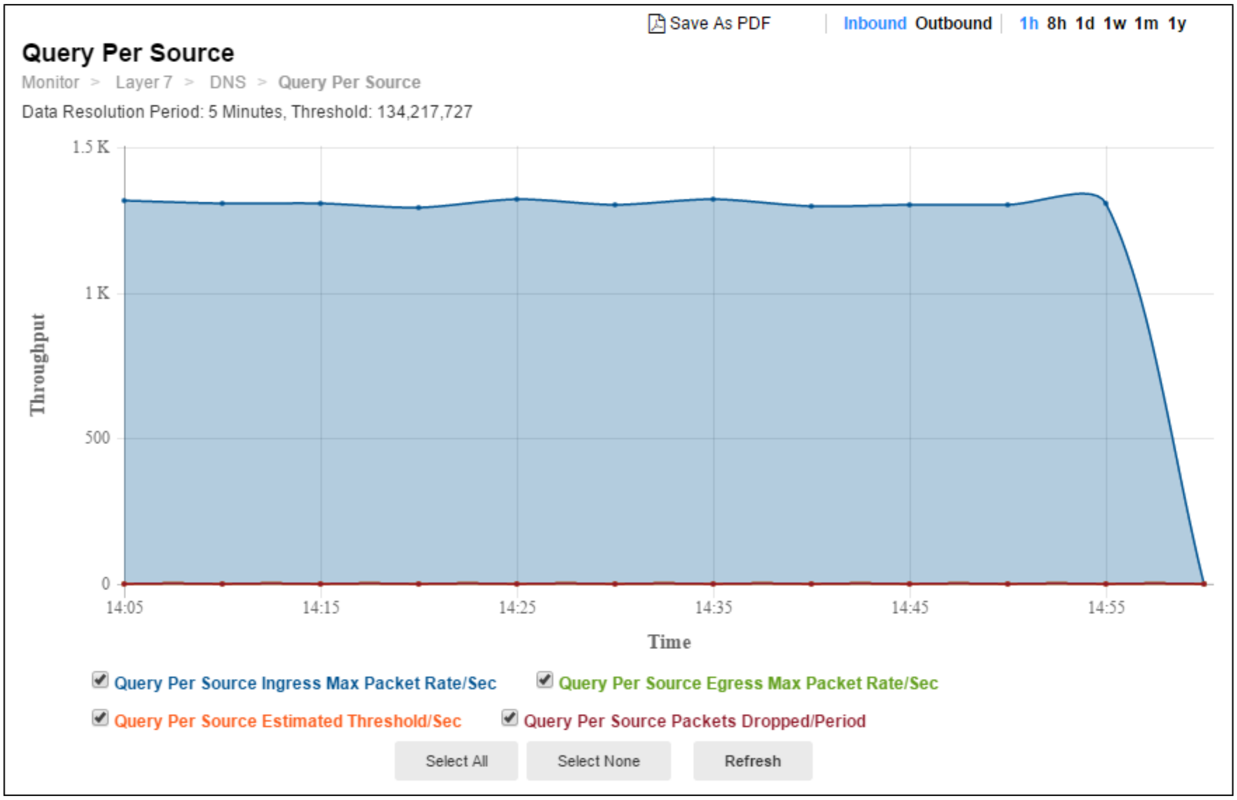
Save As PDF
Inbound Outbound
1h 8h 1d 1w 1m 1y
Query Per Source
Monitor > Layer 7 > DNS > Query Per Source
1.5 K
1 K
500
0
Throughput
14:05
14:15
14:25
14:35
14:45
14:55
Time
☑ Query Per Source Ingress Max Packet Rate/Sec
☑ Query Per Source Egress Max Packet Rate/Sec
☑ Query Per Source Estimated Threshold/Sec
☑ Query Per Source Packets Dropped/Period
Select All
Select None
Refresh

Data Resolution Period: 5 Minutes, Threshold: 134,217,727

Figure 101: DNS Suspicious Sources

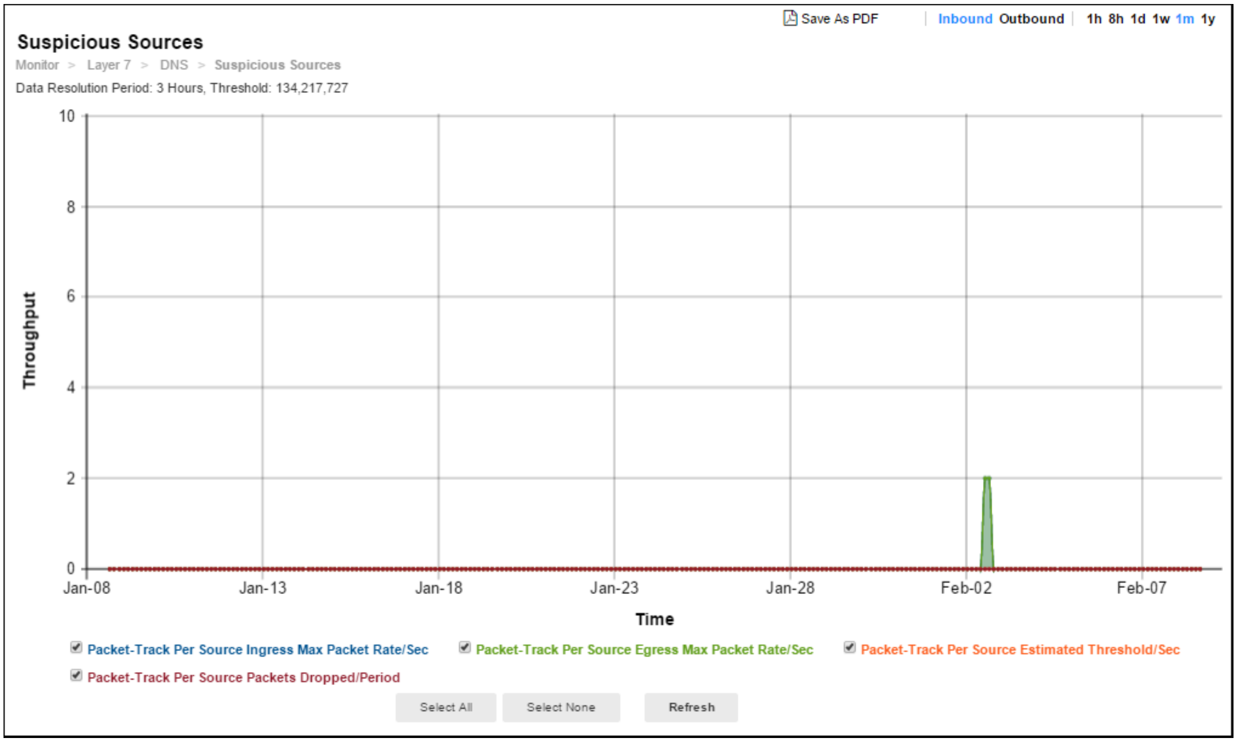

Suspicious Sources

Monitor > Layer 7 > DNS > Suspicious Sources

Data Resolution Period: 3 Hours, Threshold: 134,217,727

Save As PDF
Inbound Outbound
1h 8h 1d 1w 1m 1y
Throughput
10
8
6
4
2
0
Jan-08
Jan-13
Jan-18
Jan-23
Jan-28
Feb-02
Feb-07
Time
☑ Packet-Track Per Source Ingress Max Packet Rate/Sec
☑ Packet-Track Per Source Egress Max Packet Rate/Sec
☑ Packet-Track Per Source Estimated Threshold/Sec
☑ Packet-Track Per Source Packets Dropped/Period
Select All
Select None
Refresh

Figure 102: DNS Question Count

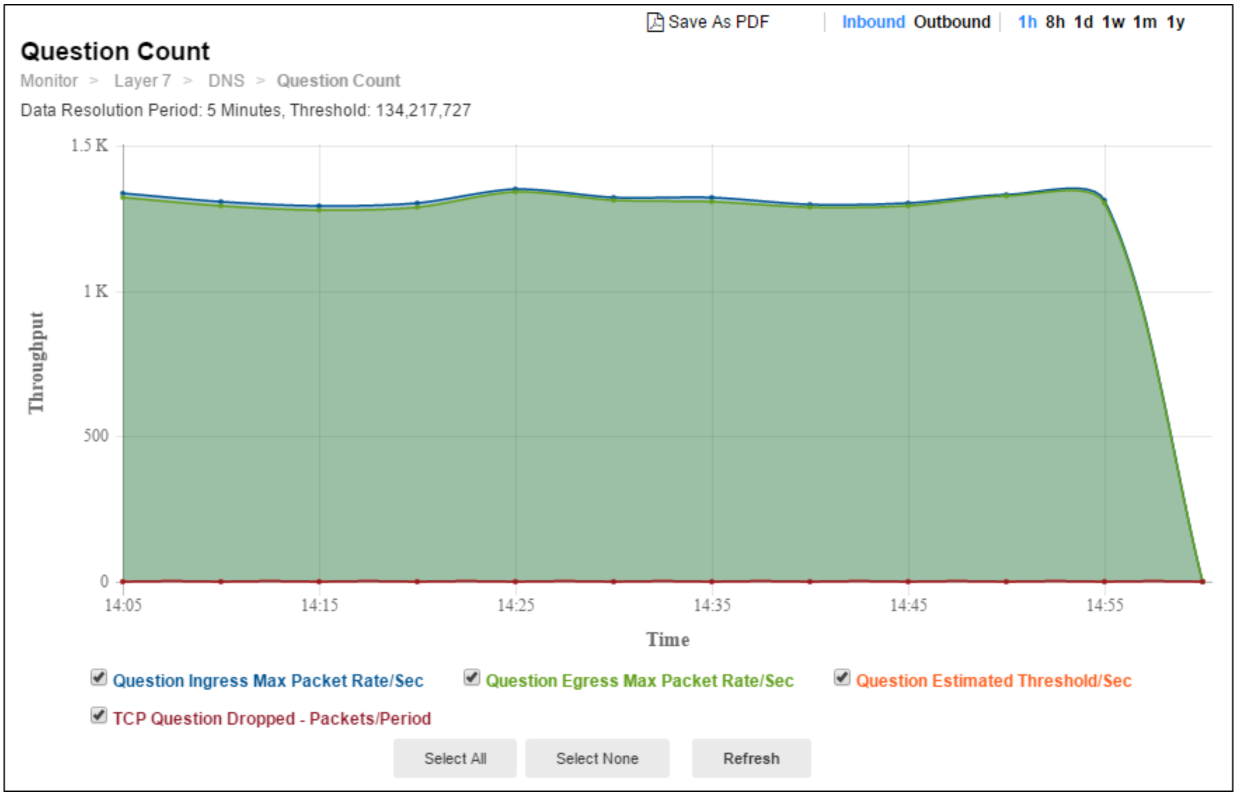

Question Count

Monitor > Layer 7 > DNS > Question Count

Data Resolution Period: 5 Minutes, Threshold: 134,217,727

Save As PDF
Inbound Outbound
1h 8h 1d 1w 1m 1y
Throughput
1.5 K
1 K
500
0
14:05
14:15
14:25
14:35
14:45
14:55
Time
☑ Question Ingress Max Packet Rate/Sec
☑ Question Egress Max Packet Rate/Sec
☑ Question Estimated Threshold/Sec
☑ TCP Question Dropped - Packets/Period
Select All
Select None
Refresh

Figure 103: DNS QType MX

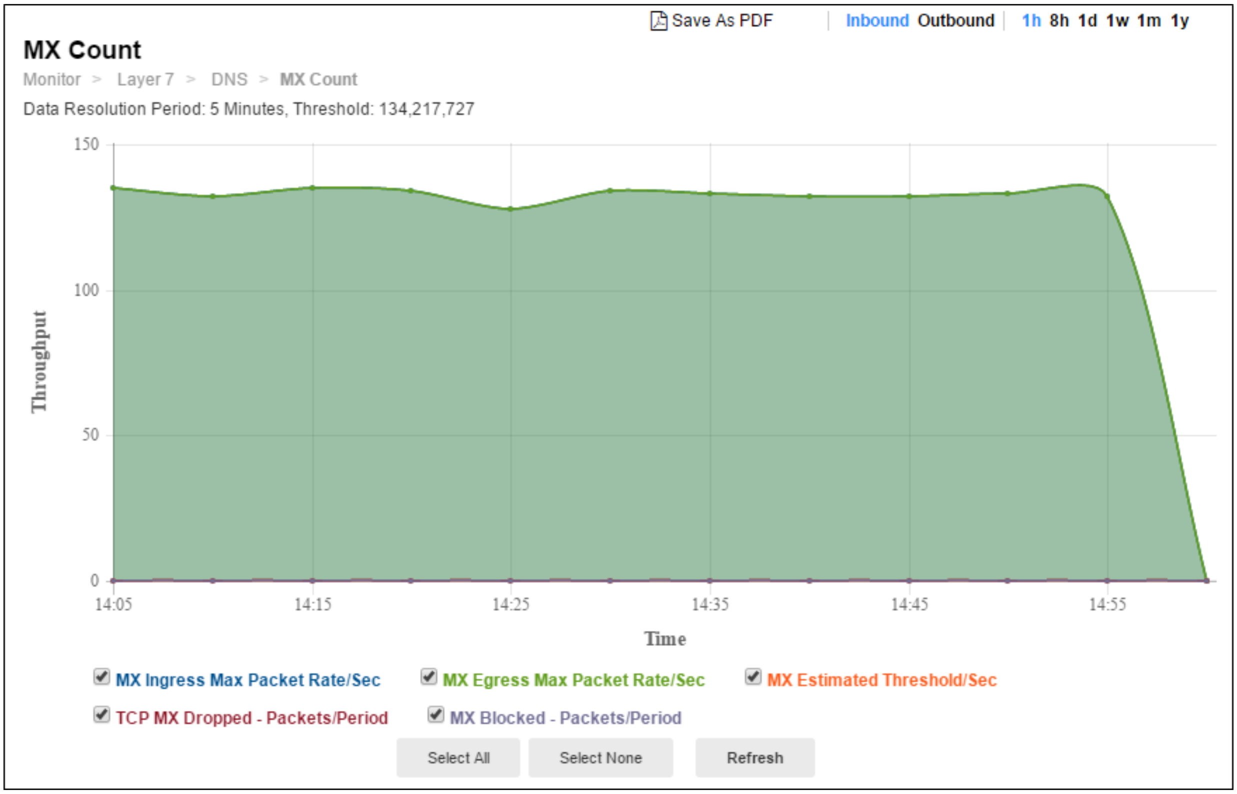

MX Count

Monitor > Layer 7 > DNS > MX Count

Data Resolution Period: 5 Minutes, Threshold: 134,217,727

Save As PDF
Inbound
Outbound
1h
8h
1d
1w
1m
1y
150
100
50
0
Throughput
14:05
14:15
14:25
14:35
14:45
14:55
Time
☑ MX Ingress Max Packet Rate/Sec
☑ MX Egress Max Packet Rate/Sec
☑ MX Estimated Threshold/Sec
☑ TCP MX Dropped - Packets/Period
☑ MX Blocked - Packets/Period
Select All
Select None
Refresh

Figure 104: DNS QType ALL

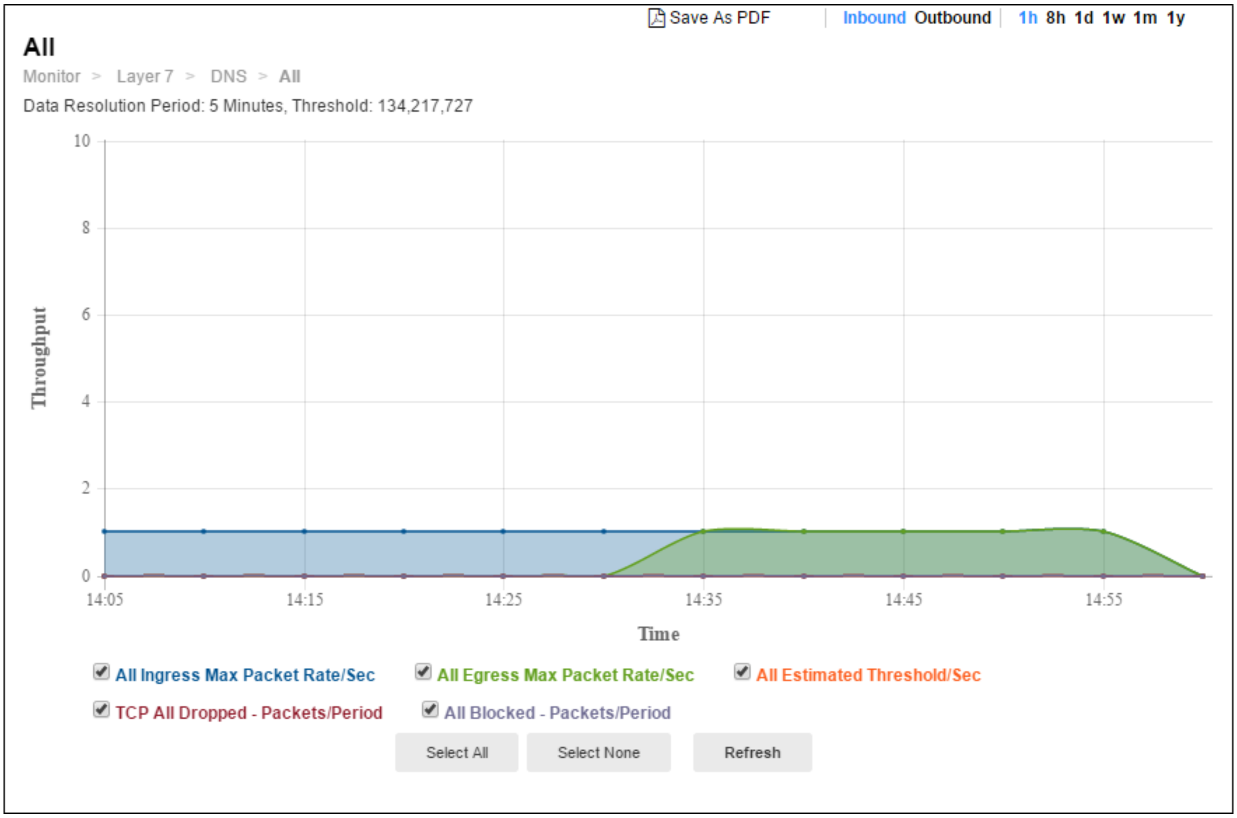

All

Monitor > Layer 7 > DNS > All

Data Resolution Period: 5 Minutes, Threshold: 134,217,727

Save As PDF
Inbound Outbound
1h 8h 1d 1w 1m 1y
Throughput
10
8
6
4
2
0
14:05
14:15
14:25
14:35
14:45
14:55
Time
☑ All Ingress Max Packet Rate/Sec
☑ All Egress Max Packet Rate/Sec
☑ All Estimated Threshold/Sec
☑ TCP All Dropped - Packets/Period
☑ All Blocked - Packets/Period
Select All
Select None
Refresh

Figure 105: DNS QType Zone Transfer

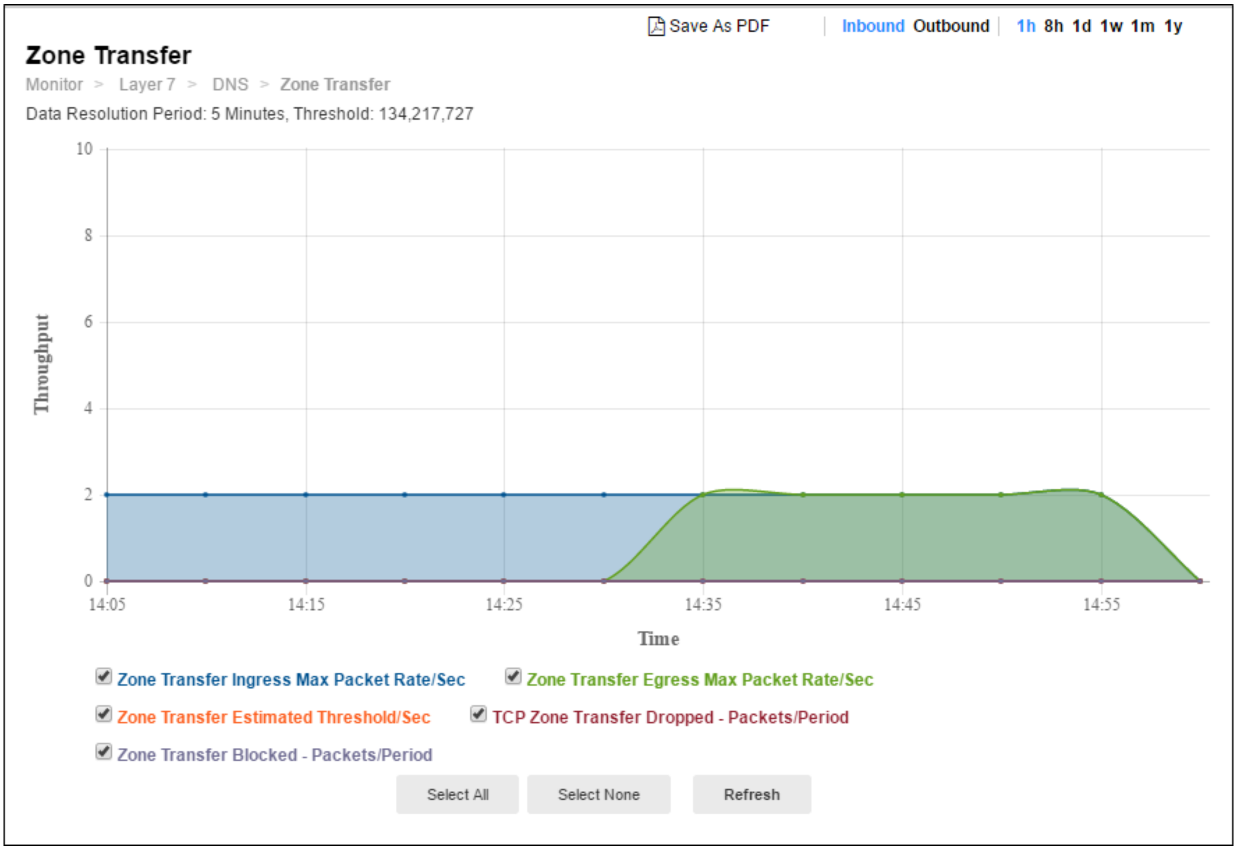

Zone Transfer

Monitor > Layer 7 > DNS > Zone Transfer

Data Resolution Period: 5 Minutes, Threshold: 134,217,727

☑ Zone Transfer Ingress Max Packet Rate/Sec ☑ Zone Transfer Egress Max Packet Rate/Sec

☑ Zone Transfer Estimated Threshold/Sec ☑ TCP Zone Transfer Dropped - Packets/Period

☑ Zone Transfer Blocked - Packets/Period

Figure 106: DNS Fragment

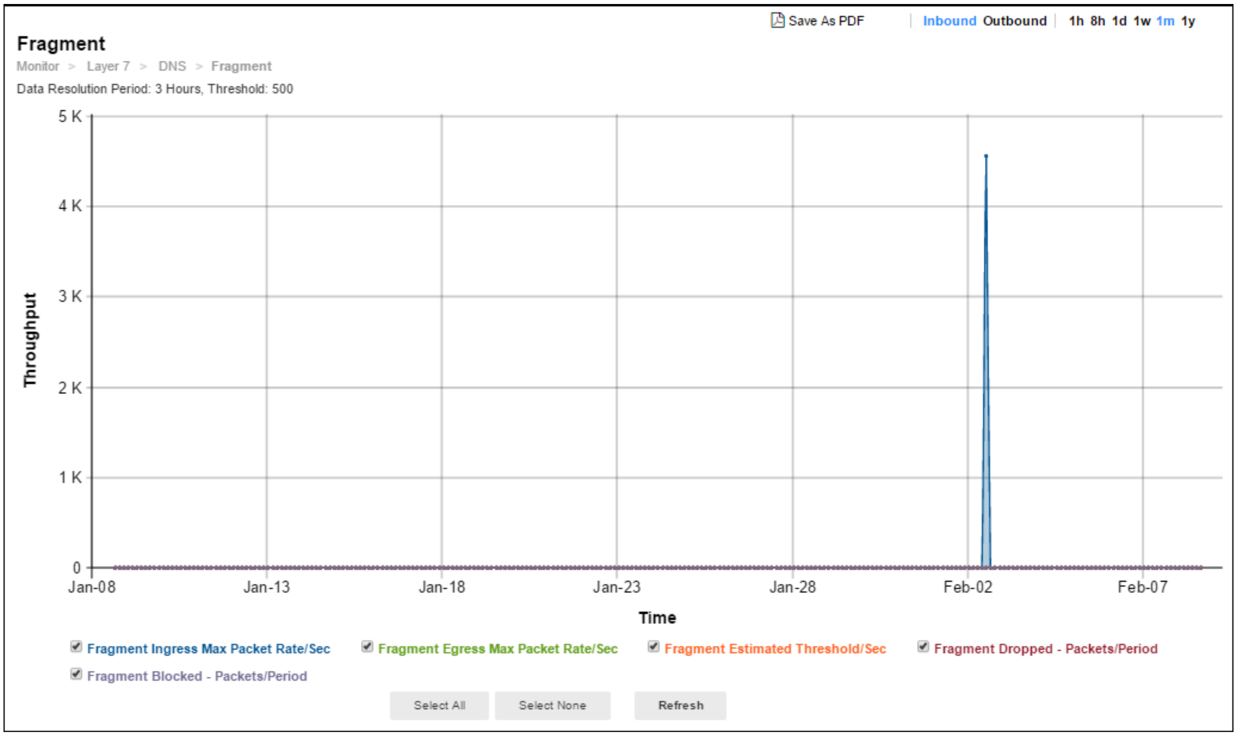

Fragment

Monitor > Layer 7 > DNS > Fragment

Data Resolution Period: 3 Hours, Threshold: 500

Save As PDF
Inbound Outbound
1h 8h 1d 1w 1m 1y
5 K
4 K
3 K
2 K
1 K
0
Throughput
Jan-08
Jan-13
Jan-18
Jan-23
Jan-28
Feb-02
Feb-07
Time
☑ Fragment Ingress Max Packet Rate/Sec
☑ Fragment Egress Max Packet Rate/Sec
☑ Fragment Estimated Threshold/Sec
☑ Fragment Dropped - Packets/Period
☑ Fragment Blocked - Packets/Period
Select All
Select None
Refresh

Figure 107: DNS Unsolicited Response

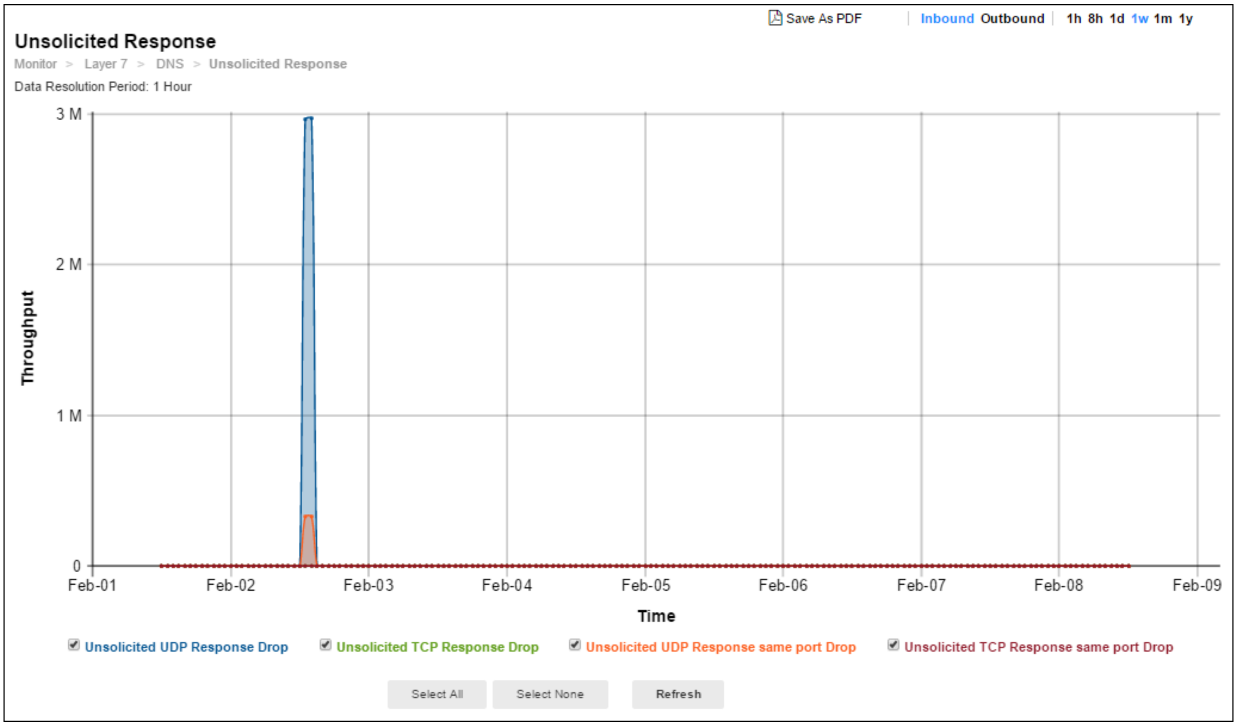

Unsolicited Response

Monitor > Layer 7 > DNS > Unsolicited Response

Data Resolution Period: 1 Hour

Figure 108: DNS Unexpected Query

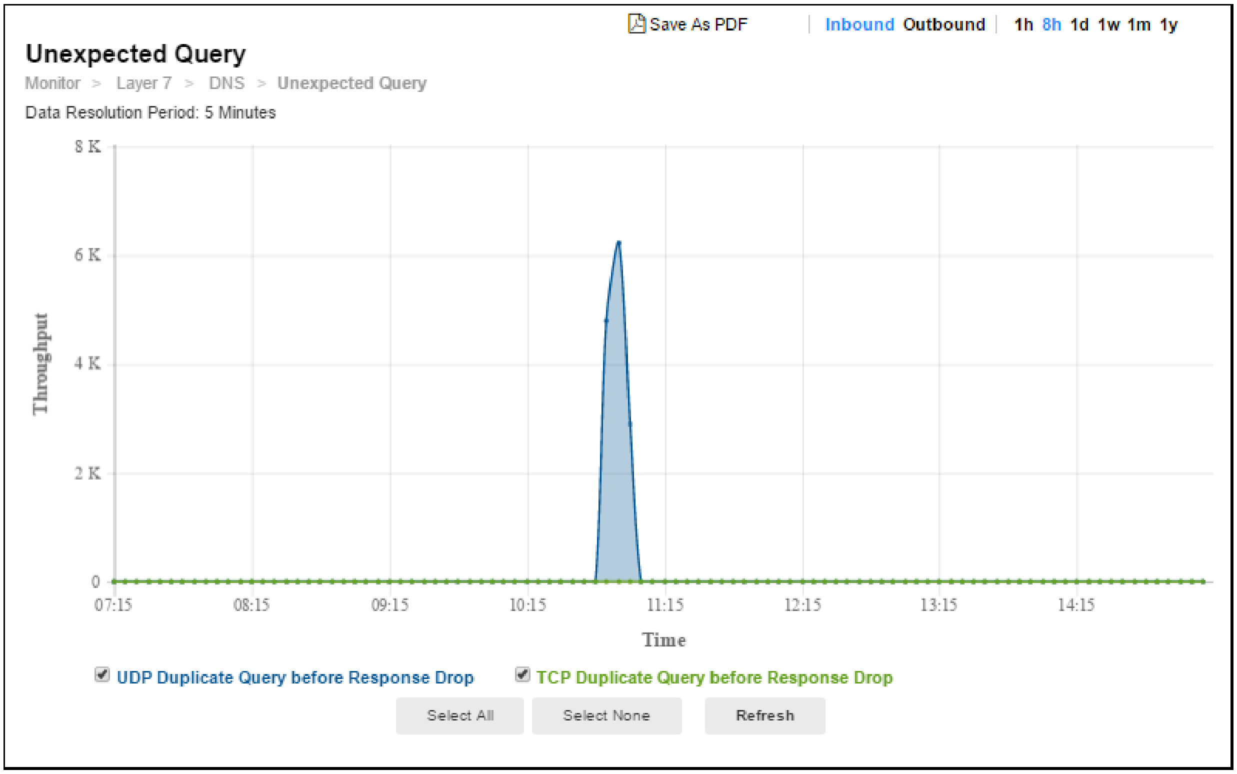

Unexpected Query

Monitor > Layer 7 > DNS > Unexpected Query

Data Resolution Period: 5 Minutes

Figure 109: DNS LQ Drop

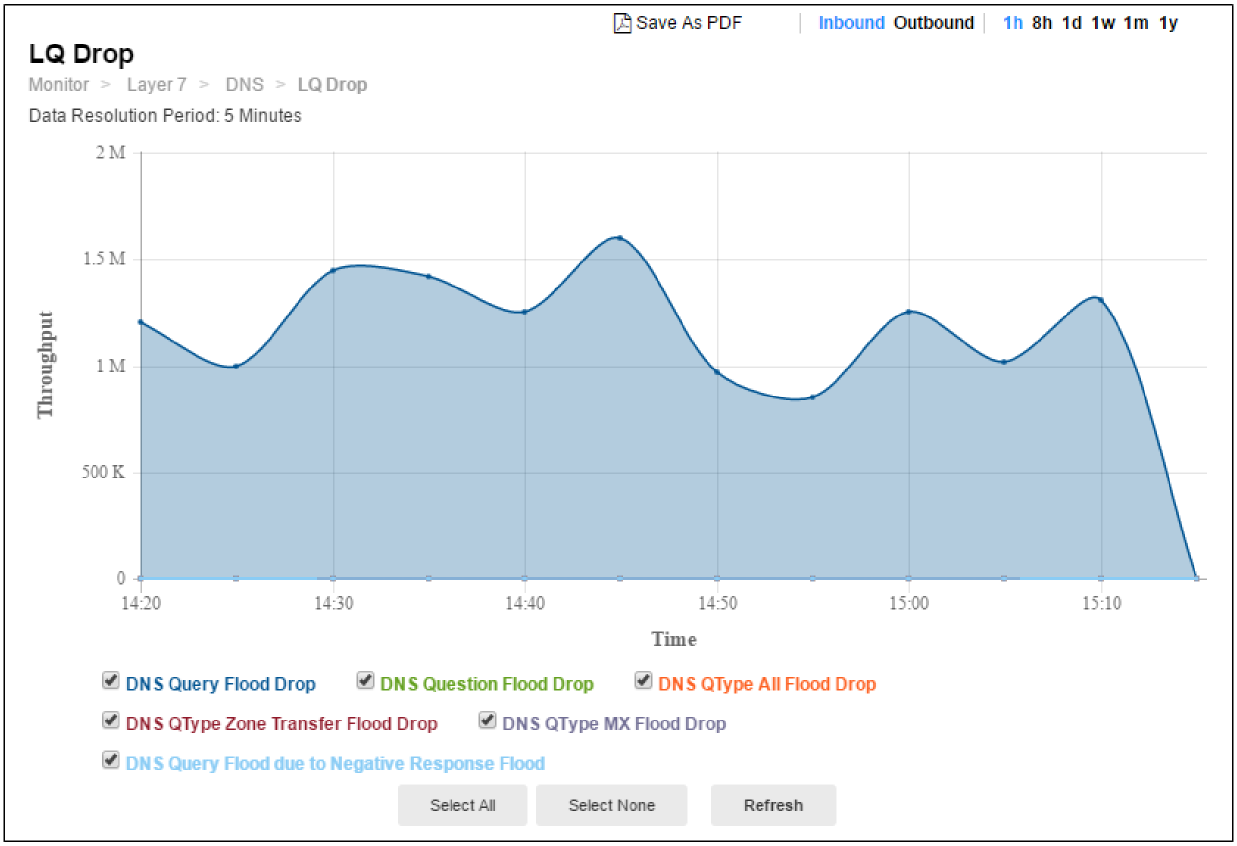

LQ Drop

Monitor > Layer 7 > DNS > LQ Drop

Data Resolution Period: 5 Minutes

Save As PDF
Inbound Outbound
1h 8h 1d 1w 1m 1y
Throughput
2 M
1.5 M
1 M
500 K
0
14:20
14:30
14:40
14:50
15:00
15:10
Time
☑ DNS Query Flood Drop
☑ DNS Question Flood Drop
☑ DNS QType All Flood Drop
☑ DNS QType Zone Transfer Flood Drop
☑ DNS QType MX Flood Drop
☑ DNS Query Flood due to Negative Response Flood
Select All
Select None
Refresh

Figure 110: DNS TTL Drop

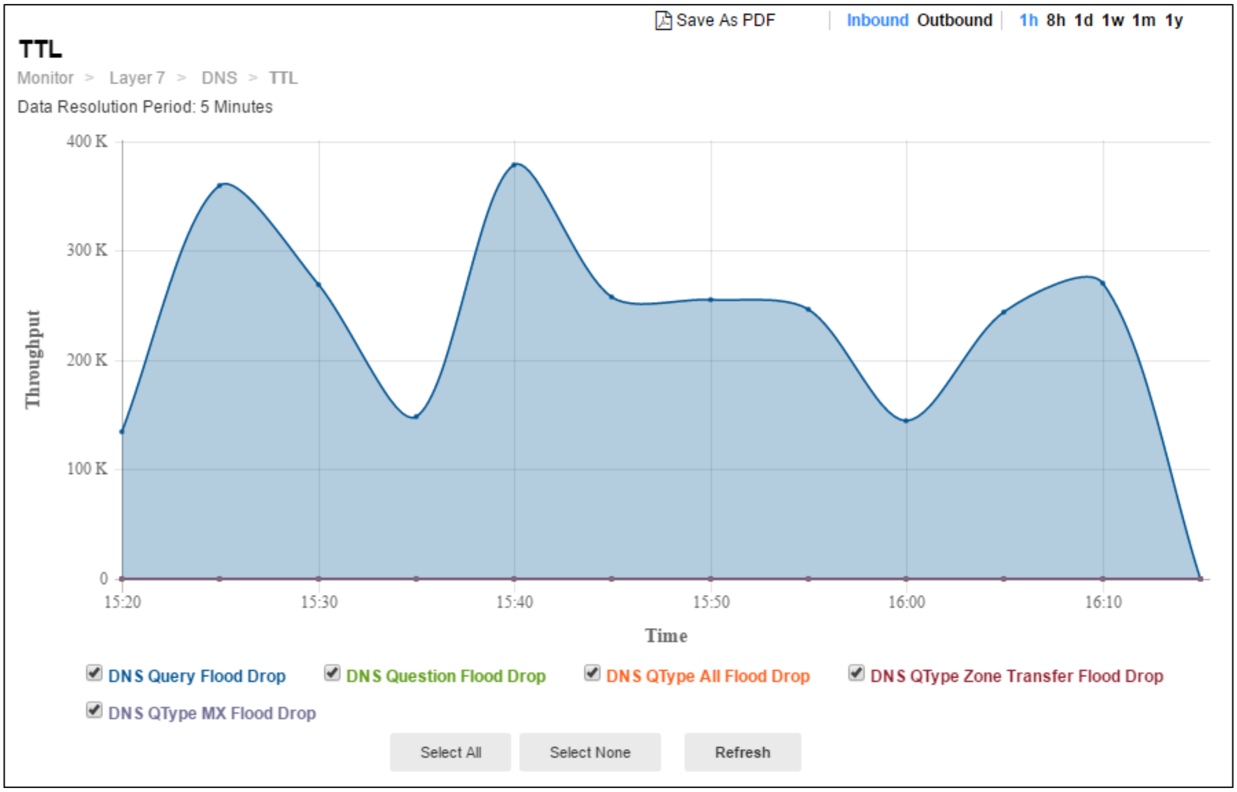
Save As PDF
Inbound Outbound
1h 8h 1d 1w 1m 1y
TTL
Monitor > Layer 7 > DNS > TTL
Data Resolution Period: 5 Minutes
400 K
300 K
200 K
100 K
0
Throughput
15:20
15:30
15:40
15:50
16:00
16:10
Time
☑ DNS Query Flood Drop
☑ DNS Question Flood Drop
☑ DNS QType All Flood Drop
☑ DNS QType Zone Transfer Flood Drop
☑ DNS QType MX Flood Drop
Select All
Select None
Refresh

Figure 111: DNS Cache Drop

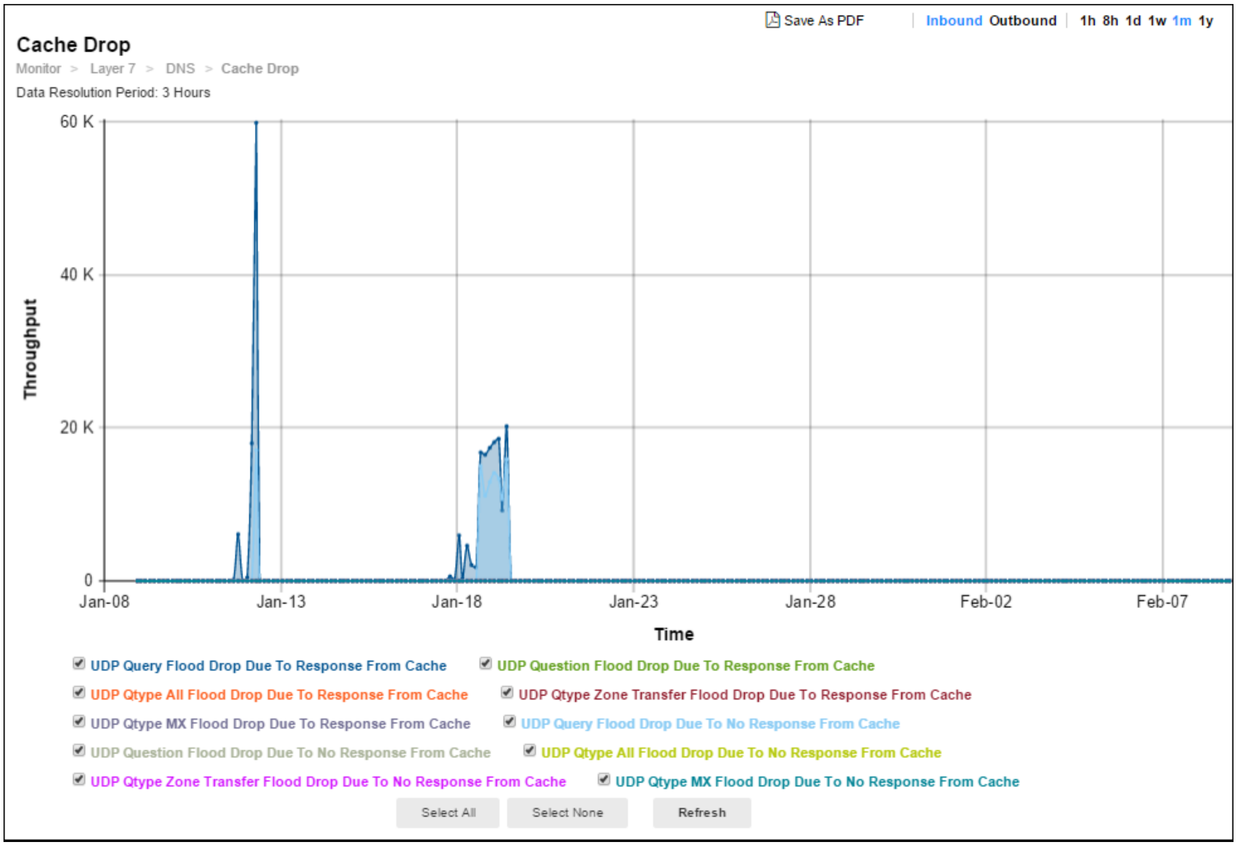
Save As PDF
Inbound Outbound
1h 8h 1d 1w 1m 1y
Cache Drop
Monitor > Layer 7 > DNS > Cache Drop
Data Resolution Period: 3 Hours
60 K
40 K
20 K
0
Throughput
Jan-08
Jan-13
Jan-18
Jan-23
Jan-28
Feb-02
Feb-07
Time
☑ UDP Query Flood Drop Due To Response From Cache
☑ UDP Question Flood Drop Due To Response From Cache
☑ UDP Qtype All Flood Drop Due To Response From Cache
☑ UDP Qtype Zone Transfer Flood Drop Due To Response From Cache
☑ UDP Qtype MX Flood Drop Due To Response From Cache
☑ UDP Query Flood Drop Due To No Response From Cache
☑ UDP Question Flood Drop Due To No Response From Cache
☑ UDP Qtype All Flood Drop Due To No Response From Cache
☑ UDP Qtype Zone Transfer Flood Drop Due To No Response From Cache
☑ UDP Qtype MX Flood Drop Due To No Response From Cache
Select All
Select None
Refresh

Figure 112: DNS Spoofed IP Drop

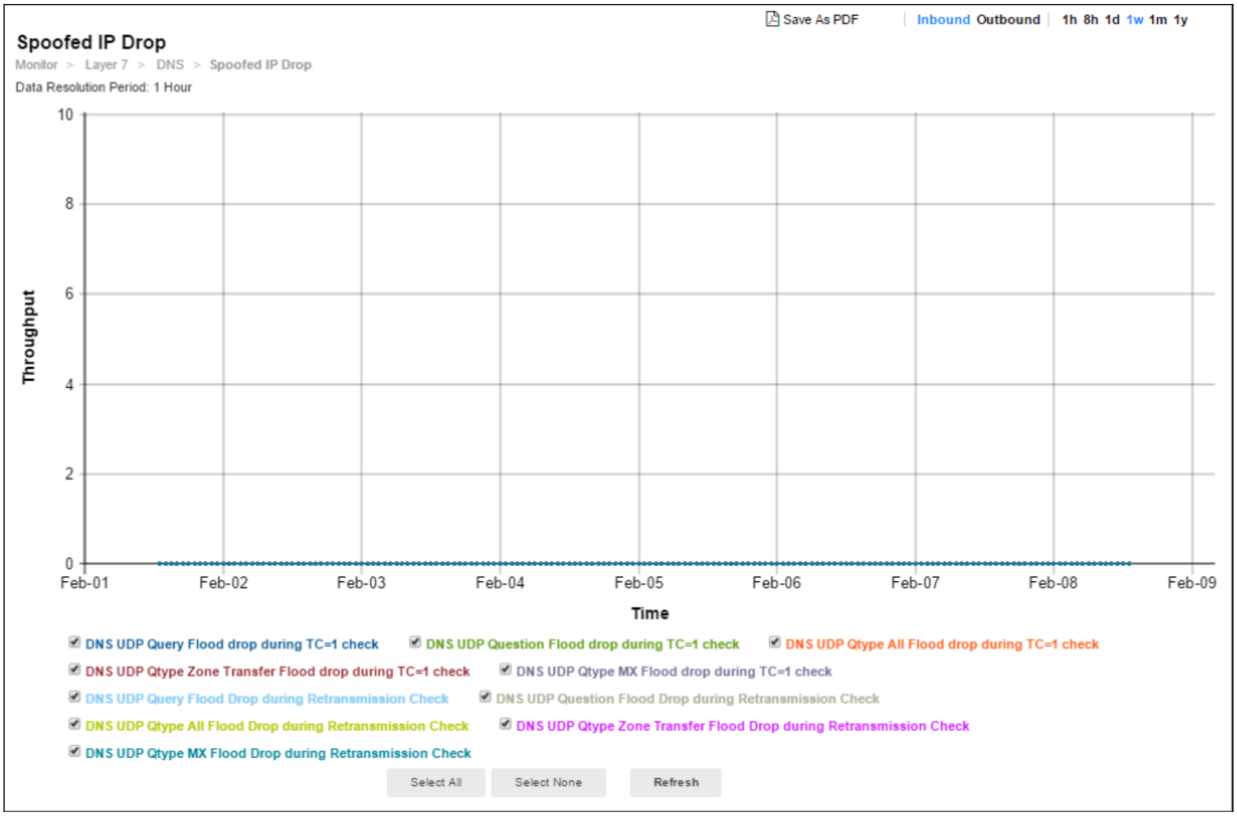
Save As PDF
Inbound Outbound
1h 8h 1d 1w 1m 1y
Spoofed IP Drop
Monitor > Layer 7 > DNS > Spoofed IP Drop
Data Resolution Period: 1 Hour
Throughput
10
8
6
4
2
0
Feb-01
Feb-02
Feb-03
Feb-04
Feb-05
Feb-06
Feb-07
Feb-08
Feb-09
Time
☑ DNS UDP Query Flood drop during TC=1 check
☑ DNS UDP Question Flood drop during TC=1 check
☑ DNS UDP Qtype All Flood drop during TC=1 check
☑ DNS UDP Qtype Zone Transfer Flood drop during TC=1 check
☑ DNS UDP Qtype MX Flood drop during TC=1 check
☑ DNS UDP Query Flood Drop during Retransmission Check
☑ DNS UDP Question Flood Drop during Retransmission Check
☑ DNS UDP Qtype All Flood Drop during Retransmission Check
☑ DNS UDP Qtype Zone Transfer Flood Drop during Retransmission Check
☑ DNS UDP Qtype MX Flood Drop during Retransmission Check
Select All
Select None
Refresh

Figure 113: DNS Rcodes

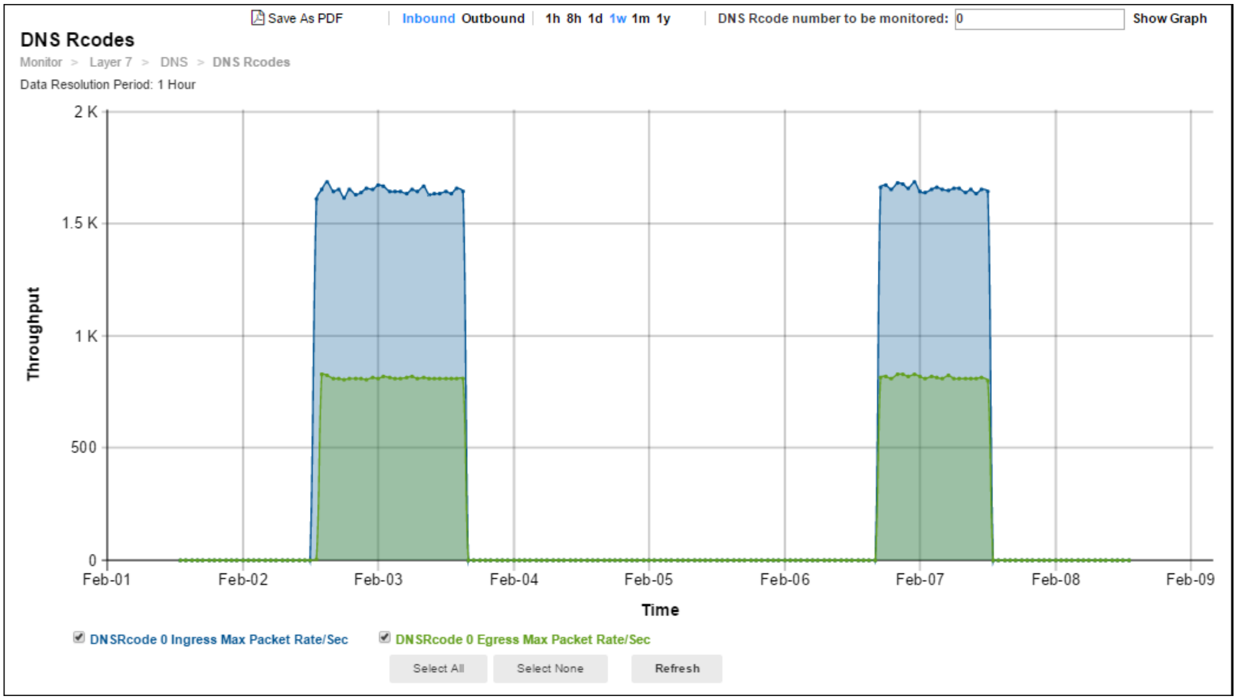
Save As PDF
Inbound Outbound
1h 8h 1d 1w 1m 1y
DNS Rcode number to be monitored: 0
Show Graph
DNS Rcodes
Monitor > Layer 7 > DNS > DNS Rcodes
Data Resolution Period: 1 Hour
Throughput
2 K
1.5 K
1 K
500
0
Feb-01
Feb-02
Feb-03
Feb-04
Feb-05
Feb-06
Feb-07
Feb-08
Feb-09
Time
☑ DNSRcode 0 Ingress Max Packet Rate/Sec
☑ DNSRcode 0 Egress Max Packet Rate/Sec
Select All
Select None
Refresh

Chapter 6: Logs and Reports

This chapter includes the following topics:

Logs and reports overview

The FortiDDoS system supports the logging and reporting features you expect in a security appliance:

- Local logging
- Remote logging (syslog)
- FortiAnalyzer support
- SNMP
- Alerts (SMTP)
- SQL support
- RESTful API support
- Realtime system status and traffic monitoring
- Configurable system event and security event logging
- Sorting and filtering of log tables
- Customizable dashboards
- Pre-defined graphic reports that can be filtered
- Customizable, scheduled reports, with multiple formats and delivery options

Configuring local log settings

The local log is a datastore hosted on the FortiDDoS system. The local log disk configuration applies to both the system event log and the DDoS attack log.

Typically, you use the local log to capture information about system health and system administration activities, to verify that your configuration and tunings behave as expected, and to understand threats in recent traffic periods. It is both standard practice and best practice to send security log data to secure remote servers where it can be stored long term and analyzed using preferred analytic tools.

Local log disk settings are configurable. You can select a subset of system events. The DDoS attack log events are not configurable.

Before you begin:

- You must have Read-Write permission for Log & Report settings.

See also: Using the event log table, Using the DDoS attack log table.

To configure local log settings:

1. Go to Log & Report > Log Configuration > Local Log Settings.
2. Complete the configuration as described in Table 54.
3. Save the configuration.

Figure 114: Local log configuration page

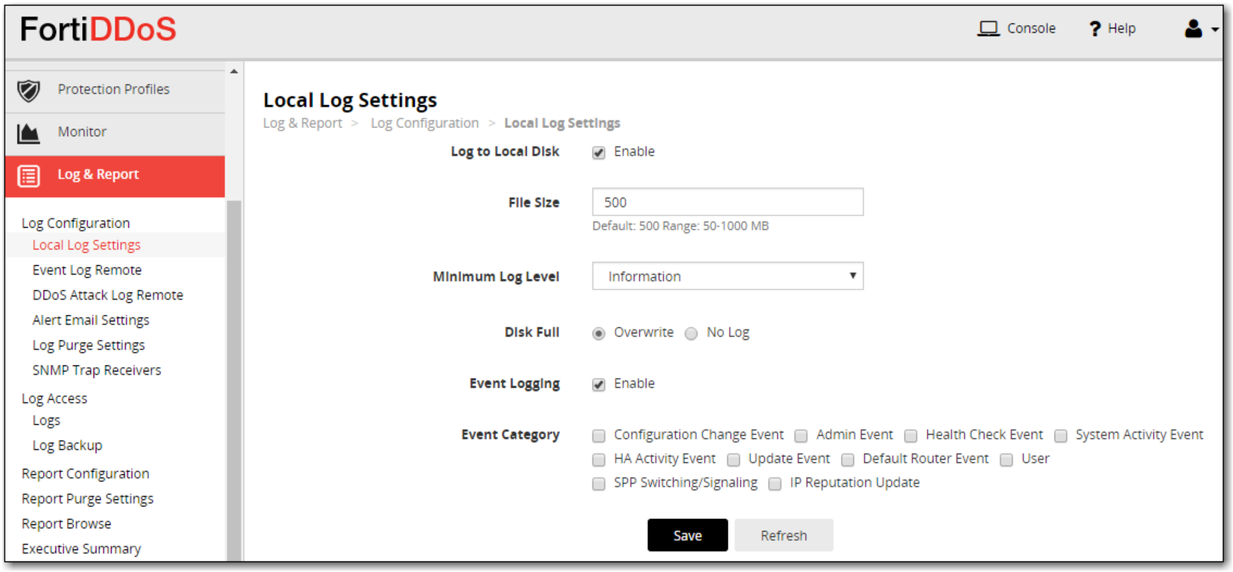

Local Log Settings

Log & Report > Log Configuration > Local Log Settings

Log to Local Disk ☑ Enable

File Size 500

Default: 500 Range: 50-1000 MB

Minimum Log Level Information

Disk Full ● Overwrite ○ No Log

Event Logging ☑ Enable

Event Category □ Configuration Change Event □ Admin Event □ Health Check Event □ System Activity Event □ HA Activity Event □ Update Event □ Default Router Event □ User □ SPP Switching/Signaling □ IP Reputation Update

Save Refresh

Table 54: Local logging configuration guidelines

Settings	Guidelines
Logging and Archiving	
Log to Local Disk	Select to display settings to manage the disk used for logging.
Minimum Log Level	Select the lowest severity to log from the following choices: • Emergency—The system has become unstable. • Alert—Immediate action is required. • Critical—Functionality is affected. • Error—An error condition exists and functionality could be affected. • Warning—Functionality might be affected. • Notification—Information about normal events. • Information—General information about system operations. • Debug—Detailed information about the system that can be used to troubleshoot unexpected behavior. For example, if you select **Error**, the system collects logs with level Error, Critical, Alert, and Emergency. If you select **Alert**, the system collects logs with level Alert and Emergency. The log level setting applies to both system events and DDoS security events. **Tip**: To prolong disk life, do not collect notification, information, and debug level logs for long periods of time.
File Size	Maximum disk space for local logs. The default is 500 MB.
Disk full	Select log behavior when the maximum disk space for local logs is reached: • Overwrite—Continue logging. Overwrite the earliest logs. • No Log—Stop logging.
Event Logging	
Event Logging	Select to enable event logging and then select the types of event category that you want included in the event log.

Configuring remote log server settings for event logs

A remote log server is a system provisioned specifically to collect logs for long term storage and analysis with preferred analytic tools. We recommend FortiAnalyzer.

The system has two configurations to support sending logs to remote log servers: remote log server settings for system event logs, and remote log server settings for DDoS logs.

The system event log configuration applies to system-wide data, such as system health indicators and system administrator activities. The DDoS log configuration applies to security data.

You can configure up to three Log Remote or Remote Event Log Servers.

Before you begin:

- You must have Read-Write permission for Log & Report settings.

See also: Configuring remote log server settings for DDoS attack log.

To configure remote event log settings:

1. Go to Log & Report > Log Configuration > Event Log Remote.
2. Click **Add.**
3. Complete the configuration as described in Table 55.
4. Save the configuration.

Figure 115: Remote log server settings

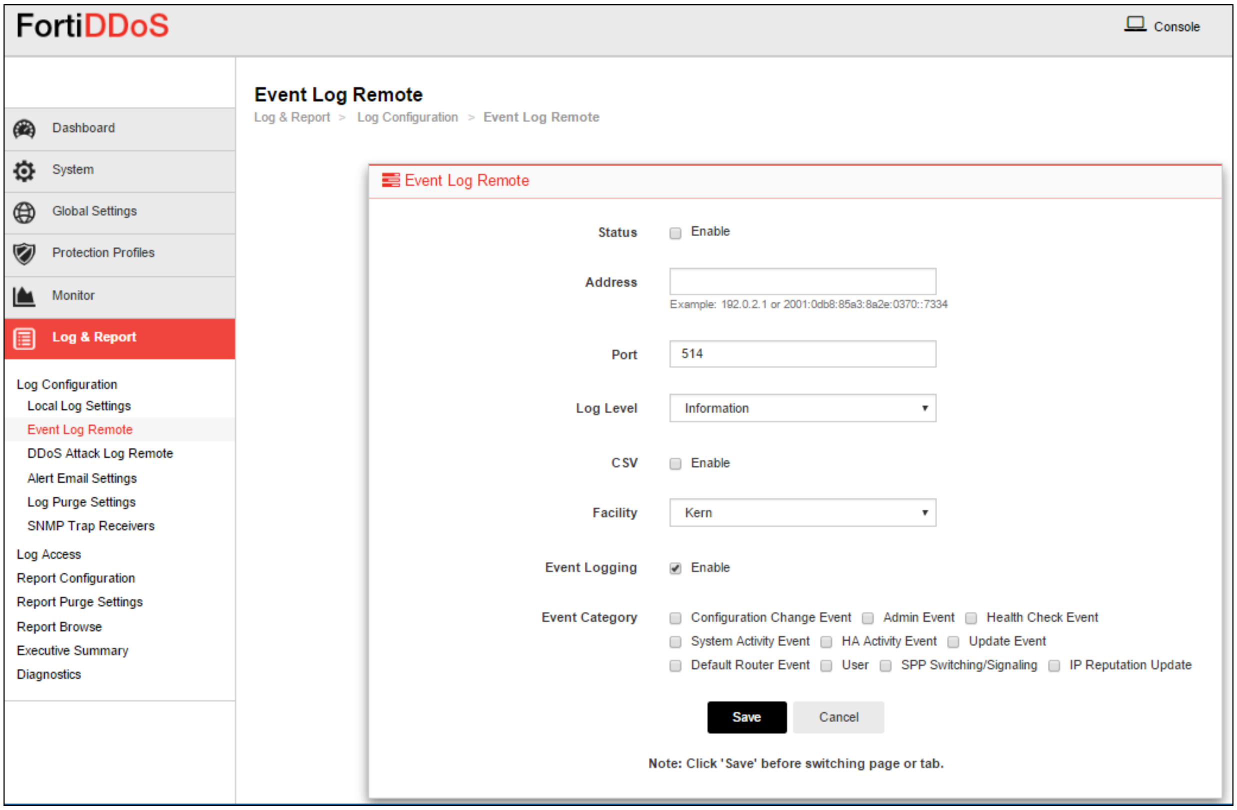

Event Log Remote

Event Log Remote

Status □ Enable

Address

Example: 192.0.2.1 or 2001:0db8:85a3:8a2e:0370::7334

Port 514

Log Level Information ▾

CSV □ Enable

Facility Kern ▾

Event Logging ☑ Enable

Event Category □ Configuration Change Event □ Admin Event □ Health Check Event
□ System Activity Event □ HA Activity Event □ Update Event
□ Default Router Event □ User □ SPP Switching/Signaling □ IP Reputation Update

Save Cancel

Note: Click 'Save' before switching page or tab.

Table 55: Remote log configuration guidelines

Settings	Guidelines
Status	Select to display settings to manage the disk used for logging.
Address	IP address of the FortiAnalyzer or syslog server.
Port	Listening port number of the FortiAnalyzer/syslog server. Usually this is UDP port 514.
Log Level	Select the severity to log from the following choices: • Emergency—The system has become unstable. • Alert—Immediate action is required. • Critical—Functionality is affected. • Error—An error condition exists and functionality could be affected. • Warning—Functionality might be affected. • Notification—Information about normal events. • Information—General information about system operations. • Debug—Detailed information about the system that can be used to troubleshoot unexpected behavior
CSV Format	Send logs in CSV format. Do not use with FortiAnalyzer.

Settings	Guidelines
Minimum Log Level	Select the lowest severity to log from the following choices: • Emergency—The system has become unstable. • Alert—Immediate action is required. • Critical—Functionality is affected. • Error—An error condition exists and functionality could be affected. • Warning—Functionality might be affected. • Notification—Information about normal events. • Information—General information about system operations. • Debug—Detailed information about the system that can be used to troubleshoot unexpected behavior. For example, if you select **Error**, the system sends the syslog server logs with level Error, Critical, Alert, and Emergency. If you select **Alert**, the system collects logs with level Alert and Emergency.
Facility	Identifier that is not used by any other device on your network when sending logs to FortiAnalyzer/syslog.
Event Logging	Select to enable event logging and then select the types of events that you want included in the event log.

The following is an example of an event syslog message:

```
device_id=SYSLOG-AC1E997F type=generic pri=information itime=1431633173 msg="date=2015-05-
    13,time=13:25:13,tz=PDT,devid=FI800B3913000032,log_
    id=0000002168,type=event,subtype=config,level=information,msg_
    id=426204,user=admin,ui=ssh
    (172.30.153.9),action=none,status=none,reason=none,msg='changed settings for 'ddos spp
    setting' on domain 'SPP-1'''"
```

Table 56 identifies the fields in the event syslog message.

Table 56: Event syslog fields

Field	Example
Syslog device ID	device_id=SYSLOG-AC1E997F
Syslog type	type=generic
Syslog log level	pri=information
Syslog time	itime=1431633173
Log datestamp	date=2015-05-13

Field	Example
Log timestamp	13:25:13
Log time zone	tz=PDT
Device ID	devid=FI800B3913000032
Log ID	log_id=0000002168
Log type	type=event
Log subtype	subtype=config
Log level	level=information
Message ID	msg_id=426204
Admin user	user=admin
Admin UI	ui=ssh(172.30.153.9)
Action	action=none
Status	status=none
Reason string	reason=none
Log message	msg='changed settings for 'ddos spp setting' on domain 'SPP-1''"

Configuring remote log server settings for DDoS attack log

The DDoS attack log remote server configuration applies to security event data. You configure individual remote log server configurations for each SPP.

You can set up two remote DDoS Attack Log Remote syslog servers per SPP.

Before you begin:

- You must have Read-Write permission for Log & Report settings.

See also: Configuring remote log server settings for event logs.

To configure remote log settings for the DDoS attack log:

1. Go to Log & Report > Log Configuration > DDoS Attack Log Remote.
2. Click **Add**.
3. Complete the configuration as described in Table 57.
4. Save the configuration.

Figure 116: DDoS Attack Log remote logging configuration page

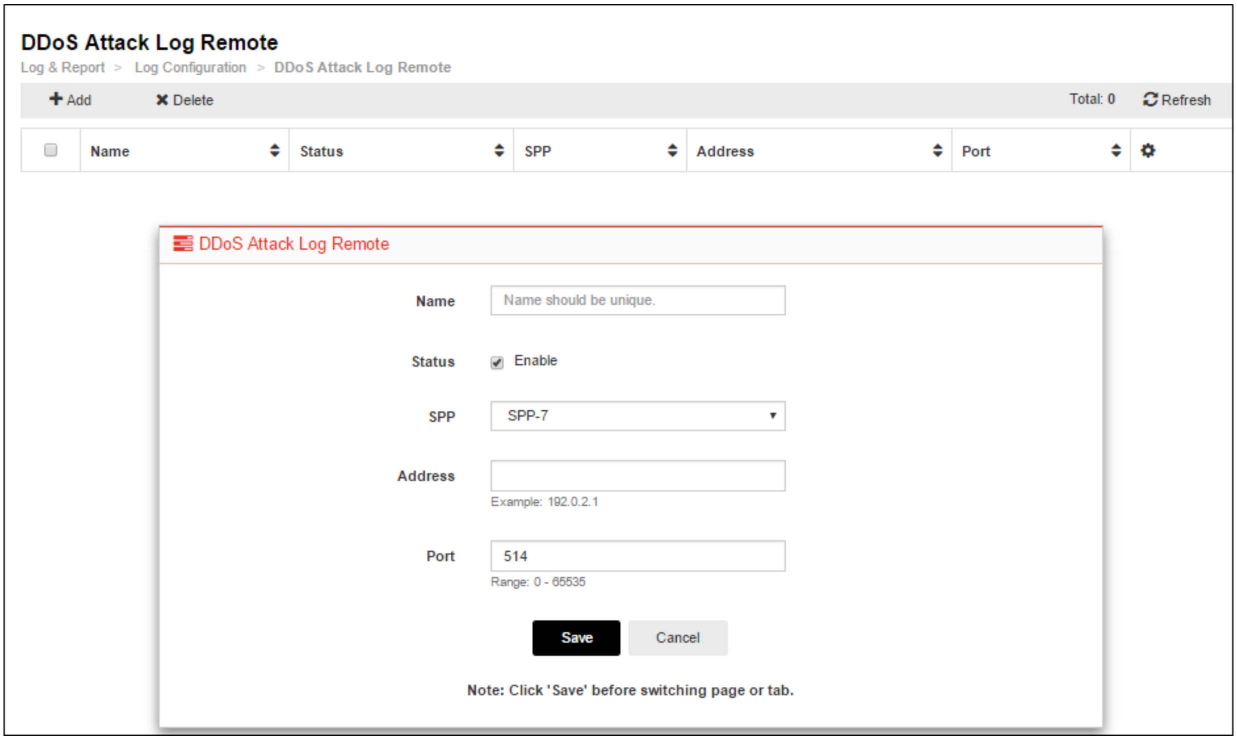

DDoS Attack Log Remote

Log & Report > Log Configuration > DDoS Attack Log Remote

DDoS Attack Log Remote

Name: Name should be unique.

Status: ☑ Enable

SPP: SPP-7

Address:

Example: 192.0.2.1

Port: 514

Range: 0 - 65535

Note: Click 'Save' before switching page or tab.

Table 57: DDoS Attack Log remote logging configuration guidelines

Settings	Guidelines
Name	Configuration name.
Status	Select to enable sending DDoS attack logs to a remote server.
SPP	Select the SPP whose logs are stored in the remote location. You can specify only one remote log server for each SPP.
Address	IP address of the FortiAnalyzer/syslog server.
Port	Listening port number of the FortiAnalyzer/syslog server. Usually this is UDP port 514.

The following is an example of a DDoS attack syslog message:

```
Apr 24 13:22:08 192.168.205.202 devid=FI800B3913000004 date=2015-04-23 time=01:10:00
    tz=PDT type=attack spp=0 evecode=1 evesubcode=14 description="IP Header checksum
    error" dir=1 sip=0.0.0.0 dip=10.10.0.1 subnet_name=VID0 subnet_comment=Dept_0
    dropcount=684138
```

Table 58 identifies the fields in the DDoS attack syslog message.

Table 58: DDoS attack syslog fields

Field	Example
Syslog send timestamp	Apr 24 13:22:08
Syslog server IP address	192.168.205.202
Device ID	devid=FI800B3913000004
Log datestamp	date=2015-04-23
Log timestamp	time=01:10:00
Log time zone	tz=PDT
Log type	type=attack
SPP ID	spp=0
Event code	evecode=1
Event subcode	evesubcode=14
Event type	description="IP Header checksum error"
Direction (1=inbound, 0=outbound)	dir=1

Field	Example
Source IP address	sip=0.0.0.0
Destination IP address	dip=10.10.0.1
Subnet name	subnet_name=VID0
Subnet comment	subnet_comment=Dept_0
Drop count	dropcount=684138

See Appendix A: DDoS Attack Log Reference for details on log categories and event types.

Using FortiAnalyzer to collect DDoS attack logs

FortiAnalyzer platforms integrate network logging, analysis, and reporting into a single system, delivering increased knowledge of security events throughout your network.

FortiAnalyzer now supports the FortiDDoS attack log. FortiAnalyzer includes the following predefined reports for FortiDDoS:

- Attacks by time period
- Attackers by time period
- Top 20 Attacks
- To 20 Attack Types

Refer to FortiAnalyzer documentation for version support details and detailed procedures on how to use FortiAnalyzer. This section describes the workflow for collecting DDoS attack logs.

To set up log collection:

1. On FortiAnalyzer, go to the System Information widget and enable Administrative Domains.
2. On FortiDDoS, use the DDos Attack Log Remote configuration to send logs to the FortiAnalyzer IP address.

   After you have saved the configuration, FortiDDoS begins sending logs to FortiAnalyzer.
3. On FortiAnalyzer, go to the Device Manager. Once FortiAnalyzer begins receiving logs from FortiDDoS, FortiDDoS appears in the Administrative Domains (ADOM).
4. Select FortiDDoS and click **Add Device** to start the Add Device wizard. Complete the wizard.
5. Go to the Device Manager and verify that the FortiDDoS device has been added.
6. Once the device has been successfully added, go to FortiView to see the attack log.

Using FortiSIEM to collect DDoS attack and event logs

FortiSIEM provides an all-in-one, seamlessly integrated and service-oriented IT infrastructure monitoring solution that covers performance, availability, change, and security monitoring aspects of network devices, servers, and applications.

FortiSIEM now supports FortiDDoS attack and event logs. FortiSIEM processes FortiDDoS events via syslog. You can configure FortiDDoS to send syslog to FortiSIEM.

Refer to FortiSIEM User Guide for version support details and detailed procedures on how to use FortiSIEM. This section describes the workflow for collecting DDoS attack logs.

To set up log collection:

1. On FortiDDoS, use **DDoS Attack Log Remote** configuration to send logs to the FortiSIEM IP address.

   Refer to section Configuring remote log server settings for DDoS attack logs and follow the procedure for configuration. Once the configuration is saved, FortiDDoS begins sending logs to FortiSIEM.

2. Use **Event Log Remote** configuration to send logs to the FortiSIEM IP address.

   Refer to section Configuring remote log server settings for event logs and follow the procedure for configuration. Once the configuration is saved, FortiDDoS begins sending event logs to FortiSIEM.

3. Go to **System** > **SNMP** and follow the steps under Configuring SNMP for system event reporting.
4. Log in to FortiSIEM and go to **Admin** > **Setup Wizard** > **Credentials** tab.
5. Click **Add** under **Step 1: Enter Credentials** and enter the details of the device in the **Access Method Definition** dialog box.

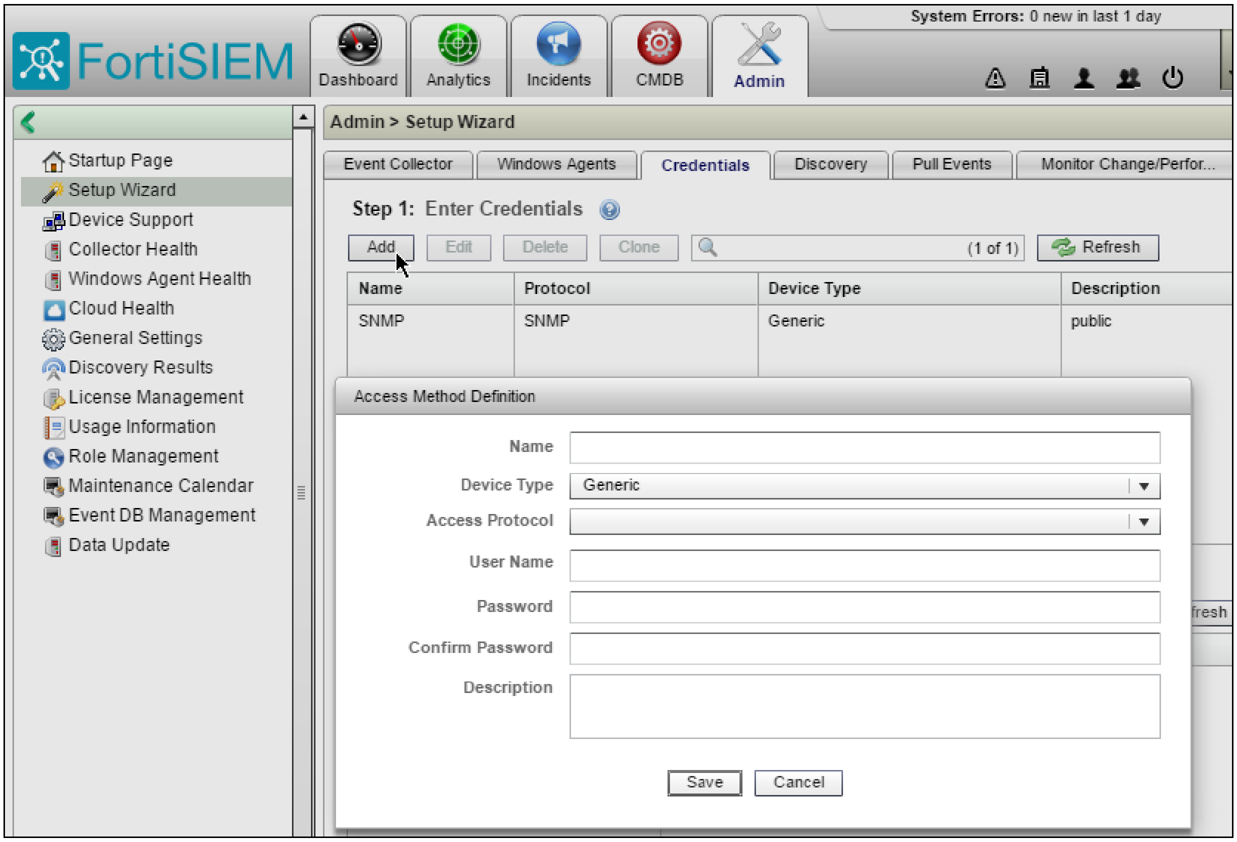

Step 1: Enter Credentials

Name	Protocol	Device Type	Description
SNMP	SNMP	Generic	public

Access Method Definition

Name	
Device Type	Generic
Access Protocol	
User Name	
Password	
Confirm Password	
Description	

6. Click **Add** under **Step 2: Enter IP Range to Credential Association** and enter the **IP/IP Range** and **Credentials** of the device in the **Device Credential Mapping Definition** dialog box.

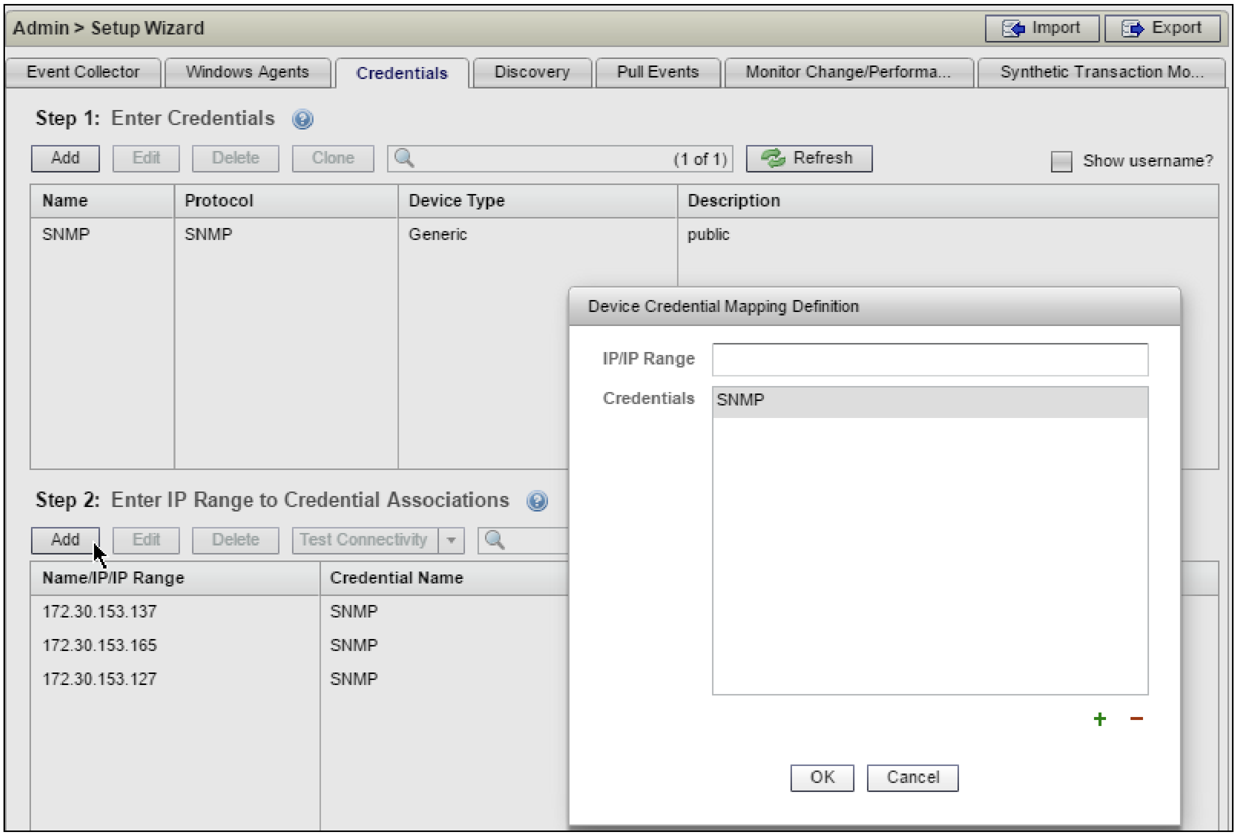

Step 1: Enter Credentials

Name	Protocol	Device Type	Description
SNMP	SNMP	Generic	public

Step 2: Enter IP Range to Credential Associations

Name/IP/IP Range	Credential Name
172.30.153.137	SNMP
172.30.153.165	SNMP
172.30.153.127	SNMP

Device Credential Mapping Definition

IP/IP Range

Credentials SNMP

7. Go to **Admin**> **Setup Wizard** > **Discovery** tab and add the **Range Definition** details.

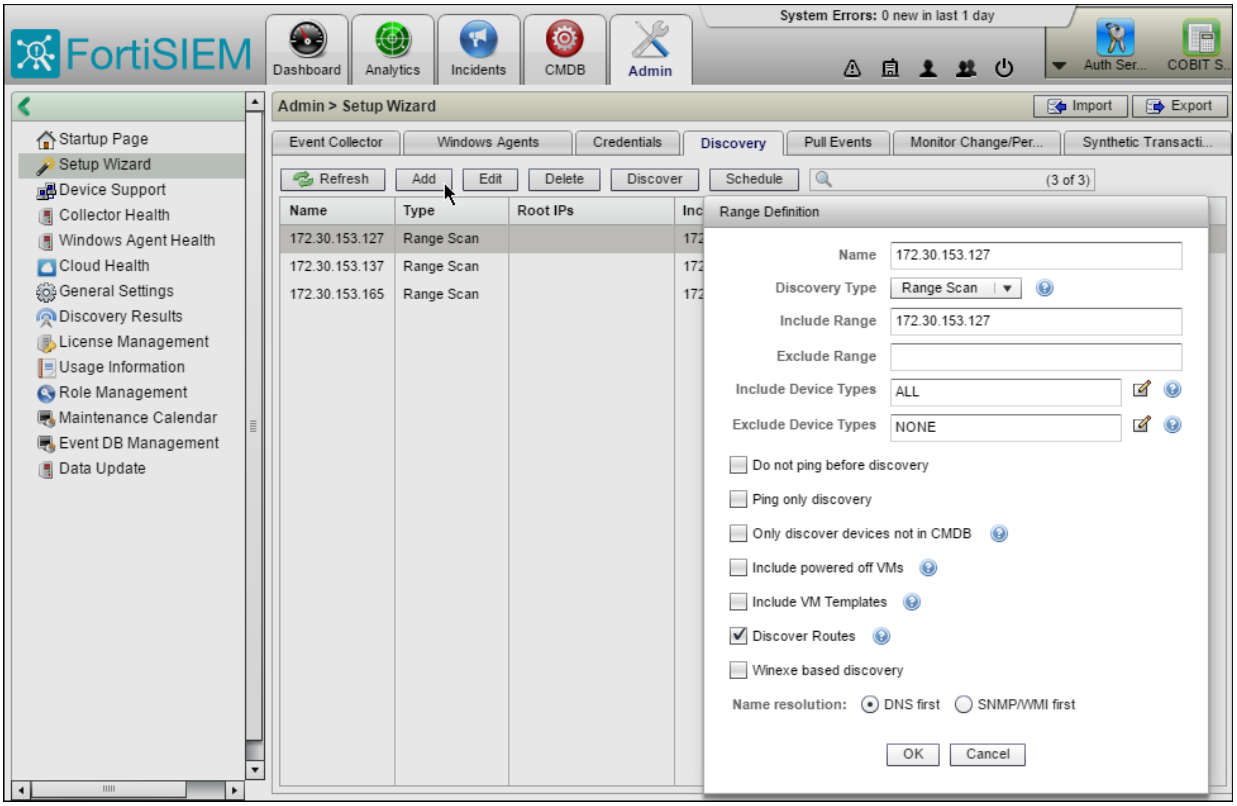

Name	Type	Root IPs	Inc
172.30.153.127	Range Scan		172
172.30.153.137	Range Scan		172
172.30.153.165	Range Scan		172

Range Definition

Name: 172.30.153.127

Discovery Type: Range Scan

Include Range: 172.30.153.127

Exclude Range:

Include Device Types: ALL

Exclude Device Types: NONE

☐ Do not ping before discovery

☐ Ping only discovery

☐ Only discover devices not in CMDB

☐ Include powered off VMs

☐ Include VM Templates

☑ Discover Routes

☐ Winexe based discovery

Name resolution: ● DNS first ○ SNMP/WMI first

8. Select the added range and run discovery by clicking **Discover**.

9. Go to **Admin** > **Discovery Results** and verify the discovered FortiDDoS devices from the list.

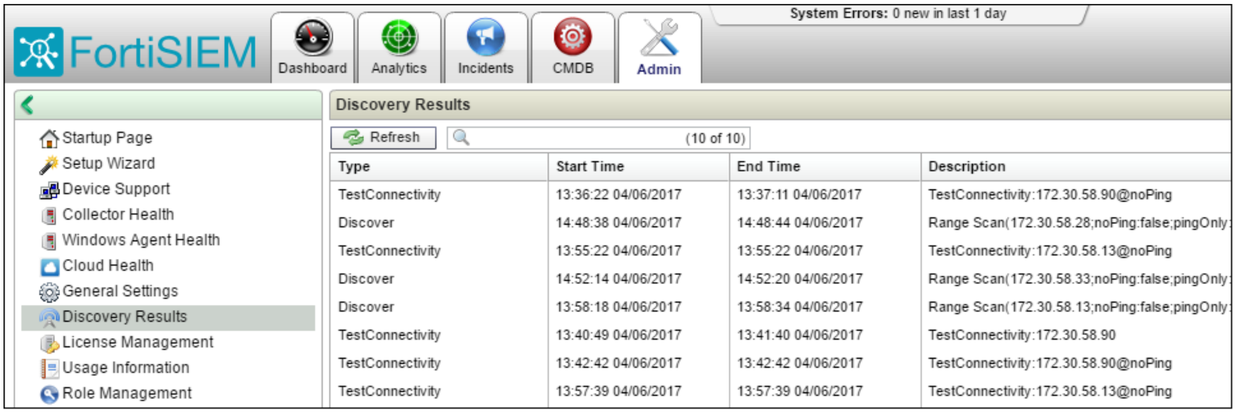

Type	Start Time	End Time	Description
TestConnectivity	13:36:22 04/06/2017	13:37:11 04/06/2017	TestConnectivity:172.30.58.90@noPing
Discover	14:48:38 04/06/2017	14:48:44 04/06/2017	Range Scan(172.30.58.28;noPing:false;pingOnly:
TestConnectivity	13:55:22 04/06/2017	13:55:22 04/06/2017	TestConnectivity:172.30.58.13@noPing
Discover	14:52:14 04/06/2017	14:52:20 04/06/2017	Range Scan(172.30.58.33;noPing:false;pingOnly:
Discover	13:58:18 04/06/2017	13:58:34 04/06/2017	Range Scan(172.30.58.13;noPing:false;pingOnly:
TestConnectivity	13:40:49 04/06/2017	13:41:40 04/06/2017	TestConnectivity:172.30.58.90
TestConnectivity	13:42:42 04/06/2017	13:42:42 04/06/2017	TestConnectivity:172.30.58.90@noPing
TestConnectivity	13:57:39 04/06/2017	13:57:39 04/06/2017	TestConnectivity:172.30.58.13@noPing

10. Go to **CMDB** > **Devices.** Select the added device from the list and click **Approve**.

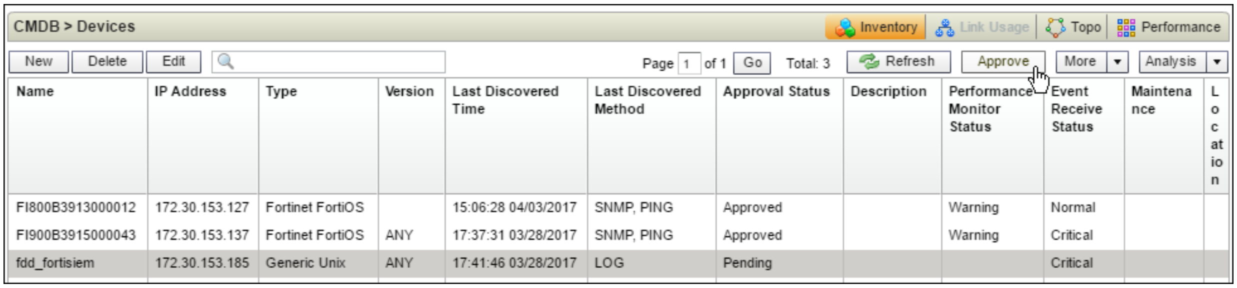

Name	IP Address	Type	Version	Last Discovered Time	Last Discovered Method	Approval Status	Description	Performance Monitor Status	Event Receive Status	Maintenance	Location
FI800B3913000012	172.30.153.127	Fortinet FortiOS		15:06:28 04/03/2017	SNMP, PING	Approved		Warning	Normal		
FI900B3915000043	172.30.153.137	Fortinet FortiOS	ANY	17:37:31 03/28/2017	SNMP, PING	Approved		Warning	Critical		
fdd_fortisiem	172.30.153.185	Generic Unix	ANY	17:41:46 03/28/2017	LOG	Pending			Critical		

11. Go to **Analytics** > **Reports** and click **New** to configure a new report.

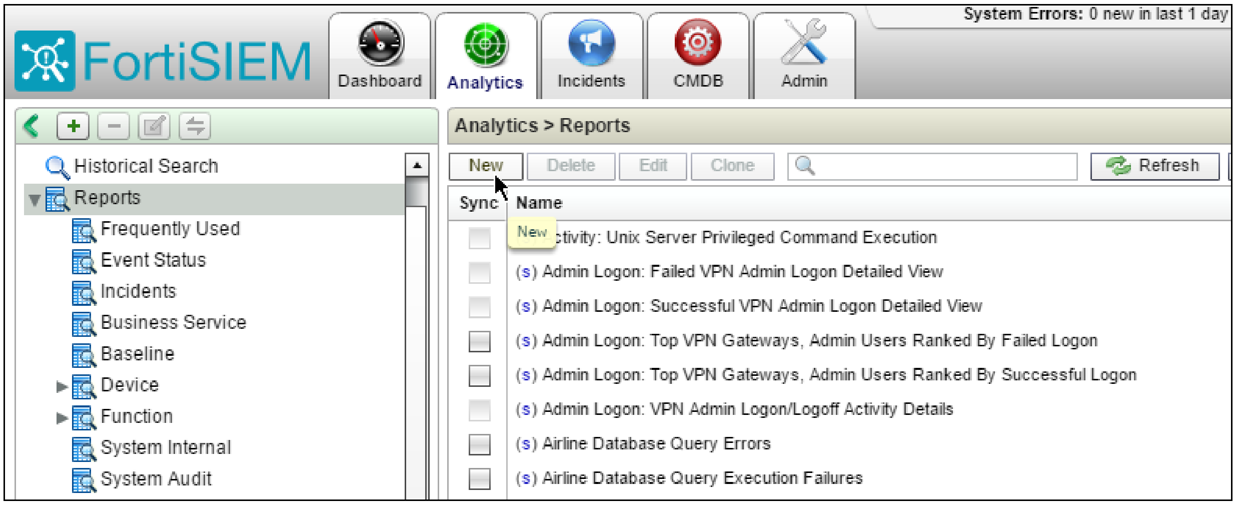

Sync	Name
□	ctivity: Unix Server Privileged Command Execution
□	(s) Admin Logon: Failed VPN Admin Logon Detailed View
□	(s) Admin Logon: Successful VPN Admin Logon Detailed View
□	(s) Admin Logon: Top VPN Gateways, Admin Users Ranked By Failed Logon
□	(s) Admin Logon: Top VPN Gateways, Admin Users Ranked By Successful Logon
□	(s) Admin Logon: VPN Admin Logon/Logoff Activity Details
□	(s) Airline Database Query Errors
□	(s) Airline Database Query Execution Failures

12. Enter the new report details in the **Add New Report** window and click **Save**.

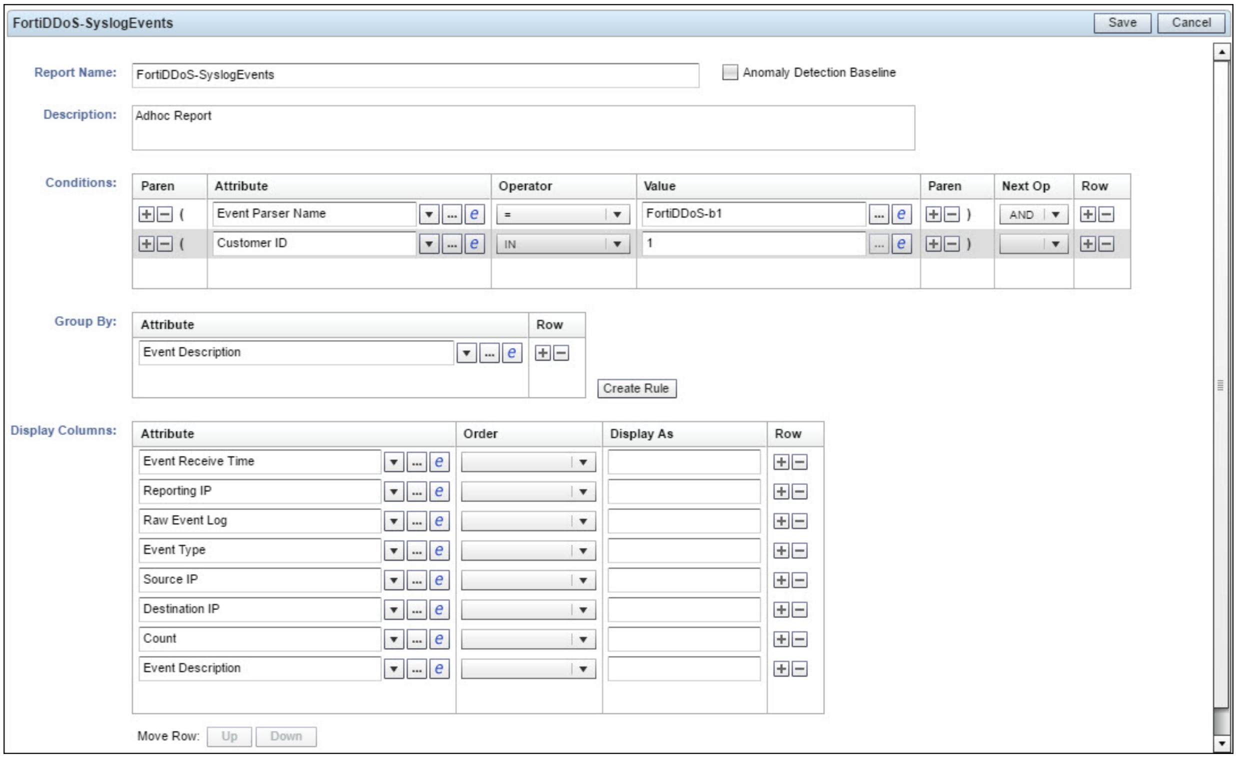

Report Name: FortiDDoS-SyslogEvents ☐ Anomaly Detection Baseline

Description: Adhoc Report

Conditions:

Paren	Attribute	Operator	Value	Paren	Next Op	Row
(	Event Parser Name	=	FortiDDoS-b1	)	AND	
(	Customer ID	IN	1	)		

Group By:

Attribute	Row
Event Description	

Display Columns:

Attribute	Order	Display As	Row
Event Receive Time			
Reporting IP			
Raw Event Log			
Event Type			
Source IP			
Destination IP			
Count			
Event Description			

13. Go to **Dashboard** > **Dashboard by Function**. Select the group and click **Add Reports to Dashboard**.

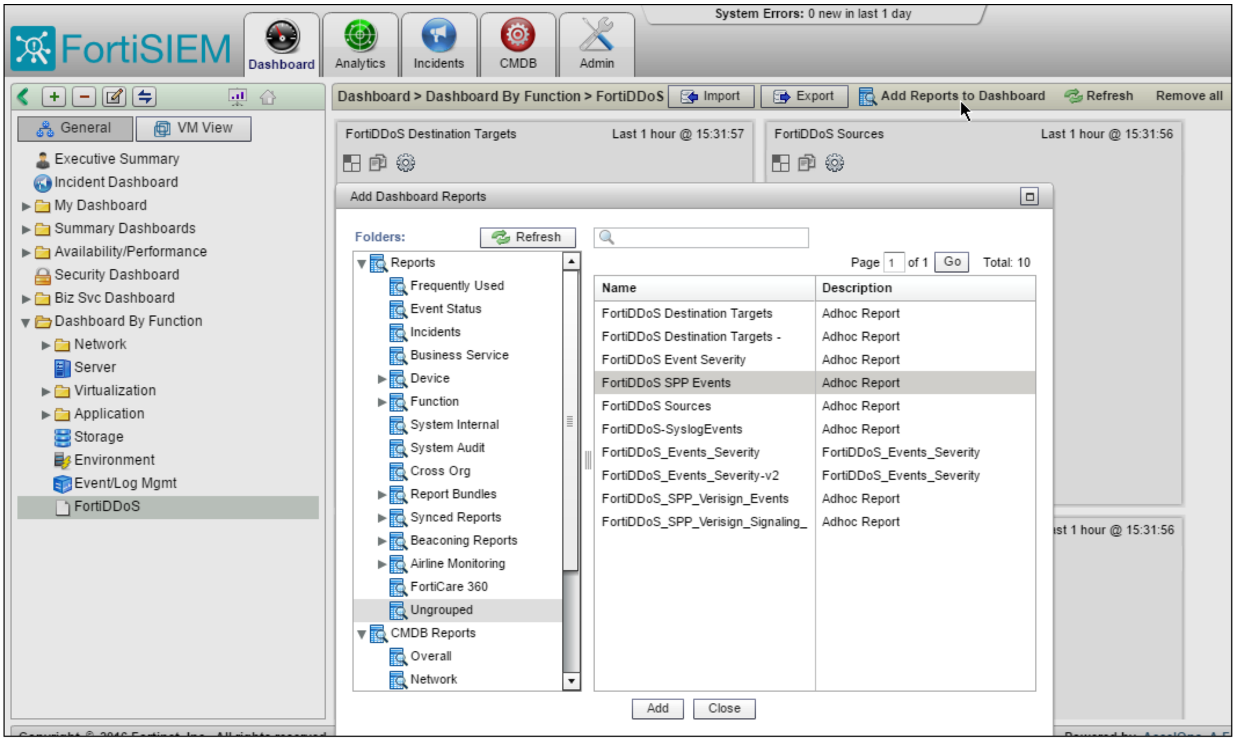

Name	Description
FortiDDoS Destination Targets	Adhoc Report
FortiDDoS Destination Targets -	Adhoc Report
FortiDDoS Event Severity	Adhoc Report
FortiDDoS SPP Events	Adhoc Report
FortiDDoS Sources	Adhoc Report
FortiDDoS-SyslogEvents	Adhoc Report
FortiDDoS_Events_Severity	FortiDDoS_Events_Severity
FortiDDoS_Events_Severity-v2	FortiDDoS_Events_Severity
FortiDDoS_SPP_Verisign_Events	Adhoc Report
FortiDDoS_SPP_Verisign_Signaling_	Adhoc Report

14. Select the required reports from the list and click **Add**.

15. Go to **Dashboard** > **Executive Summary** to see the selected reports. The following figures show the sample dashboard reports.

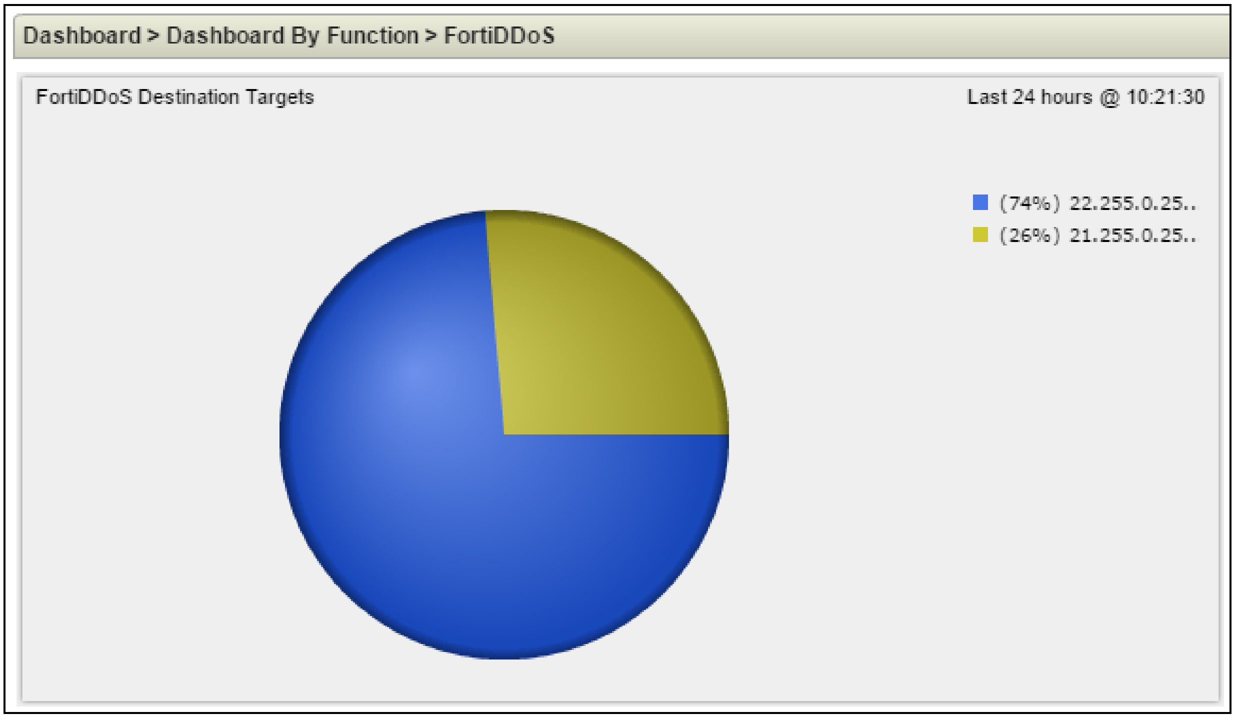
Dashboard > Dashboard By Function > FortiDDoS
FortiDDoS Destination Targets
Last 24 hours @ 10:21:30
(74%) 22.255.0.25..
(26%) 21.255.0.25..

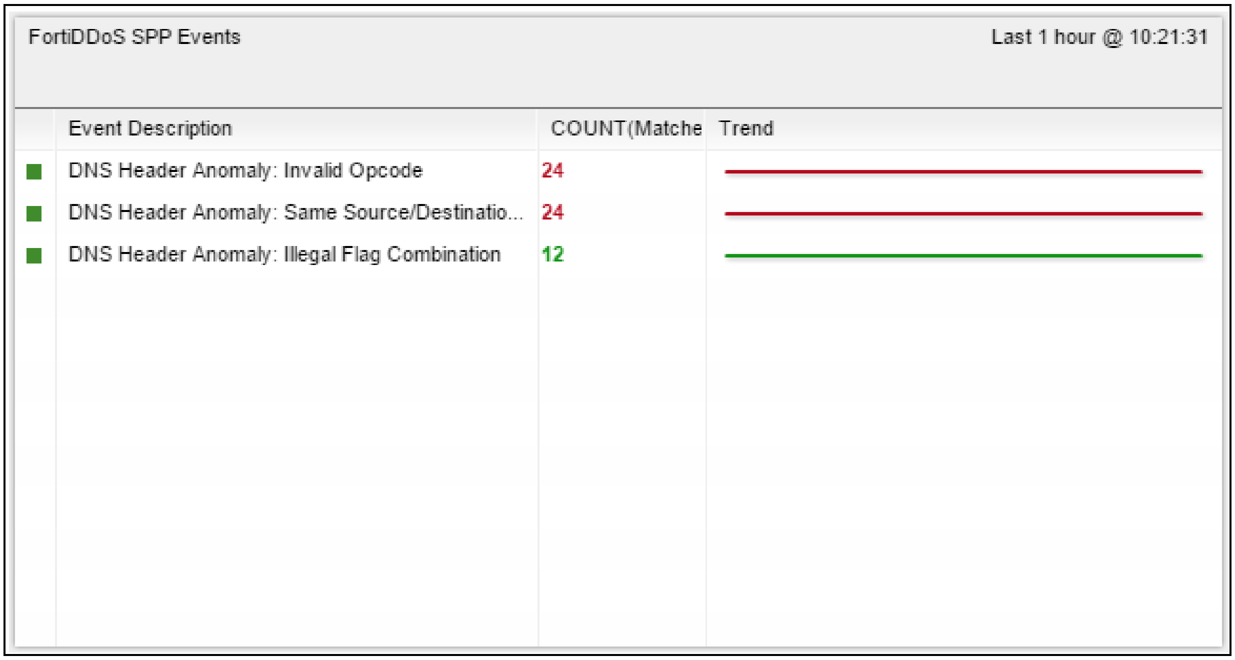

FortiDDoS SPP Events

Last 1 hour @ 10:21:31

	Event Description	COUNT(Matche	Trend
🟩	DNS Header Anomaly: Invalid Opcode	24	
🟩	DNS Header Anomaly: Same Source/Destinatio...	24	
🟩	DNS Header Anomaly: Illegal Flag Combination	12	

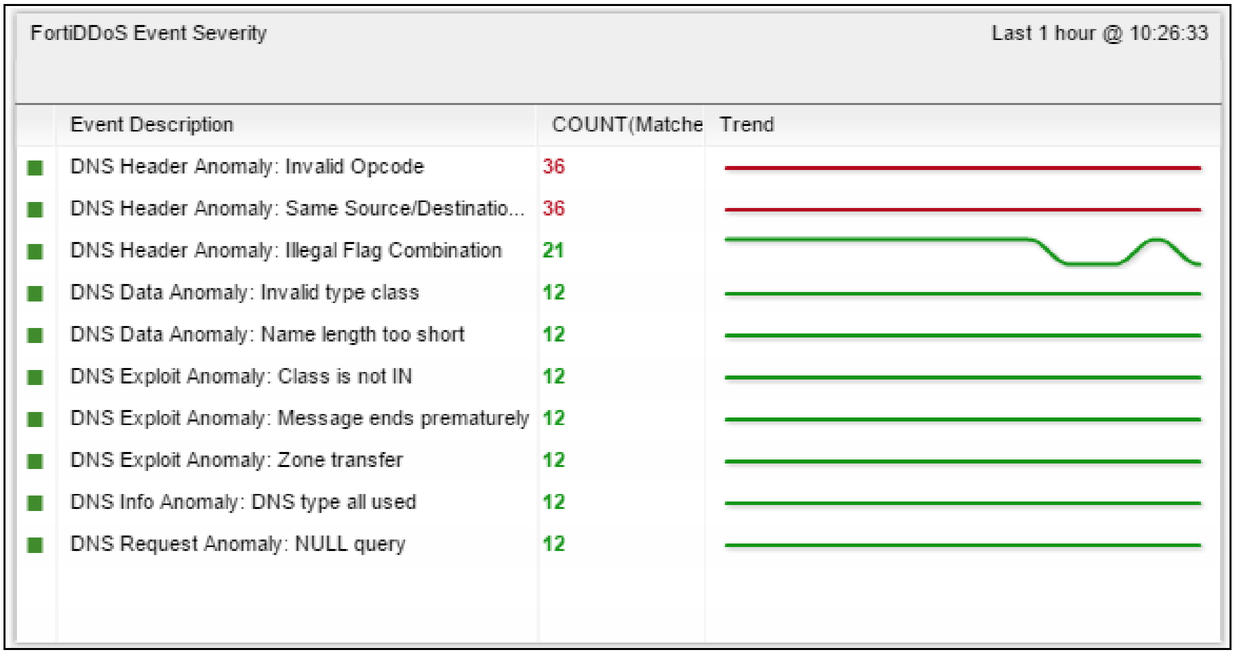

FortiDDoS Event Severity Last 1 hour @ 10:26:33

	Event Description	COUNT(Matche	Trend
🟩	DNS Header Anomaly: Invalid Opcode	36	
🟩	DNS Header Anomaly: Same Source/Destinatio...	36	
🟩	DNS Header Anomaly: Illegal Flag Combination	21	
🟩	DNS Data Anomaly: Invalid type class	12	
🟩	DNS Data Anomaly: Name length too short	12	
🟩	DNS Exploit Anomaly: Class is not IN	12	
🟩	DNS Exploit Anomaly: Message ends prematurely	12	
🟩	DNS Exploit Anomaly: Zone transfer	12	
🟩	DNS Info Anomaly: DNS type all used	12	
🟩	DNS Request Anomaly: NULL query	12	

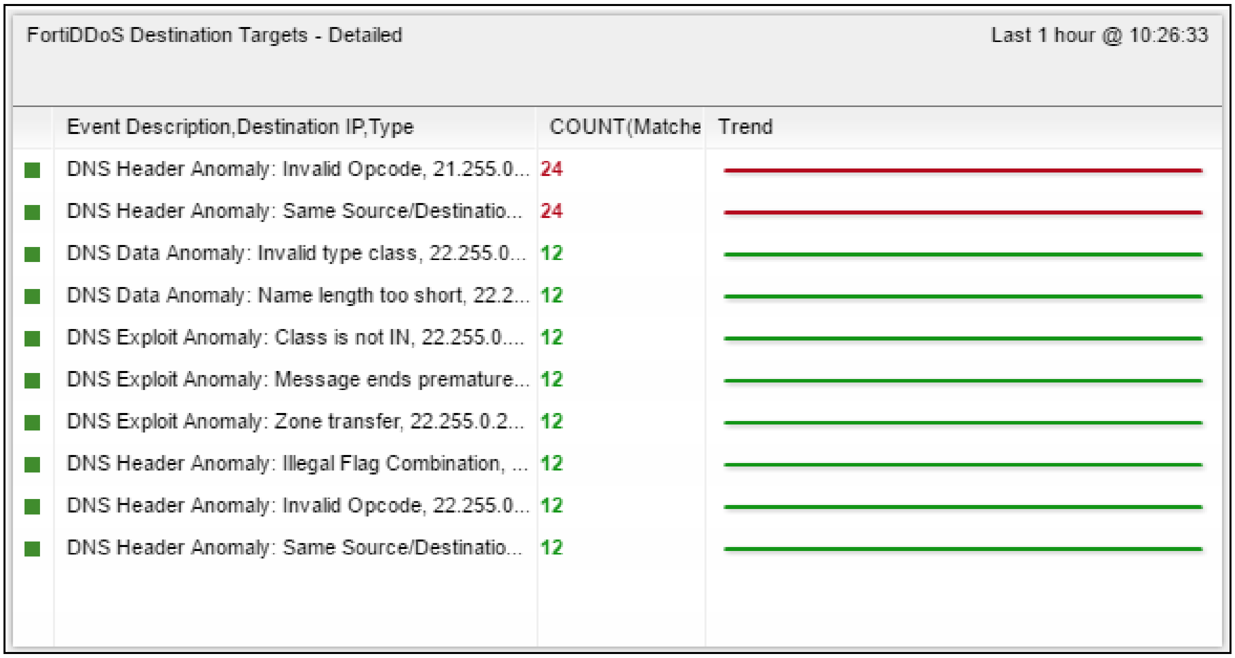

FortiDDoS Destination Targets - Detailed Last 1 hour @ 10:26:33

	Event Description,Destination IP,Type	COUNT(Matche	Trend
🟩	DNS Header Anomaly: Invalid Opcode, 21.255.0...	24	
🟩	DNS Header Anomaly: Same Source/Destinatio...	24	
🟩	DNS Data Anomaly: Invalid type class, 22.255.0...	12	
🟩	DNS Data Anomaly: Name length too short, 22.2...	12	
🟩	DNS Exploit Anomaly: Class is not IN, 22.255.0....	12	
🟩	DNS Exploit Anomaly: Message ends premature...	12	
🟩	DNS Exploit Anomaly: Zone transfer, 22.255.0.2...	12	
🟩	DNS Header Anomaly: Illegal Flag Combination, ...	12	
🟩	DNS Header Anomaly: Invalid Opcode, 22.255.0...	12	
🟩	DNS Header Anomaly: Same Source/Destinatio...	12	

Configuring alert email settings

Alerts are emails sent to specified addresses when specified events are triggered.

You can specify whether event severity or event category is the basis for your alerts configuration.

Before you begin:

- You must have Read-Write permission for Log & Report settings.

To configure alert email settings:

1. Go to Log & Report > Log Configuration > Alert Email Settings.
2. Complete the configuration under the tabs: Mail Server, Settings and Recipients as described in Table 59.
3. Save the configuration.

Figure 117: Alert email settings

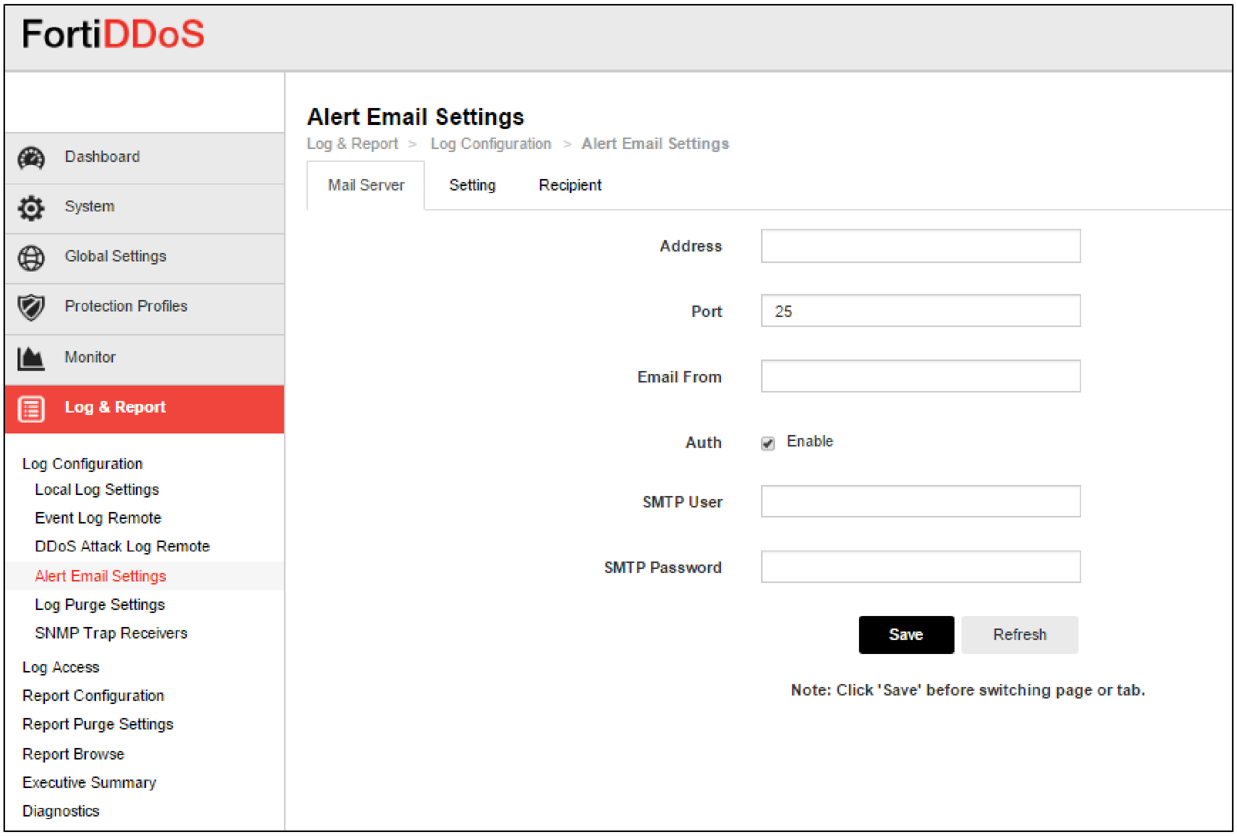

Address	
Port	25
Email From	
Auth	☑ Enable
SMTP User	
SMTP Password	

Note: Click 'Save' before switching page or tab.

Table 59: Alert mail configuration guidelines

Settings	Guidelines
Mail Server	
SMTP Server	IP address or FQDN of an SMTP server (such as FortiMail) or email server that the appliance can connect to in order to send alerts and/or generated reports.
Port	Listening port number of the server. Usually, SMTP is 25.
Email From	Sender email address used in alert email.
Authentication	Enable or disable authentication.
SMTP Username	Username for authentication to the SMTP server.
SMTP Password	Password for authentication to the SMTP server.
Settings	
By Category	Select to enable email alerts based on category.
Category	Select the categories for receiving alerts.
Interval time (min)	If identical alerts are occurring continuously, select the interval between each email that will be sent while the event continues.
Recipient	
Name	Name of the recipient.
Mail To	Up to three recipient email addresses, one per field. **Tip**: To temporarily disable alert emails, delete all recipients. This allows you to preserve the other SMTP settings in case you want to enable alert emails in the future.

Configuring log purge settings

Log purging is the deleting of logs to preserve log space and maintain log system performance.

By default, DDoS attack event logs are purged on a first-in, first-out basis when the log reaches 1,000,000 entries. Log purge settings are configurable. You can specify a different threshold, and you can purge logs manually.

Before you begin:

- You must have Read-Write permission for Log & Report settings.

To configure purge settings:

1. Go to Log & Report > Log Configuration > Log Purge Settings.
2. Complete the configuration as described in Table 60.
3. Save the configuration.

Figure 118: Log purge settings

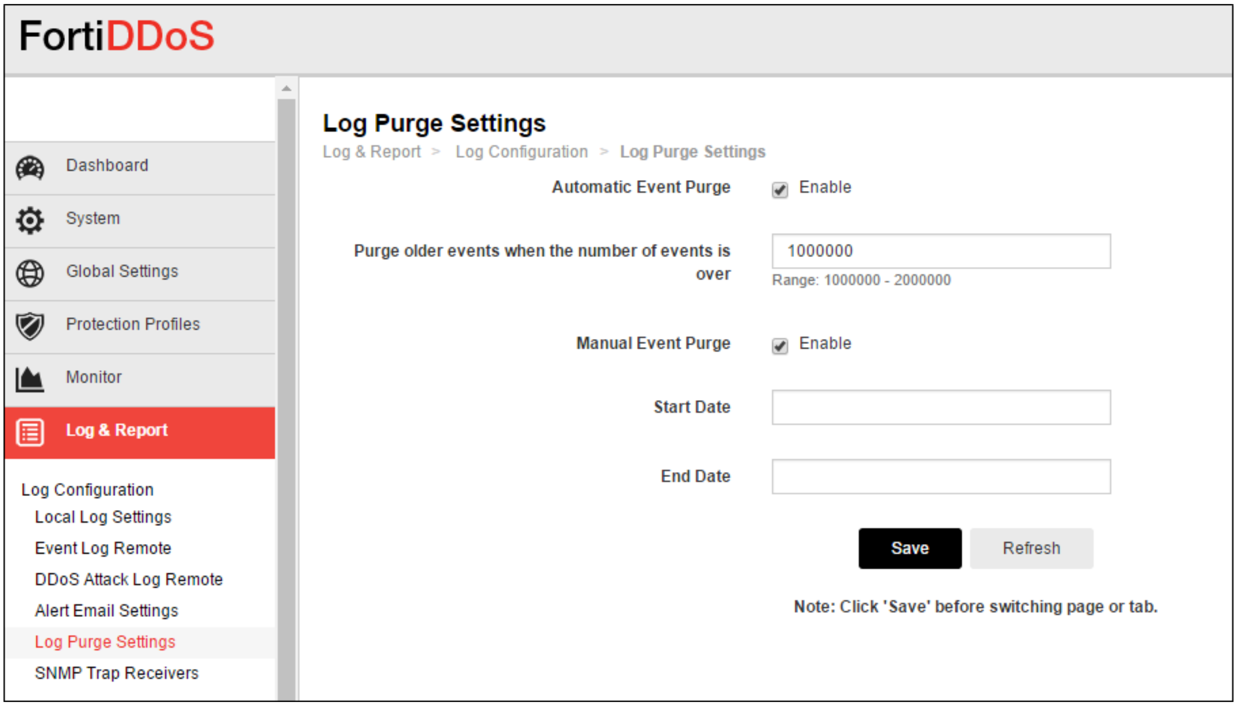

Log Purge Settings

Log & Report > Log Configuration > Log Purge Settings

Automatic Event Purge ☑ Enable

Purge older events when the number of events is over 1000000

Range: 1000000 - 2000000

Manual Event Purge ☑ Enable

Start Date

End Date

Save Refresh

Note: Click 'Save' before switching page or tab.

Table 60: Log purge settings configuration guidelines

Settings	Guidelines
Automatic Event Purge	Select to automatically purge logs after the max number of entries is reached.
Purge older events when the number of events is over	Purge the earliest events from the attack log when this threshold is reached. The default is 1,000,000 entries.
Manual Event Purge	Select to purge entries logged during the specified period.
Start Date / End Date	Specify a period when purging logs manually. The period begins at 0:00 on the start date and ends at 23:59 on the end date.

CLI commands:

```
config ddos global attack-event-purge
   [set automatic-event-purge {enable | disable}
   [set purge-older-events-when-the-number-of-events-is-
       over <int>]
   [set manual-event-purge {enable | disable}]
   [set purge-start-date <purge_date_str>]
   [set purge-end-date <purge_date_str>]
end
```

Configuring SNMP for system event reporting

Many organizations use SNMP (simple network management protocol) to track the health of their systems. SNMP is a commonly used protocol for communication between SNMP agents that reside on network nodes and an SNMP manager that resides on a management host.

An SNMP community is a grouping of equipment for network monitoring purposes. The FortiDDoS SNMP agent does not respond to SNMP managers whose query packets do not contain a matching community name. Similarly, trap packets from the FortiDDoS agent include community name, and an SNMP manager might not accept the trap if its community name does not match.

⚠ Fortinet strongly recommends that you do *not* add FortiDDoS to the community named `public`. This popular default name is well-known, and attackers that gain access to your network will often try this name first.

The FortiDDoS SNMP agent can be configured to enable both queries and traps (alarms or event messages). You configure the SNMP settings for FortiDDoS system events and FortiDDoS attack events separately.

Basic steps:

1. Add the Fortinet and FortiDDoS MIBs to your SNMP manager. See Appendix B: Management Information Base (MIB).
2. Go to System > SNMP and configure the SNMP agent and traps for system events. See below.
3. Go to Log & Report > Log Configuration > SNMP Trap Receivers and configure SNMP traps for DDoS security events. See Configuring SNMP trap receivers for DDoS attack reporting.

Before you begin:

- On the SNMP manager, you must verify that the SNMP manager is a member of the community to which the FortiDDoS system belongs, and compile the necessary Fortinet-proprietary management information blocks (MIBs) and Fortinet-supported standard MIBs.
- In the FortiDDoS interface settings, you must enable SNMP access on the network interface through which the SNMP manager connects.
- You must have Read-Write permission for System settings.

To configure SNMP:

1. Go to System > SNMP.
2. Complete the configuration as described in .
3. Save the configuration.
4. Verify the SNMP configuration and network connectivity between your SNMP manager and this system.

Test both traps and queries (assuming you have enabled both). Traps and queries typically occur on different port numbers, and therefore verifying one does not necessarily verify that the other is also functional.

To test queries, from your SNMP manager, query the FortiDDoS appliance. To test traps, cause one of the events that should trigger a trap.

Table 61: SNMP settings for system event reporting

Settings	Guidelines
SNMP Information	
SNMP Agent	Enable to activate the SNMP agent, so that the system can send traps and receive queries.
Description	A description or comment about the system, such as `dont-reboot`. The description can be up to 35 characters long, and can contain only letters (a-z, A-Z), numbers, hyphens (-) and underscores (_).
Contact	Contact information for the administrator or other person responsible for this system, such as a phone number (`555-5555`) or name (`jdoe`). The contact information can be up to 35 characters long, and can contain only letters (a-z, A-Z), numbers, hyphens (-) and underscores (_).
Location	Physical location of the appliance, such as `floor2`. The location can be up to 35 characters long, and can contain only letters (a-z, A-Z), numbers, hyphens (-) and underscores (_).
Threshold	
CPU	• Trigger—The default is 80% utilization. • Threshold—The default is 3, meaning the event is reported when the condition has been triggered 3 times in a short period. • Sample Period—The default is 600 seconds. • Sample Frequency—The default is 30 seconds.
Memory	• Trigger—The default is 80% utilization. • Threshold—The default is 3, meaning the event is reported when the condition has been triggered 3 times in a short period. • Sample Period—The default is 600 seconds. • Sample Frequency—The default is 30 seconds.
Log disk usage	• Trigger—The default is 90% utilization. • Threshold—The default is 1, meaning the event is reported each time the condition is triggered. • Sample Period—The default is 7200 seconds. • Sample Frequency—The default is 3600 seconds.
Community	

Settings	Guidelines
Name	Name of the SNMP community to which the FortiDDoS system and at least one SNMP manager belongs, such as `management`. You must configure the FortiDDoS system to belong to at least one SNMP community so that community's SNMP managers can query system information and receive SNMP traps. You can add up to three SNMP communities. Each community can have a different configuration for queries and traps, and the set of events that trigger a trap. You can also add the IP addresses of up to eight SNMP managers to each community to designate the destination of traps and which IP addresses are permitted to query the FortiDDoS system.
Status	Select to enable the configuration.
Community Hosts	IP address of the SNMP manager to receive traps and be permitted to query the FortiDDoS system. SNMP managers have read-only access. You can add up to 8 SNMP managers to each community. To allow any IP address using this SNMP community name to query the FortiDDoS system, enter `0.0.0.0`. For security best practice reasons, however, this is not recommended. **Caution:** The system sends security-sensitive traps, which should be sent only over a trusted network, and only to administrative equipment. **Note:** If there are no other host IP entries, entering only `0.0.0.0` effectively disables traps because there is no specific destination for trap packets. *If you do not want to disable traps, you must add at least one other entry*.
Queries	Port number on which the system listens for SNMP queries from the SNMP managers in this community. The default is 161. Enable queries for SNMP v1, SNMP v2c, or both.
Traps	Source (**Local**) port number and destination (**Remote**) port number for trap packets sent to SNMP managers in this community. The default is 162. Enable traps for SNMP v1, SNMP v2c, or both.
SNMP Event	Select to enable SNMP event reporting for the following thresholds: • CPU—CPU usage has exceeded 80%. • Memory—Memory (RAM) usage has exceeded 80%. • Log disk usage—Disk space usage for the log partition or disk has exceeded 90%.

Configuring SNMP trap receivers for DDoS attack reporting

You configure SNMP trap receivers for FortiDDoS attack events separately from the system traps. This enables you to have separate configurations for each SPP, if necessary. You can configure multiple SNMP trap receivers—multiple managers that receive traps from the same SPP, or managers that receive traps from multiple SPPs.

Before you begin:

- You must have Read-Write permission for Log & Report settings.

To configure SNMP trap receivers:

1. Go to Log & Report > Log Configuration > SNMP Trap Receivers.
2. Click **Add** to display the configuration editor.
3. Complete the configuration as described in Table 62.
4. Save the configuration.

Table 62: SNMP Trap Receivers configuration guidelines

Settings	Guidelines
Name	Identifies this SNMP trap receiver in the list of receivers.
Enable	Enable the configuration.
SPP	Select the SPP for the configuration.
IP Address	IP address of the SNMP manager that receives attack log traps.
Port	Listening port of the SNMP manager. The default value is 162.
Community User-name	String that specifies the SNMP community to which the FortiDDoS system and the SNMP manager at the specified address belong.
SNMP Version	• v2c • v3
SNMPv3	

Settings	Guidelines
Engine ID	ID that uniquely identifies the SNMP agent. If the Engine ID is not specified, the MAC address of the management port is used to generate the Engine ID. For example, if the MAC address is: 08:5b:0e:9f:05:f0, the Engine ID will be: 8000304403085b0e9f05f0 **Note** Mac address will be prefixed with Fortinet prefix string: 8000304403 For Auth and Privacy modes, the following are supported: -a PROTOCOL set authentication protocol SHA -x PROTOCOL set privacy protocol AES
v3 Access Type	• No authentication • Authentication • Privacy and Authentication - You need to enter BOTH Authentication and Privacy as configured on the SNMP Manager.
Authentication Passphrase	If authentication is required, specify the authentication passphrase configured on the SNMP manager.

Downloading collected logs

You can download the DDoS Attack Log collection. You might do this if you are following manual procedures for storing log data or a manual process for purging the local log.

The download file is a MySQL export. You can import it into a MySQL database server to rebuild the flg database, including the dlog table.

Before you begin:

- You must have Read-Write permission for Log & Report settings.

To download collected logs:

1. Go to Log & Report > Log Access > Log Backup.
2. Enable **DDoS Attack Log Backup**.
3. Click **Save** to start the backup process.
4. Click **Refresh** to check whether the backup is complete.
5. Click **Download**.

Figure 119: Log backup

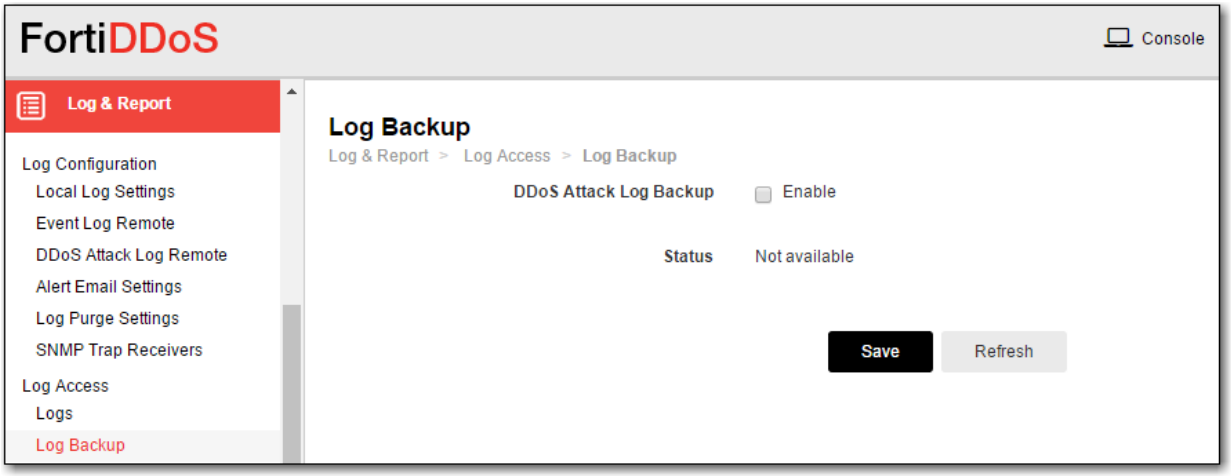

DDoS Attack Log Backup □ Enable

Status Not available

Using SQL to query the DDoS Attack Log

You can use SQL to query the DDoS Attack Log using a third-party tool such as the MySQL command-line tool or MySQL Workbench. Access to the log database is read-only.

This feature allows you to view attack log information in a report format other than the one provided by the web UI. For example, to generate consolidated reports when FortiDDoS is integrated with other appliances in your network.

You access the log using the user `root` and the password is the serial number of the appliance.

SQL connections are not secure, and can be intercepted by a third party. If possible, enable this option only for network interfaces connected to a trusted private network, or directly to your management computer. Failure to restrict administrative access through this protocol could compromise the security of your FortiDDoS appliance.

To enable SQL access:

1. Go to System > Network > Interface.
2. Double-click either **mgmt1** or **mgmt2**.
3. Under Administrative Access, select **SQL**.

   To allow other types of access to FortiDDoS (for example, HTTPS or SSH), ensure other types of access are selected.

CLI commands:

```
config system interface
   edit {mgmt1|mgmt2}
      set ip <address_ipv4> <netmask_ipv4mask>
      set allow access sql
   next
end
```

To access the DDoS attack log database with a GUI tool:

The following workflow gives steps for getting started with MySQL Workbench. You can download MYSQL Workbench for Windows from the following location:

http://dev.mysql.com/downloads/tools/workbench/

1. Open the workbench and log into the IP address of the appropriate management network interface.

   Use the user `root`; the password is the serial number of the appliance.

2. Open a connection to start querying.
3. In the SQL Editor, select database flg and table dlog.

   The following is an example query:

   ```
   select dropcount from dlog where dropcount>10000 order by dropcount desc
   ```

To access the DDoS attack log database with a CLI tool:

The following example illustrates accessing the log using MySQL on a Linux terminal:

```
mysql -h 172.30.153.122 -u root -p
Enter password:
Welcome to the MySQL monitor. Commands end with ; or \g.
Your MySQL connection id is 1393
Server version: 5.5.23-MariaDB Source distribution

...
mysql> show databases;
+--------------------+
| Database |
+--------------------+
| information_schema |
| flg |
+--------------------+
2 rows in set (0.00 sec)

mysql> use flg

Reading table information for completion of table and column names
You can turn off this feature to get a quicker startup with -A

Database changed

mysql> select timestamp, inet_ntoa(ip_src4), dropcount from dlog where dropcount > 1000
    order by timestamp desc limit 10;

+---------------------+--------------------+-----------+
| timestamp | inet_ntoa(ip_src4) | dropcount |
+---------------------+--------------------+-----------+
| 2014-03-13 10:03:07 | 12.0.0.2 | 7471 |
| 2014-03-13 09:47:31 | 10.0.0.2 | 3571 |
| 2014-03-13 09:40:35 | 12.0.0.2 | 3991 |
| 2014-03-13 09:07:29 | 12.0.0.2 | 5649 |
| 2014-03-13 08:38:19 | 10.0.0.2 | 7557 |
| 2014-03-13 07:38:49 | 10.0.0.2 | 2418 |
| 2014-03-13 06:57:48 | 12.0.0.2 | 3425 |
| 2014-03-13 06:57:25 | 10.0.0.2 | 3610 |
| 2014-03-13 06:46:00 | 10.0.0.2 | 1051 |
| 2014-03-13 06:39:12 | 10.0.0.2 | 4853 |
+---------------------+--------------------+-----------+
10 rows in set (0.00 sec)
```

Using the DDoS attack log table

The DDoS Attack Log table under Log & Report > Log Access > Logs displays attack event records for the selected SPP. The DDoS Attack Log table is updated every few seconds. It contains a maximum of 1 million events. If the number of events exceeds 1 million, the system deletes the 200,000 oldest events.

Table 63 describes the columns in the DDoS attack log.

Table 63: DDoS attack log

Column	Example	Description
Event ID	462380959	Log ID.
Timestamp	2015-05-05 16:31:00	Log timestamp.
SPP ID	0	SPP ID.
Source IP	-	Source IP address. Reported only for drops due to per-source thresholds (see Source tracking table).
Protected IP	74.255.0.253	Protected IP address. • For outbound traffic, Protected IP is the Source IP. • For inbound traffic, Protected IP is the Destination IP. The reported IP address varies, depending on the protection feature that dropped the packet. See FAQ: Logs and Reports.
Direction	Inbound	Direction: Inbound, Outbound.
Protocol	6/tcp	Protocol number.
ICMP type/code	-	ICMP type/code number.
Event Type	SYN flood	Event type.
Destination port	-	Destination port number.
Drop Count	14	Packets dropped per this event.
Event Detail	-	Reason string.
Subnet ID	0	Subnet ID.

Column	Example	Description
SPP Policy	-	SPP policy: SPP-0, SPP-1, SPP-2, SPP-3, SPP-4, SPP-5, SPP-6, SPP-7.
SPP Policy Comment	-	SPP policy comment.

Note: In the DDoS attack log, a table cell displays ”-” (hyphen) when data was not collected or is invalid. Protocol, Destination Port, and ICMP Type/Code use a number/name format. If the name cannot be determined, a hyphen is displayed.

Figure 120 shows the DDoS Attack Log page. By default, the table displays most recent records first and the columns Timestamp, SPP, DIR (direction), Event Type, and Drop Count.

You click a row to select a record. Log details for the selected event are displayed below the table.

You can filter the rows displayed in the table based on timestamp, direction (inbound or outbound), category, source IP address, destination IP address, destination port, protocol, ICMP type/code, or SPP policy.

See Appendix A: DDoS Attack Log Reference for details on log categories and event types.

Before you begin:

- You must have an administrator account with the System Admin option enabled.

To view and filter the log:

1. Go to Log & Report > Log Access > Logs and select DDoS Attack log tab.
2. Click **Filter Settings** to display the filter tools.
3. Use the tools to create filter logic.
4. Click **OK** to apply the filter and redisplay the log.
5. Click Preview icon to view the details of the attack.

Note: The Protocol field may show a blank value if there is traffic from multiple protocols since FPGA does not report a specific protocol in this scenario.

Figure 120: DDoS Attack Log table

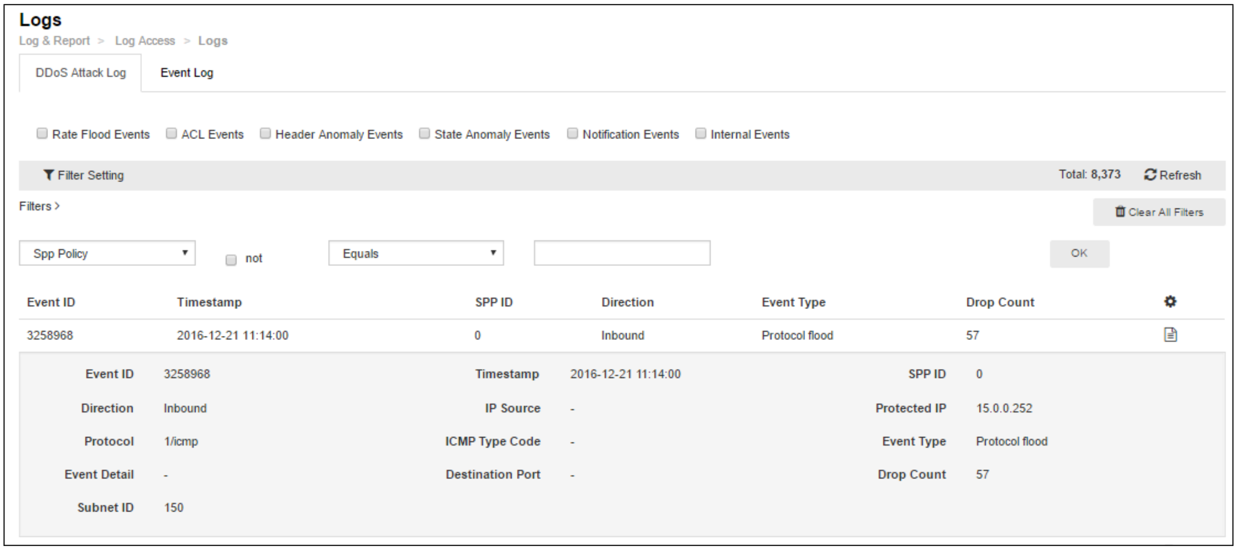

Logs

Log & Report > Log Access > Logs

DDoS Attack Log | Event Log

□ Rate Flood Events □ ACL Events □ Header Anomaly Events □ State Anomaly Events □ Notification Events □ Internal Events

Filter Setting Total: 8,373 Refresh

Filters > Clear All Filters

Spp Policy ▾ □ not Equals ▾ OK

Event ID	Timestamp	SPP ID	Direction	Event Type	Drop Count
3258968	2016-12-21 11:14:00	0	Inbound	Protocol flood	57

Event ID	3258968	Timestamp	2016-12-21 11:14:00	SPP ID	0
Direction	Inbound	IP Source	-	Protected IP	15.0.0.252
Protocol	1/icmp	ICMP Type Code	-	Event Type	Protocol flood
Event Detail	-	Destination Port	-	Drop Count	57
Subnet ID	150				

Using the event log table

The Event Log table displays logs related to system-wide status and administrator activity.

Table 64 describes the columns in the event log.

Table 64: Event log

Column	Example	Description
Date	2015-05-04	Log date.
Time	15:50:37	Log time.
Log ID	1005081	Log ID.
Type	event	Log type: event
Sub Type	config	Log subtype: config, admin, system, ha, update, healthcheck, vserver, router, user, anti-dos.
Priority	information	Log level.
Msg ID	36609	Message ID.
User	admin	User that performed the operation.
UI	GUI(172.30.153.4)	User interface from which the operation was performed.
Action	none	Administrator action.
Status	success	Status of the event.
Message	"changed settings for 'ddos spp threshold-adjust' on domain 'SPP-0'"	Log message.

Figure 121 shows the Event Log page. By default, the table displays most recent records first and all columns. You can click a column heading to display controls to sort the rows or show/hide columns.

You click a row to select a record. Log details for the selected event are displayed below the table.

You can use the Filter Settings controls to filter the rows displayed in the table based on event type, severity, action, status, and other values. shows the Filter Settings controls.

Before you begin:

- You must have enabled local logging. See Configuring local log settings.
- You must have Read-Write permission for Log & Report settings.

To view and filter the log:

1. Go to Log & Report > Executive Summary > Event Log to display the event log.
2. Click **Filter Settings** to display the filter tools.
3. Use the tools to create filter logic.
4. Click **Apply** to apply the filter and redisplay the log.

Figure 121: Event log

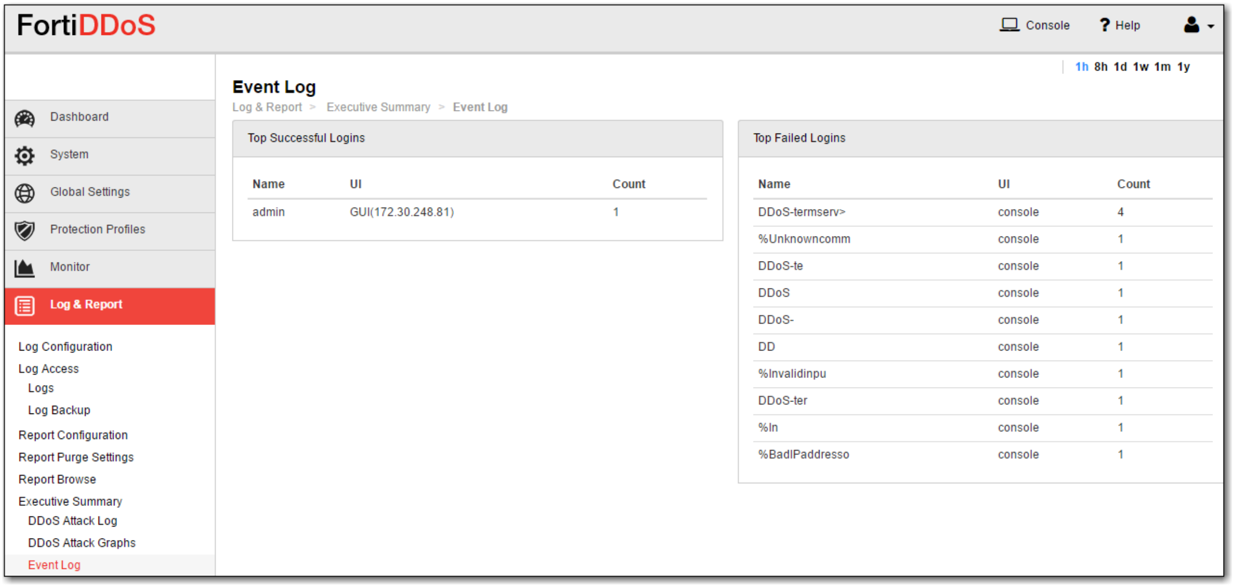

Event Log

Log & Report > Executive Summary > Event Log

Top Successful Logins

Name	UI	Count
admin	GUI(172.30.248.81)	1

Top Failed Logins

Name	UI	Count
DDoS-termserv>	console	4
%Unknowncomm	console	1
DDoS-te	console	1
DDoS	console	1
DDoS-	console	1
DD	console	1
%Invalidinpu	console	1
DDoS-ter	console	1
%In	console	1
%BadIPaddresso	console	1

Configuring reports

The report generator enables you to configure report profiles that can be run on demand or automatically according to a schedule you specify. The report generator is typically used to generate reports that can be distributed to subscribers or similar stakeholders who do not have administrative access to the FortiDDoS system. You can configure profiles that include system event data, DDoS attack data, or both.

Top attack reports are ranked by drop count (highest to lowest).

The following attack reports are available:

- Top Attacks—Drop count by DDoS attack type.
- Top ACL Attacks—Drop count by ACL rules.
- Top Attackers—Drop count by Source IP address.
- Top Attacked Subnets—Drop count by subnet.
- Top ACL Subnets—Drop count by ACL subnets.
- Top Attacked Protocols—Drop count by protocol.
- Top Attacked TCP Ports—Drop count by TCP port.
- Top Attacked UDP Ports—Drop count by UDP port.
- Top Attacked ICMP Type Codes—Drop count by ICMP type code.
- Top Attacked HTTP URLs—Drop count by HTTP URL (hash index).
- Top Attacked HTTP Methods—Drop count by HTTP method.
- Top Attacked HTTP Hosts—Drop count by Host header (hash index).
- Top Attacked HTTP Referers—Drop count by Referer header (hash index).
- Top Attacked HTTP Cookies—Drop count by Cookie header (hash index).
- Top Attacked HTTP User Agents—Drop count by User-Agent header (hash index).
- Top Attacked HTTP Servers—Drop count by HTTP server IP address.
- Top Attacked Destinations—Drop count by Destination IP address.
- Top Attacked SPPs—Drop count by SPPs.
- Top Attacked ACL SPPs—Drop count by ACL SPPs.
- Top Attacked DNS Servers—Drop count by DNS server IP address.
- Top Attacked DNS Anomalies—Drop count due to anomalies by DNS server IP address.

Before you begin:

- You must have enabled local logging for system events if you want to generate system event reports.
- You must have Read-Write permission for Log & Report settings.

To configure alert email settings:

1. Go to Log & Report > Report Configuration.
2. Click **Add** to display the configuration editor.
3. Complete the configuration as described in Table 65.
4. Save the configuration.

After you save the configuration, the profile is added to the report profile list. You can edit and delete profiles, and you can select them to generate on demand reports.

Table 65: Report configuration guidelines

Settings	Guidelines
Name	Name for the configuration. Spaces are not valid.
Report Title	Title displayed at the top of the report.
Report Type	• Global • SPP • Subnet • Default Subnet
DDoS Event Subtype	Select the type of attack report required.
Event Subtype	Select the event subtype based on the following: • Top Successful logins • Top Failed logins
Format	• HTML • PDF • Microsoft Word
Direction	• Inbound • Outbound **Note**: Shift-click to select both inbound and outbound.
Period	Select a time period. **Not Used** means all available data is included in the report, regardless of time period. **Absolute** means you specify precise dates and hours. The other options are self-explanatory.
On Schedule	If configuring a scheduled report, enable this option to specify when to run the report and set the following values: • Schedule Type • Schedule Hour
Email settings	To receive the reports via email, you can use the following fields: • Email Subject • Email Body • Email Attachment • Recipients 1, 2 and 3

Configuring report purge settings

Report purging is the deleting of report files to preserve log space and maintain log system performance.

By default, DDoS report files are purged on a first-in, first-out basis when the disk allocation for reports reaches 10 GB. Report purge settings are configurable. You can specify a different threshold, and you can purge reports manually.

Before you begin:

- You must have Read-Write permission for Log & Report settings.

To configure purge settings:

1. Go to Log & Report > Report Purge Settings.
2. Complete the configuration as described in Table 66.
3. Save the configuration.

Figure 122: Purge Settings

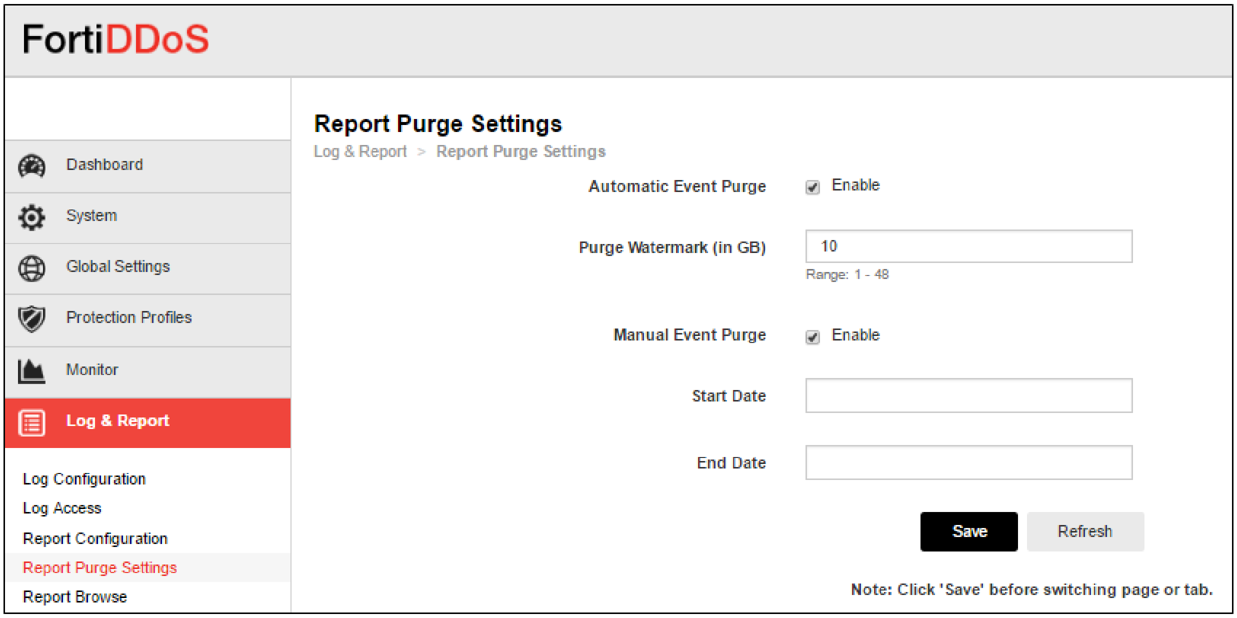

Report Purge Settings

Automatic Event Purge ☑ Enable

Purge Watermark (in GB) 10

Range: 1 - 48

Manual Event Purge ☑ Enable

Start Date

End Date

Note: Click 'Save' before switching page or tab.

Table 66: Report purge settings configuration guidelines

Settings	Guidelines
Automatic	Select to automatically purge reports when the disk allocation is reached.
Purge Watermark (in GB)	Purge the earliest reports when this limit is reached. The default is 10 GB. The valid range is 1-48 GB.

Settings	Guidelines
Manual Event Purge	Select to purge reports that were generated during the specified period manually.
Start Date / End Date	Specify a period when purging reports manually. The period begins at 0:00 on the start date and ends at 23:59 on the end date.

CLI commands:

```
config ddos global report-purge
   [set automatic-report-purge {enable | disable}]
   [set report-purge-watermark <watermark_int>]
   [set manual-report-purge {enable | disable}]
   [set purge-start-date <purge_date_str>]
   [set purge-end-date <purge_date_str>]
end
```

Using Report Browse

The report browse is a list of generated reports (scheduled or on demand). You can use the report browser to view the reports or delete them from the system. The attack categories and types reported correspond with the DDoS Attack log categories and event types. Refer to for descriptions.

Before you begin:

- You must have Read-Write permission for Log & Report settings.

To view or delete reports:

1. Go to Log & Report > Report Browse.
2. Click the report title to view it or the delete link to delete it.

Figure 123: Report Browse

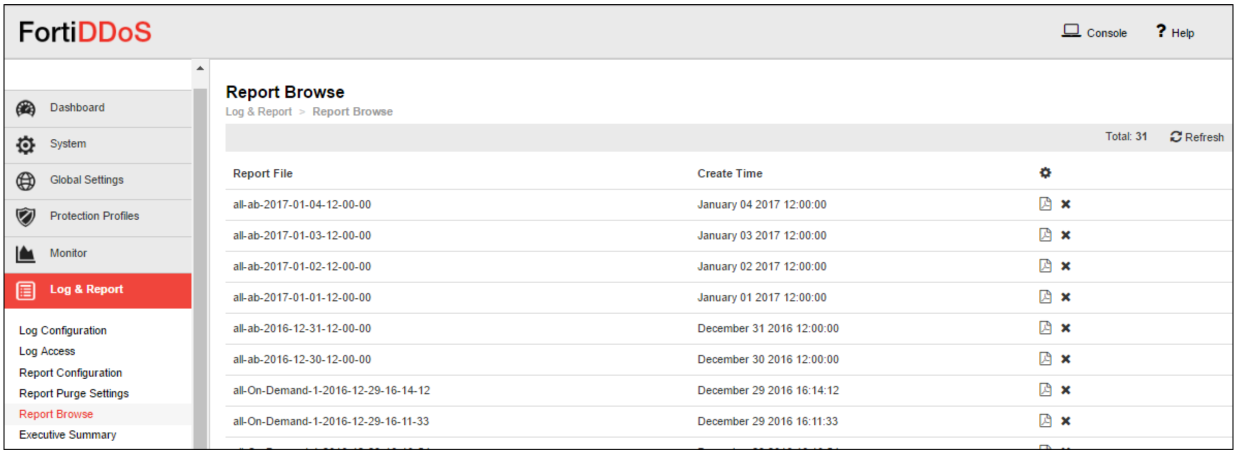

Report Browse

Total: 31

Report File	Create Time	
all-ab-2017-01-04-12-00-00	January 04 2017 12:00:00	✖
all-ab-2017-01-03-12-00-00	January 03 2017 12:00:00	✖
all-ab-2017-01-02-12-00-00	January 02 2017 12:00:00	✖
all-ab-2017-01-01-12-00-00	January 01 2017 12:00:00	✖
all-ab-2016-12-31-12-00-00	December 31 2016 12:00:00	✖
all-ab-2016-12-30-12-00-00	December 30 2016 12:00:00	✖
all-On-Demand-1-2016-12-29-16-14-12	December 29 2016 16:14:12	✖
all-On-Demand-1-2016-12-29-16-11-33	December 29 2016 16:11:33	✖

Using the DDoS Attack Log dashboard

Figure 124 shows the DDoS Attack dashboards. Each table summarizes top attacks ranked by drop count (highest to lowest). The data is filtered by SPP, so this dashboard gives you insight into the attacks that have been thwarted by that SPP's security posture.

The following attack reports are available:

- Top Attacked SPPs—Drop and Event counts by SPP.
- Top SPPs with Denied Packets—ACL drop count by SPP.
- Top Attacks—Drop count by DDoS attack type.
- Top ACL Drops—Drop count by ACL rules.
- Top Attacked Subnets (SPP Policies)—Drop count by SPP Policy.
- Top Attacked Subnets with Denied Packets—ACL drop count by subnet ID.
- Top Attacked Destinations—Drop count by Destination IP address.
- Top Attacked HTTP Servers—Drop count by HTTP server IP address.
- Top Attackers—Drop count by Source IP address.
- Top Attacked Protocols—Drop count by protocol.
- Top Attacked TCP Ports—Drop count by TCP port.
- Top Attacked UDP Ports—Drop count by UDP port.
- Top Attacked ICMP Type Codes—Drop count by ICMP type code.
- Top Attacked URLs—Drop count by HTTP URL (hash index).
- Top Attacked HTTP Methods—Drop count by HTTP method.
- Top Attacked HTTP Hosts—Drop count by Host header (hash index).
- Top Attacked HTTP User Agents—Drop count by User-Agent header (hash index).
- Top Attacked HTTP Referers—Drop count by Referer header (hash index).
- Top Attacked HTTP Cookies—Drop count by Cookie header (hash index).
- Top Attacked DNS Servers—Drop count by DNS server IP address.
- Top Attacked DNS Anomalies—Drop count by DNS server IP address for packets dropped by DNS anomaly rules.

To display the DDoS Attack Log dashboard:

1. Go to Log & Report > Executive Summary > DDoS Attack Log.
2. Select the SPP of interest, time period, and traffic direction.

Figure 124: DDoS Attack Log

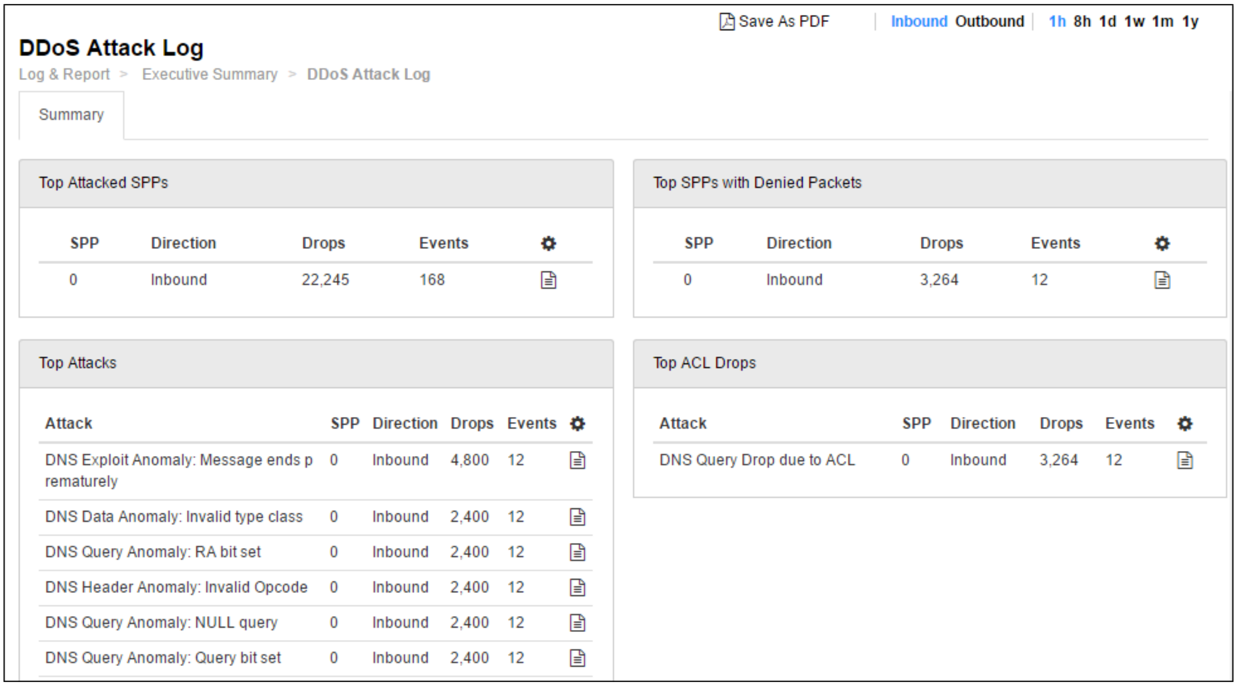

DDoS Attack Log

Top Attacked SPPs

SPP	Direction	Drops	Events	⚙
0	Inbound	22,245	168	🗎

Top SPPs with Denied Packets

SPP	Direction	Drops	Events	⚙
0	Inbound	3,264	12	🗎

Top Attacks

Attack	SPP	Direction	Drops	Events	⚙
DNS Exploit Anomaly: Message ends p rematurely	0	Inbound	4,800	12	🗎
DNS Data Anomaly: Invalid type class	0	Inbound	2,400	12	🗎
DNS Query Anomaly: RA bit set	0	Inbound	2,400	12	🗎
DNS Header Anomaly: Invalid Opcode	0	Inbound	2,400	12	🗎
DNS Query Anomaly: NULL query	0	Inbound	2,400	12	🗎
DNS Query Anomaly: Query bit set	0	Inbound	2,400	12	🗎

Top ACL Drops

Attack	SPP	Direction	Drops	Events	⚙
DNS Query Drop due to ACL	0	Inbound	3,264	12	🗎

Using the DDoS Attack Graph dashboard

shows the DDoS Attack Graph dashboard. This dashboard contains graphs that summarizes top attacks. The data is filtered by SPP, so this dashboard gives you insight into the attacks that have been thwarted by that SPP’s security posture.

The following attack graphs are available:

- Top Attacked SPPs—Drop and Event counts by SPP.
- Top SPPs with Denied Packets—ACL drop count by SPP.
- Top Attacks—Drop count by DDoS attack type.
- Top ACL Drops—Drop count by ACL rules.
- Top Attacked Subnets—Drop count by subnet ID.
- Top Attacked Subnets with Denied Packets—ACL drop count by subnet ID.
- Top Attacked Destinations—Drop count by Destination IP address.
- Top Attacked HTTP Servers—Drop count by HTTP server IP address.
- Top Attackers—Drop count by Source IP address.
- Top Attacked Protocols—Drop count by protocol.
- Top Attacked TCP Ports—Drop count by TCP port.
- Top Attacked UDP Ports—Drop count by UDP port.
- Top Attacked ICMP Type Codes—Drop count by ICMP type code.
- Top Attacked URLs—Drop count by HTTP URL (hash index).
- Top Attacked HTTP Methods—Drop count by HTTP method.
- Top Attacked HTTP Hosts—Drop count by Host header (hash index).
- Top Attacked HTTP User Agents—Drop count by User-Agent header (hash index).
- Top Attacked HTTP Referers—Drop count by Referer header (hash index).
- Top Attacked HTTP Cookies—Drop count by Cookie header (hash index).
- Top Attacked DNS Servers—Drop count by DNS server IP address.
- Top Attacked DNS Anomalies—Drop count by DNS server IP address for packets dropped by DNS anomaly rules.

To display the DDoS Attack Graphs dashboard:

- Go to Log & Report > Executive Summary > DDoS Attack Graphs.

Figure 125: DDoS Attack Graphs Dashboard

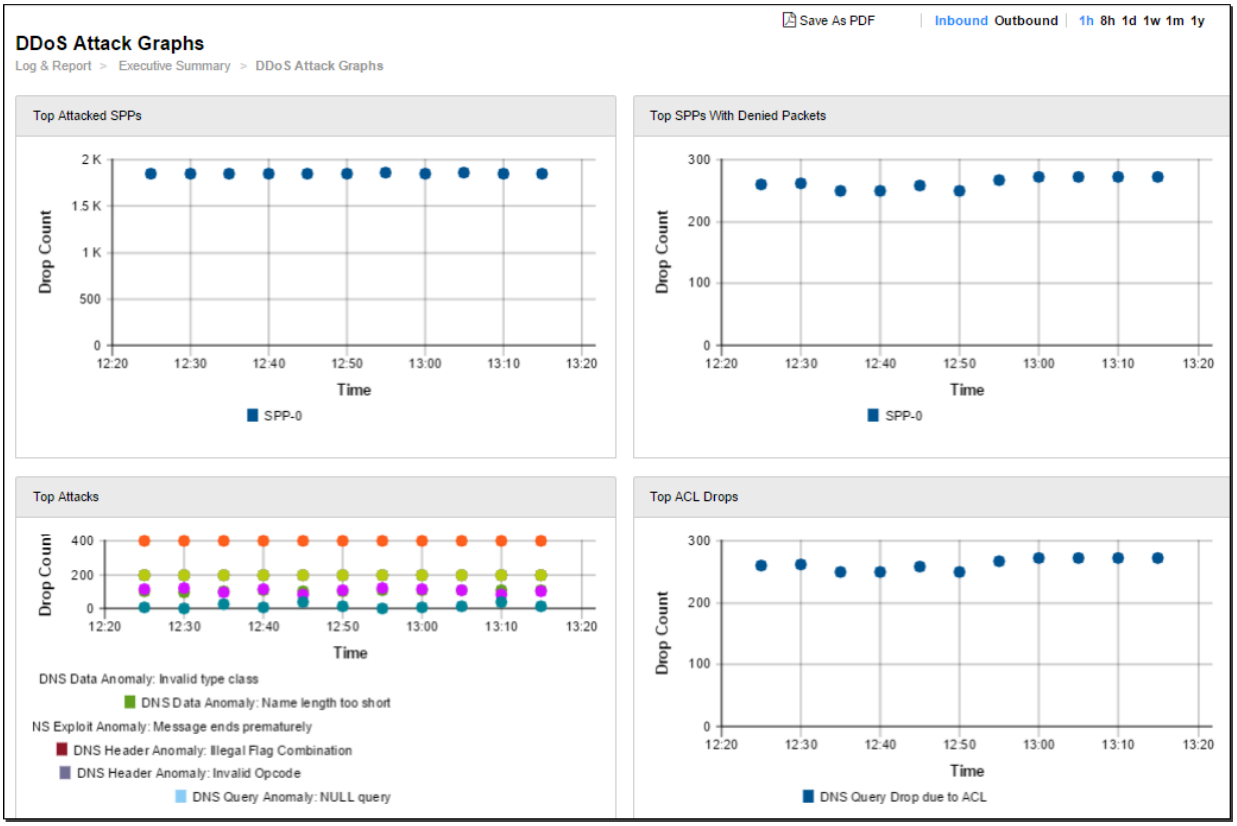
Save As PDF
Inbound
Outbound
1h
8h
1d
1w
1m
1y
DDoS Attack Graphs
Log & Report > Executive Summary > DDoS Attack Graphs
Top Attacked SPPs
Drop Count
2 K
1.5 K
1 K
500
0
12:20
12:30
12:40
12:50
13:00
13:10
13:20
Time
SPP-0
Top SPPs With Denied Packets
Drop Count
300
200
100
0
12:20
12:30
12:40
12:50
13:00
13:10
13:20
Time
SPP-0
Top Attacks
Drop Count
400
200
0
12:20
12:30
12:40
12:50
13:00
13:10
13:20
Time
DNS Data Anomaly: Invalid type class
🟩 DNS Data Anomaly: Name length too short
NS Exploit Anomaly: Message ends prematurely
🟥 DNS Header Anomaly: Illegal Flag Combination
▨ DNS Header Anomaly: Invalid Opcode
▨ DNS Query Anomaly: NULL query
Top ACL Drops
Drop Count
300
200
100
0
12:20
12:30
12:40
12:50
13:00
13:10
13:20
Time
DNS Query Drop due to ACL

Using the Event Log dashboard

The Event Log dashboard which shows information about the following:

- Top Successful Logins
- Top Failed Logins

Figure 126: Event Log

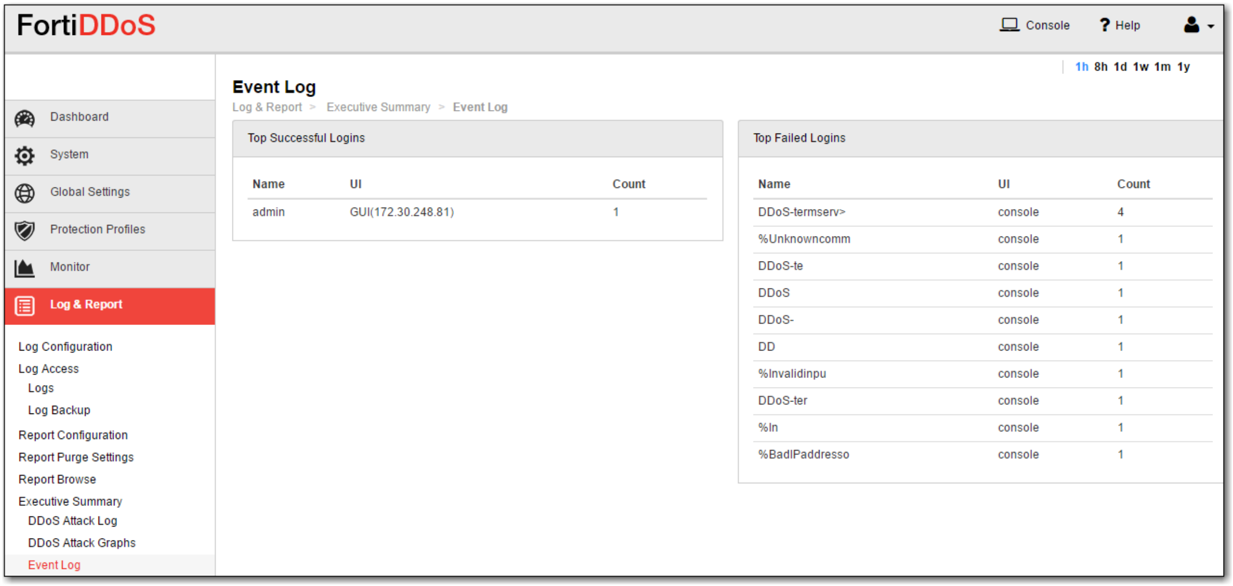

Event Log

Log & Report > Executive Summary > Event Log

Top Successful Logins

Name	UI	Count
admin	GUI(172.30.248.81)	1

Top Failed Logins

Name	UI	Count
DDoS-termserv>	console	4
%Unknowncomm	console	1
DDoS-te	console	1
DDoS	console	1
DDoS-	console	1
DD	console	1
%Invalidinpu	console	1
DDoS-ter	console	1
%In	console	1
%BadIPaddresso	console	1

Using the Session Diagnostic report

Figure 127 shows the Session Diagnostic report. You can use the Session Diagnostic report to check the session counters. The count is for current traffic. You can correlate the count with source IP address, destination IP address, destination port, or TCP state. You can also filter the records to include or exclude matching expressions.

To display the Session Diagnostic report

1. Go to Log & Report > Diagnostics > Sessions.
2. Select the SPP of interest.
3. Select a Group By option.

Figure 127: Session Diagnostic report

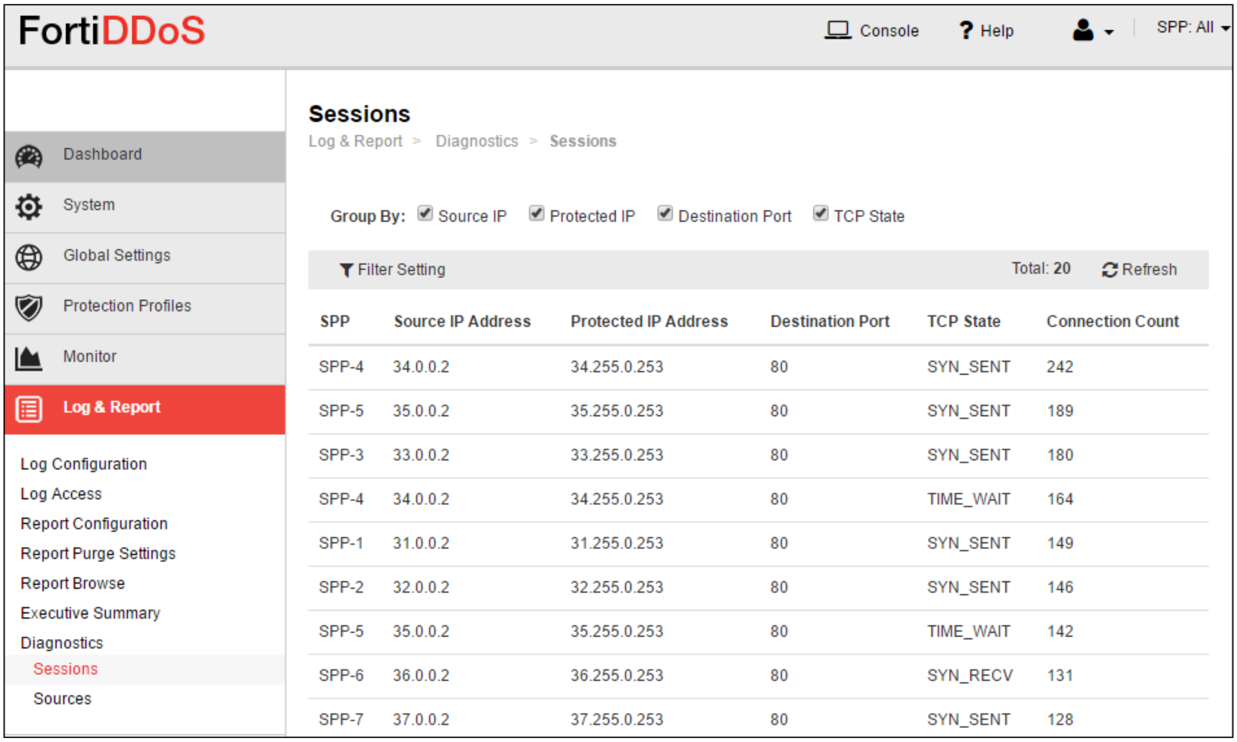

Sessions

Log & Report > Diagnostics > Sessions

Group By: ☑ Source IP ☑ Protected IP ☑ Destination Port ☑ TCP State

Total: 20

SPP	Source IP Address	Protected IP Address	Destination Port	TCP State	Connection Count
SPP-4	34.0.0.2	34.255.0.253	80	SYN_SENT	242
SPP-5	35.0.0.2	35.255.0.253	80	SYN_SENT	189
SPP-3	33.0.0.2	33.255.0.253	80	SYN_SENT	180
SPP-4	34.0.0.2	34.255.0.253	80	TIME_WAIT	164
SPP-1	31.0.0.2	31.255.0.253	80	SYN_SENT	149
SPP-2	32.0.0.2	32.255.0.253	80	SYN_SENT	146
SPP-5	35.0.0.2	35.255.0.253	80	TIME_WAIT	142
SPP-6	36.0.0.2	36.255.0.253	80	SYN_RECV	131
SPP-7	37.0.0.2	37.255.0.253	80	SYN_SENT	128

Using the Source Diagnostics report

Figure 128 shows the Source Diagnostic report. You can use the Source Diagnostic report to check the connection and drop counters per source IP address. The count is for current traffic.

You can select one the following options to filter the results:

- Source IP
- Direction

You can also filter the records to include or exclude matching expressions.

To display the Source Diagnostic report

1. Go to Log & Report > Diagnostics > Source.
2. Select Source or Destination IP.
3. Click **OK**.

If desired, use the Filter Settings controls to filter records to include or exclude matching expressions.

Figure 128: Source Diagnostics report

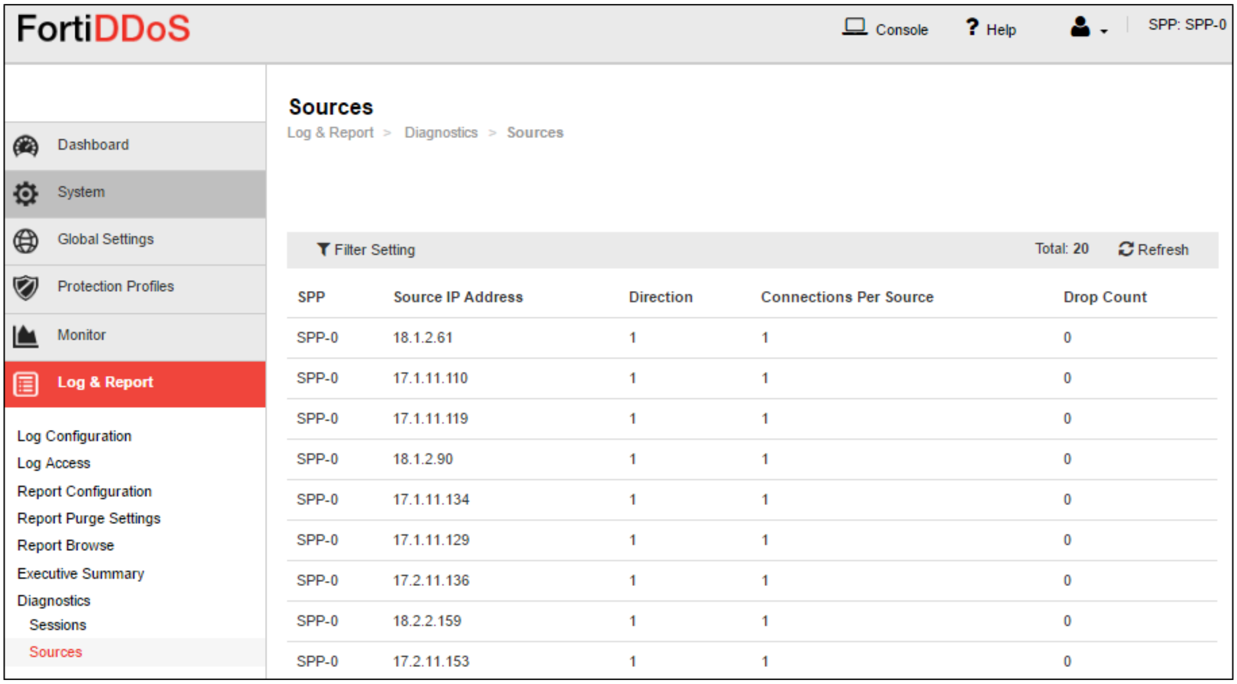

SPP	Source IP Address	Direction	Connections Per Source	Drop Count
SPP-0	18.1.2.61	1	1	0
SPP-0	17.1.11.110	1	1	0
SPP-0	17.1.11.119	1	1	0
SPP-0	18.1.2.90	1	1	0
SPP-0	17.1.11.134	1	1	0
SPP-0	17.1.11.129	1	1	0
SPP-0	17.2.11.136	1	1	0
SPP-0	18.2.2.159	1	1	0
SPP-0	17.2.11.153	1	1	0

FAQ: Logs and Reports

Attack log

This section discusses some of the questions that users often have about the attack log events.

Why is source IP address not reported for a SYN flood?

During a SYN flood attack, the system reports a SYN flood, identifies the SPP that is under attack, and reports how many packets it has dropped. SYN floods are spoofed—the reported source of the packets is not their true source. Trying to determine the source IP address or report potentially millions of source IP addresses that have no consistent pattern is resource-intensive and does not help you to determine the identity of the attacker.

When is source IP address reported?

Source IP address is reported only for drops due to per-source thresholds (see Source tracking table).

Why is destination not reported for some types of attacks?

To keep its reporting processes manageable, the system does not always report a source IP address or destination port in the DDoS attack log. For example, for an HTTP GET flood, the system reports the protocol of the dropped packets but not their source IP address or destination port.

Why are some attack events not reported in real time?

For some types of attacks, such as TCP and ICMP checksum errors, the system collects aggregate data and reports every 5 minutes only. If the appliance reported each dropped packet as soon as it dropped it and generated some kind of alert every time it dropped a packet, it would log events and generate alerts continuously.

Reports

This section discusses some of the questions that users often have about reports.

Where can I find information about the attack types listed in reports?

Reports are presentations of DDoS attack log database queries. The attack categories and types reported correspond with the DDoS Attack log categories and event types. Refer to Appendix A: DDoS Attack Log Reference for descriptions.

Why do I see records for SPP-0 in a report filtered by SPP-1?

If you change the SPP policy configuration or the resources it monitors, the data can become skewed. For example, if you remove a subnet from the profile, or change the servers that are deployed in the subnet, or change the services offered by those servers, the traffic history becomes less relevant.

Fortinet strongly recommends that you reset the traffic history for a profile before you make any significant changes to its configuration. Go to Protection Profiles > Factory Reset > Factory Reset.

If you do not reset traffic statistics, changes to an SPP policy can result in counter-intuitive data accumulated in the longer reporting periods (year, month). For example, if a subnet belonged to the default SPP-0 before you assigned it to SPP-1, a report filtered by SPP-1 includes the SPP-0 traffic history for that subnet.

What is the difference between the link status reported in the web UI and the link status reported with CLI commands?

The link status reported on the Dashboard page is the detected link state.

The link status shown in the `show system interface` and `get system interface` commands is the configured status.

To display the detected link state with the CLI, use the following command:

```
FI-2K# diagnose hardware get deviceinfo data-port
port1 down 10G FD SW No Forward TX RX None F XGMII 16356
port2 down 10G FD SW No Forward TX RX None F XGMII 16356
port3 down 10G FD SW No Forward TX RX None F XGMII 16356
```

Chapter 7: Using the Dashboard

The Dashboard contains tables or graph summary of system information or system status.

You can use the dashboard to check system status at-a-glance or to quickly find system information, like the hardware serial number, firmware version, license status, or interface status. For a deeper look at attack traffic, use the Monitor and Log & Report menus.

Before you begin:

- You must have Read-Write permission for Log & Report settings.

To display the Dashboard:

- Go to Dashboard.

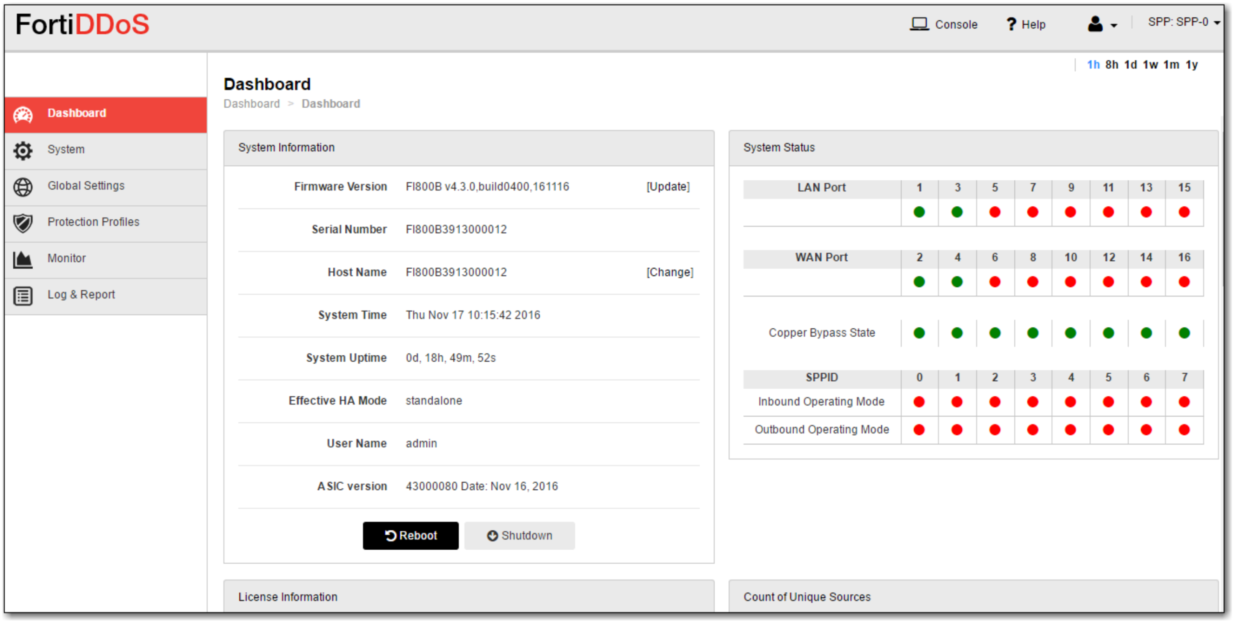

System Information

Firmware Version	FI800B v4.3.0,build0400,161116	[Update]
Serial Number	FI800B3913000012	
Host Name	FI800B3913000012	[Change]
System Time	Thu Nov 17 10:15:42 2016	
System Uptime	0d, 18h, 49m, 52s	
Effective HA Mode	standalone	
User Name	admin	
ASIC version	43000080 Date: Nov 16, 2016	

System Status

LAN Port	1	3	5	7	9	11	13	15
	🟢	🟢	🔴	🔴	🔴	🔴	🔴	🔴

WAN Port	2	4	6	8	10	12	14	16
	🟢	🟢	🔴	🔴	🔴	🔴	🔴	🔴
Copper Bypass State	🟢	🟢	🟢	🟢	🟢	🟢	🟢	🟢

SPPID	0	1	2	3	4	5	6	7
Inbound Operating Mode	🔴	🔴	🔴	🔴	🔴	🔴	🔴	🔴
Outbound Operating Mode	🔴	🔴	🔴	🔴	🔴	🔴	🔴	🔴

The default dashboard setup includes the following tables/graphs:

- System Information
- System Status
- License Information
- Count of Unique Sources
- Top Attacked SPPs
- System Resources
- Top SPPs with Denied Packets
- Recent Event Logs

System Information

It displays basic information, such as the firmware version, serial number, host name, and system time.

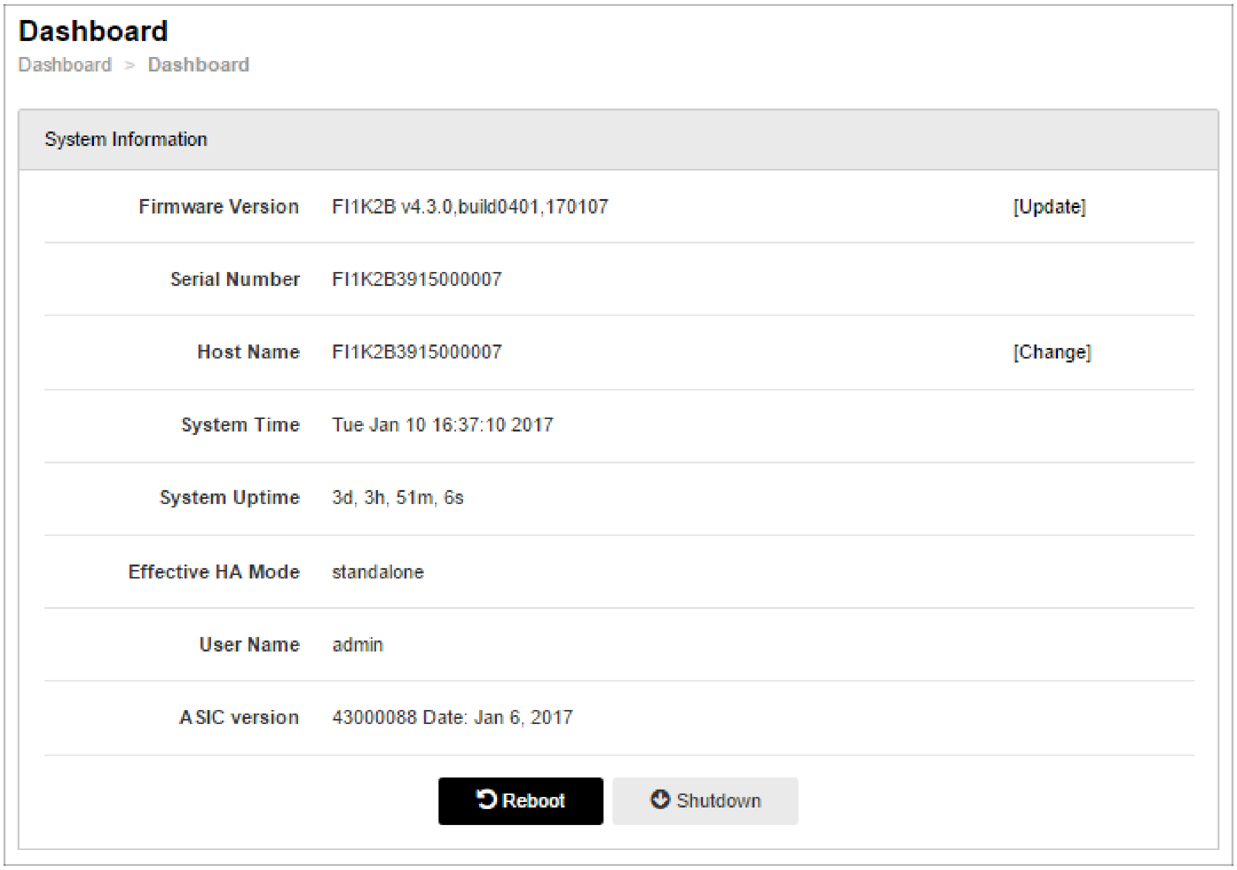

Dashboard

Dashboard > Dashboard

System Information

Firmware Version	FI1K2B v4.3.0,build0401,170107	[Update]
Serial Number	FI1K2B3915000007	
Host Name	FI1K2B3915000007	[Change]
System Time	Tue Jan 10 16:37:10 2017	
System Uptime	3d, 3h, 51m, 6s	
Effective HA Mode	standalone	
User Name	admin	
ASIC version	43000088 Date: Jan 6, 2017	

System Status

It displays status for network interfaces, bypass, and SPPs.

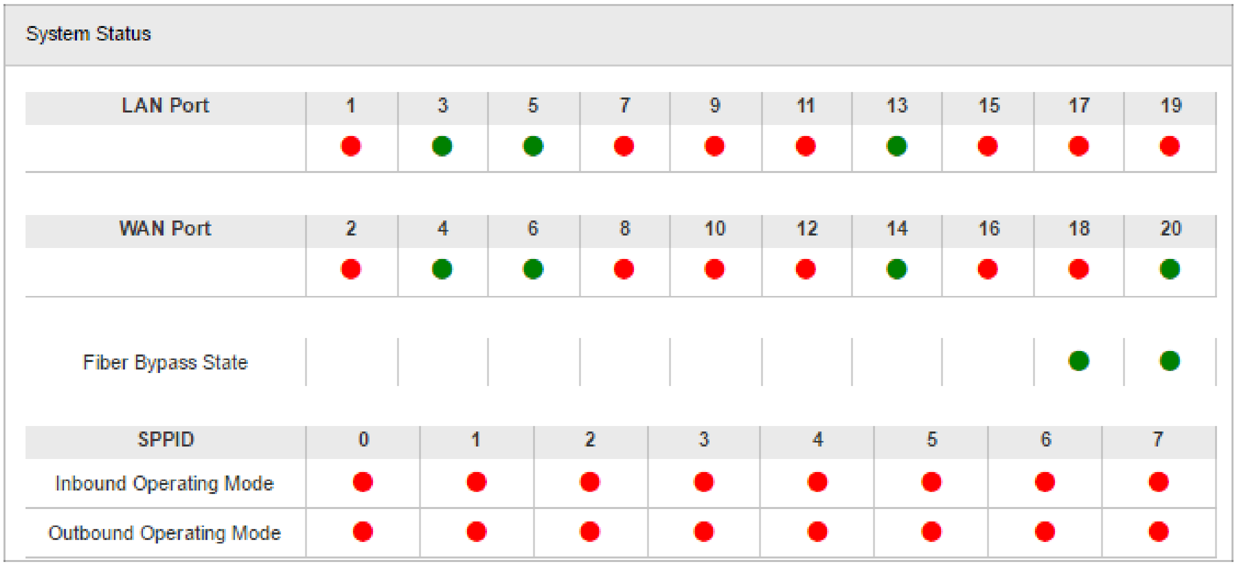

System Status

LAN Port	1	3	5	7	9	11	13	15	17	19
	🔴	🟢	🟢	🔴	🔴	🔴	🟢	🔴	🔴	🔴

WAN Port	2	4	6	8	10	12	14	16	18	20
	🔴	🟢	🟢	🔴	🔴	🔴	🟢	🔴	🔴	🟢
Fiber Bypass State									🟢	🟢

SPPID	0	1	2	3	4	5	6	7
Inbound Operating Mode	🔴	🔴	🔴	🔴	🔴	🔴	🔴	🔴
Outbound Operating Mode	🔴	🔴	🔴	🔴	🔴	🔴	🔴	🔴

Table 67: System Status

Category	Status Indicators
Port	Green icon - The port is physically connected the network. Red icon - The port has no physical connection to the network. Odd-numbered ports are LAN connections (Protected Network) that have a corresponding even-numbered port, which is the associated WAN connection. (For example, Port 1 connected to the Ethernet and Port 2 connected to the Internet). Hover over the status icons to see additional information: port number, link status, speed, auto-negotiation, and medium (copper or fiber).
Bypass State	Used for copper-based (RJ-45) Ethernet connections only. Green icon - Normal operation (no bypass). Red icon - Link is operating in bypass mode. **Note**: For FortiDDoS 1200B and 2000B: ports 17, 18, 19 and 20 (Fiber 10G ports) have Bypass states.
SPP ID	Green icon - Profile is operating in Detection Mode (monitoring traffic to generate statistics). Red icon - Profile is operating in Prevention Mode (dropping or blocking attacks as well as generating statistics). Gray icon - Profile is not configured. Hover over the status icons to see additional information: port number, ID, name, mode.

License Information

It displays license and registration status, including status for the FortiGuard IP Reputation Service.

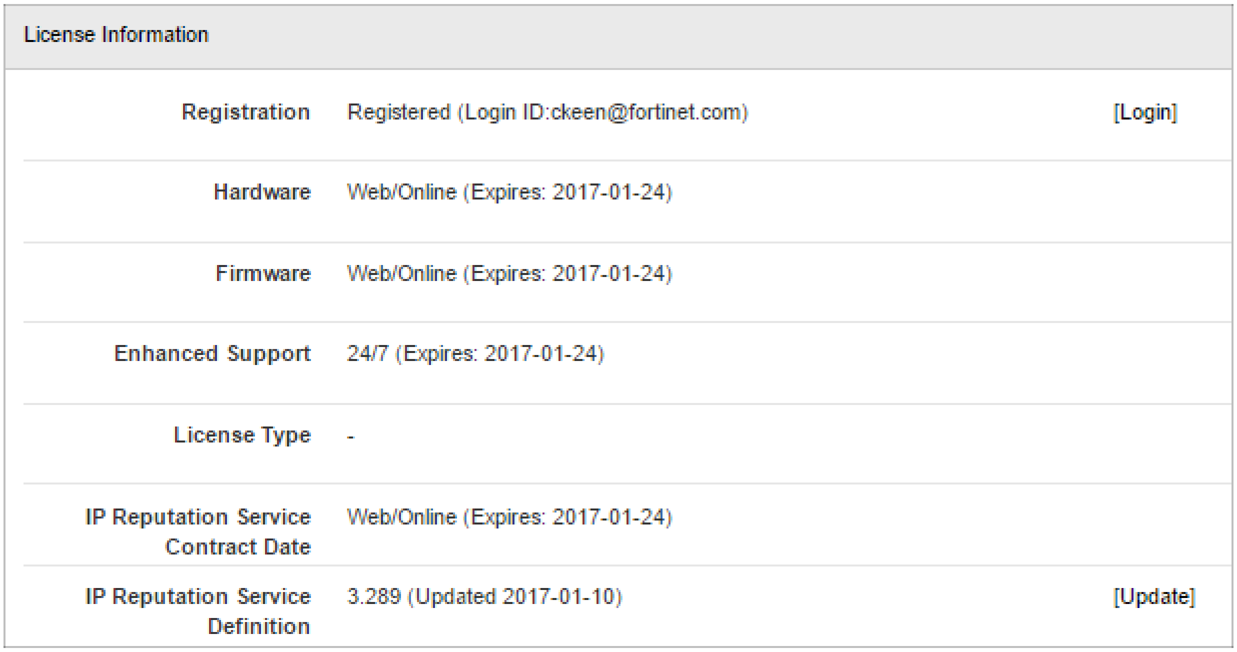

License Information

Registration	Registered (Login ID:ckeen@fortinet.com)	[Login]
Hardware	Web/Online (Expires: 2017-01-24)	
Firmware	Web/Online (Expires: 2017-01-24)	
Enhanced Support	24/7 (Expires: 2017-01-24)	
License Type	-	
IP Reputation Service Contract Date	Web/Online (Expires: 2017-01-24)	
IP Reputation Service Definition	3.289 (Updated 2017-01-10)	[Update]

Count of Unique Sources

It displays the trend in the unique Source IP count for the selected SPP and time period. A spike in this graph indicates a possible DDoS attack.

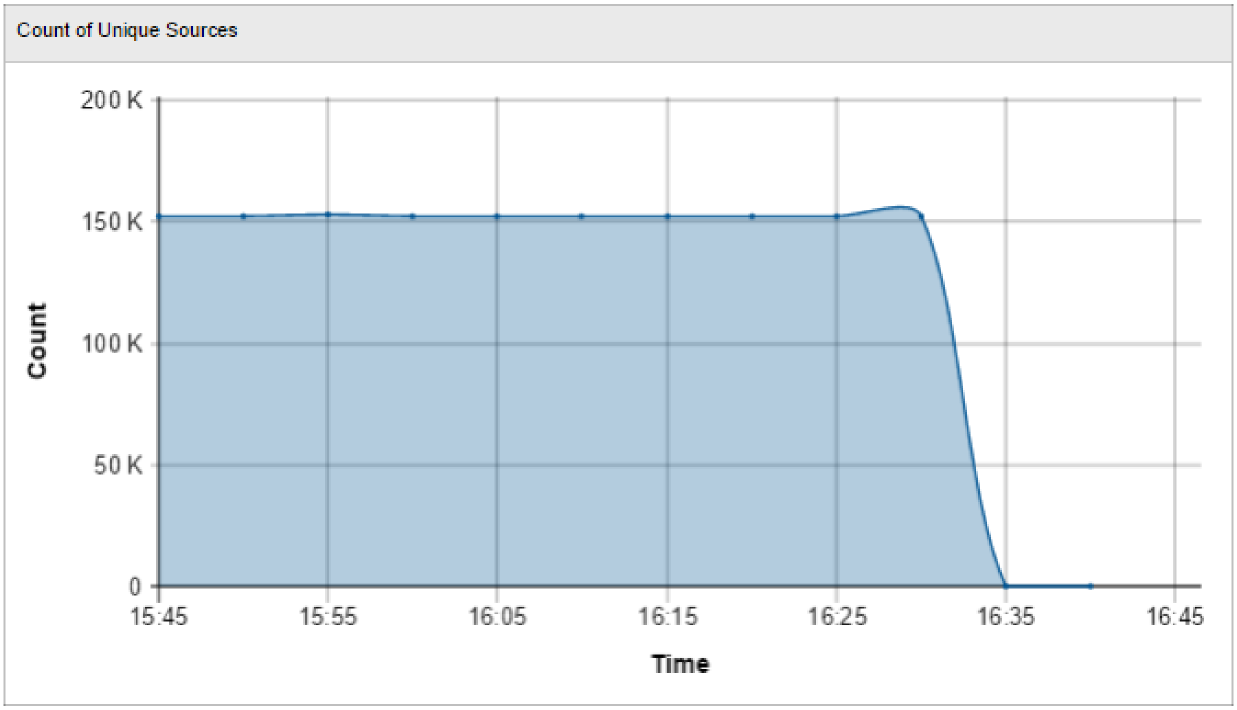

Top Attacked SPPs

Figure 129 shows the Top Attacked SPPs graph. It displays the trend in drop count for traffic that has been dropped by SPP threshold rules.

To hide or display the count for an SPP, click its name.

To display details, hover over a point on the graph.

Figure 129: Top Attacked SPPs

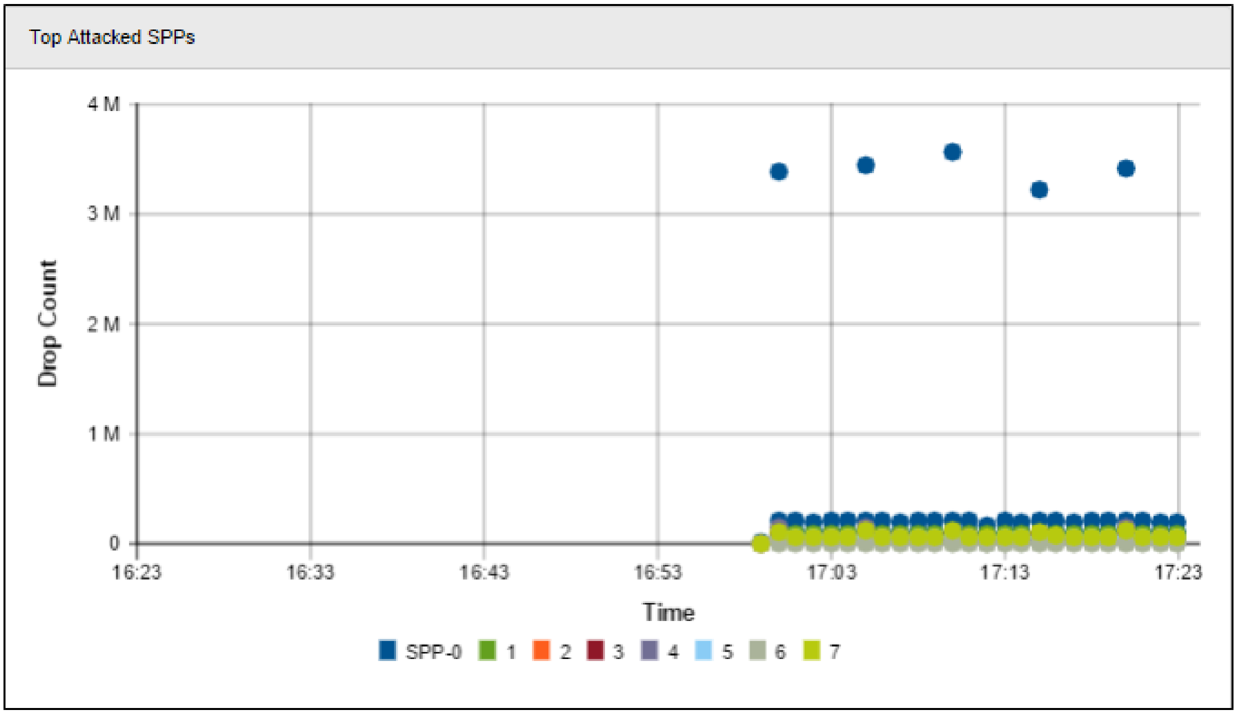

System Resources

It displays CPU and memory usage as a dial gauge and as a percentage of the usage for core processes only. CPU and memory usage for management processes (for example, for HTTPS connections to the web UI) is excluded.

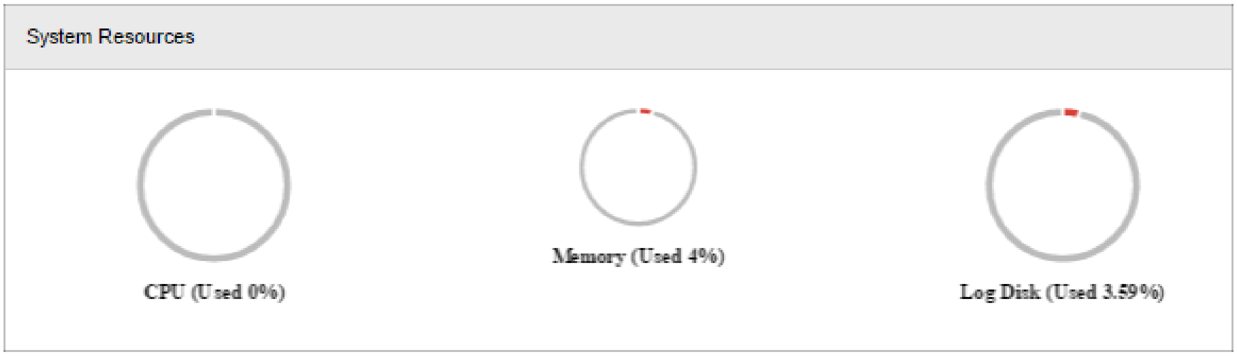

This does not display disk usage.To view disk space information, connect to the CLI and enter the following command:

```
diagnose hardware get sysinfo df
```

Top SPPs with Denied Packets

Figure 130 shows the Top SPPs with Denied Packets. It displays the trend in drop count for traffic that has been dropped by Global ACL rules or SPP ACL rules.

Figure 130: Top SPPs with Denied Packets

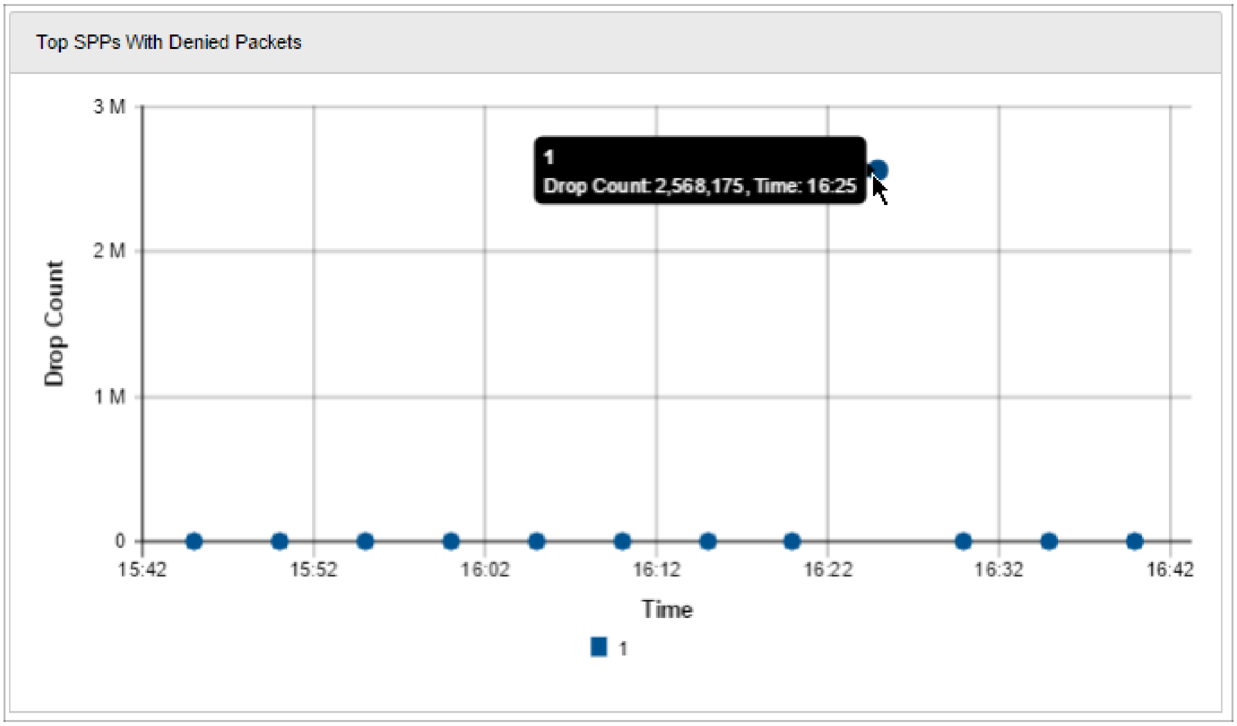

Recent Event Logs

Figure 131 shows the Recent Event Log information. Event logs help you track system events, such as administrator logins and firmware changes.

Figure 131: Event Log Console

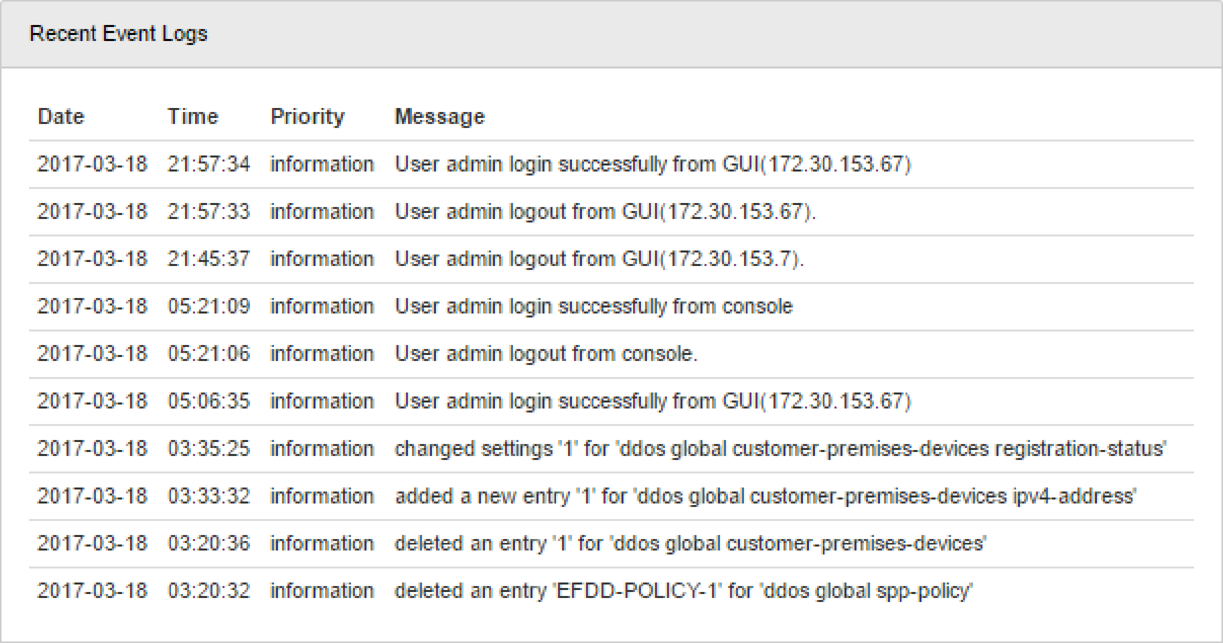

Recent Event Logs

Date	Time	Priority	Message
2017-03-18	21:57:34	information	User admin login successfully from GUI(172.30.153.67)
2017-03-18	21:57:33	information	User admin logout from GUI(172.30.153.67).
2017-03-18	21:45:37	information	User admin logout from GUI(172.30.153.7).
2017-03-18	05:21:09	information	User admin login successfully from console
2017-03-18	05:21:06	information	User admin logout from console.
2017-03-18	05:06:35	information	User admin login successfully from GUI(172.30.153.67)
2017-03-18	03:35:25	information	changed settings '1' for 'ddos global customer-premises-devices registration-status'
2017-03-18	03:33:32	information	added a new entry '1' for 'ddos global customer-premises-devices ipv4-address'
2017-03-18	03:20:36	information	deleted an entry '1' for 'ddos global customer-premises-devices'
2017-03-18	03:20:32	information	deleted an entry 'EFDD-POLICY-1' for 'ddos global spp-policy'

CLI Console

Figure 132 shows the CLI Console . It enables you to enter CLI commands through the web UI, without making a separate Telnet, SSH, or local console connection.

To use the console, click **Console** on the FortiDDoS UI header. You are logged in as the same admin account you used to access the web UI.

Figure 132: CLI Console

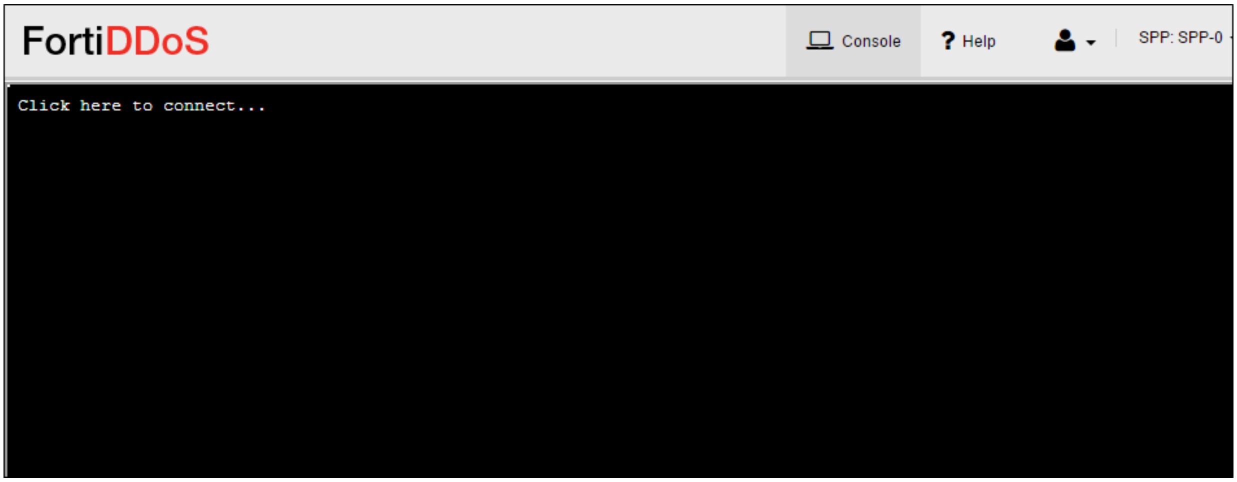

```
Click here to connect...
```

Chapter 8: System Management

This chapter includes the following topics:

Configuring network interfaces

The network interfaces that are bound to physical ports have three uses:

- Management—Ports mgmt1 and mgmt2 are management interfaces. Management interfaces are used for administrator connections and to send management traffic, like syslog and SNMP traffic. Typically, administrators use mgmt1 for the management interface.
- HA—If you plan to deploy HA, you must reserve a physical port for HA heartbeat and synchronization traffic. Do *not* configure a network interface for the port that will be used for HA; instead, leave it unconfigured or “reserved” for HA. Typically, administrators use mgmt2 for the HA interface.
- Traffic—The remaining physical ports can be used for your target traffic—these are your “traffic interfaces.” The FortiDDoS system is deployed inline (between the Internet and your local network resources). Consecutively numbered ports belong to port pairs: Use an odd port numbers (1, 3, 5, and so on) for the LAN-side connection and an even port number (2, 4, 6, and so on) for the WAN-side connection. For example, port1 and port2 are a pair. The port1 interface is connected to a switch that connects servers in the local network; the port2 interface is connected to the network path that receives traffic from the Internet.

By default, ports use auto-negotiation to determine the connection speed. In general, you change the speed if the interface is connected to a device that does not support auto-negotiation. If the other device uses a fixed speed/duplex setting, you use the configuration page to set the FortiDDoS network interface speed/duplex to the appropriate matching values.

The interface modules for FortiDDoS 900B/1000B and FortiDDoS 1200B/2000B models have special guidelines. To avoid issues with speed/duplex for these interface modules, please disregard the possible choices and use the required settings shown in Table 68.

Table 68: Speed/Duplex settings

Transceiver/Interface Module	Possible Choices	Required Settings
SFP (1 Gbps)	Auto, 1000Mbps Full Duplex	1000Mbps Full Duplex
SFP+ (10 Gbps)	Auto, 1000Mbps Full Duplex	Auto
LC 850nm optical (10 Gbps)*	Auto, 1000Mbps Full Duplex	Auto
*Available on FortiDDOS 1200B/2000B only.		

Before you begin:

- You must have Read-Write permission for System settings.

To configure a network interface:

1. Go to System > Network > Interface.
2. Double-click the row of the port you want to configure to display the configuration editor.
3. Complete the configuration as described in Table 69.
4. Save the configuration.

Figure 133: Network interface status page

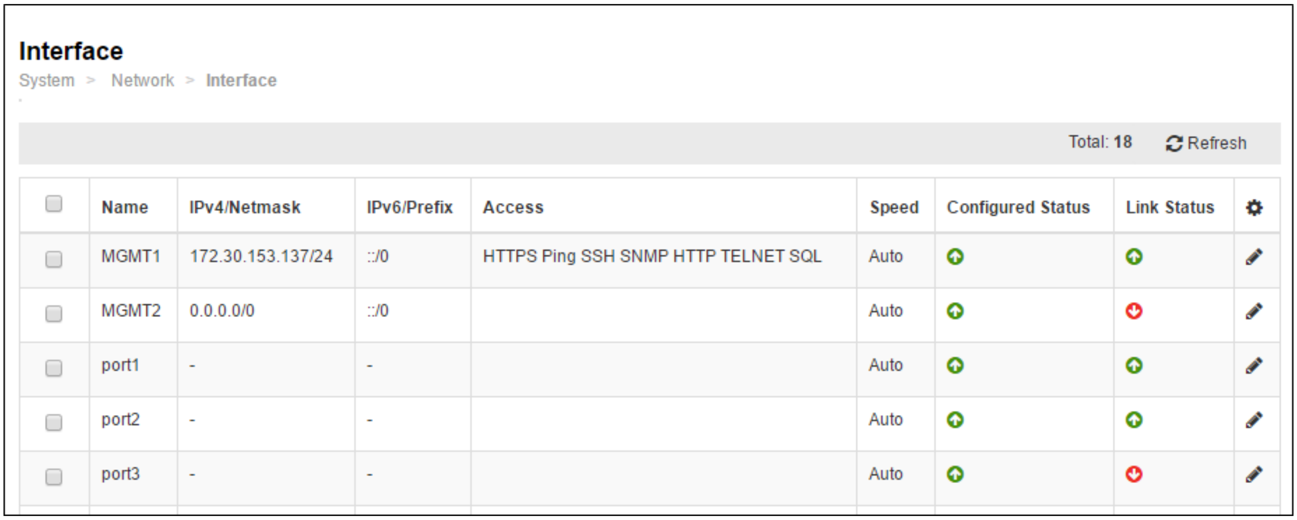

Interface

System > Network > Interface

Total: 18 Refresh

□	Name	IPv4/Netmask	IPv6/Prefix	Access	Speed	Configured Status	Link Status	⚙
□	MGMT1	172.30.153.137/24	::/0	HTTPS Ping SSH SNMP HTTP TELNET SQL	Auto	🟢	🟢	✎
□	MGMT2	0.0.0.0/0	::/0		Auto	🟢	🔴	✎
□	port1	-	-		Auto	🟢	🟢	✎
□	port2	-	-		Auto	🟢	🟢	✎
□	port3	-	-		Auto	🟢	🔴	✎

The Status indicators on the Interface Configuration page display the connectivity status. A green indicator means that the link is connected and negotiation was successful. A red indicator means that the link is not connected or is down.

Figure 134: Network interface speed/duplex settings page

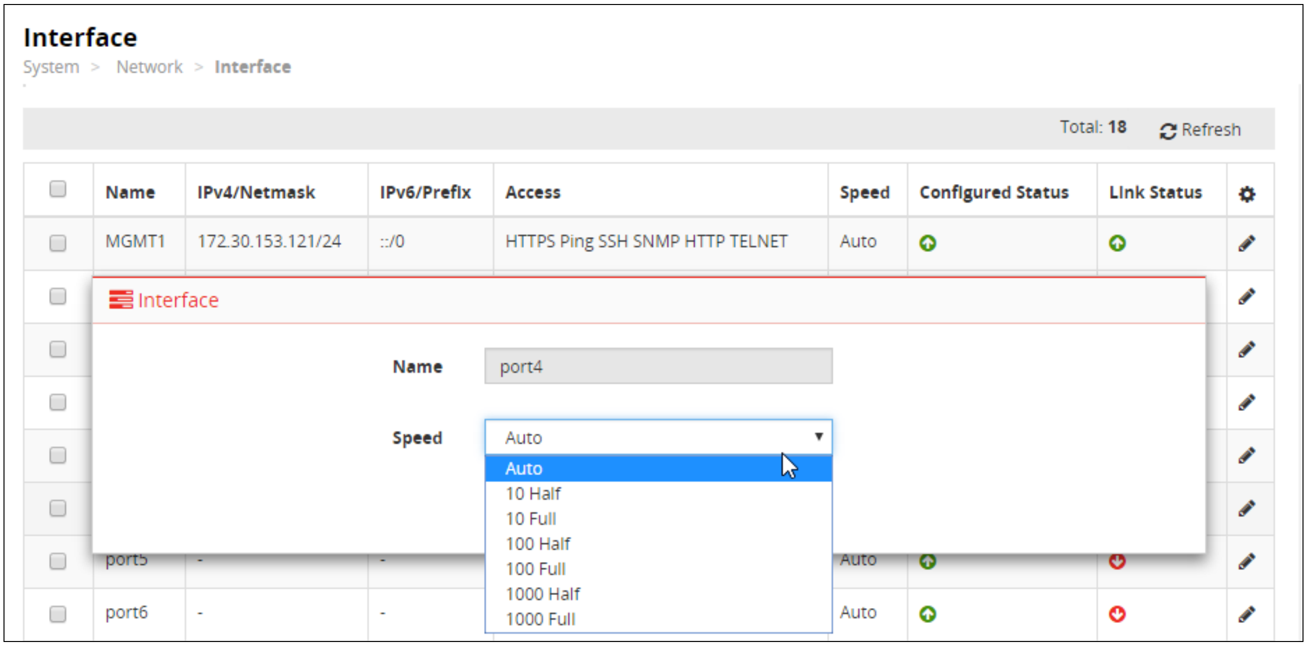

Interface

System > Network > Interface

Total: 18 Refresh

□	Name	IPv4/Netmask	IPv6/Prefix	Access	Speed	Configured Status	Link Status	⚙
□	MGMT1	172.30.153.121/24	::/0	HTTPS Ping SSH SNMP HTTP TELNET	Auto	🟢	🟢	✎
□	port5	-	-		Auto	🟢	🔴	✎
□	port6	-	-		Auto	🟢	🔴	✎

Interface

Name: port4

Speed: Auto

Figure 135: Management interfaces settings page

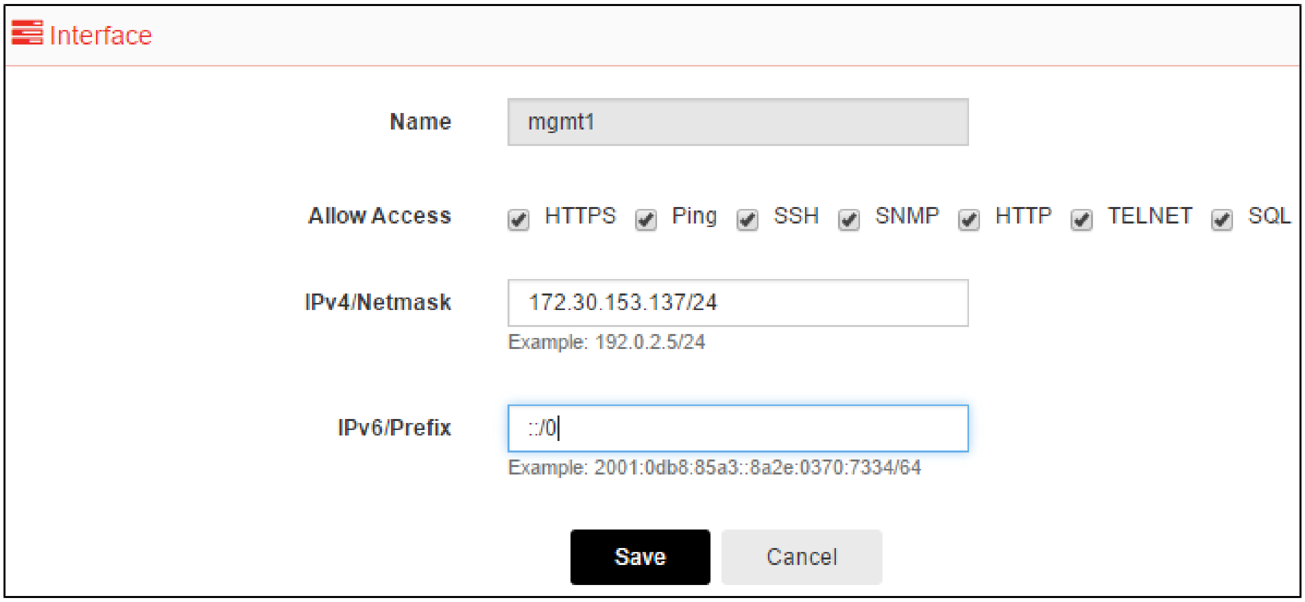

Name: mgmt1

Allow Access: ☑ HTTPS ☑ Ping ☑ SSH ☑ SNMP ☑ HTTP ☑ TELNET ☑ SQL

IPv4/Netmask: 172.30.153.137/24

Example: 192.0.2.5/24

IPv6/Prefix: ::/0

Example: 2001:0db8:85a3::8a2e:0370:7334/64

Table 69: Network interface configuration guidelines

Settings	Guidelines
Speed	Select one of the following speed/duplex settings: • Auto—Speed and duplex are negotiated automatically. Recommended. • 10half—10 Mbps, half duplex. • 10full—10 Mbps, full duplex. • 100half—100 Mbps, half duplex. • 100full—100 Mbps, full duplex. • 1000half—1000 Mbps, half duplex. • 1000full—1000 Mbps, full duplex.
IPv4/Netmask	Management interfaces only. Specify the IP address and CIDR-formatted subnet mask, separated by a forward slash (/), such as 192.0.2.5/24. Dotted quad formatted subnet masks are not accepted.
IPv6/Netmask	Management interfaces only. Specify the IP address and CIDR-formatted subnet mask, separated by a forward slash (/), such as 2001:0db8:85a3:::8a2e:0370:7334/64. Dotted quad formatted subnet masks are not accepted.

Settings	Guidelines
Administrative Access	Management interfaces only. Allow inbound service traffic. Select from the following options: • HTTP—Enables connections to the web UI. We recommend this option only for network interfaces connected to a trusted private network, or directly to your management computer. • HTTPS—Enables secure connections to the web UI. We recommend this option instead of HTTP. • Ping—Enables ping and traceroute to be received on this network interface. When it receives an ECHO_REQUEST (“ping”), the FortiDDOS system replies with ICMP type 0 (ECHO_RESPONSE or “pong”). • SNMP—Enables SNMP queries to this network interface. • SSH—Enables SSH connections to the CLI. We recommend this option instead of Telnet. • Telnet—Enables Telnet connections to the CLI. We recommend this option only for network interfaces connected to a trusted private network, or directly to your management computer. • SQL—Enables SQL queries. **Note**: We recommend that you enable administrative access only on network interfaces connected to trusted private networks or directly to your management computer. If possible, enable only secure administrative access protocols such as HTTPS or SSH.

CLI commands:

```
config system interface
  edit <interface>
    set speed
        {auto|10half|10full|100half|100full|1000half|1000full}
    set status {up|down}
    set ip <address_ipv4> <netmask_ipv4mask>
    set allowaccess {http https ping snmp ssh telnet
        sql}
  end
```

Configuring DNS

The system must be able to contact DNS servers to resolve IP addresses and fully qualified domain names.

Before you begin:

- You must know the IP addresses of the DNS servers used in your network.
- Your Internet service provider (ISP) might supply IP addresses of DNS servers, or you might want to use the IP addresses of your own DNS servers. You must provide unicast, non-local addresses for your DNS servers. Local host and broadcast addresses are not accepted.
- Incorrect DNS settings or unreliable DNS connectivity can cause issues with other features, such as FortiGuard services and NTP system time.
- You must have Read-Write permission for System settings.

To configure DNS:

1. Go to System > Network > DNS.
2. Complete the configuration as described in Table 70.
3. Save the configuration.

Figure 136: DNS configuration page

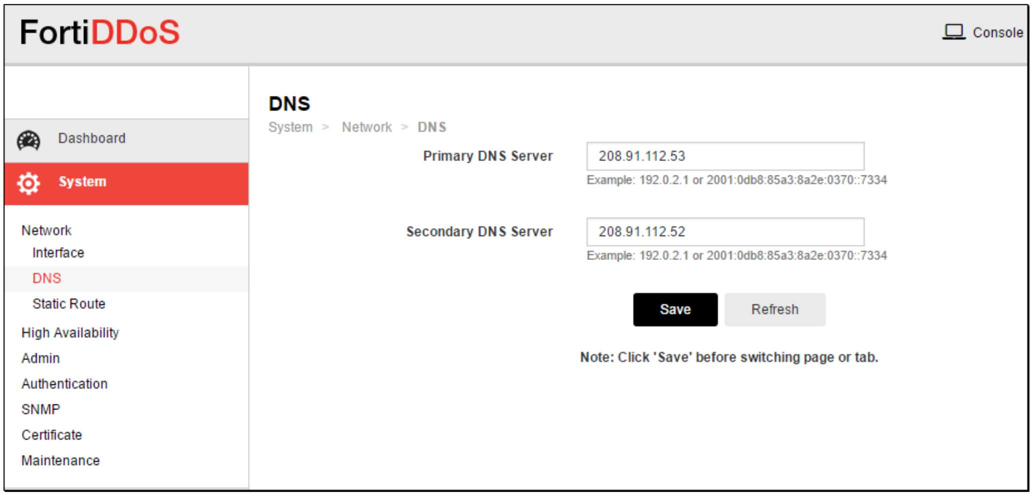

DNS

System > Network > DNS

Primary DNS Server 208.91.112.53

Example: 192.0.2.1 or 2001:0db8:85a3:8a2e:0370::7334

Secondary DNS Server 208.91.112.52

Example: 192.0.2.1 or 2001:0db8:85a3:8a2e:0370::7334

Note: Click 'Save' before switching page or tab.

Table 70: DNS configuration guidelines

Settings	Guidelines
Primary DNS Server	IPv4 address of the primary DNS server. For best performance, use a DNS server on your local network.
Secondary DNS Server	IPv4 address of the secondary DNS server for your local network.

CLI commands:

```
config system dns
   set primary <ip address>
   set secondary <ip address>
end
```

To verify DNS:

```
execute traceroute <server_fqdn>
```

where `<server_fqdn>` is a domain name such as www.example.com.

Configuring static routes

You configure a static route to enable you to connect to the web UI and CLI from a remote location, like your desk.

Before you begin:

- You must have Read-Write permission for System settings.

To configure a static route:

1. Go to System > Network > Static Route.
2. Click **Add** to display the configuration editor.
3. Complete the configuration as described in Table 71.
4. Save the configuration.

Figure 137: Static route configuration page

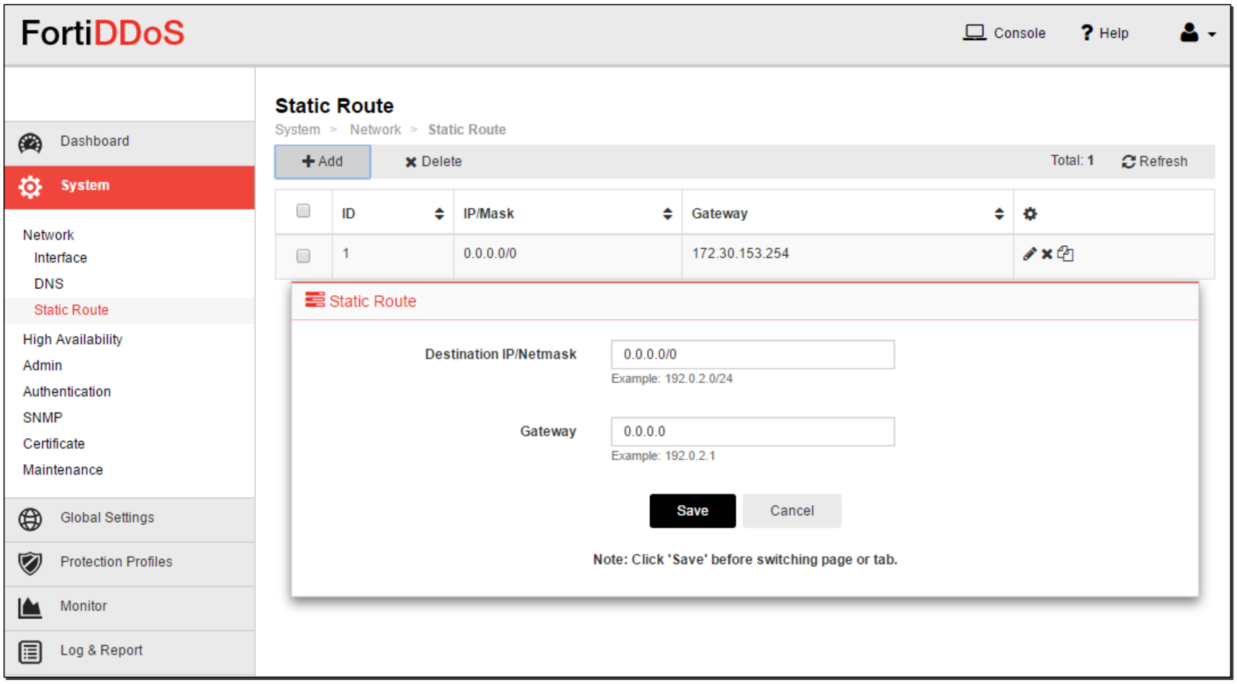

ID	IP/Mask	Gateway
1	0.0.0.0/0	172.30.153.254

Static Route

Destination IP/Netmask: 0.0.0.0/0
Example: 192.0.2.0/24

Gateway: 0.0.0.0
Example: 192.0.2.1

Note: Click 'Save' before switching page or tab.

Table 71: Static route configuration guidelines

Settings	Guidelines
Interface	Select the network interface that uses the static route.

Settings	Guidelines
Destination IP/mask	Destination IP address and network mask of packets that use this static route, separated by a slash (/) or space. The value `0.0.0.0/0` is a default route, which matches all packets
Gateway	IP address of the next-hop router for the FortiDDoS management computer.

To configure a static route using the CLI:

```
config system default-gateway
  edit <route_number>
    set destination <destination_ipv4/mask>
    set gateway <gateway_ipv4>
    set interface {mgmt1|mgmt2}
  end
```

Configuring RADIUS authentication

You can configure administrator authentication against a RADIUS server.

After you have completed the RADIUS server configuration and enabled it, you can select it when you create an administrator user on the System > Admin > Administrators page. On that page, you specify the username but not the password. You also specify the SPP assignment, trusted host list, and access profile for that user.

If RADIUS is enabled, when a user logs in, an authentication request is made to the remote RADIUS server. If authentication succeeds, and the user has a configuration on the System > Admin > Administrators page, the SPP assignment, trusted host list, and access profile are applied. If the user does not have a configuration on the System > Admin > Administrators page, these assignments are obtained from the Default Access Strategy settings described in Table 72.

Before you begin:

- You must have Read-Write permission for System settings.

To configure a RADIUS server:

1. Go to System > Authentication > RADIUS.
2. Complete the configuration as described in Table 72.
3. Save the configuration.

Figure 138: RADIUS server configuration page

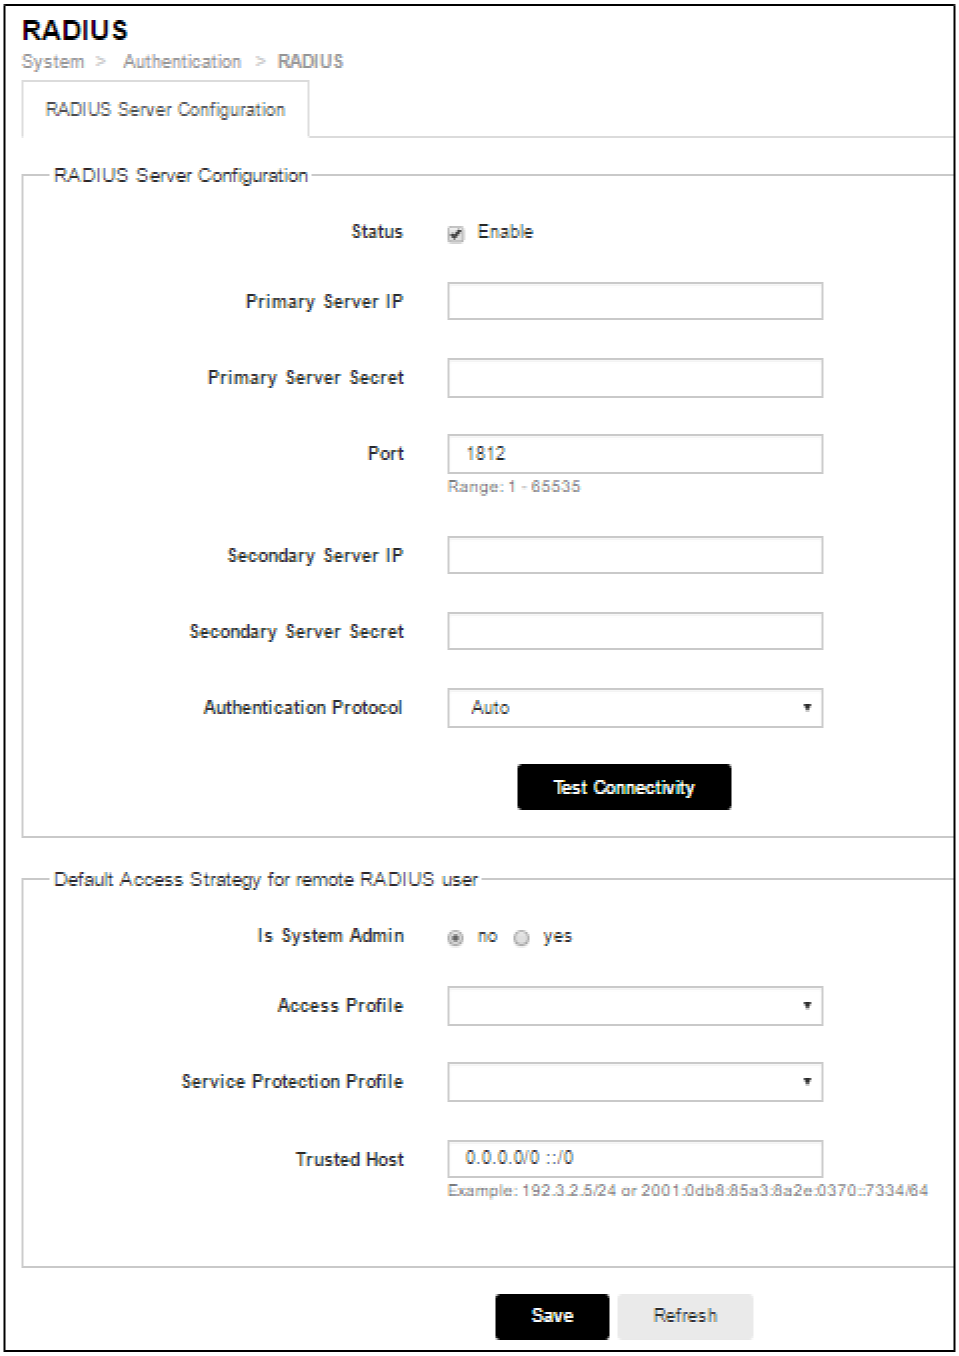

RADIUS

System > Authentication > RADIUS

RADIUS Server Configuration

RADIUS Server Configuration

Status ☑ Enable

Primary Server IP

Primary Server Secret

Port 1812

Range: 1 - 65535

Secondary Server IP

Secondary Server Secret

Authentication Protocol Auto ▾

Default Access Strategy for remote RADIUS user

Is System Admin ● no ○ yes

Access Profile ▾

Service Protection Profile ▾

Trusted Host 0.0.0.0/0 ::/0

Example: 192.3.2.5/24 or 2001:0db8:85a3:8a2e:0370::7334/64

Table 72: RADIUS server configuration guidelines

Settings	Guidelines
Enable	Unique name. No spaces or special characters.
Primary Server Name/IP	IP address of the primary RADIUS server.

Settings	Guidelines
Primary Server Secret	RADIUS server shared secret.
Secondary Server Name/IP	Optional. IP address of a backup RADIUS server.
Secondary Server Secret	Optional. RADIUS server shared secret.
Port	RADIUS port. Usually, this is 1812.
Auth Protocol	• Auto—If you leave this default value, the system uses MSCHAP2. • PAP—Password Authentication Protocol • CHAP—Challenge Handshake Authentication Protocol (defined in RFC 1994) • MSCHAP—Microsoft CHAP (defined in RFC 2433) • MSCHAP2—Microsoft CHAP version 2 (defined in RFC 2759)
Test Connectivity	
Test Connectivity	Select to test connectivity using a test username and password specified next. Click the **Test** button before you save the configuration.
Username	Username for the connectivity test.
Password	Corresponding password.
Default Access Strategy for remote RADIUS user	
System Admin	If enabled, the user is regarded as a system administrator with access to all SPPs.
Service Protection Profile	If this administrator is not a system administrator, select the profile that this account manages.

Settings	Guidelines
Trusted Hosts	Source IP address and netmask from which the administrator is allowed to log in. For multiple addresses, separate each entry with a space. You can specify up to three trusted areas. They can be single hosts, subnets, or a mixture. Configuring trusted hosts hardens the security of the system. In addition to knowing the password, an administrator can connect only from the computer or subnets you specify. Trusted host definitions apply both to the web UI and to the CLI when accessed through Telnet, SSH, or the CLI console widget. Local console access is *not* affected by trusted hosts, as the local console is by definition not remote, and does not occur through the network. If ping is enabled, the address you specify here is also a source IP address to which the system will respond when it receives a ping or traceroute signal. To allow logins only from *one* computer, enter only its IP address and 32- or 128-bit netmask:`192.0.2.2/32` `2001:0db8:85a3:::8a2e:0370:7334/128` To allow login attempts from any IP address (not recommended), enter:`0.0.0.0/0.0.0.0`. **Caution:** If you restrict trusted hosts, do so for *all* administrator accounts. Failure to do so means that all accounts are still exposed to the risk of brute force login attacks. This is because if you leave even *one* administrator account unrestricted (i.e. `0.0.0.0/0`), the system must allow login attempts on all network interfaces where remote administrative protocols are enabled, and wait until *after* a login attempt has been received in order to check that user name's trusted hosts list. **Tip:** If you allow login from the Internet, set a longer and more complex password, and enable only secure administrative access protocols. We also recommend that you restrict trusted hosts to IPs in your administrator's geographical area. **Tip:** For improved security, restrict all trusted host addresses to single IP addresses of computer(s) from which *only* this administrator will log in.
Default Access Profile	Select a user-defined or predefined profile. The predefined profile named **super_admin_prof** is a special access profile used by the **admin** account. However, selecting this access profile will *not* confer all permissions of the **admin** account. For example, the new administrator would not be able to reset lost administrator passwords. **Note**: This option does not appear for the **admin** administrator account, which by definition always uses the **super_admin_prof** access profile.

```
config system authentication radius
  set state {enable|disable}
  set primary-server <ip>
  set primary-secret <string>
  set backup-server <ip>
  set backup-secret <string>
  set port <port>
  set authprot {auto|chap|mschap|mschapv|pap}
  set is-system-admin {yes|no}
  set dft-domain <SPP>
  set dft-accprofile <profile>
  set dft-trusted-hosts <CIDR list>
end
```

Configuring LDAP authentication

You can configure administrator authentication against an LDAP server.

Before you begin:

- You must have Read-Write permission for System settings.
- You must work with your LDAP administrator to determine an appropriate DN for FortiADC access. The LDAP administrator might need to provision a special group.

After you have completed the LDAP server configuration and enabled it, you can select it when you create an administrator user on the System > Admin > Administrators page. On that page, you specify the username but not the password. You also specify the SPP assignment, trusted host list, and access profile for that user.

If LDAP is enabled, when a user logs in, an authentication request is made to the remote LDAP server. If authentication succeeds, and the user has a configuration on the System > Admin > Administrators page, the SPP assignment, trusted host list, and access profile are applied. If the user does not have a configuration on the System > Admin > Administrators page, these assignments are obtained from the Default Access Strategy settings described in Table 73.

To configure an LDAP server:

1. Go to System > Authentication > LDAP.
2. Complete the configuration as described in Table 73.
3. Save the configuration.

Figure 139: LDAP server configuration page

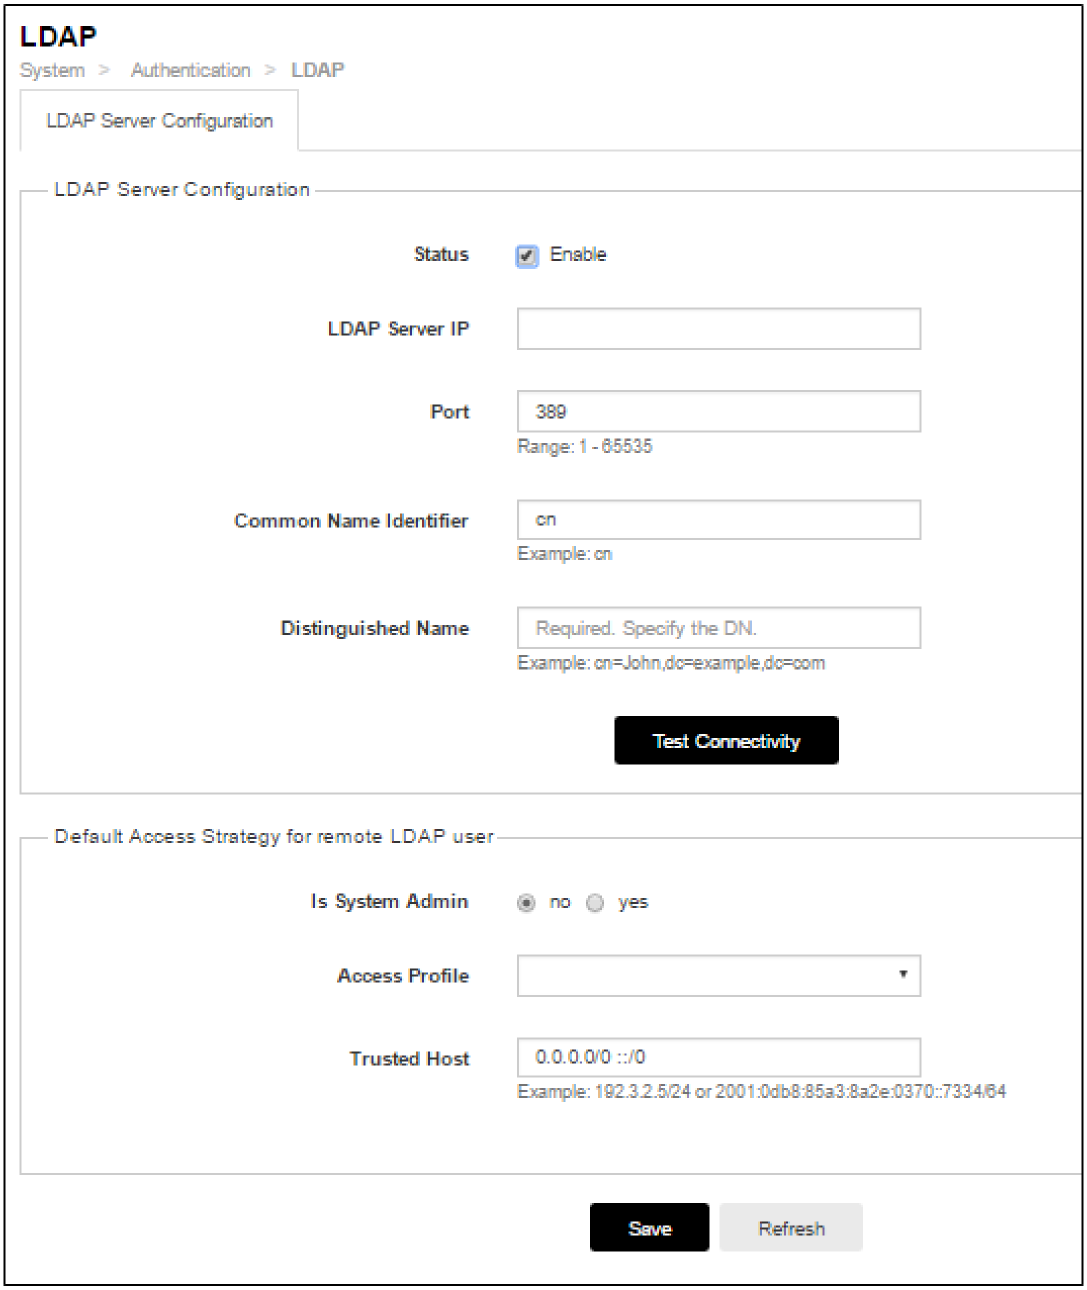

LDAP

System > Authentication > LDAP

LDAP Server Configuration

LDAP Server Configuration

Status ☑ Enable

LDAP Server IP

Port 389
Range: 1 - 65535

Common Name Identifier cn
Example: cn

Distinguished Name Required. Specify the DN.
Example: cn=John,dc=example,dc=com

Default Access Strategy for remote LDAP user

Is System Admin ● no ○ yes

Access Profile

Trusted Host 0.0.0.0/0 ::/0
Example: 192.3.2.5/24 or 2001:0db8:85a3:8a2e:0370::7334/64

Table 73: LDAP server configuration guidelines

Settings	Guidelines
Enable	Unique name. No spaces or special characters.
LDAP Server Name/IP	IP address of the LDAP server.
Port	LDAP port. Usually, this is 389.

Settings	Guidelines
Common Name Identifier	Common name (cn) attribute for the LDAP record. For example: `cn` or `uid`.
Distinguished Name	Distinguished name (dn) attribute for the LDAP record. The dn uniquely identifies a user in the LDAP directory. For example: `cn=John%20Doe,dc=example,dc=com` Most likely, you must work with your LDAP administrator to know the appropriate DN to use for FortiADC access. The LDAP administrator might need to provision a special group.
Test Connectivity	
Test Connectivity	Select to test connectivity using a test username and password specified next. Click the **Test** button after you have saved the configuration.
User DN	User DN for the connectivity test.
Password	Corresponding password.
Default Access Strategy for remote LDAP user	
System Admin	If enabled, the user is regarded as a system administrator with access to all SPPs.
Service Protection Profile	If this administrator is not a system administrator, select the profile that this account manages.

Settings	Guidelines
Trusted Hosts	Source IP address and netmask from which the administrator is allowed to log in. For multiple addresses, separate each entry with a space. You can specify up to three trusted areas. They can be single hosts, subnets, or a mixture. Configuring trusted hosts hardens the security of the system. In addition to knowing the password, an administrator can connect only from the computer or subnets you specify. Trusted host definitions apply both to the web UI and to the CLI when accessed through Telnet, SSH, or the CLI console widget. Local console access is *not* affected by trusted hosts, as the local console is by definition not remote, and does not occur through the network. If ping is enabled, the address you specify here is also a source IP address to which the system will respond when it receives a ping or traceroute signal. To allow logins only from *one* computer, enter only its IP address and 32- or 128-bit netmask:`192.0.2.2/32` `2001:0db8:85a3:::8a2e:0370:7334/128` To allow login attempts from any IP address (not recommended), enter:`0.0.0.0/0.0.0.0`. **Caution:** If you restrict trusted hosts, do so for *all* administrator accounts. Failure to do so means that all accounts are still exposed to the risk of brute force login attacks. This is because if you leave even *one* administrator account unrestricted (i.e. `0.0.0.0/0`), the system must allow login attempts on all network interfaces where remote administrative protocols are enabled, and wait until *after* a login attempt has been received in order to check that user name's trusted hosts list. **Tip:** If you allow login from the Internet, set a longer and more complex password, and enable only secure administrative access protocols. We also recommend that you restrict trusted hosts to IPs in your administrator's geographical area. **Tip:** For improved security, restrict all trusted host addresses to single IP addresses of computer(s) from which *only* this administrator will log in.
Default Access Profile	Select a user-defined or predefined profile. The predefined profile named **super_admin_prof** is a special access profile used by the **admin** account. However, selecting this access profile will *not* confer all permissions of the **admin** account. For example, the new administrator would not be able to reset lost administrator passwords. **Note**: This option does not appear for the **admin** administrator account, which by definition always uses the **super_admin_prof** access profile.

```
config system authentication LDAP
  set state {enable|disable}
  set server <ip>
  set port <port>
  set cnid <cn>
  set dn <dn>
  set is-system-admin {yes|no}
  set dft-domain <SPP>
  set dft-accprofile <profile>
  set dft-trusted-hosts <CIDR list>
end
```

If you initially set is-system-admin to no, but later change your mind, you must first change dft-domain to SPP-0 and commit it; then configure the system admin setting. For example:

```
config system authentication LDAP
   set dft-domain SPP-0
end
config system authentication LDAP
   set is-system-admin yes
end
```

Managing administrator users

This topic includes the following information:

- Administrator user overview
- Configuring access profiles
- Creating administrator users
- Changing user passwords
- Configuring administration settings

Administrator user overview

In its factory default configuration, FortiDDoS has one administrator account named **admin**. This administrator has permissions that grant Read-Write access to all system functions.

Unlike other administrator accounts, the administrator account named admin exists by default and cannot be deleted. The **admin** account is similar to a root administrator account. This account always has full permission to view and change all system configuration options, including viewing and changing *all* other administrator accounts. Its name and permissions cannot be changed. It is the only administrator account that can reset another administrator’s password without being required to enter that administrator’s existing password.

To prevent accidental changes to the configuration, it is best if only network administrators—and if possible, only a single person—use the **admin** account. You can use the **admin** account to configure more administrator accounts for other people. Accounts can be made with different scopes of access. You can associate each of these accounts with either all SPPs or a single SPP, and you can specify the type of profile settings that each account can access. If you require such role-based access control (RBAC) restrictions, or if you simply want to harden security or prevent inadvertent changes to other administrators’ areas, you can do so with access profiles. For example, you can create an account for a security auditor who must only be able to view the configuration and logs, but *not* change them.

Basic steps

1. Configure profiles to provision permissions to roles.
2. Optional. Create RADIUS or LDAP server configurations if you want to use a RADIUS or LDAP server to authenticate administrators. Otherwise, you can use local authentication.
3. Create administrator user accounts with permissions provisioned by the profiles.

Configuring access profiles

Access profiles provision permissions to roles. The following permissions can be assigned:

- Read (view access)
- Read-Write (view, change, and execute access)
- No access

When a profile includes only read access to a category, the user can access the web UI page for that category, and can use the `get` and `show` CLI command for that category, but cannot make changes to the configuration.

When a profile includes no categories with read-write permissions, the user can log into the web UI but not the CLI.

In larger companies where multiple administrators share the workload, access profiles often reflect the specific job that each administrator does ("role"), such as account creation or log auditing. Access profiles can limit each administrator account to their assigned role. This is sometimes called role-based access control (RBAC).

Table 74 lists the administrative areas that can be provisioned. If you provision read access, the role can view the web UI menu (or issue a CLI get command). If you provision read-write access, the role can save configuration changes (or issue a CLI set command).

For complete access to *all* commands and abilities, you must log in with the administrator account named **admin**.

Table 74: Areas of control in access profiles

Web UI Menus	CLI Commands
System	`config system` ... `show full-configuration` `diagnose` ... `execute` ...
Global Settings	`config ddos global` ...
Protection Profiles	`config spp` ...

Monitor	`get system status`
	`get system performance`
	`show system status`
	`show system performance`
	`show full-configuration`

Log & Report	`config log` ... `config system`

* For each `config` command, there is an equivalent `get`/`show` command, unless otherwise noted. `config` commands require write permission. `get`/`show` commands require read permission.

Before you begin:

- You must have Read-Write permission for System settings.

To configure administrator profiles:

1. Go to System > Admin > Access Profile.
2. Click **Add** to display the configuration editor.
3. Complete the configuration as described in Table 75.
4. Save the configuration.

Figure 140: Admin profile configuration page

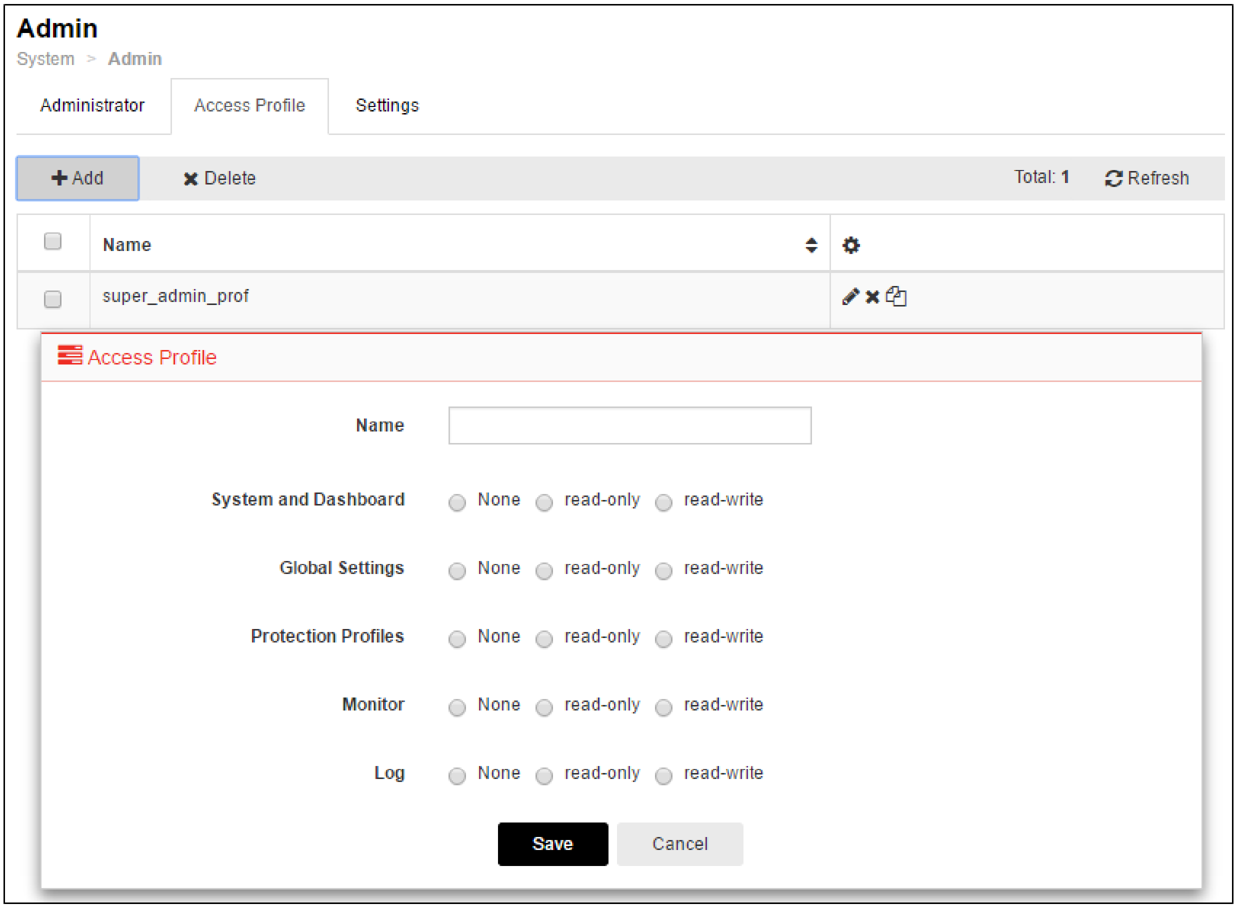

Admin

System > Admin

□	Name
□	super_admin_prof

Access Profile

Name:

System and Dashboard: ○ None ○ read-only ○ read-write

Global Settings: ○ None ○ read-only ○ read-write

Protection Profiles: ○ None ○ read-only ○ read-write

Monitor: ○ None ○ read-only ○ read-write

Log: ○ None ○ read-only ○ read-write

Table 75: Admin profile configuration guidelines

Settings	Guidelines
Profile name	Unique name. No spaces or special characters.
Access Control	• None—Do not provision access for the menu. • Read Only—Provision ready-only access. • Read-Write—Enable the role to make changes to the configuration.

The **super_admin_prof** access profile, a special access profile assigned to the **admin** account and required by it, appears in the list of access profiles. It exists by default and cannot be changed or deleted. The profile has permissions similar to the UNIX root account.

Creating administrator users

We recommend that only network administrators—and if possible, only a single person—use the **admin** account. You can configure accounts that provision different scopes of access. For example, you can create an account for a security auditor who must only be able to view the configuration and logs, but *not* change them.

Before you begin:

- You must have Read-Write permission for System settings.

To create administrator users:

1. Go to System > Admin > Administrators.
2. Click **Add** to display the configuration editor.
3. Complete the configuration as described in Table 76.
4. Save the configuration.

Figure 141: Administrator user configuration page

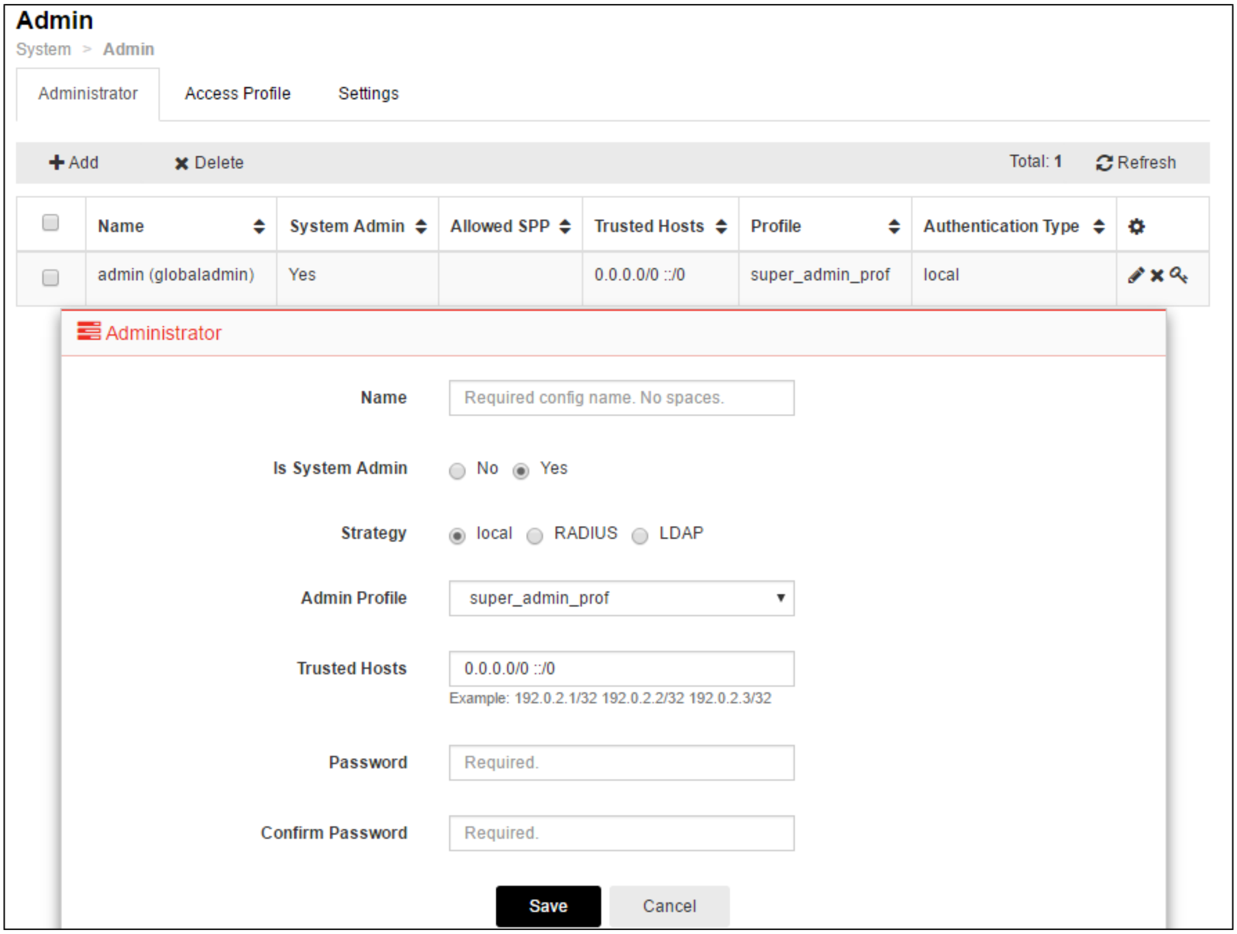

Admin

System > Admin

□	Name	System Admin	Allowed SPP	Trusted Hosts	Profile	Authentication Type	⚙
□	admin (globaladmin)	Yes		0.0.0.0/0 ::/0	super_admin_prof	local	

Administrator

Name	Required config name. No spaces.
Is System Admin	○ No ● Yes
Strategy	● local ○ RADIUS ○ LDAP
Admin Profile	super_admin_prof ▾
Trusted Hosts	0.0.0.0/0 ::/0 Example: 192.0.2.1/32 192.0.2.2/32 192.0.2.3/32
Password	Required.
Confirm Password	Required.

Table 76: Administrator user configuration guidelines

Settings	Guidelines
Name	Name of the administrator account, such as `admin1` or `admin@example.com`, that can be referenced in other parts of the configuration. Do not use spaces or special characters except the 'at' symbol (`@`). The maximum length is 35 characters. If you use LDAP or RADIUS authentication, this is the username stored in the LDAP or RADIUS authentication server. **Note:** This is the user name that the administrator must provide when logging in to the CLI or web UI. If using an external authentication server such as RADIUS or Active Directory, this name will be passed to the server via the remote authentication query.
System Admin	• Yes—Administrator for all SPPs. • No—Administrator for the selected SPP only.
Service Protection Profile	If this administrator is not a system administrator, select the profile that this account manages.
Auth Strategy	• Local—Use the local authentication server. When you use the local authentication server, you also configure a password. • LDAP—Authenticate against an LDAP server. When you use LDAP, you do not configure a password. The system authenticates against the username and password stored in the LDAP server. • RADIUS—Authenticate against a RADIUS server. When you use RADIUS, you do not configure a password. The system authenticates against the username and password stored in the RADIUS server.
New Password	Type a password for the administrator account. Passwords may have a maximum of 16 characters, may include numbers, upper and lowercase characters, and the following special characters: _ (underscore), - (hyphen), !, @, #, $, %, ^, &, *
Confirm Password	Type the password again to confirm its spelling.

Settings	Guidelines
Trusted Hosts	Source IP address and netmask from which the administrator is allowed to log in. For multiple addresses, separate each entry with a space. You can specify up to three trusted areas. They can be single hosts, subnets, or a mixture. Configuring trusted hosts hardens the security of the system. In addition to knowing the password, an administrator can connect only from the computer or subnets you specify. Trusted host definitions apply both to the web UI and to the CLI when accessed through Telnet, SSH, or the CLI console widget. Local console access is *not* affected by trusted hosts, as the local console is by definition not remote, and does not occur through the network. If ping is enabled, the address you specify here is also a source IP address to which the system will respond when it receives a ping or traceroute signal. To allow logins only from *one* computer, enter only its IP address and 32- or 128-bit netmask:`192.0.2.2/32` `2001:0db8:85a3:::8a2e:0370:7334/128` To allow login attempts from any IP address (not recommended), enter:`0.0.0.0/0.0.0.0`. **Caution:** If you restrict trusted hosts, do so for *all* administrator accounts. Failure to do so means that all accounts are still exposed to the risk of brute force login attacks. This is because if you leave even *one* administrator account unrestricted (i.e. `0.0.0.0/0`), the system must allow login attempts on all network interfaces where remote administrative protocols are enabled, and wait until *after* a login attempt has been received in order to check that user name’s trusted hosts list. **Tip:** If you allow login from the Internet, set a longer and more complex password, and enable only secure administrative access protocols. We also recommend that you restrict trusted hosts to IPs in your administrator’s geographical area. **Tip:** For improved security, restrict all trusted host addresses to single IP addresses of computer(s) from which *only* this administrator will log in.
Admin Profile	Select a user-defined or predefined profile. The predefined profile named **super_admin_prof** is a special access profile used by the **admin** account. However, selecting this access profile will *not* confer all permissions of the **admin** account. For example, the new administrator would not be able to reset lost administrator passwords. **Note**: This option does not appear for the **admin** administrator account, which by definition always uses the **super_admin_prof** access profile.

CLI commands:

```
config system admin
edit admin
set access-profile super_admin_prof
next
edit admin-spp1
     set is-system-admin no
     set domain SPP-1
     set password ENC $1$0b721b38$vk7GoO147JXXqy5B3ag8z/
     set access-profile admin
  end
```

Changing user passwords

By default, this administrator account has no password. Set a strong password for the `admin` administrator account. Change the password regularly.

Before you begin:

- You must have Read-Write permission for System settings.

To change the password:

1. Go to System > Admin > Administrator.
2. Click **Change Password** icon.
3. Complete the configuration as described in Table 77.
4. Save the configuration.

Figure 142: Administrator settings page

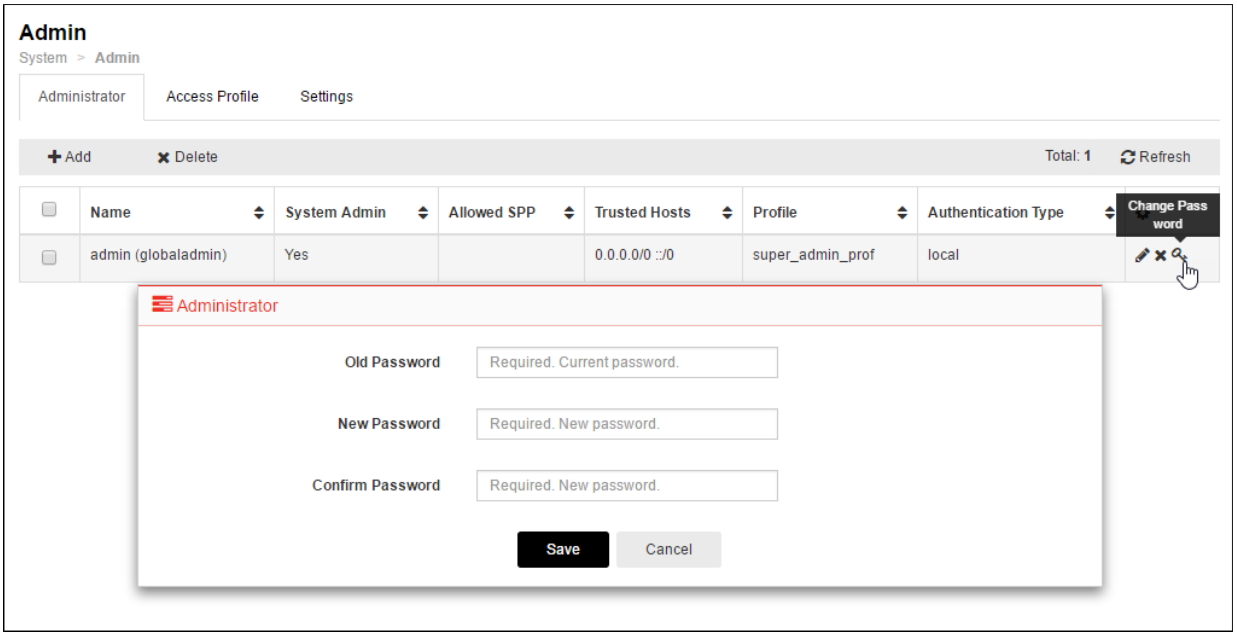

□	Name	System Admin	Allowed SPP	Trusted Hosts	Profile	Authentication Type	
□	admin (globaladmin)	Yes		0.0.0.0/0 ::/0	super_admin_prof	local	

Old Password	Required. Current password.
New Password	Required. New password.
Confirm Password	Required. New password.

Table 77: Password configuration

Settings	Guidelines
Old Password	Type the current password.
New Password	Type a password for the administrator account. Passwords may have a maximum of 16 characters, may include numbers, upper and lowercase characters, and the following special characters: _ (underscore), - (hyphen), !, @, #, $, %, ^, &, *
Confirm Password	Type the password again to confirm its spelling.

CLI commands:

```
config system admin
   edit admin
      set password <new-password_str>
   end
```

Configuring administration settings

Before you begin:

- You must have Read-Write permission for System settings.

To change the administration settings:

1. Go to System > Admin > Settings.
2. Complete the configuration as described in Table 78.
3. Save the configuration.

Figure 143: Administration settings page

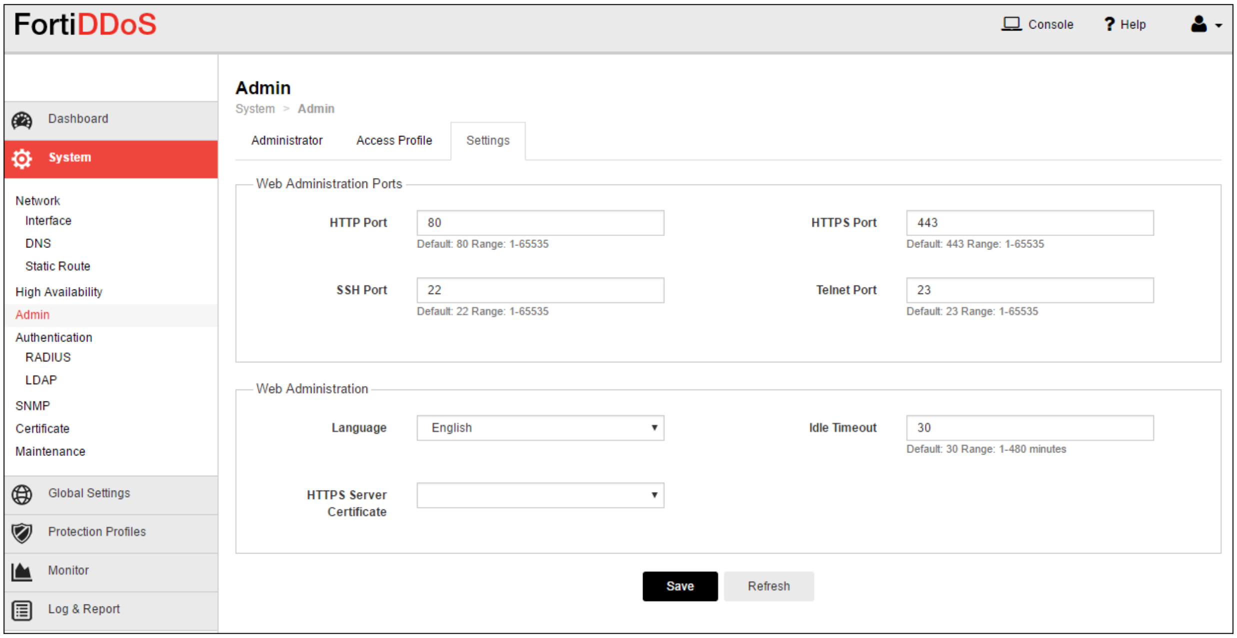

Admin

System > Admin

Web Administration Ports

HTTP Port: 80
Default: 80 Range: 1-65535

HTTPS Port: 443
Default: 443 Range: 1-65535

SSH Port: 22
Default: 22 Range: 1-65535

Telnet Port: 23
Default: 23 Range: 1-65535

Web Administration

Language: English

Idle Timeout: 30
Default: 30 Range: 1-480 minutes

HTTPS Server Certificate:

Table 78: Administration settings guidelines

Settings	Guidelines
Web Administration Ports	
HTTP	Specify the port for the HTTP service. Usually, HTTP uses port 80.
HTTPS	Specify the port for the HTTPS service. Usually, HTTPS uses port 443.
Telnet	Specify the port for the Telnet service. Usually, Telnet uses port 25.
SSH	Specify the port for the SSH service. Usually, SSH uses port 22.
Web Administration	

Settings	Guidelines
Language	Language of the web UI. The following languages are supported: • English • Simplified Chinese • Korean • Japanese The display’s web pages use UTF-8 encoding, regardless of which language you choose. UTF-8 supports multiple languages, and allows them to display correctly, even when multiple languages are used on the same web page. For example, your organization could have websites in both English and simplified Chinese. Your FortiDDoS administrators prefer to work in the English version of the web UI. They could use the web UI in English while writing rules to match content in both English and simplified Chinese *without* changing this setting. Both the rules and the web UI will display correctly, as long as all rules were input using UTF-8. Usually, your text input method or your management computer’s operating system should match the display by also using UTF-8. If they do not, your input and the web UI may not display correctly at the same time. For example, your web browser’s or operating system’s default encoding for simplified Chinese input may be GB2312. However, you *usually* should switch it to be UTF-8 when using the web UI, *unless* you are writing regular expressions that must match HTTP client’s requests, and those requests use GB2312 encoding. **Note:** This setting does *not* affect the display of the CLI.
Timeout	Number of minutes that a web UI connection can be idle before the administrator must log in again. The maximum is 480 minutes (8 hours). The default is 30 minutes.

Managing local certificates

This section includes the following information:

- Overview
- Generating a Certificate Signing Request (CSR)
- Importing certificates
- Using certificates
- Viewing certificates

Overview

While requesting secure administrator access to a FortiDDoS device via HTTPS, the device uses SSL protocol to ensure that all communication between the device and the HTTP browser is secure no matter which client application is used. Regarding basic authentication made by an HTTP client, the device will use its self-signed security certificate to allow authentication whenever HTTPS is initiated by the client.

Note: The self-signed certificate proposal is the default setting on the device.

The HTTP browser notices the following discrepancies:

- The 'issuer' of the certificate offered by the device is unknown.
- The 'subject' of the certificate doesn't match the FQDN of the HTTP request a.b.c.d.

To avoid the triggering of these messages in the scenario where you don't require your HTTP browser to 'Permanently store this exception':

- Always ensure that the certificate of the CA signed by the device certificate is stored in the browser repository.
- Always ensure that the device is accessed with a correct FQDN.

Once the security exception is confirmed, the login page will be displayed. All the data sent to the device is encrypted and a HTTPS connection is created without reading the self-signed certificate proposal. Once the HTTP browser has permanently stored this exception, the exception prompt is not shown again. If the HTTP client declines the certificate, then the device does not allow the connection.

If you want to avoid these warnings and have a custom certificate, you must assign a host name to the appliance, generate a key pair and certificate request and import the certificate from a signing authority.

Generating a Certificate Signing Request (CSR)

FortiDDoS allows you to generate CSRs that you can send to a CA to sign and give you a signed certificate. FortiDDoS creates a key pair that it keeps in a protected storage and is later used for SSL.

Before you begin:

• You must have Read-Write permission for System settings.

To generate a certificate request:

1. Go to System > Certificate > Local Certificates.
2. Click **Generate** to display the configuration editor.
3. Complete the configuration as described in the Table 94.

4. Save the configuration.
   The system creates a private and public key pair. The generated request includes the public key of the FortiDDoS appliance and information such as the IP address, domain name, or email address. The FortiDDoS appliance private key remains confidential in the FortiDDoS appliance. The Status column of the new CSR entry is Pending.
5. Select the row that corresponds to the certificate request.
6. Click **Download**.
   Standard dialogs appear with buttons to save the file to the location you select. Your web browser downloads the certificate request (.csr) file.
7. Upload the certificate request to your CA.
   After you submit the request to a CA, the CA will verify the information in the certificate, give it a serial number, an expiration date, and sign it with the public key of the CA.
8. If you are not using a commercial CA whose root certificate is already installed by default on web browsers, download your CA's root certificate, then install it on all computers that will be connecting to your appliance. (If you do not install these, those computers might not trust your new certificate.)
9. When you receive the signed certificate from the CA, you can import the certificate into the FortiDDoS system.

Table 79: CSR configuration

Settings	Guidelines
Generate Certificate	**Signing Request**
Certification Name	Configuration name. Valid characters are `A-Z`,`a-z`,`0-9`,`_`, and -. No spaces. The maximum length is 35 characters. **Note**: This is the name of the CSR file, not the host name/IP contained in the certificate's `Subject:` line.
Subject Information	

Settings	Guidelines
ID Type	Select the type of identifier to use in the certificate to identify the virtual server: • Host IP—The static public IP address of the FortiDDoS virtual server in the **IP Address** field. If the FortiDDoS appliance does not have a static public IP address, use the email or domain name options instead. **Note**: If your network has a dynamic public IP address, you should not use this option. An “Unable to verify certificate” or similar error message will be displayed by users’ browsers when your public IP address changes. • Domain Name—The fully qualified domain name (FQDN) of the FortiDDoS virtual server, such as `www.example.com`. This does not require that the IP address be static, and may be useful if, for example, your network has a dynamic public IP address and therefore clients connect to it via dynamic DNS. Do not include the protocol specification (http://) or any port number or path names. • E-Mail—The email address of the owner of the FortiDDoS virtual server. Use this if the virtual server does not require either a static IP address or a domain name. Depending on your choice for **ID Type**, related options appear.
IP Address	Type the static IP address of the FortiDDoS appliance, such as `10.0.0.1`.The IP address should be the one that is visible to clients. Usually, this should be its public IP address on the Internet, or a virtual IP that you use NAT to map to the appliance’s IP address on your private network. This option appears only if **ID Type** is **Host IP**.
Domain Name	Type the FQDN of the FortiDDoS appliance, such as `www.example.com`. The domain name must resolve to the IP address of the FortiDDoS appliance or backend server according to the DNS server used by clients. (If it does not, the clients’ browsers will display a Host name mismatch or similar error message.) This option appears only if **ID Type** is **Domain Name**.
E-mail	Type the email address of the owner of the FortiDDoS appliance, such as `admin@example.com`. This option appears only if **ID Type** is **E-Mail.**
Distinguished Information	
Organization Unit	Name of organizational unit (OU), such as the name of your department. This is optional. To enter more than one OU name, click the **+** icon, and enter each OU separately in each field

Settings	Guidelines
Organization	Legal name of your organization.
Locality (City)	City or town where the FortiDDoS appliance is located.
State/Province	State or province where the FortiDDoS appliance is located.
Country/Region	Country where the FortiDDoS appliance is located.
Email	Email address that may be used for contact purposes, such as `admin@example.com`.
Key Information	
Key Type	RSA
Key Size	Select a secure key size. Larger keys use more computing resources, but provide better security. For RSA, select one of the following: • 1024 Bit • 1536 Bit • 2048 Bit
Enrollment Information	
Enrollment Method	File Based—You must manually download and submit the resulting certificate request file to a CA for signing. Once signed, upload the local certificate. Online SCEP—The FortiDDoS appliance automatically uses HTTP to submit the request to the simple certificate enrollment protocol (SCEP) server of a CA, which will validate and sign the certificate. For this selection, two options appear. Enter the **CA Server URL** and the **Challenge Password**.

Importing certificates

You can import or upload the following types of server certificates and private keys to the FortiDDoS system:

- local
- PKCS12
- certificate

Before you begin:

- You must have Read-Write permission for System settings.
- You must have downloaded the certificate and key files to browse and upload.

To import a local certificate:

1. Go to System > Certificate > Local Certificate.
2. Click Import to display the configuration editor.

Figure 144: Importing a Local Certificate

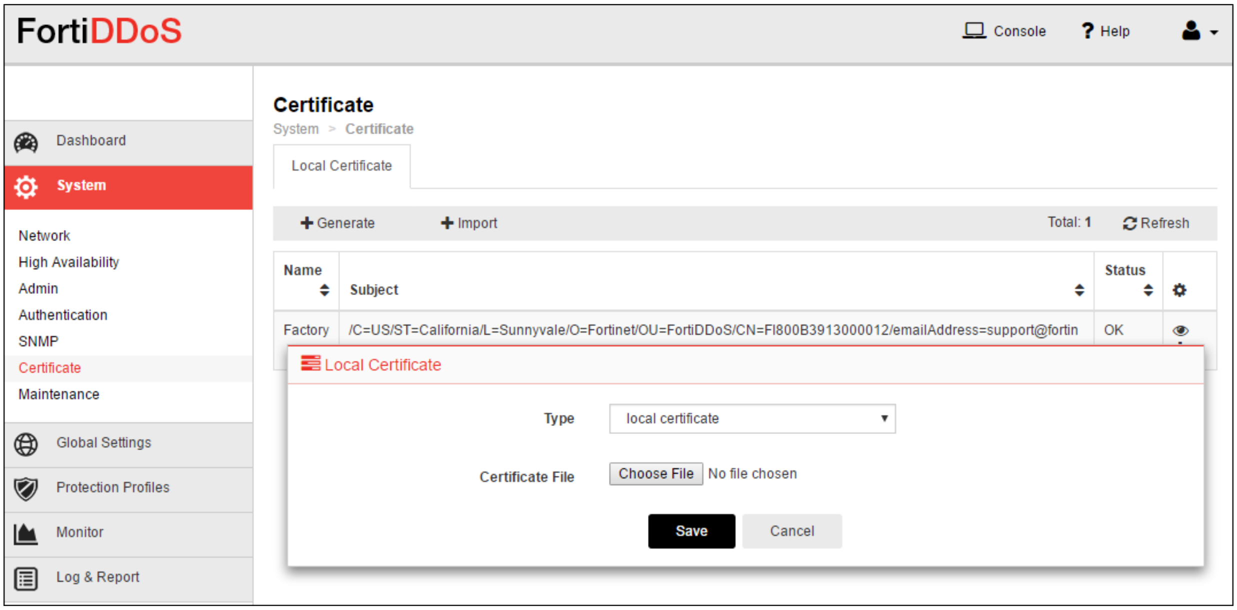

3. Complete the configuration based on the Type selected, as described in Table 95.
4. Save the configuration.

Table 80: Local certificate import configuration

Settings	Guidelines
Type	• Local Certificate: An unencrypted certificate in PEM format. • PKCS12 Certificate: A PKCS #12 password-encrypted certificate with key in the same file. • Certificate: An unencrypted certificate in PEM format. The key is in a separate file. Additional fields are displayed depending on your selection.
Local Certificate	
Certificate File	Browse and locate the certificate file that you want to upload.
PKCS12 Certificate	
Certificate Name	Name that can be referenced by other parts of the configuration, such as `www_example_com`. • Do not use spaces or special characters. • Maximum length is 35 characters.
Certificate File	Browse and locate the certificate file that you want to upload.
Password	Password to encrypt the file in local storage.

Settings	Guidelines
Certificate	
Certificate Name	Name that can be referenced by other parts of the configuration, such as `www_example_com`. • Do not use spaces or special characters. • Maximum length is 35 characters.
Certificate File	Browse and locate the certificate file that you want to upload.
Password	Password to encrypt the files in local storage.

After the certificate is imported, status shows OK.

Using certificates

1. Go to System > Admin > Settings.
2. Select the certificate from the dropdown under Under Web Administration > HTTPS Server Certificate.
3. Save the configuration.

Viewing certificates

The system has its own default “Factory” certificate that it presents to establish secure connections with the administrator client computer.

To view the local certificate:

1. Go to System > Certificates.
2. Click the Local Certificate tab.
3. Double-click the row corresponding to the Factory Certificate.

Figure 145: Factory Local Certificate

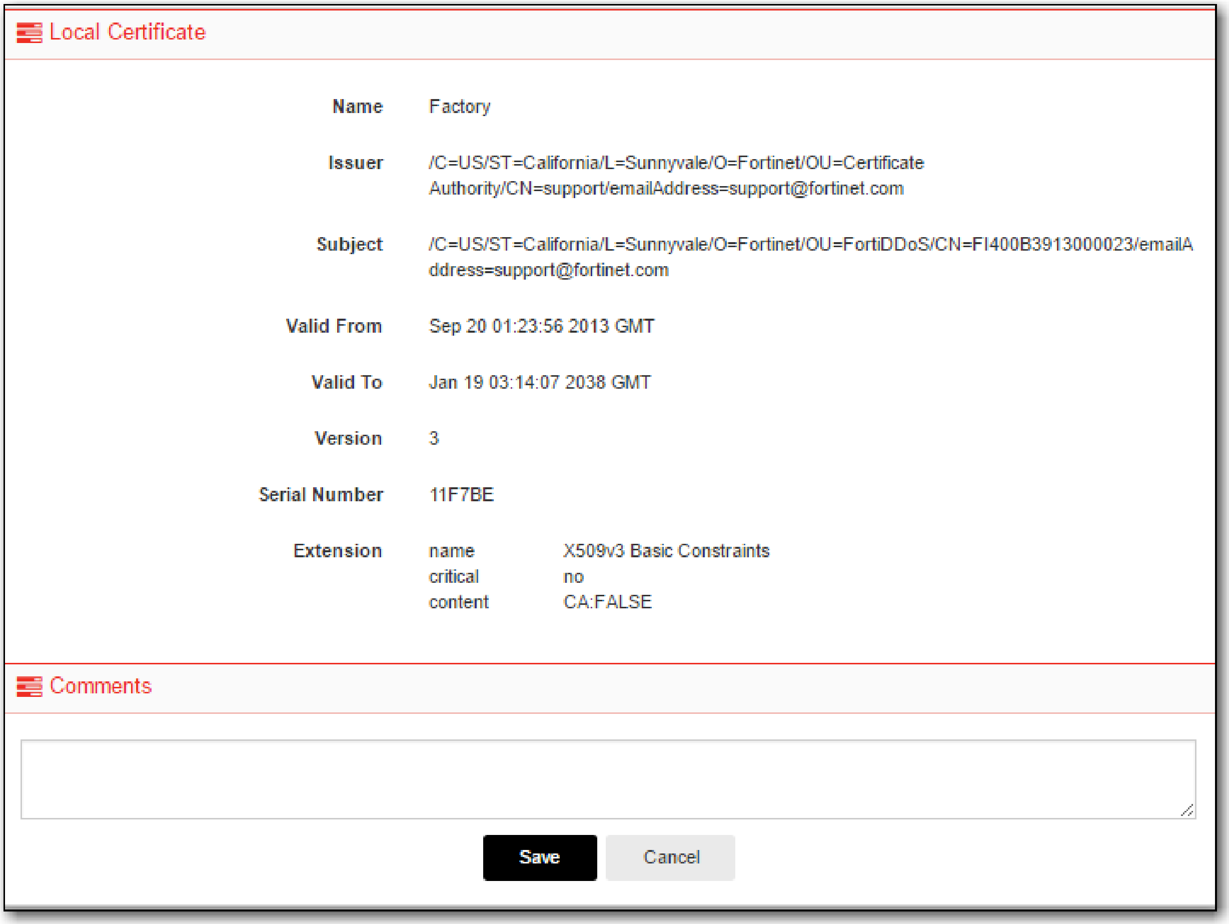

Local Certificate

Name	Factory	
Issuer	/C=US/ST=California/L=Sunnyvale/O=Fortinet/OU=Certificate Authority/CN=support/emailAddress=support@fortinet.com	
Subject	/C=US/ST=California/L=Sunnyvale/O=Fortinet/OU=FortiDDoS/CN=FI400B3913000023/emailAddress=support@fortinet.com	
Valid From	Sep 20 01:23:56 2013 GMT	
Valid To	Jan 19 03:14:07 2038 GMT	
Version	3	
Serial Number	11F7BE	
Extension	name critical content	X509v3 Basic Constraints no CA:FALSE

Comments

Backing up and restoring the configuration

You use the backup procedure to save a copy of the configuration. You can create a backup of a specific SPP configuration or the whole system configuration (including all SPPs). The backup file created by the web UI is a text file with the following naming convention: FDD-<serialnumber>-<YYYY-MM-DD>[-SPP<No>]. If you use the CLI to create a backup, you specify the filename.

The backup feature has a few basic uses:

- Restoring the system to a known functional configuration.
- Creating an SPP template configuration that you can edit and then import. You must carefully edit the SPP name and ID to avoid issues, and the SPP must exist on the running system before you can import a configuration for it. For example, if you want to import a configuration with the name SPP-2, and ID 2, you must first create an SPP-2 configuration (name and ID) on the running system.
- Saving the configuration as CLI commands that a co-worker or Fortinet support can use to help you resolve issues with misconfiguration.

Note: When you restore an SPP configuration, the SPP traffic statistics and counters are reset.

Before you begin:

- If you are restoring a system configuration, you must know its management interface configuration in order to access the web UI after the restore procedure is completed. Open the configuration file and make note of the IP address and network requirements for the management interface. You also must know the administrator username and password.
- You must have Read-Write permission for System settings.

To backup or restore the system configuration:

1. Go to System > Maintenance > Backup & Restore.
2. Complete the actions described in Table 81.

Figure 146: Backup and restore configuration page

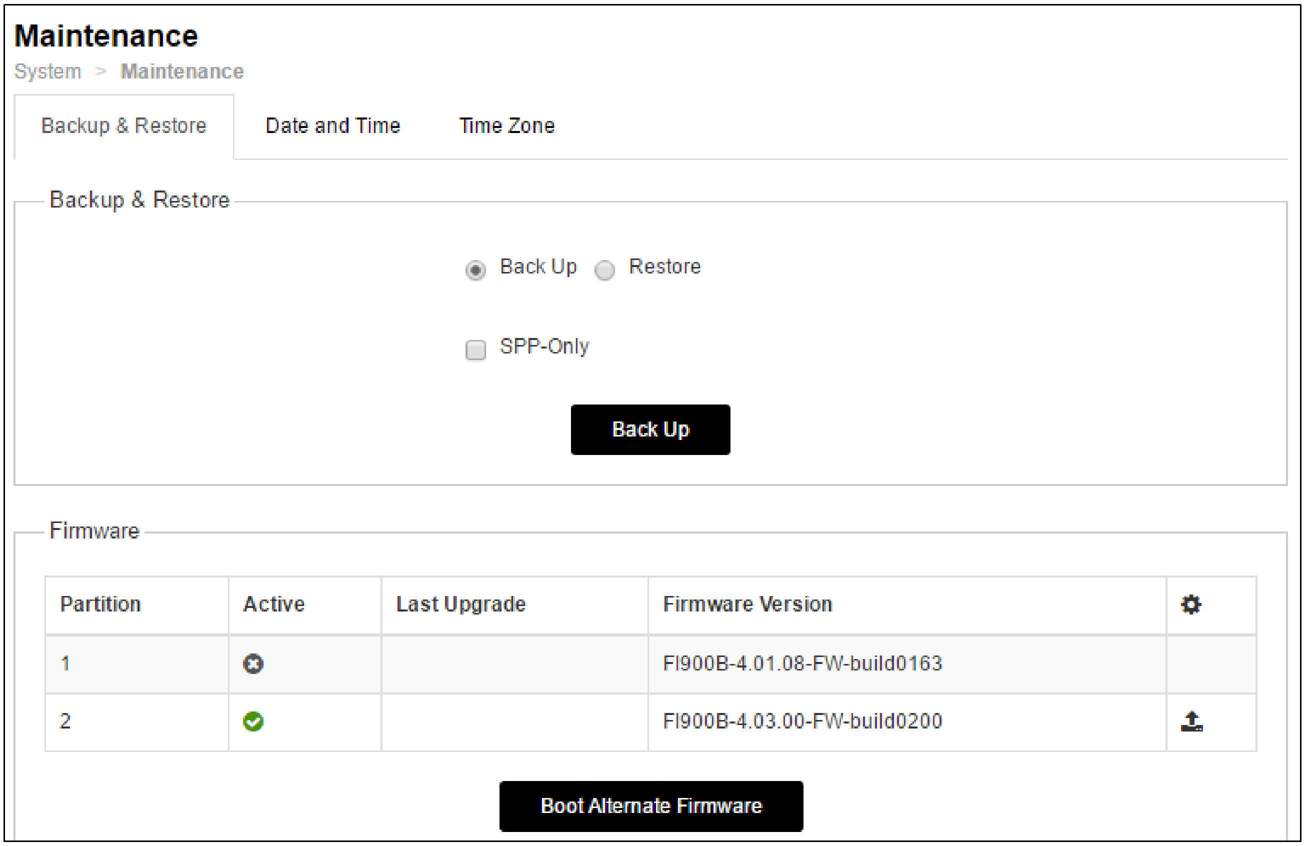

● Back Up ○ Restore

□ SPP-Only

Partition	Active	Last Upgrade	Firmware Version	⚙
1	☒		FI900B-4.01.08-FW-build0163	
2	✓		FI900B-4.03.00-FW-build0200	⇪

Table 81: Backup and restore configuration guidelines

Actions	Guidelines
Backup	
SPP-Only	To create a backup of a single SPP configuration, select this option and then select the SPP. If this option is not selected, the system creates a backup of the complete configuration.
Backup (button)	Click the **Backup** button to start the backup.
Restore	
SPP-Only	To restore the configuration for a single SPP configuration, select this option and then select the SPP. If this option is not selected, the system processes the update as a complete restore.
From File	Type the path and backup file name or click **Browse** to locate the file.

Actions	Guidelines
Restore (button)	Click the **Restore** button to start the restore procedure. Your web browser uploads the configuration file and the system reboots with the new configuration. The time required to restore varies by the size of the file and the speed of your network connection. Your web UI session is terminated when the system reboots. To continue using the web UI, refresh the web page and log in again. If the restored system has a different management interface configuration than the previous configuration, you must access the web UI using the new management interface IP address. **WARNING**: Restoring a configuration (full system) results in a system REBOOT which can interrupt traffic if your traffic links do not have fail-open capability.

To back up the configuration using the CLI to a TFTP server:

1. If necessary, start your TFTP server.
2. Log into the CLI as the `admin` administrator using either the local console, the **CLI Console** widget in the web UI, or an SSH or Telnet connection.

   Other administrator accounts do not have the required permissions.
3. Use the following command:

```
execute backup config tftp <filename> <ipaddress> [spp_name]
```

`<filename>`	Name of the file to be used for the backup file, such as `Backup.conf`.
`<ipaddress>`	IP address of the TFTP server.
`[spp_name]`	Optional. SPP configuration name, for example, SPP-0 or SPP-1. Use this option to back up only the SPP configuration. If you do not specify this option, a backup is created for the complete system configuration.

The following command creates a backup of the SPP-1 configuration:

```
exec backup config tftp Backup-SPP-1.conf 192.0.2.1 SPP-1
```

To restore a configuration:

```
execute restore config tftp <filename> <ipaddress> [spp_name]
```

`<filename>`	Name of the file, such as `Backup.conf`.
`<ipaddress>`	IP address of the TFTP server.
`[spp_name]`	Optional. SPP configuration name, for example, SPP-0 or SPP-1. Use this option to restore only the SPP configuration. If you do not specify this option, the imported file is regarded as a complete system configuration.

For example:

```
execute restore config tftp Backup-SPP-1.conf 192.0.2.1 SPP-1
```

⚠ TFTP is not secure, and it does not support authentication. You should run it only on trusted administrator-only networks, and never on computers directly connected to the Internet. Turn off tftpd off immediately after completing this procedure.

Updating firmware

This topic includes the following information:

- Upgrade considerations
- Updating firmware using the web UI
- Updating firmware using the CLI
- Downgrading firmware

Upgrade considerations

The following considerations help you determine whether to follow a standard or non-standard upgrade procedure:

- HA—Updating firmware on an HA cluster requires some additions to the usual steps for a standalone appliance. See Updating firmware on an HA cluster
- Downgrades—Special guidelines apply when you downgrade firmware to an earlier version. See Downgrading firmware. In some cases, the downgrade path requires reimaging. Take care to study the release notes for each version in your downgrade path.
- Re-imaging—If you are installing a firmware version that requires a different size of system partition, you might be required to re-image the boot device.

Important: Read the release notes for release-specific upgrade considerations.

Updating firmware using the web UI

Figure 147 shows the user interface for managing firmware. Firmware can be loaded on two disk partitions. You can use the web UI to boot the firmware version stored on the alternate partition or to upload and boot firmware updates (either upgrades or downgrades).

Figure 147: Firmware update page

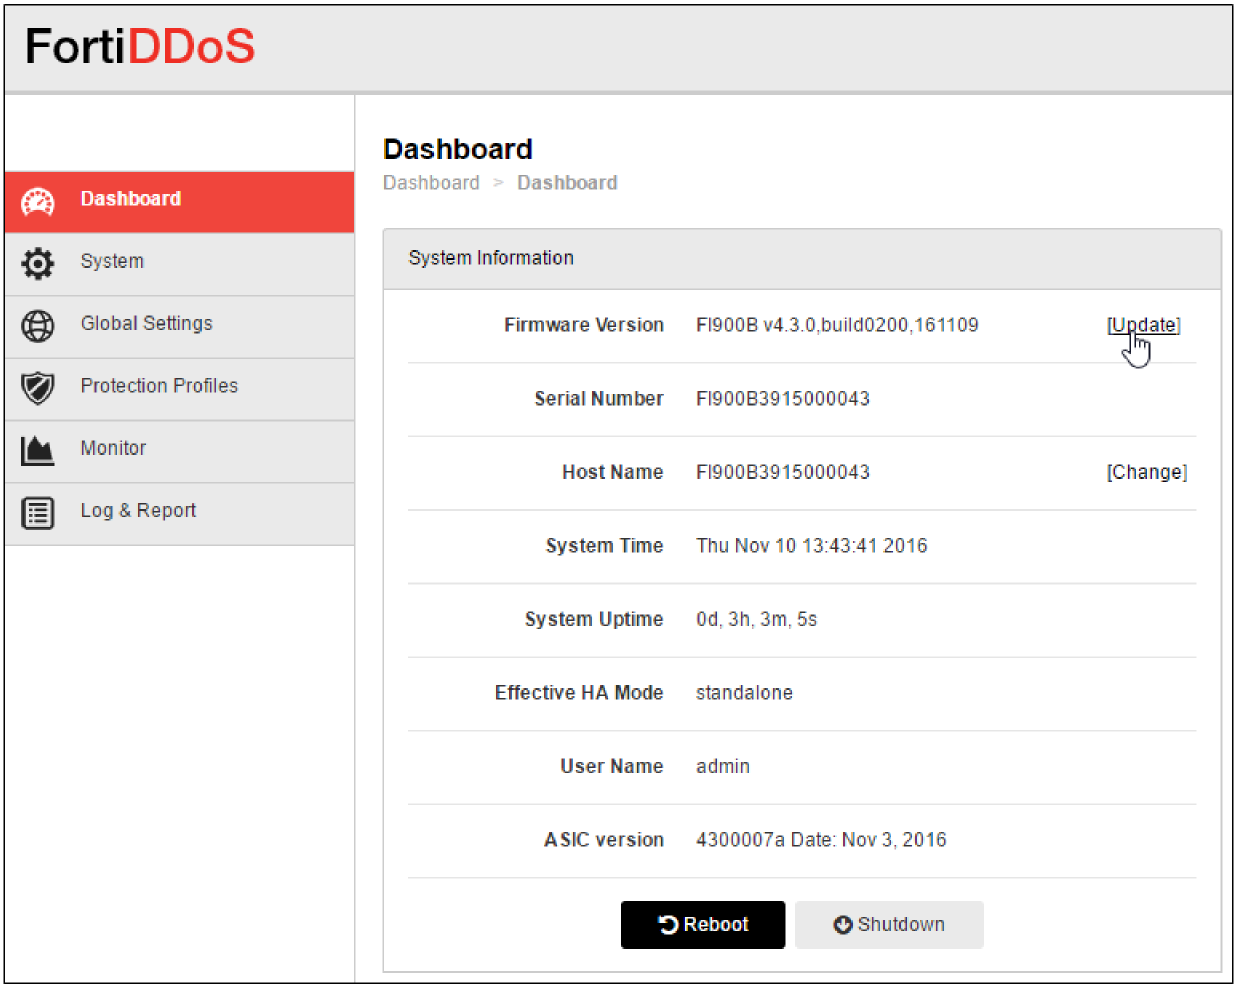

Firmware Version	FI900B v4.3.0,build0200,161109	[Update]
Serial Number	FI900B3915000043	
Host Name	FI900B3915000043	[Change]
System Time	Thu Nov 10 13:43:41 2016	
System Uptime	0d, 3h, 3m, 5s	
Effective HA Mode	standalone	
User Name	admin	
ASIC version	4300007a Date: Nov 3, 2016	

Before you begin:

- Download the firmware file from the Fortinet Technical Support website.
- Read the release notes for the version you plan to install.
- **Important**: Back up your configuration before beginning this procedure. If you revert to an earlier firmware version, the running configuration is erased, and you must restore a saved configuration. We recommend you restore a configuration you knew to be working effectively on the firmware version you revert to. Some 4.2 settings are incompatible with 4.1.x, so we recommend you not restore a 4.2 configuration to a 4.1.x system.
- Make a note of configurations that are disabled in your active configuration. Configurations that are not enabled are not preserved in the upgrade. For example, if a custom HTTP service port, log remote port, or event log port have been configured and then disabled in 4.1.11, the port information is not preserved in the upgrade to 4.2.1.
- You must have super user permission (user **admin)** to upgrade firmware.

To install firmware:

1. Go to System > Maintenance > Backup & Restore tab.
2. Under Firmware Upgrade/Downgrade, use the controls to select the firmware file that you want to install and click Update and Reboot icon.

Clear the cache of your web browser and restart it to ensure that it reloads the web UI.

Updating firmware using the CLI

This procedure is provided for CLI users.

Before you begin:

- Read the release notes for the version you plan to install. If information in the release notes is different from this documentation, follow the instructions in the release notes.
- You must be able to use TFTP to transfer the firmware file to the FortiDDoS system. If you do not have a TFTP server, download and install one, like `tftpd`, on a server located on the same subnet as the FortiDDoS system.
- Download the firmware file from the Fortinet Technical Support website.
- Copy the firmware image file to the root directory of the TFTP server.
- Back up your configuration before beginning this procedure. Reverting to an earlier firmware version could reset settings that are not compatible with the new firmware.
- Make a note of configurations that are disabled in your active configuration. Configurations that are not enabled are not preserved in the upgrade. For example, if a custom HTTP service port, log remote port, or event log port have been configured and then disabled in 4.1.11, the port information is not preserved in the upgrade to 4.2.1.
- You must have super user permission (user **admin)** to upgrade firmware.

To install firmware via the CLI:

1. Connect your management computer to the FortiDDoS console port using an RJ-45-to-DB-9 serial cable or a null-modem cable.
2. Initiate a connection to the CLI and log in as the user **admin**.
3. Use an Ethernet cable to connect FortiDDoS port1 to the TFTP server directly, or connect it to the same subnet as the TFTP server.
4. If necessary, start the TFTP server.
5. Enter the following command to transfer the firmware image to the FortiDDoS system:

```
execute restore image tftp <filename_str> <tftp_ipv4>
```

where <filename_str> is the name of the firmware image file and <tftp_ipv4> is the IP address of the TFTP server. For example, if the firmware image file name is image.out and the IP address of the TFTP server is 192.168.1.168, enter:

```
execute restore image tftp image.out 192.168.1.168
```

One of the following message appears:

```
This operation will replace the current firmware version!
   Do you want to continue? (y/n)
```

or:

```
Get image from tftp server OK.
   Check image OK.
   This operation will downgrade the current firmware version!
   Do you want to continue? (y/n)
```

6. Type `y`.

The system installs the firmware and restarts:

```
MAC:00219B8F0D94
###########################
Total 28385179 bytes data downloaded.
Verifying the integrity of the firmware image.
Save as Default firmware/Backup firmware/Run image without saving:[D/B/R]?
```

7. To verify that the firmware was successfully installed, use the following command:

   ```
   get system status
   ```

The firmware version number is displayed.

If the download fails after the integrity check with the error message `invalid compressed format (err=1`, but the firmware matches the integrity checksum on the Fortinet Technical Support website, try a different TFTP server.

TFTP is not secure, and it does not support authentication. You should run it only on trusted administrator-only networks, and never on computers directly connected to the Internet. Turn off tftpd off immediately after completing this procedure.

Downgrading firmware

You can use the web UI or CLI to downgrade to a previous software image. The commands are the same as for upgrading. However, special guidelines apply:

- When the image version is different from than the existing version, the initializing code also updates the TP2 ASIC image to match the correct version in the image. In other words, if you upgrade or downgrade software, the procedure also upgrades or downgrades the TP2 ASIC image.
- Always keep a back up of the configuration before you change the software image (upgrade or downgrade). For example, before you upgrade to 4.2.1, save a copy of the 4.1.11 configuration and label it 4.1.11.
- Upgrades preserve the running configuration, but downgrades do not. When you downgrade, the running configuration is erased, including the management IP address.
- You must use a console port connection to reconfigure the management interface.
- After you have configured the management interface, you can restore the earlier configuration. We recommend you restore a configuration you knew to be working effectively on the firmware version you revert to. Some 4.2 settings are incompatible with 4.1.x, so we recommend you not restore a 4.2 configuration to a 4.1.x system. Instead, if you downgrade from 4.2.1 to 4.1.11, we recommend you restore the 4.1.11 configuration.
- After the configuration is restored, the system reboots, and the restored configuration will be in effect.

Configuring system time

Accurate system time is critical to the correct FortiDDoS operation including all graphs, logs, and scheduling.

Changing the time of a system that is operational may have extreme consequences for the data already collected by the system and the system's ability to detect and mitigate attacks. All saved traffic graphs, drop data and logs may be lost along with traffic statistics and system Thresholds.

We strongly recommend that you use Network Time Protocol (NTP) to maintain the system time and that you configure date and time NTP settings before you do any other configurations.

As an alternative when NTP is not available or is impractical, you can set the system time manually, but this time will drift and changing times later can have serious consequences on existing data and mitigation.

You can change the system time with the web UI or the CLI.

Before you begin:

- You must have Read-Write permission for System settings.
- Ensure you have:
  - DNS settings set in System > DNS.
  - Gateway settings so that NTP queries can exit the network: System > Static Route.
  - Verify that your firewalls or routers do not block or proxy UDP port 123 (NTP).
- We recommend that — before you change system time settings —If the system is new or has already been factory reset, proceed with the instructions below.
- If the system has saved Traffic Statistics for SPPs and/or graphs and logs from existing traffic:
  - If you need to set the time ahead (from 8 am to 9 am, for example):
    - You may proceed with the instructions below.
  - If you need to set the time back (from 9 am to 8 am, for example):
    - You must perform command execute `formatlogdisk` to clear data that already exists within the time period that will be overwritten. This removes ALL graph, drop and log data from the system.

To configure the system time:

1. Navigate to the system time settings page in one of the following ways:
   - Go to System > Maintenance > Time Zone.
2. Complete the configuration as described in Table 82.

   Important: You can change settings for only one group at a time: Time Zone or Time Setting. You must save your changes after each group before making changes in the next.
3. Save your changes. The system will reboot.

Note: If the time or time zone is changed, you need to reset the system to its factory state using the command `# execute formatlogdisk` so that the graphs are updated accordingly.

4. Change tabs to Date and Time.
5. Complete NTP or manual time settings.
6. Save your changes. The system will reboot.
7. Verify your changes:

- Success — If you manually configured the system time, or if you enabled NTP and the NTP query for the current time succeeds, the new clock time appears in the System Time field at the top of the page shown in Figure 147. If the NTP query reply is slow, you might need to wait a couple of seconds, and then click **Refresh** to update the time displayed in the System Time field.
- Failure — If the NTP query fails, the system clock continues without adjustment. For example, if the system time had been 3 hours late, the system time is still 3 hours late. To troubleshoot the issue, check settings for your DNS server IP addresses, your NTP server IP address or name, and routing addresses; verify that your firewalls or routers do not block or proxy UDP port 123.

Figure 148: Time settings

Table 82: System time configuration

Setting	Guidelines
Time Zone	
Time zone	1. Ensure the 'Daylight Saving Time' checkbox is unchecked. 2. Select the time zone where the appliance is located. Check that the GMT offset is correct for the location. In recent years, Time Zone changes have been frequent. If the City/Country-Time Zone pair is inaccurate, use the correct GMT offset and ignore the text information.
Automatically adjust clock for daylight saving changes	Enable if you want the system to adjust its own clock when its time zone changes between daylight saving time (DST) and standard time. When enabled, you will see that the Time Zone GMT offset immediately changes, no matter if you are in DST or not. This is for display only and will not affect the system time.
Synchronize with NTP Server	
Server	Specify the IP address or domain name of an NTP server or pool, such as `pool.ntp.org`. To find an NTP server, go to http://www.ntp.org. You may enter more than one IP address or domain name with a space between them.
Sync Interval	Specify how often the system should synchronize its time with the NTP server, in minutes. For example, to synchronize once a day, type 1440.
OR	
Set Time	
Hour, Minute, Second, Date	This is not required if you have set NTP. Use the controls to set the time manually. The clock is initialized with the time you specified when you click **Save**. **NOTE: Manual time setting is NOT recommended.**

To configure NTP using the CLI:

```
config system time ntp
   set ntpsync enable
   set ntpserver {<server_fqdn> | <server_ipv4>}
   set syncinterval <minutes_int>
end
```

To configure the system time manually:

```
config system time ntp
   set ntpsync disable
end
config system time manual
   set zone <timezone_index>
   set daylight-saving-time {enable | disable}
end
execute date <mm:dd:yyyy> <hh:mm:ss>
```

Configuring the hostname

You can configure a hostname to facilitate system management. If you use SNMP, for example, the SNMP system name is derived from the configured hostname.

Before you begin:

- You must have Read-Write permission for System settings.

To configure the hostname:

1. Go to Dashboard.
2. Find the Host Name in the System Information table and click **Change**.
3. Complete the configuration as described in Table 83.
4. Save the configuration.

Table 83: Hostname configuration

Settings	Guidelines
New Name	The hostname can be up to 35 characters in length. It can include US-ASCII letters, numbers, hyphens, and underscores, but not spaces and special characters. The System Information widget and the `get system status` CLI command display the full hostname. If the hostname is longer than 16 characters, the name is truncated and ends with a tilde (~) to indicate that additional characters exist, but are not displayed.

CLI commands:

```
config system global
   set hostname <name>
end
```

Rebooting, shutting down, and resetting the system

This section includes the following information:

- Rebooting the system
- Shutting down the system
- Resetting the system

Rebooting the system

To reboot the operating system:

1. Go to Dashboard.
2. In the System Information widget, click the **Reboot** button.

CLI commands:

```
execute reboot
```

Shutting down the system

Always properly shut down the system before turning off the power switch or unplugging it. This causes it to finish writing any buffered data, and to correctly spin down and park the hard disks.

Do not unplug or switch off the FortiDDoS appliance without first halting the operating system. Failure to do so could cause data loss and hardware problems.

To power off the system:

1. Go to Dashboard.
2. In the System Information widget, click the **Shutdown** button.

The system does not emit disk activity noise when shutdown is complete.

CLI commands:

```
execute shutdown
```

To completely power off:

- Press the power button if there is one.

Power supplies and switches vary by hardware model. On some, you press the power button; on others, you flip the switch to either the off (O) or on (I) position.

Resetting the system

Table 84 summarizes “factory reset” options.

Table 84: “Factory reset” options

Task	Menu
Reset the threshold configuration for an SPP.	Protection Profiles > Thresholds > Factory Defaults
Reset the threshold configuration and clear traffic history for an SPP.	Protection Profiles > Factory Reset > Factory Reset
Reset the system to its factory state. All SPPs, statistics, and logs will be deleted.	See below.

To reset the system to its factory state:

Use the command based on the following:

- `# execute factoryreset`: Deletes all the configuration without deleting any data.
- `# execute formatlogdisk`: Saves the configuration, but deletes all the data, including mysql database (attack log, event log) and RRDs (graphs).
- `# execute factoryreset`: Deletes the attack log database.

Chapter 9: Deployment Topologies

This chapter provides guidelines for basic and advanced deployments. It includes the following sections:

Basic enterprise deployment

Figure 149 shows a basic deployment. The FortiDDoS appliance is positioned ‘inline’, meaning it is installed between the Internet and the protected network.

FortiDDoS is stateful and bidirectional. The data packet traffic is described as either incoming (inbound) and outgoing (outbound).

Figure 149: Basic topology

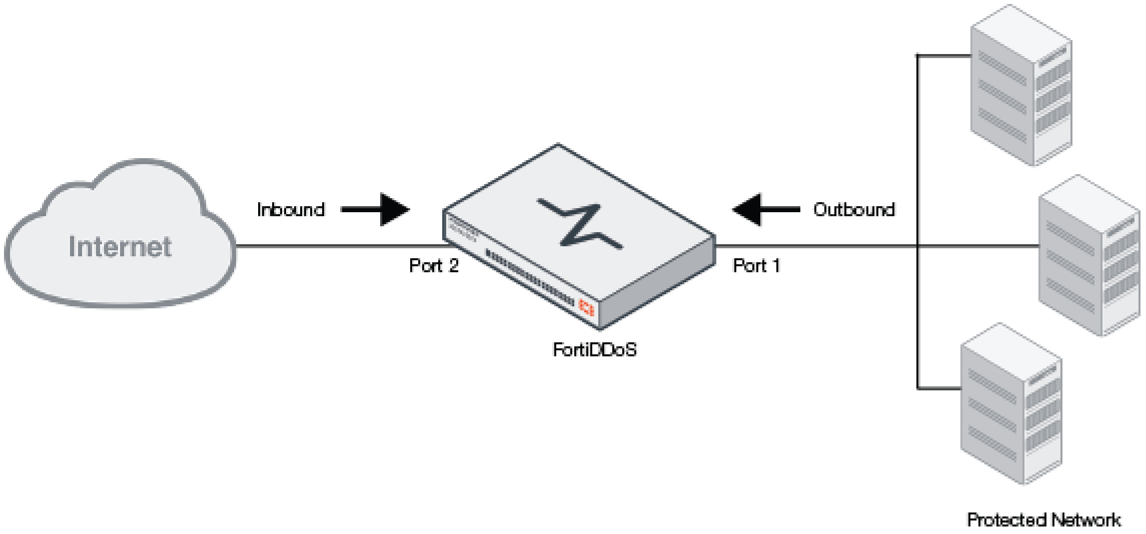

For networks with multiple servers, you can provision a port pair to each server. For example, to protect 8 servers with connections with a total throughput of up to 12 Gbps, connect each server to a port pair on a FortiDDoS 800B.

Basic multi-tenant deployment

Figure 150 shows a basic multi-tenant deployment. A web hosting company leases FortiDDoS services to its customers. You can provision individual SPPs for up to seven customers.

Figure 150: Basic web hosting deployment of FortiDDoS appliances

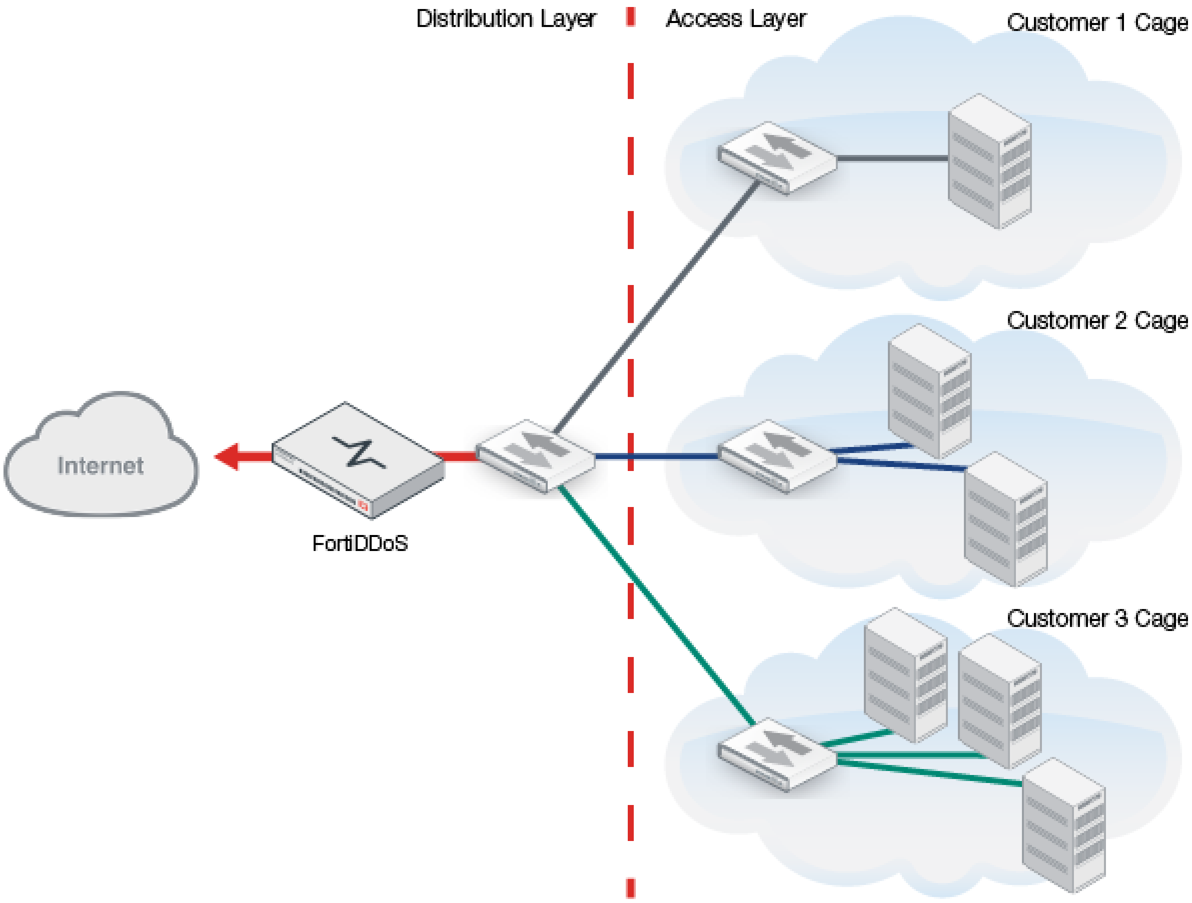

Built-in bypass

The following FortiDDoS network interface connections have a built-in bypass mechanism:

- Any copper (RJ-45) network connections (for example, the RJ-45 connections for ports 1-16 on FortiDDoS 400B or 600B/800B)
- Ports 17-20 on the FortiDDoS 1200B/2000B, which are fixed LC connectors

This automatic bypass functionality is not available for the other fiber-optic connections on the FortiDDoS 1200B/2000B (ports 1-17) or for any of the fiber-optic connections found on other models.

Bypass is activated under the following conditions:

- The appliance is not powered up or is starting up or rebooting
- The appliance's FortiASIC processor or integrated switch fabric fail

You can use the Global Settings > Settings page to configure the internal bypass mechanism to fail open or fail closed.

By default, the interfaces are configured to fail open. This means that interfaces pass traffic through without performing any monitoring or prevention tasks. Packets that arrive at ingress ports are simply transferred to the corresponding egress ports, just like a wire.

If you use an external bypass solution, configure the interfaces to fail closed. This means traffic is not forwarded through the interfaces. An external bypass system can detect the outage and forward traffic around the FortiDDoS.

If you deploy an active-passive cluster, configure the interfaces on the primary node to fail closed so the adjacent switches can select the secondary node. The secondary unit can be set to fail closed or fail open, depending on how you want to handle the situation if both FortiDDoS nodes are down.

Table 85 summarizes bypass behavior for a sequence of system states. During boot up, daemons and drivers are started. When boot up is complete and all memory tables are clean, the TP2-ASIC is ready for packet processing, and the appliance exits the bypass state. Traffic is routed through the TP2-ASIC, it is monitored, and policies enforced. In the event of failure, manual reboot, or graceful shutdown, system services are unavailable because they are either being restarted or shut down, and the appliance enters the bypass state.

Table 85: System state and bypass

User Option	State 1 Power Off	State 2 Just Powered Up	State 3 Boot Up Process	State 4 System Ready	State 5 Failure, Reboot, or Graceful Shutdown	State 6 Power Off
Fail Open	Bypass	Bypass	Bypass	Bypass off	Bypass	Bypass
Fail Closed	Closed	Closed	Closed	Bypass off	Closed	Closed

In addition to the automatic bypass settings, FortiDDoS 200B, 400B, and 600B/800B support manual bypass (for copper ports) with the following CLI command:

```
execute bypass-traffic {enable|disable}
```

This command forces the appliance interfaces to fail open. This command does not have an option to fail closed.

Note that if you use the CLI command to initiate bypass, you must use the CLI command to disable that state.

After you have executed this command, go to the System Dashboard to confirm the bypass state for the interfaces. If not all of the interfaces have gone to bypass state or returned from bypass state, execute the command a second time.

External bypass

FortiDDoS can be deployed with an external bypass mechanism, such as a bypass switch. When both the FortiDDoS appliance and the failover switch share the same power supply, external connectivity is maintained during a power failure.

Figure 151 shows a bypass deployment when bypass is not active. The inline traffic flows through the FortiDDoS appliance.

Figure 151: Bypass ready but not active

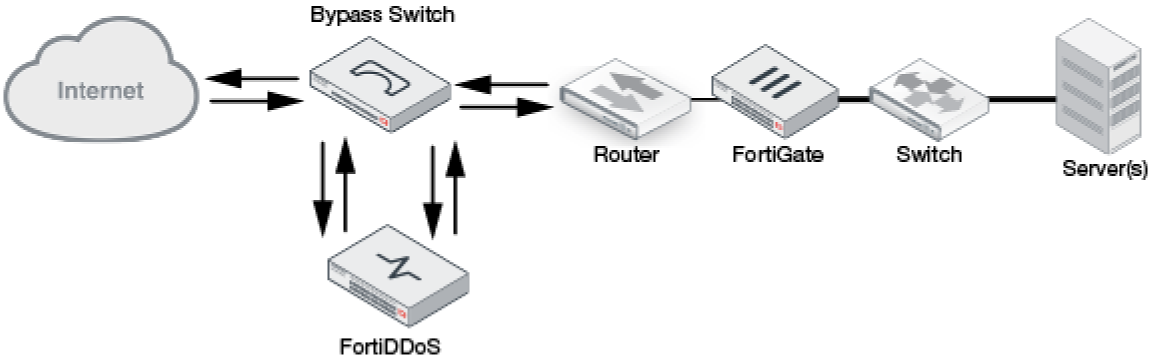

Figure 152 shows a bypass deployment when bypass is active. All inline traffic is routed through the switch until FortiDDoS is back online.

Figure 152: Active bypass

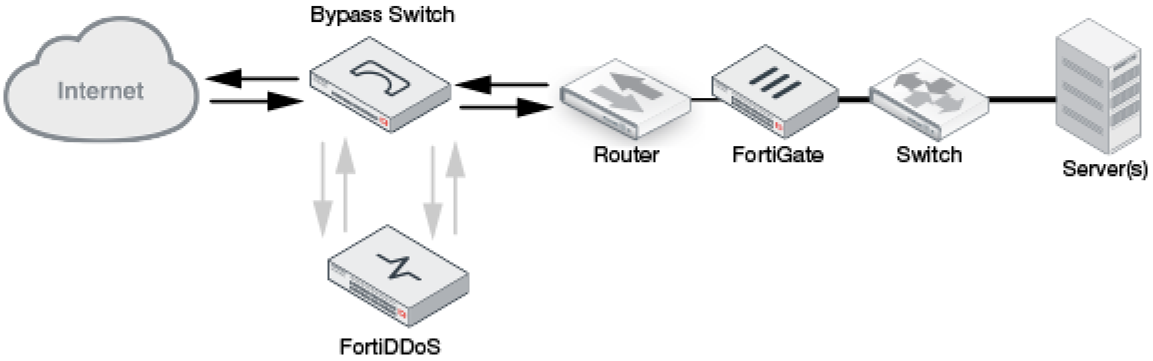

Either the automatic bypass mechanism or a bypass switch can maintain data traffic when there is a power or appliance failure. However, it is recommended that you automate failover behavior using a bypass switch with heartbeat. A bypass switch with heartbeat detects the failure of the FortiDDoS appliance (and the failure of traffic monitoring and mitigation) even when the appliance maintains the copper-based data link.

Using an optical bypass switch

Fortinet recommends using a FortiBridge device as your optical bypass switch. Contact your Sales Engineer to see if other bypass switches can support your FortiDDoS deployment.

Figure 153 shows a deployment with an optical bypass switch that monitors the link to the attached FortiDDoS appliance by sending a heartbeat packet to the appliance once every second. If the optical bypass switch does not receive the heartbeat back, it automatically switches network traffic to bypass the unresponsive FortiDDoS appliance, even if the appliance is still receiving power. The optical bypass continues to send the heartbeat and restores the traffic through the FortiDDoS appliance as soon as the link is restored.

Figure 153: Optical bypass device

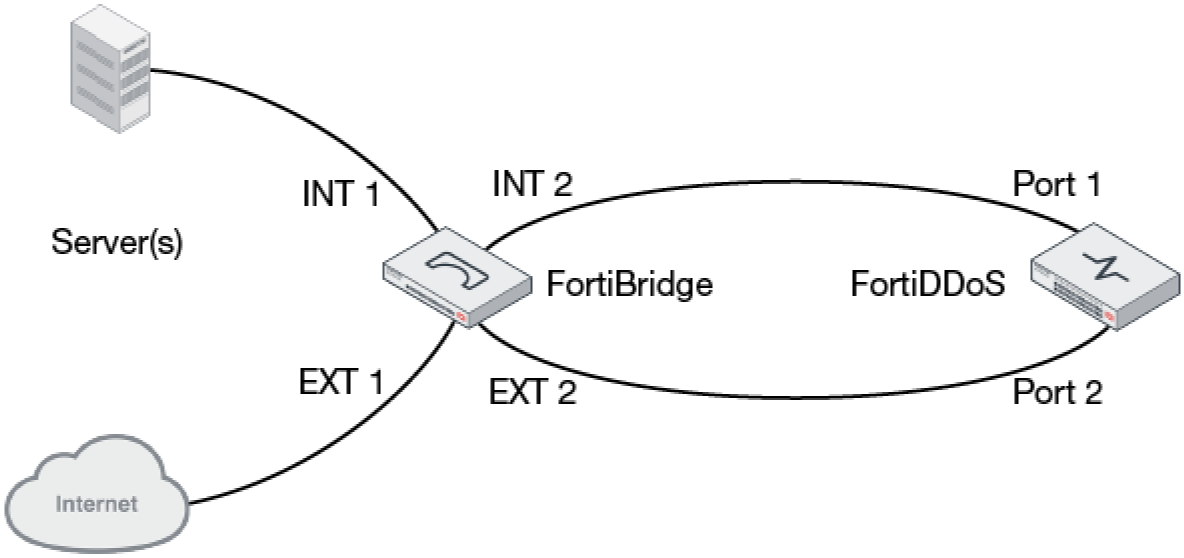

Configuring the optical bypass switch

Refer to the *FortiBridge QuickStart Guide* and *FortiGate Hardware Guide* to set the following parameters:

- Input timeout period
- Input retry count

Connecting the optical bypass switch to the network and FortiDDoS

1. Connect the INT 1 port to the Ethernet segment.
2. Connect the EXT 1 port to the Internet side.
3. Connect the INT 2 port to the FortiDDoS server port (for example, Port 1).
4. Connect the EXT 2 port to the FortiDDoS Internet port (for example, Port 2).

Configuring MAC addresses for bypass switch heartbeat packets

When a FortiDDoS appliance is used in conjunction with a bypass switch such as FortiBridge, ensure that FortiDDoS allows heartbeat packets from the bypass switch in all possible cases.

Typical bypass switches use heartbeat packets to check if the data path is connected. If the data path is broken for some critical reason, the bypass switch switches to bypass mode from normal mode.

To ensure that it passes on the heartbeat packets, FortiDDoS allows you to specify the MAC addresses that the bypass switch uses for the packets.

You can view these MAC addresses in the FortiBridge status page.

Every FortiDDoS link pair can be connected via a FortiBridge link pair. For example, you can use a FortiBridge link to bridge the Port 1/Port 2 link pair and another FortiBridge link to bridge the Port 3/Port 4 link pair. Each link pair is associated with a pair of MAC addresses. Therefore, if you are using two links, you configure four MAC addresses. If you are using one link, specify two MAC addresses.

You can program up to 16 MAC addresses.

If the bypass switches are from the same vendor, the most significant 24-bits of their MAC addresses are the same.

To configure bypass MAC addresses:

1. Go to Global Settings > Bypass MAC > Bypass MAC.
2. Click **Add**, and then enter a name for the MAC address and the address.
3. Save the configuration.

Tap Mode deployments

This section provides the following information about FortiDDoS Tap Mode deployments:

Overview

The FortiDDoS appliance is a transparent Layer 2 bridge that could become a point-of-failure without proper bypass mechanisms. It is common to deploy a Layer 2 bridge in-path and the FortiDDoS appliance in an out-of-path segment so that you are never faced with outages due to failure, maintenance, or replacement of a FortiDDoS appliance.

FortiBridge appliances support inline, bypass, and recovery features. FortiBridge 3000 series appliances also support Tap Mode—a mode in which the Layer 2 bridge can simultaneously perform bypass through its network ports and mirroring through its monitor ports.

FortiDDoS appliances have a complementary Tap Mode setting that turns off the transmit (Tx) component of the FortiDDoS network interface cards. This ensures the FortiDDoS is a passive listener that cannot disrupt traffic or cause an outage.

In a Tap Mode deployment, FortiDDoS can use the mirrored packets to build the traffic history it uses to establish rate thresholds, and it can detect volumetric attacks (rate anomalies), but it does not take actions, like dropping traffic, blocking identified source attackers, or aggressively aging connections.

When an attack is detected, you can turn off Tap Mode on FortiDDoS and the FortiDDoS interfaces resume packet transmission. FortiBridge probes will then pass through FortiDDoS successfully, FortiBridge will detect that the out-of-path segment is available, and it will switch to Inline Mode.

Deployment Topology

Figure 154 illustrates how FortiBridge deployment modes are used in a deployment with FortiDDoS. FortiBridge is deployed in-path and FortiDDoS is deployed out-of-path.

In Inline Mode, FortiBridge passes heartbeat packets through its monitor ports to detect whether the out-of-path segment is available. When the health probes indicate the path is available, inbound traffic that is received by the FortiBridge Net0 interface is forwarded through the Mon0 interface to the FortiDDoS WAN port. FortiDDoS processes the traffic, takes action on attacks and passes non-attack traffic through its LAN port to the FortiBridge Mon1 interface. The traffic is passed through the FortiBridge Net1 interface towards its destination.

Figure 154: Inline Mode

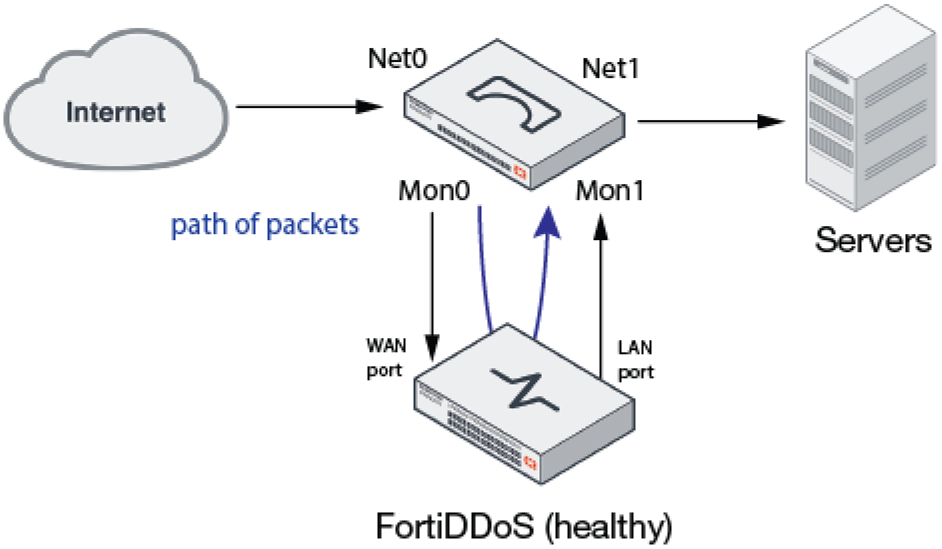

If the heartbeat probe fails due to FortiDDoS failure or maintenance, FortiBridge can be set up to switch from Inline mode to Bypass mode. In Bypass Mode, traffic is not forwarded through the monitor ports. Instead, it is forwarded from Net0 to Net1, bypassing the out-of-path segment.

Figure 155: Bypass Mode

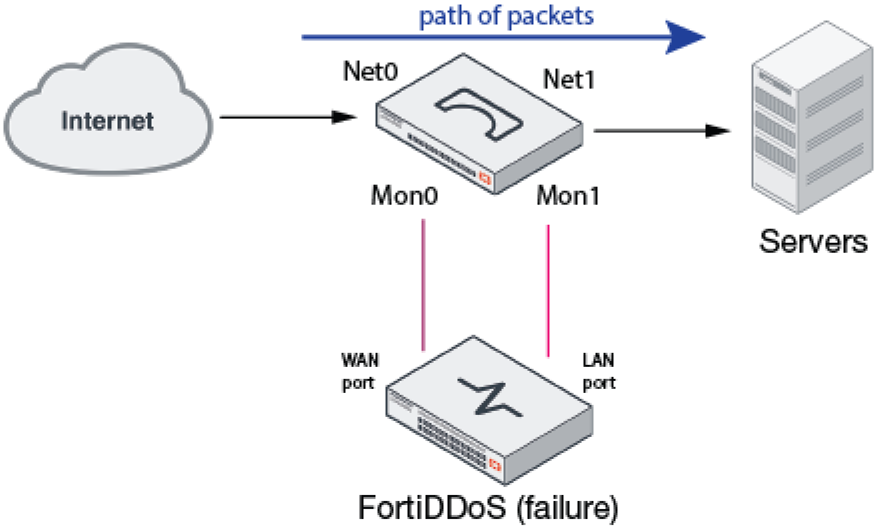

Alternatively, you can set up FortiBridge to switch from Inline Mode to Tap Mode when probes fail. In FortiBridge Tap Mode, traffic is forwarded from Net0 to Net1, and it is also mirrored to Mon0. This is what you want when you want to deploy FortiDDoS as a passive listener.

Figure 156: Tap Mode

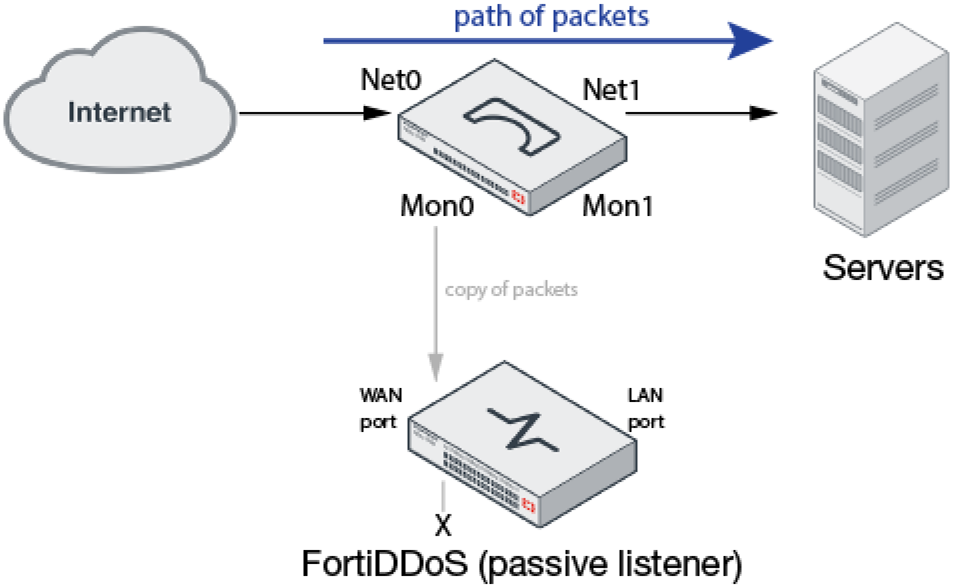

Although not shown in the illustrations, the reverse paths are processed the same way.

Note: When in Tap Mode, FortiDDoS discards packets after processing (noted by an X in Figure 156). You should not expect to see egress traffic on the Monitor > Port Statistics graphs.

Requirements

This solution has been verified with FortiBridge 3000 Series appliances running OS version 4.2.x or later. Contact your Fortinet sales engineer to learn more about FortiBridge appliances.

Fortinet does not support Tap Mode deployments with other bridge or tap devices. If you attempt a deployment with other devices, consider the following Tap Mode requirements:

- The bridge device must be deployed and configured to forward traffic along the data path and send mirrored traffic towards FortiDDoS on both its monitor ports (inbound traffic on one port and outbound on the other).
- The bridge must block any transmit packets from FortiDDoS on its monitor ports so that any traffic sent by FortiDDoS is blocked.
- The bridge device should have the ability to set inline/bypass/tap mode manually so that administrators take direct action when there is an attack.
- FortiDDoS passes heartbeat packets from its ingress to egress ports, so the bridge must not be affected by seeing these heartbeat packets (it will not switch to inline mode).
- Passive optical TAPs will generally not work since the TAPs usually have a single duplex monitor port output on 1 pair of fiber ports. FortiDDoS requires 2 separate monitor ports for inbound and outbound traffic on 2 separate fiber pairs. Custom cabling can support this, but FortiDDoS can never be switched inline using passive TAPs.

Limitations

In Tap Mode, FortiDDoS is a passive listener. It records actions it would have taken were it placed inline, so ACL, anomaly, rate threshold drops, source blocking, and aggressive aging events and statistics are just simulations.

However, some features cannot be simulated when FortiDDoS is a passive listener. The following Prevention Mode features depend on being deployed in-path and interacting with clients and servers to work correctly:

- SYN flood mitigation—With SYN validation enabled, FortiDDoS performs antispoofing tests to determine whether the source is legitimate. In Tap Mode, if the source was not already in the legitimate IP table, it will fail the test. As a result, the simulation is skewed, and the reports will show an inordinate spike in blocked sources.
- TCP state anomaly detection—With Foreign Packet Validation enabled, FortiDDoS drops unexpected packets (for example, if there is a sequence of events in which FortiDDoS drops inbound packets, it does not expect to receive corresponding outbound packets, so a foreign packet drop event is triggered).
- Aggressive aging—Aggressive aging resets are not actually sent when slow connection attacks and Layer 7 floods are detected, but the connections are cleared from the TCP state table. As a result, subsequent packets for the connection are treated as foreign packets.

Talk with your Fortinet CSE to make sure you thoroughly understand your choices, which include:

- Disabling TCP session feature control when FortiDDoS is deployed in Tap Mode. (But remember to enable it if you want its protections when you FortiDDoS is deployed inline.)
- Interpreting or disregarding the logs and graphs for these anomalies.

Tap Mode is not a perfect deployment simulation, but it does enable you prepare for volumetric attacks by building traffic history without risk of disruption or outage.

Configuration

This section includes the following information:

- FortiBridge configuration guidelines
- FortiDDoS configuration guidelines

FortiBridge configuration guidelines

FortiBridge must be deployed physically in the network where it can forward traffic along the data path and send mirrored traffic to the FortiDDoS.

The following key settings are related to a deployment with FortiDDoS:

1. Configure probes. Probes are used when FortiBridge is in Inline mode. Heartbeat packets are sent between monitor interfaces, through the out-of-path device. Set the heartbeats to **bidirectional** (both Mon ports send to each other) and set the heartbeat fail to **unidirectional** so if either direction fails, the system is bypassed.
2. Set Operation Mode to Inline Mode.
3. Set action on failure to Tap Mode.
4. Set action on recovery to Inline Mode.

You can find FortiBridge documentation on the Fortinet web site:

http://docs.fortinet.com/fortibridge/admin-guides

We recommend you set up FortiBridge to Inline Mode with action on failure set to Tap Mode; and then force a failure by turning on FortiDDoS Tap Mode.

If you enable the FortiBridge operation mode called Tap Mode manually, you must turn off probes first. If you do this, there are no probes, so there is no automatic recovery to Inline Mode.

FortiDDoS configuration guidelines

This section gives pointers for FortiDDoS configuration.

Before you begin:

- Physically connect FortiDDoS to FortiBridge.

You must add the MAC addresses for the FortiBridge Monitor ports so that FortiDDoS accepts heartbeats from them. Heartbeats are used when FortiBridge is in Inline Mode.

To configure bypass MAC addresses:

1. Go to Global Settings > Bypass MAC > Bypass MAC.
2. Click **Add**, and then enter a name for the MAC address and the address.
3. Save the configuration.

To enable Tap Mode:

1. Go to Global Settings > Settings > Settings > Deployement tab.
2. Enable **Tap Mode**.
3. Save the configuration.

The system reboots when you enable/disable Tap Mode.

```
config ddos global setting
   set tap-mode {enable|disable}
end
```

Best Practices

The following best practices are recommended by Fortinet CSEs:

- Do not set FortiBridge Tap Mode manually. Set it up as the action on failure for FortiBridge Inline Mode, and then force a failure of the out-of-path segment by turning on FortiDDoS Tap Mode.
- In a FortiDDoS Tap Mode deployment, you can set SPPs in Detection Mode or Prevention Mode. Set it to whichever mode you want enabled when you toggle off Tap Mode and put FortiDDoS inline.

Load balancing

Many data center and server farm architectures require network infrastructure to protect them. However, traffic volumes on some networks can exceed the capabilities of a single link pair on a FortiDDoS appliance or even the maximum throughput of a single appliance. To increase the overall throughput, some topologies require some type of load-balancing solution using multiple link pairs or multiple FortiDDoS appliances.

The capacity of the load-balancing device must exceed the combined throughput of the multiple FortiDDoS appliances. For details on the capacity of FortiDDoS appliances, refer to the product datasheet.

The load-balancing device intercepts all traffic between the server side and the Internet side and dynamically distributes the load among the available FortiDDoS appliances, based on the device's configuration. Load balancing utilizes all the appliances concurrently, providing overall improved performance, scalability and availability.

The FortiDDoS appliance is a Layer 2 bridge and therefore does not have either a MAC address or an IP address in the data path. For transparent bridges, the load-balancing device receives a packet, makes a load-balancing decision, and forwards the packet to a FortiDDoS appliance. The FortiDDoS appliance does not perform NAT on the packets; the source and destination IP addresses are not changed.

The load-balancing device performs the following tasks:

- Balances traffic across two or more FortiDDoS appliances in your network, allowing them to work in parallel.
- Maintains state information about the traffic that flows through it and ensures that all traffic between specific IP address source and destination pairs flows through the same FortiDDoS appliance.
- Performs health checks on all paths through the FortiDDoS appliances. If any path is not operational, the load balancer maintains connectivity by diverting traffic away from that path.

You can use an external load balancer such as Linux Virtual Server (LVS), Cisco Content Switching Module (CSM), or Avaya Load Balancing Manager.

Load Balancing allows you to:

- Maximize FortiDDoS productivity
- Scale FortiDDoS performance
- Eliminate the FortiDDoS appliance as a single point of failure

Load balancing for FortiDDoS appliances requires a sandwich topology.

Sandwich topology for load balancing

Figure 157 shows a sandwich topology. In this example, load-balancing devices are deployed before and after a pair of FortiDDoS appliances. For example, two 400B appliances to support a total throughput of 12 Gbps. This same topology and throughput is possible using a single 800B appliance.

This type of design ensures the highest level of security because it physically separates the FortiDDoS interfaces using multiple switches.

Each load-balancing device balances traffic between IP address interfaces of the peer device behind the FortiDDoS appliance. Each FortiDDoS appliance resides in a different VLAN and subnet and the physical ports connected to the FortiDDoS appliance are also on different VLANs. In addition, for each VLAN, both load-balancing devices are in the same subnet. Each load balancer interface and the FortiDDoS appliance connected to it reside in a separate VLAN. This configuration ensures persistency because all the traffic through a particular FortiDDoS appliance is contained in the appliance’s VLAN.

In a typical load-balancing device, there are two hash predictors:

- **Bidirectional hash** requires both load-balancing devices to share a common hash value that ultimately produces the same route. You create bidirectional hashing by hashing the source and destination IP address along with the destination port of the given flow. The load-balancing devices ensure that all packets belonging to a session pass through the same FortiDDoS appliance in both directions. The devices select a FortiDDoS appliance based on a symmetric hash function of the source and destination IP addresses. This ensures that packets traveling between the same source and destination IP addresses traverse the same FortiDDoS appliance.
- **Unidirectional hash** produces the route in the same fashion as a bidirectional hash and also creates a TCP connection table with the reverse flow path defined. This allows you to match return path traffic against this connection table rather than being hashed.

Figure 157: Sandwich topology for load balancing

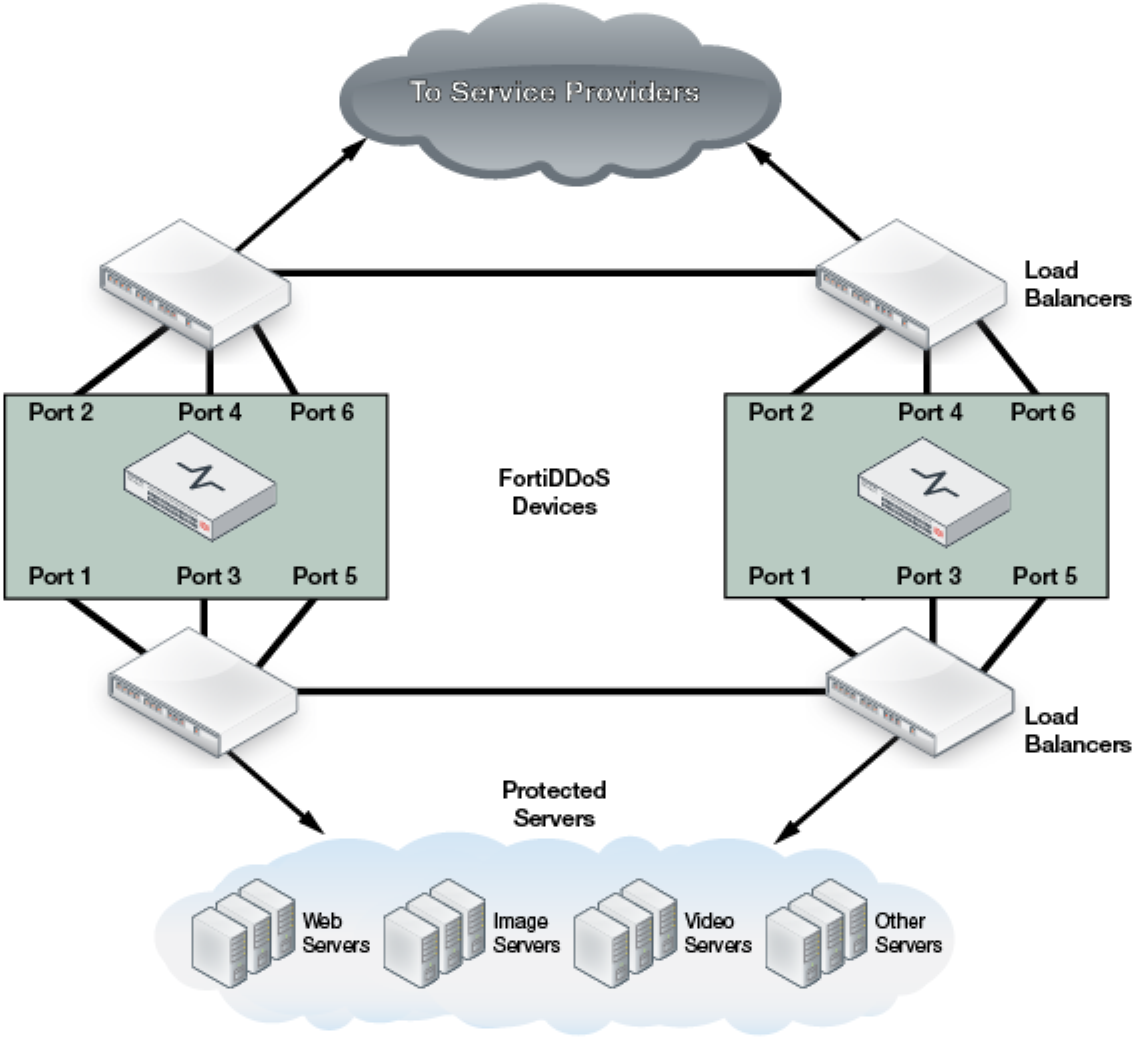

Switch configuration for load balancing using FortiSwitch

For an example configuration for the FortiSwitch 248-B DPS Ethernet switch, see Appendix D: Switch and Router Configuration.

Traffic diversion deployment

In some environments, such as a service provider environment, the total bandwidth is more than what the FortiDDoS appliance supports. However, the attack traffic to a specific subnet or server is within the appliance's capacity. You can route normal traffic through its regular path and manually divert the attack traffic. The FortiDDoS cleanses the diverted traffic and injects it back to the network.

The FortiDDoS appliance is a Layer 2 bridge and therefore does not have either a MAC address or an IP address in the data path. To allow traffic to be diverted, connect the appliance to interfaces on the routers or switches that have a routable IP address.

Figure 158 shows an example topology. The topology uses the following terminology:

- **Divert-from router:** Router from which the FortiDDoS appliance diverts the attacked customer traffic.
- **Inject-to router:** Router to which the FortiDDoS appliance forwards legitimate traffic.

Traffic diversion using separate divert-from and inject-to routers

Figure 158 shows a basic topology. In this deployment, Router 1 forwards traffic through the FortiDDoS appliance.

An additional interface on Router 1 Divert-from Router diverts the traffic that is destined for the attacked destination. This traffic passes through the FortiDDoS appliance. The traffic is then forwarded to Router 2 Inject-to Router. These two interfaces are in the same network (192.168.1.x) and therefore an ARP request from Router 1 for 192.168.1.2 passes through the FortiDDoS appliance and reaches Router 2 and Router 2 can respond back with an ARP reply and vice versa.

A static route is added on Router 1 for addresses for the attacked customer network. Because it has the longest matching prefix, the rule matches first and therefore all traffic to the attacked customer network is diverted from Router 1 to Router 2 through the FortiDDoS appliance network rather than going straight from Router 1 to Router 2. Preferably, the return path for traffic is also through the FortiDDoS appliance. Although this solution works even if the traffic is unidirectional through the FortiDDoS appliance, bidirectional traffic helps the appliance determine the statefulness within connections.

Figure 158: Traffic diversion and a FortiDDoS appliance

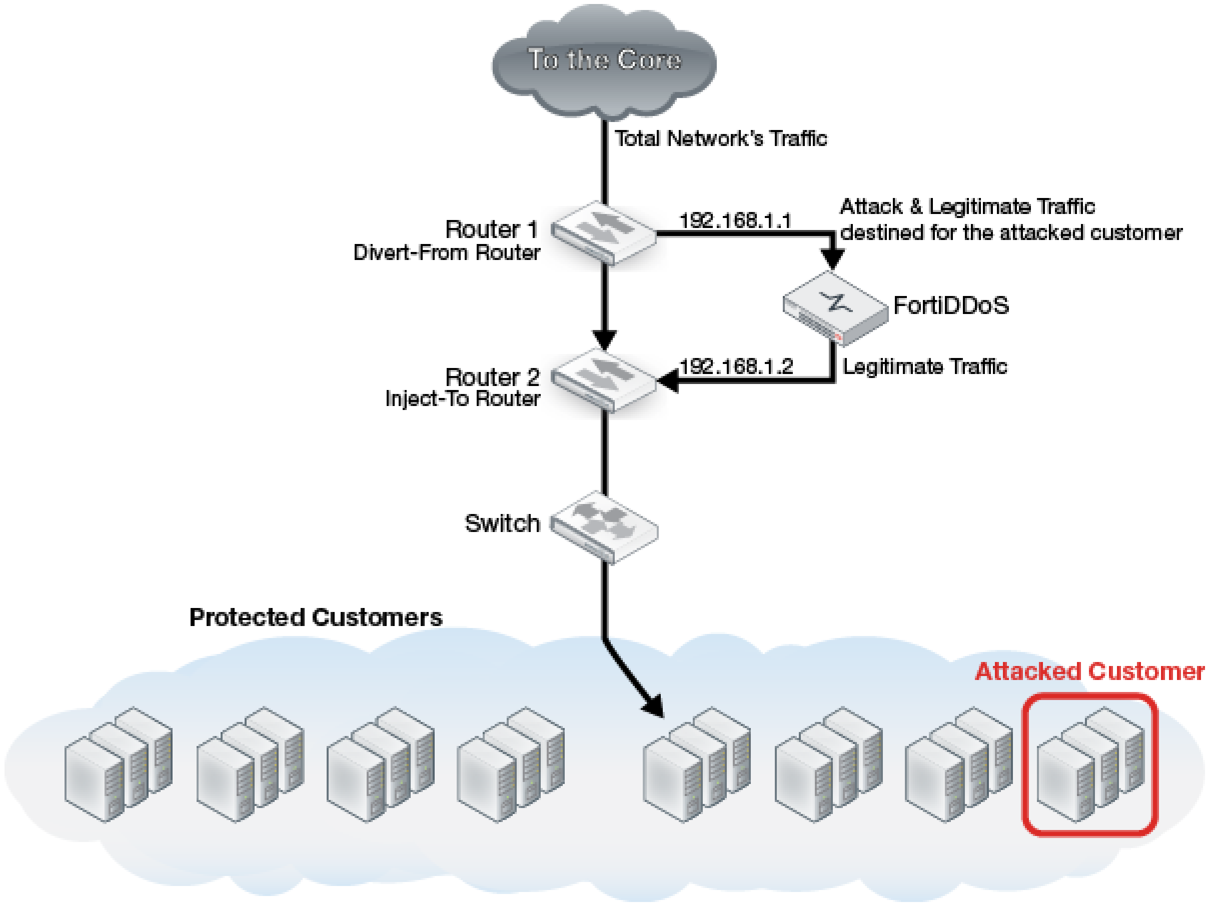

Traffic diversion using a single divert-from and inject-to router and a switch

Figure 159 shows a single router that is acting as both a divert-from and inject-to router. Layer 2 forwards through the FortiDDoS appliance.

One interface on the Internet side of the router diverts traffic to the attacked destination. This traffic passes through the FortiDDoS appliance through a switch. The traffic is then forwarded to the inject-to interface on the router through the same switch.

To ensure that the traffic is symmetric and both incoming and outgoing traffic to and from the attacked destination go through the FortiDDoS appliance, the LAN interface of the router diverts the traffic from the attacked destination. This traffic passes through the FortiDDoS appliance through a switch. The traffic is then forwarded to the inject-to interface on the same router through the same switch.

A static route is added on the router for addresses for the attacked customer network. Because it has the longest matching prefix, the rule matches first and therefore all traffic to the attacked customer network is diverted to the Layer 3 switch through the FortiDDoS appliance rather than going straight from the router to the distribution switch.

Preferably, the return path for traffic is through a FortiDDoS appliance. Although the solution works even if the traffic is unidirectional through the FortiDDoS appliance, bidirectional traffic helps the appliance determine the statefulness within connections.

To ensure that the return traffic passes through the FortiDDoS appliance, use the Policy Based Routing (PBR) that is available in most routers. PBR allows you to base routing on the source address of the packets and interface.

Figure 159: Traffic diversion using a single divert-from and inject-to router

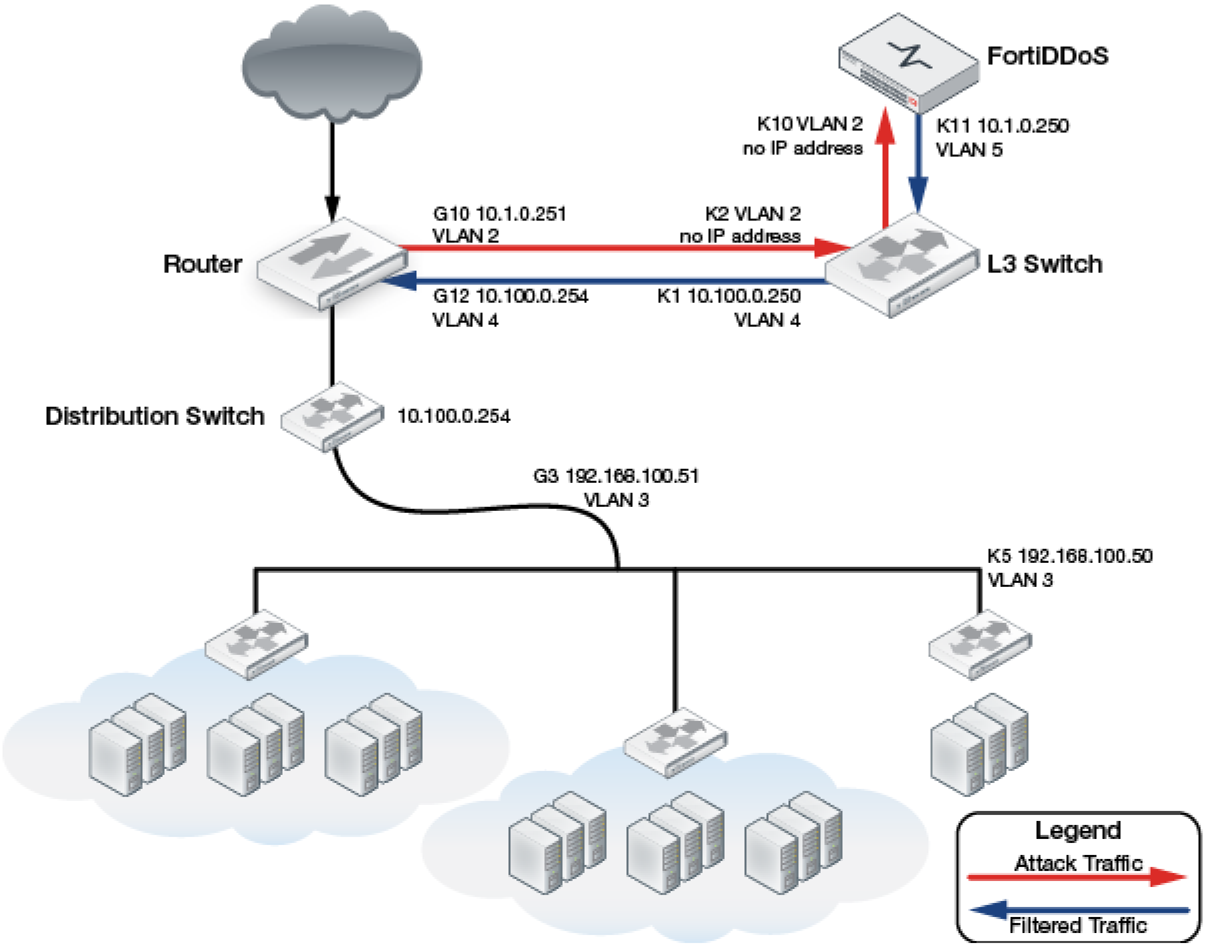

Router and switch configuration for diversion

For an example router and switch configuration for traffic diversion, see Appendix D: Switch and Router Configuration.

Setting thresholds for diverted traffic

In some cases, when traffic for a customer network is diverted through the FortiDDoS appliance during attacks, the appliance does not have traffic thresholds that correspond to the diverted network's normal traffic level or characteristics.

To solve this issue, do one of the following:

- **Archive a learning period:** During a time of normal traffic activity, divert the customer network traffic to a FortiDDoS appliance. Then, archive the configuration file created during this learning period using System > Maintenance > Backup & Restore. During an attack, restore the configuration and divert the affected traffic to the appliance with the restored configuration.

- **Create predefined profiles:** Create a series of backup configurations for different traffic levels. For example, define normal traffic levels for 1 Mbps, 10 Mbps, 20 Mbps, 100 Mbps, and so on. In addition, predefine profile parameters such as SYN/second, SYNs/Src, Concurrent Connections/Source, and so on. During an attack, restore the configuration that corresponds to the customer traffic level based on past historical knowledge of the data, and then divert the affected traffic to the appliance with the restored configuration.

Chapter 10: High Availability Deployments

This chapter includes the following sections:

HA feature overview

FortiDDoS appliances can be deployed as standalone appliances or as members of a high availability (HA) pair. FortiDDoS supports *active-passive* cluster pairs. In an HA pair, one node is the *master node*, and the other is called the *slave node*.

Figure 160 shows an active-passive deployment. The cluster uses the connection of MGMT2 ports for two types of HA communication:

- *Heartbeats*. A cluster node indicates to other nodes in the cluster that it is up and available. The absence of heartbeat traffic indicates the node is not up and is unavailable.
- *Synchronization*. During initialization and periodically thereafter, the master node pushes its configuration (with noted exceptions) to the secondary nodes.

You can log into the management interface (MGMT1) of either node, but you actively manage the configuration of the master node only.

Figure 160: Active-passive cluster

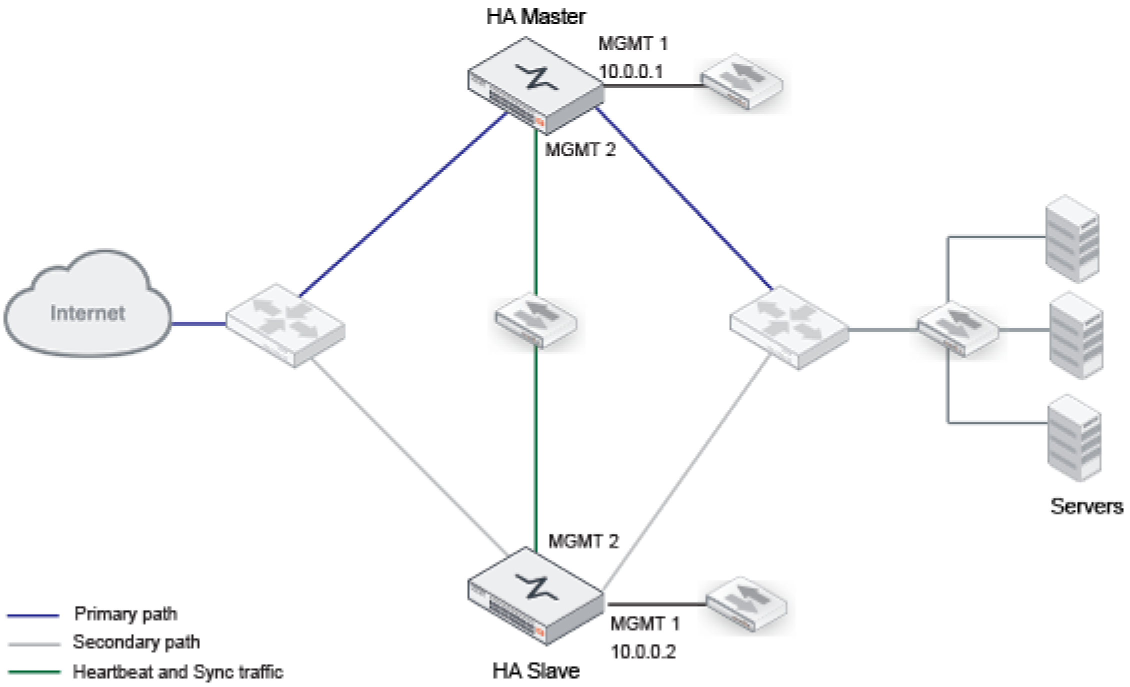

Although one appliance is deemed active (the master) and one passive (the slave), the ports are not turned off on the passive node. It can receive traffic, mitigate it, and forward it. This can cause issues monitoring traffic rates for inbound and outbound traffic.

You must use the adjacent routers to ensure that traffic is forwarded through only the active path. For example, you can set a path priority or costing to set a high priority (low cost) path that goes through the primary node, ignoring the secondary, even if it can pass traffic. If the primary fails, its interfaces can be configured to "fail closed"; the router can detect this and switch to the alternative path.

If that secondary node fails as well (double failure), you do not want the traffic to fail, so can configure the secondary system to "fail open" (copper fail open, 1200B LC ports fail open, or you need a FortiBridge if you are using SFP/SFP+s).

HA system requirements

- Two identical appliances (the same hardware model and same firmware version).
- By default, you use MGMT2 port to connect the HA appliances directly or through a Layer 2 switch. The HA port can be changed but be aware of the settings on the System > Network > Interface page before changing from default.
- Heartbeat and synchronization traffic between cluster nodes occur over the physical network ports you specify. If switches are used to connect heartbeat interfaces between nodes, the heartbeat interfaces must be reachable by Layer 2 multicast. HA traffic uses multicast UDP on port numbers 6065 (heartbeat) and 6056 (synchronization). The HA multicast IP address is 239.0.0.1; it is hard-coded, and cannot be configured.

HA synchronization

The master node pushes the following configuration elements to the slave nodes. This is known as synchronization.

Table 86: HA Synchronization

Configuration elements		HA Synchronization status
System	- Network - High Availability - Admin - Authentication - SNMP - Certificate	No
	- Maintenance	Yes
Global settings	- Service Protection Profiles - Settings - IP reputation - Proxy IP Address - Do Not Track Policy - Access Control List - Bypass MAC	Yes
	- Blacklisted Domains - Blacklisted IPv4 Address	No
	Note: Power fail bypass mode and Tap Mode settings are not synchronized.	
Log and Report	- Log configuration - Log Access - Report Configuration - Report Purge settings - Report Browse - Executive Summary - Diagnostics	No

Synchronization occurs immediately when an appliance joins the cluster, and thereafter every 30 seconds. In an active-passive cluster, these settings are read-only on the slave node.

All other system configuration, network and interface configuration, HA configuration, and log/report configuration are *not* synchronized.

Note the following:

It is not recommended to perform the below actions on a master node. You need to switch to standalone mode to modify these settings:

  - Time zone change
  - Configuration restore
  - TAP mode change

- HA slave does not synchronize time/date from HA master. To synchronize the new time/date from master, manually reboot the slave machine.
- Settings that are not synchronized should be configured before the HA active-passive setting is enabled because all of these except HA Settings become read-only on the slave node when HA is enabled.
- HA settings are read-write on all nodes in all modes so that you can switch from HA to standalone mode as needed.

Collected data is also *not* synchronized. The following data is not synchronized:

- Session data—It does not synchronize session information or any other element of the data traffic.
- Estimated thresholds—Configured thresholds are part of the configuration and are synchronized, but estimated thresholds that are shown in Monitor graphs are based on the history of traffic processed by the local system.
- Log messages—These describe events that happened on that specific appliance. After a failover, you might notice that there is a gap in the original active appliance's log files that corresponds to the period of its down time. Log messages created during the time when the standby was acting as the active appliance (if you have configured local log storage) are stored there, on the original standby appliance.
- Generated reports—Like the log messages that they are based upon, PDF, HTML, RTF, and plain text reports also describe events that happened on that specific appliance. As such, report settings are synchronized, but report output is not.

In an HA deployment, avoid using the following CLI commands:

```
config ddos spp threshold-report
config ddos spp threshold-adjust
```

These commands generate other commands and a command context, and could lead to unexpected behavior when synchronized to the secondary node. In an HA deployment, be sure to use the GUI or REST API to configure these particular settings.

Configuring HA settings

Before you begin:

- You must have Read-Write permission to items in the System category.

To configure HA settings:

1. Go to System > High Availability.
2. Complete the configuration as described in Table 87.
3. Save the configuration.

After you have saved the configuration, cluster members begin to send heartbeat traffic to each other. Members with the same Group ID join the cluster. They send synchronization traffic directly through the HA connection.

Table 87: High availability configuration

Settings	Guidelines
Configured HA Mode	• Standalone • Active-passive
Group Name	Name to identify the HA cluster if you have more than one.This setting is optional, and does not affect HA function. The maximum length is 35 characters.
Device Priority	Number indicating priority of the member node when electing the cluster master node. This setting is optional. The smaller the number, the higher the priority. The valid range is 0 to 9. The default is 5.
Override	Enable/disable to make Device Priority a more important factor than uptime when selecting the master node. Enabled by default and recommended.
Group ID	Number that identifies the HA cluster. Nodes with the same group ID join the cluster. If you have more than one HA cluster on the same network, each cluster must have a different group ID. The valid range is 0 to 63. The default is 0.
Detection Interval	Number of 100-millisecond intervals at which heartbeat packets are sent. This is also the interval at which a node expects to receive heartbeat packets. The valid range is 1 to 20 (that is, between 100 and 2,000 milliseconds). The default is 2.
Heartbeat Lost Threshold	Number of times a node retries the heartbeat and waits to receive HA heartbeat packets from the other node before concluding the other node is down. The valid range is from 1 to 60. The default is 6.

Settings	Guidelines
ARP Packet Numbers	Not used.
ARP Packet Interval	Not used.
Monitor	Mark the checkboxes for the network interface to be used for port monitoring and heartbeat packets. The standard practice is to use mgmt2 for the port monitoring and heartbeat packets. Use the same port number for all cluster members. For example, if you select mgmt2 on the primary node, select mgmt2 as the heartbeat interface on the other member nodes.

```
CLI commands:
         config system ha
            set mode <standalone | active-passive>
            set group-name <group_name_str>
            set priority <priority_int>
            set override <enable | disable>
            set group-id <group_id_integer>
            set hb-interval <hb_interval_int>
            set hb-lost-threshold <hb_lost_thresh_int>
            set hbdev <mgmt1 | mgmt2>
            set arps <arps_int>
            set arps-interval <arps_interval_int>
         end
```

Monitoring an HA cluster

You can use SNMP, log messages, and alert email to monitor HA events, such as when failover has occurred. The system logs HA node status changes as follows:

- When a member joins a group: `Member (FDD2HD3A12000003) join to the HA group`
- When the HA configuration is changed from standalone to an active-passive: `HA device into Slave mode`
- When HA synchronization is initialized: `HA device Init`

Updating firmware on an HA cluster

Installing firmware on an HA cluster is similar to installing firmware on a single, standalone appliance.

To ensure minimal interruption of service to clients, use the following steps.

If *downgrading* to a previous version, do *not* use this procedure. The HA daemon on a member node might detect that the primary node has older firmware, and attempt to upgrade it to bring it into sync, undoing your downgrade.

Instead, switch out of HA, downgrade each node individually, then switch them back into HA mode.

To update the firmware of an HA cluster:

1. Verify that the cluster node members are powered on and available.
2. Log into the web UI of the master node as the `admin` administrator.

   Alternatively, log on with an administrator account whose access profile contains **Read** and **Write** permissions in the **Maintenance** category.
3. Install the firmware on the master node. For details, see Updating firmware.

   After the new firmware has been installed, the system reboots. When the system is rebooting, the standby node assumes master status.
4. Log into the standby node as soon as it assumes master status and upgrades its firmware. HA requires the same firmware version on all nodes.

After both nodes have finished rebooting and indicate via the HA heartbeat that they are up again, the system determines whether the original node becomes the master node, according to the HA Override setting:

- If Override is *enabled* (default), the cluster considers the Device Priority setting first. Both nodes usually make a second failover in order to resume their original roles.
- If Override is *disabled*, the cluster considers uptime first. The original master node will have a smaller uptime due to the order of reboots during the firmware upgrade. Therefore it will *not* resume its active role; instead, the node with the greatest uptime will remain the new master node. A second failover will *not* occur.

Reboot times vary by the appliance model, and also by differences between the original firmware version and the firmware version you are installing.

Deploying an active-passive cluster

This topic includes the following information:

- Overview
- Basic steps
- Deploying an active-passive cluster

Overview

Figure 161 shows an active-passive deployment. When HA is enabled, the system sends heartbeat packets between the pair to monitor availability, and the master node pushes its configuration to the slave node.

Figure 161: Active-passive cluster

When the primary node goes down, the secondary becomes the master node. When the primary node comes back online, the system selects the master based on the following criteria:

- Most available ports
- Lowest device priority number (1 has greater priority than 2)
- Highest uptime value (if you disable the HA setting "Overide", then uptime will have precedence over device priority)
- Highest-sorting serial number—Serial numbers are sorted by comparing each character from left to right, where 9 and z are the greatest values. The system gives preference to higher values over lower values.

Basic steps

To deploy an active-passive cluster:

1. License all FortiDDoS appliances in the HA cluster, and register them, including FortiGuard services, with the Fortinet Technical Support website:

   https://support.fortinet.com/

2. Physically link the FortiDDoS appliances that make up the HA cluster.

   You must link at least one of their ports (for example, mgmt2 to mgmt2) for heartbeat and synchronization traffic between members of the cluster. You can do either of the following:

   - Connect the two appliances directly with an Ethernet cable.
   - Link the appliances through a switch. If connected through a switch, the HA interfaces must be reachable by Layer 2 multicast.

3. Configure the secondary node:
   a. Log into the secondary appliance as the **admin** user.
   b. Configure settings that are not subject to synchronization, including network and interface settings, Power Failure Bypass Mode (in Global Settings), SNMP, syslog, and email notification servers. You must do this when the appliance is in standalone mode because these settings are read-only on an HA secondary node.

c. Go to Global Settings > Settings and set the Power Failure Bypass Mode to **Fail Open** or **Fail Closed**, according to your preference on how to handle traffic when both nodes fail. If you use an external bypass unit, you configure **Fail Closed**.

d. Complete the HA settings as described in Configuring HA settings.

Important: Set the Device Priority to a higher number than the primary appliance; for example, set Device Priority to 2.

4. Configure the primary node:

a. Log into the primary appliance as the **admin** user.

b. Go to Global Settings > Settings and set the Power Failure Bypass Mode to **Fail Closed**.

c. Complete the configuration for all features, as well as the HA configuration.

Important: Set the Device Priority to a lower number than the secondary appliance; for example, set Device Priority to 1.

Note: After you have saved the HA configuration changes, cluster members might join or rejoin the cluster. After you have saved configuration changes on the master node, it automatically pushes its Global Settings and Protection Profiles configuration to the slave node.

Chapter 11: Service Provider Signaling Deployments

This chapter includes the following sections:

Overview

When a subnet or destination server in a customer premises network (CPN) is the target of a DDoS attack, a FortiDDoS appliance deployed in the customer premises network can detect the attack and enforce the SPP threshold policies to protect the destination servers; however, at this point, the WAN uplink connecting the customer premises network to the Internet might have already become saturated with attack traffic, resulting in legitimate traffic being dropped and the destination being unreachable.

The Service Provider Signaling feature enables small/medium businesses and enterprises to work with participating service providers to route traffic through a "scrubbing station" in the service provider network (SPN) before it is forwarded through the WAN link to the customer premises network (CPN). The scrubbing station is a large-scale FortiDDoS appliance or a third-party device that enforces its own Global Settings policies and the SPP policy assigned to the subnet. Traffic that is not dropped at the scrubbing station is forwarded to the customer premises network.

The feature uses REST API communication to support:

- Registration of CPN FortiDDoS appliances with an SPN FortiDDoS appliance or third-party appliance.
- Status checks (every 1 minute) for the connection from the CPN to the SPN.
- Signaling from the CPN FortiDDoS to the SPN FortiDDoS when traffic volume reaches the configured threshold.
- Export of the CPN FortiDDoS SPP policy and settings to the SPN FortiDDoS appliance so that a security policy based on normal CPN SPP baseline traffic rates can be enforced.

The SPN FortiDDoS administrator can be alerted of the attack through the event log, SNMP, or alert email notification. The SPN administrator must then use the BGP routing policy to divert traffic to the attack destination through the SPN scrubbing station.

When the attack is over, the SPN administrator should remove the SPP policy that was installed for the CPN (if not, a subsequent signaling from that SPP will fail).

The SPN FortiDDoS must be a model that supports DNS protection. FortiDDoS 600B and 900B do not support DNS protection.

Topology

Figure 162 shows the network topology for a Service Provider Signaling deployment. Under normal conditions, the network traffic through the service provider network WAN link to the customer premises network follows the path of the green arrow.

When the volume of traffic exceeds the high volume threshold, the FortiDDoS appliance in the CPN signals the FortiDDoS in the SPN. The SPN administrator can then route traffic through the scrubbing station.

Figure 162: Service Provider Signaling topology

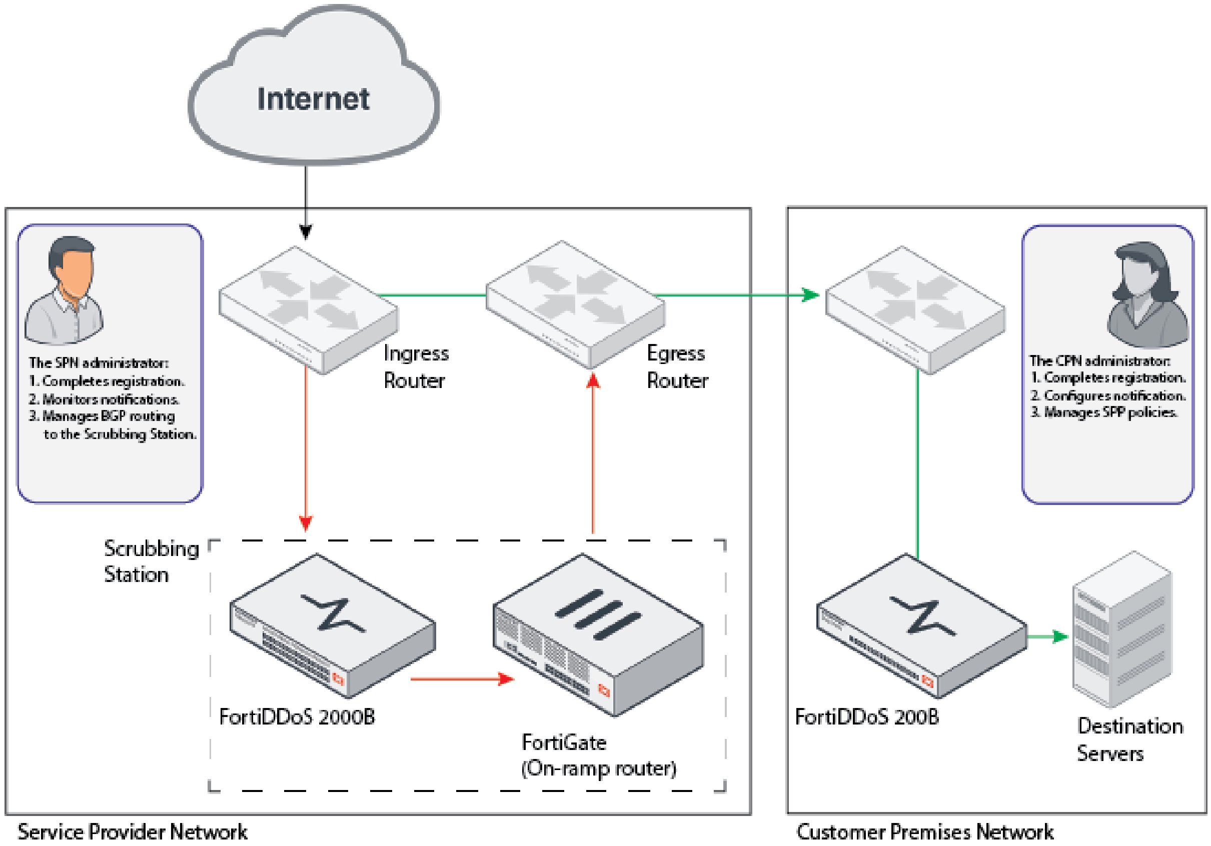
Internet
The SPN administrator:
1. Completes registration.
2. Monitors notifications.
3. Manages BGP routing to the Scrubbing Station.
Ingress Router
Egress Router
The CPN administrator:
1. Completes registration.
2. Configures notification.
3. Manages SPP policies.
Scrubbing Station
FortiDDoS 2000B
FortiGate (On-ramp router)
FortiDDoS 200B
Destination Servers
Service Provider Network
Customer Premises Network

Notes:

- Ingress Router—When an attack is signaled, the SPN administrator updates the BGP routing policy so that traffic to the destination that is under attack is forwarded to the scrubbing station. Once the attack is over, the SPN administrator can update the BGP routing policy again so the traffic can go directly to the CPN.
- FortiGate (on-ramp router)—The on-ramp router is responsible for receiving the traffic that passes through the SPN FortiDDoS and injecting it back to the service provider network. We recommend a FortiGate model that supports routing and EtherChannel aggregation features. An on-ramp router is required because the FortiDDoS deployment is transparent to the routers; it has no IP address in the path of packets.
- Egress Router—The service provider router closest to the customer premises network. During an attack, this link to the CPN edge router can become saturated, hence this signaling solution.

Registration

This section includes the following information:

- Overview
- CPN registration tasks
- SPN registration tasks

Overview

Participation is coordinated in advance by the SPN and CPN administrators. An SPN appliance can accept registration from up to 8 CPN appliances. A CPN appliance can register with only one SPN appliance.

1. On the SPN appliance, the administrator creates a configuration for the connection with the CPN appliance. The configuration includes the CPN appliance serial number, IP address, and a shared secret.
2. On the CPN appliance, the administrator creates a configuration for the connection with the SPN appliance. The configuration includes the SPN appliance serial number, IP address, and a shared secret.
3. When the configuration on the CPN appliance is saved, the CPN appliance initiates a registration request to the SPN appliance. Thee registration request includes its serial number, IP address, and the shared secret.
4. The SPN appliance attempts to register the CPN appliance and returns a status message indicating "registered" if the registration information matches or "mismatch" if the registration information does not match. The SPN appliance sends a "declined" massage if the SPN administrator actively declines the registration.

Figure 163: Registration

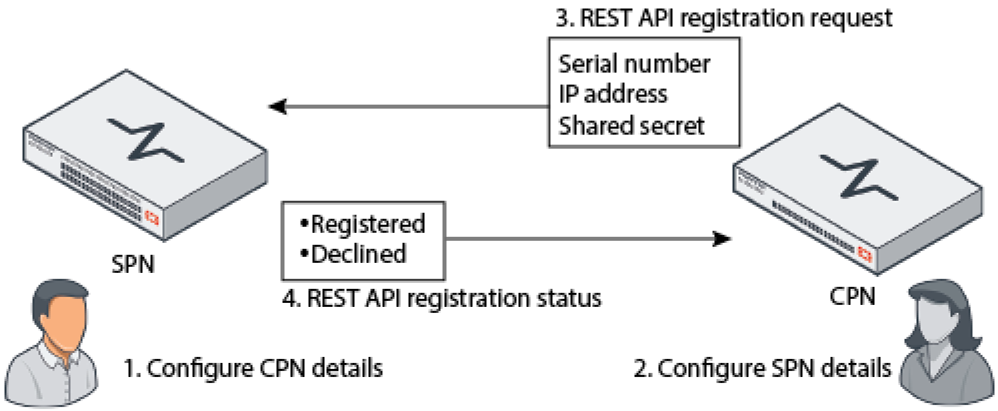

SPN registration tasks

The SPN administrator completes the registration.

Basic steps

1. Go to Global Settings > Deployment and select **Service Provider**.
2. Go to Global Settings > Signaling and review details for the connection with the CPN FortiDDoS.
3. Change the pending registration status to registered.

Figure 164: Global Settings > Deployment tab

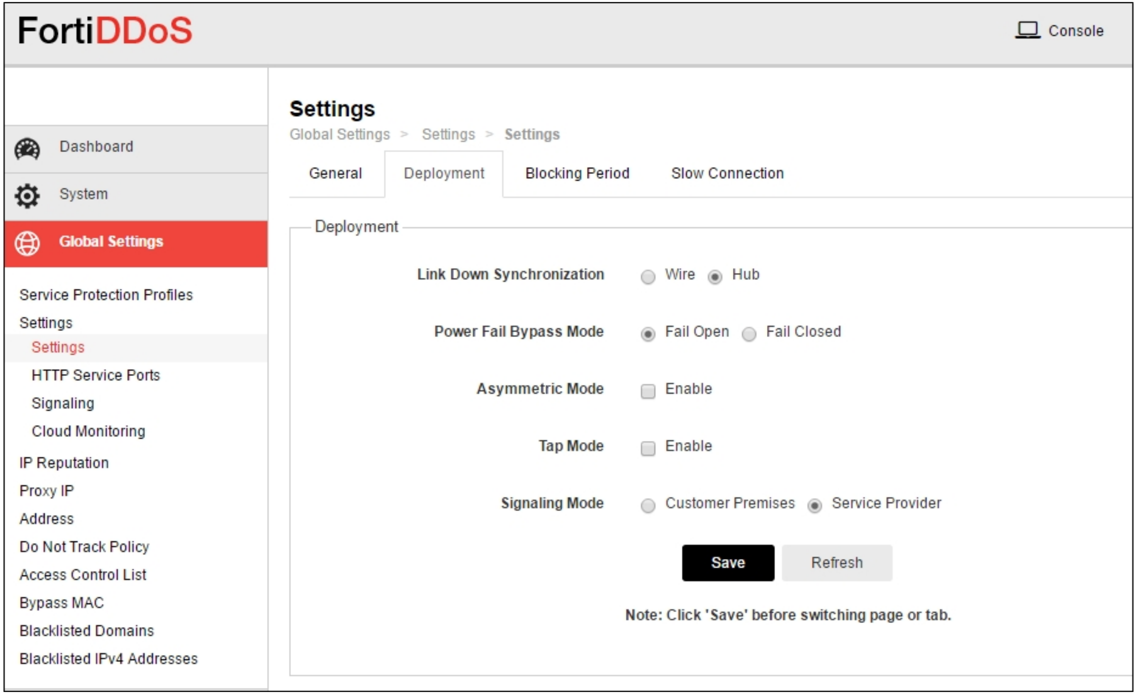

Settings

Global Settings > Settings > Settings

Deployment

Link Down Synchronization ○ Wire ● Hub

Power Fail Bypass Mode ● Fail Open ○ Fail Closed

Asymmetric Mode □ Enable

Tap Mode □ Enable

Signaling Mode ○ Customer Premises ● Service Provider

Note: Click 'Save' before switching page or tab.

CLI commands:

```
SP-FDD # config ddos global setting
SP-FDD (setting) # set signaling-mode service-provider
SP-FDD (setting) # end
```

Figure 165: Global Settings > Signaling page

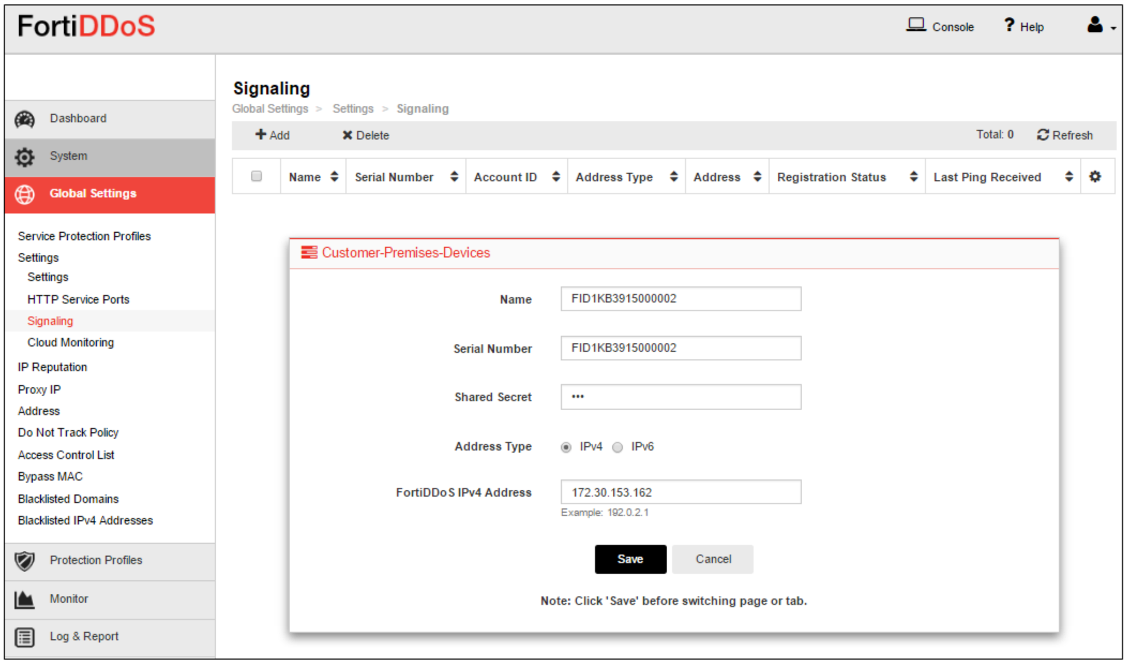

Signaling

Global Settings > Settings > Signaling

Total: 0

□	Name	Serial Number	Account ID	Address Type	Address	Registration Status	Last Ping Received	⚙

Customer-Premises-Devices

Name: FID1KB3915000002

Serial Number: FID1KB3915000002

Shared Secret: •••

Address Type ● IPv4 ○ IPv6

FortiDDoS IPv4 Address: 172.30.153.162
Example: 192.0.2.1

Note: Click 'Save' before switching page or tab.

CLI commands:

```
SP-FDD # config ddos global signaling-devices
SP-FDD (signaling-devi~e) # edit FI800B3913800021
SP-FDD (FI800B3913800024) # get
serial-number : FI800B3913800021
shared-secret : test1
address-type : ipv4
ipv4-address : 172.30.153.121
registration-status : pending-registration
```

CPN registration tasks

The CPN administrator initiates registration of the CPN FortiDDoS to the SPN FortiDDoS.

Basic steps

1. Go to Global Settings > Settings > Settings > Deployment tab and select **Customer Premises**.
2. Go to Global Settings > Signaling and provide details for the connection with the SPN FortiDDoS.

Figure 166: Global Settings > Deployment tab

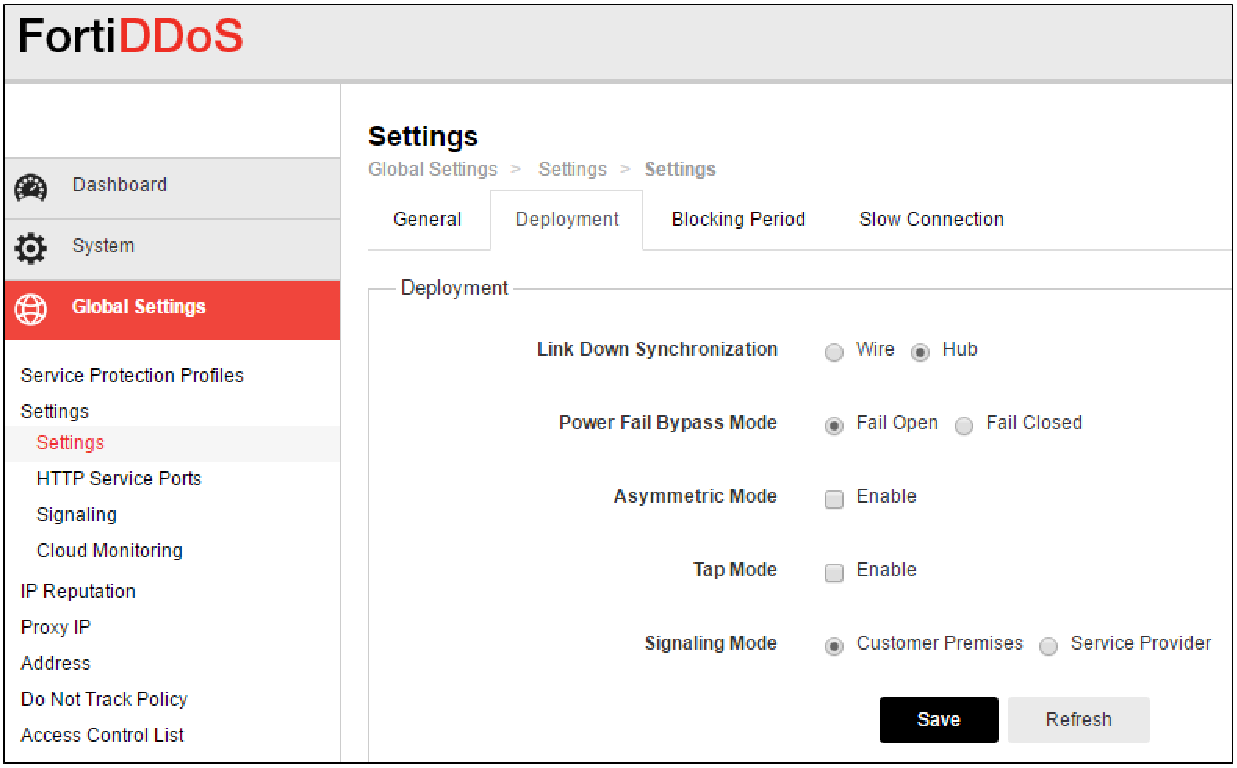

Settings

Global Settings > Settings > Settings

Deployment

Link Down Synchronization ○ Wire ● Hub

Power Fail Bypass Mode ● Fail Open ○ Fail Closed

Asymmetric Mode □ Enable

Tap Mode □ Enable

Signaling Mode ● Customer Premises ○ Service Provider

CLI commands:

```
CP-FDD # config ddos global setting
CP-FDD (setting) # set signaling-mode customer-premises
CP-FDD (setting) # end
```

Figure 167: Global Settings > Signaling page

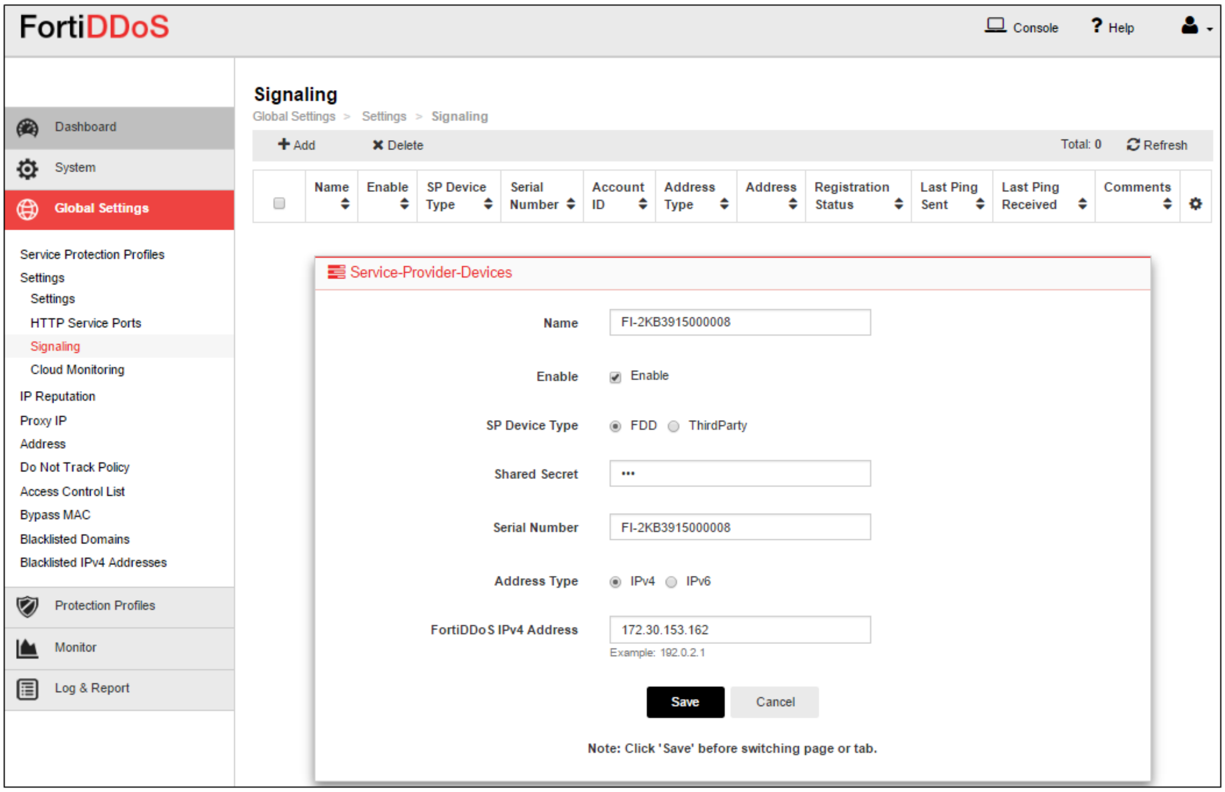

Name: FI-2KB3915000008

Enable: ☑ Enable

SP Device Type: ● FDD ○ ThirdParty

Shared Secret: •••

Serial Number: FI-2KB3915000008

Address Type: ● IPv4 ○ IPv6

FortiDDoS IPv4 Address: 172.30.153.162

Example: 192.0.2.1

Note: Click 'Save' before switching page or tab.

CLI commands:

```
CP-FDD # config ddos global service-provider-devices
CP-FDD (service-provid~r) # edit SP-FDD
CP-FDD (SP-FDD) # set enable-sp-device enable
CP-FDD (SP-FDD) # set serial-number FI800B3913800024
CP-FDD (SP-FDD) # set shared-secret test1
CP-FDD (SP-FDD) # set ipv4-address 172.30.153.125
CP-FDD (SP-FDD) # end
```

Signaling

The CPN FortiDDoS sends the SPP policy settings and SPP settings configuration to the SPN FortiDDoS to signal the volume threshold for the subnet has been reached. The SPN FortiDDoS checks registration status and availability of an SPP "slot" to import the SPP configuration. If those checks pass, the SPP configuration from the CPN is installed automatically in the SPN FortiDDoS.

Figure 168: Signaling

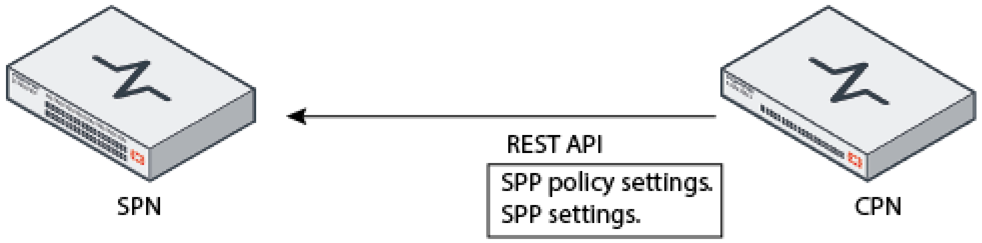

The CPN SPP policy settings determine the volume threshold at which to signal the SPN.

Basic Steps

1. On the CPN appliance, go to Global Settings > Switching Policy and enable the feature.
2. Go to Global Settings > SService Protection Profiles > Switching Policy. When you configure the SPP policy, follow these guidelines:
   - SPP Switching—Enable.
   - Alternate Service Protection Profile—Specify the same SPP name. For example, if you are configuring SPP-1, specify SPP-1 as the alternate as well.
   - Threshold—Packet rate at which signaling occurs.

Figure 169: SPP Switching Policy

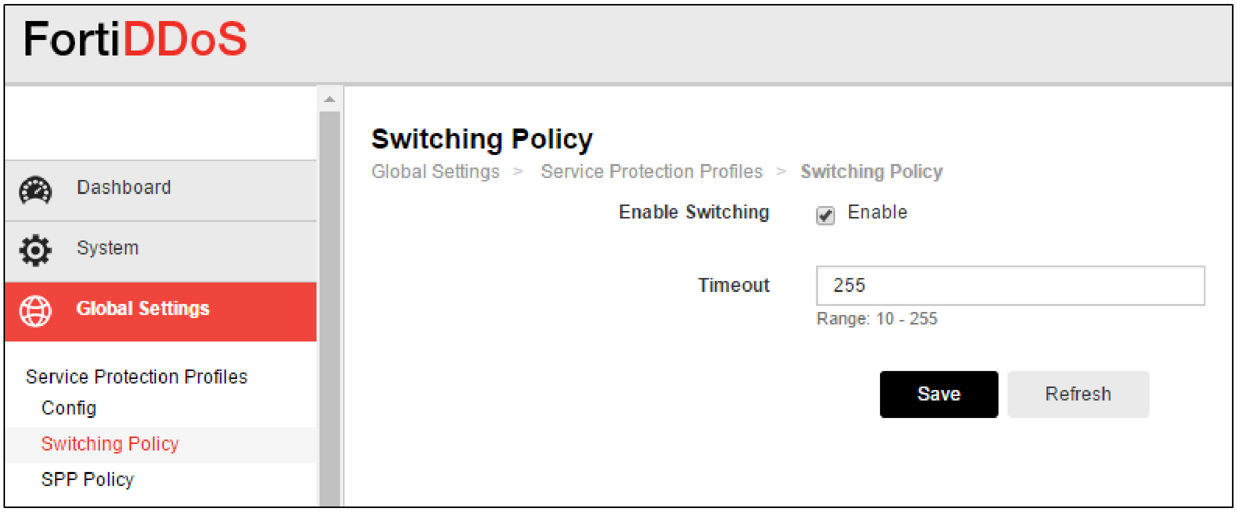

Switching Policy

Global Settings > Service Protection Profiles > Switching Policy

Enable Switching ☑ Enable

Timeout 255

Range: 10 - 255

Figure 170: SPP Policy

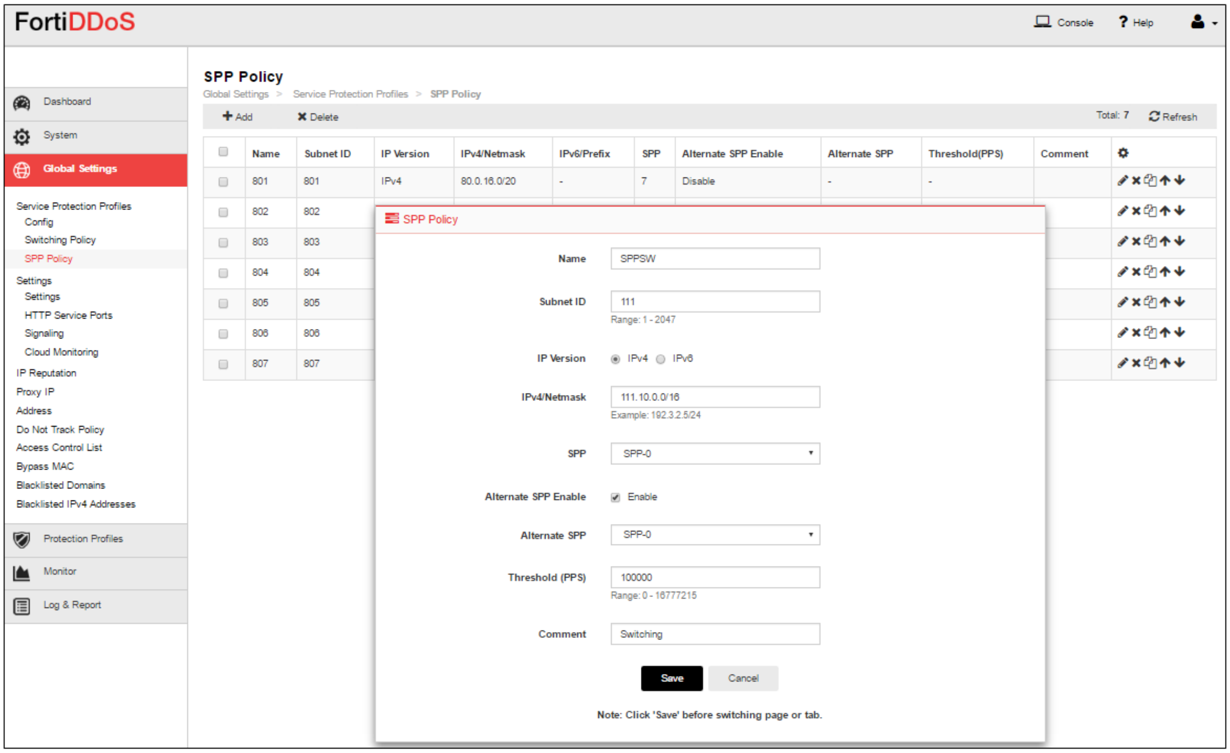

SPP Policy

Global Settings > Service Protection Profiles > SPP Policy

	Name	Subnet ID	IP Version	IPv4/Netmask	IPv6/Prefix	SPP	Alternate SPP Enable	Alternate SPP	Threshold(PPS)	Comment
□	801	801	IPv4	80.0.16.0/20	-	7	Disable	-	-	
□	802	802								
□	803	803								
□	804	804								
□	805	805								
□	806	806								
□	807	807								

SPP Policy

Name: SPPSW

Subnet ID: 111
Range: 1 - 2047

IP Version: ● IPv4 ○ IPv6

IPv4/Netmask: 111.10.0.0/16
Example: 192.3.2.5/24

SPP: SPP-0

Alternate SPP Enable: ☑ Enable

Alternate SPP: SPP-0

Threshold (PPS): 100000
Range: 0 - 16777215

Comment: Switching

Save Cancel

Note: Click 'Save' before switching page or tab.

Notification

The CPN and SPN FortiDDoS administrators must work out the details on how the SPN administrator (or team) is notified when signaling occurs. The following features support notification:

- SNMP traps
- Email alerts
- Event logs

SNMP traps

On the CPN FortiDDoS, the following SNMP traps are sent:

- The SPP switching policy threshold has been reached, and therefore the signaling has been initiated.
- CPN FortiDDoS attempt to signal SPN FortiDDoS failed.

On the SPN FortiDDoS, an SNMP trap is sent when the signaling CPN SPP policy is loaded into one of the 8 SPP slots. The SPN administrator must first create 8 empty SPP slots and configure SNMP trap receivers (Log & Report > Log Configuration > SNMP Trap Receivers) for each. Then, if a CPN signals and its SPP policy is installed in slot SPP-3, for example, an SNMP trap is sent to the SNMP trap receiver configured for slot SPP-3.

Email alerts

You can configure email alerts. Go to Log & Report > Alert Email Settings and configure alert email settings. Include alerts when SPP Switching/Signaling occurs.

On the CPN FortiDDoS, the following alerts are sent:

- Registration has been attempted, and it indicates success or failure.
- Signaling has been attempted, and it indicates success or failure.

On the SPN FortiDDoS, an alert is sent when the SPN FortiDDoS attempts to install the SPP policy and SPP settings, and it indicates success or failure.

Figure 171: Alert notifications

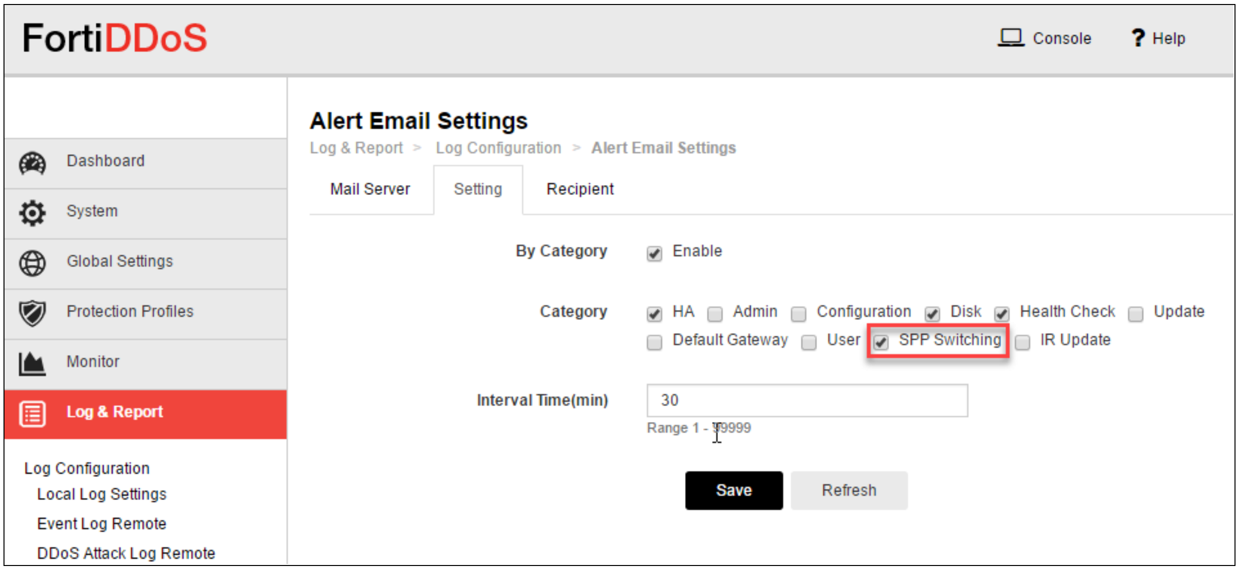

By Category ☑ Enable

Category ☑ HA □ Admin □ Configuration ☑ Disk ☑ Health Check □ Update
□ Default Gateway □ User ☑ SPP Switching □ IR Update

Interval Time(min) 30
Range 1 - 99999

Event logs

You can enable event logs. Go to Log & Report > Log Configuration > Local Log Settings and configure event log settings. Include events when SPP Switching/Signaling occurs.

On the CPN FortiDDoS, the following logs are generated:

- Registration has been attempted, and it indicates success or failure.
- Signaling has been attempted, and it indicates success or failure.

On the SPN FortiDDoS, logs are generated when the SPN FortiDDoS attempts to install the SPP policy and SPP settings, and it indicates success or failure.

Figure 172: Event logging

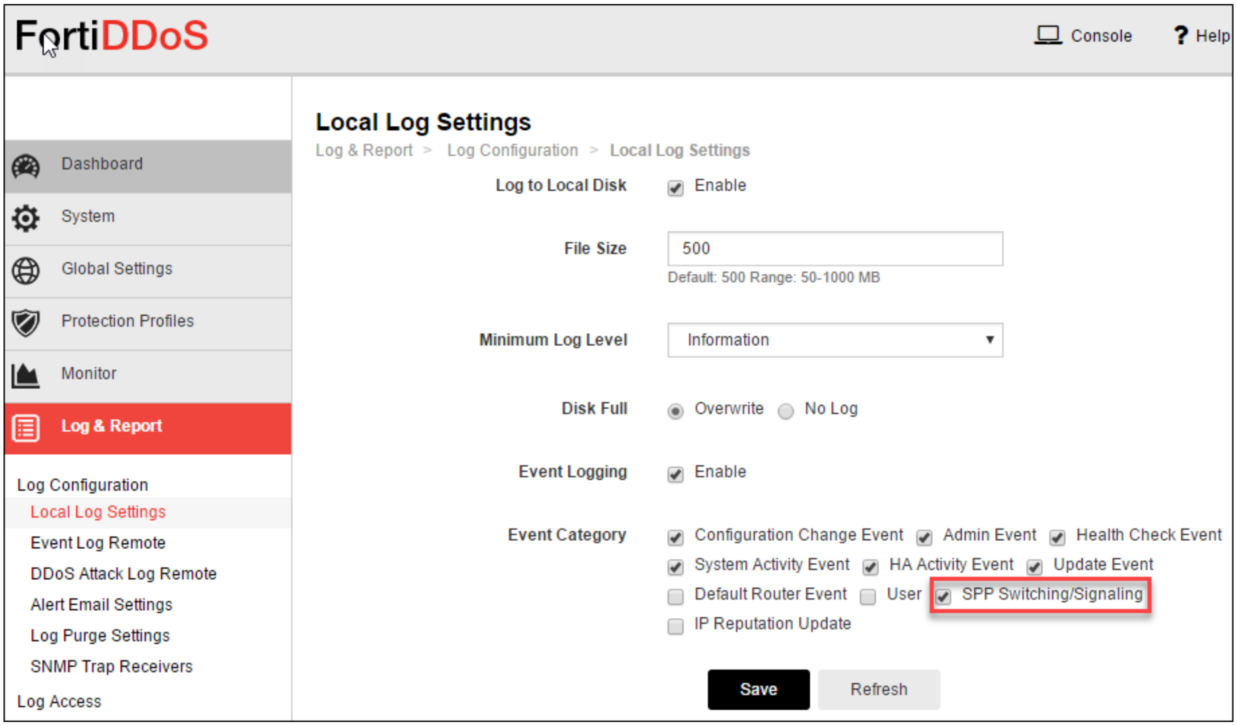

Local Log Settings

Log & Report > Log Configuration > Local Log Settings

Log to Local Disk ☑ Enable

File Size 500
Default: 500 Range: 50-1000 MB

Minimum Log Level Information ▾

Disk Full ● Overwrite ○ No Log

Event Logging ☑ Enable

Event Category ☑ Configuration Change Event ☑ Admin Event ☑ Health Check Event
☑ System Activity Event ☑ HA Activity Event ☑ Update Event
□ Default Router Event □ User ☑ SPP Switching/Signaling
□ IP Reputation Update

Chapter 12: Troubleshooting

This chapter includes the following information:

Logs

Log messages often contain clues that can aid you in determining the cause of a problem.

Depending on the type, log messages may appear in either the system event logs or the DDoS attack logs. To enable logging of different categories of system events, go to Log & Report > Log Configuration > Local Log Settings. All DDoS attack log categories are enabled automatically and cannot be disabled.

During troubleshooting, you might find it useful to lower the logging severity threshold for more verbose logs, to include more information on less severe events. To configure the log level, go to Log & Report > Log Configuration > Local Log Settings.

Tools

This section describes the following troubleshooting tools:

- execute commands
- diagnose commands
- Special Fortinet Support commands

execute commands

You can use the command-line interface (CLI) execute commands to run diagnostic utilities, such as nslookup, ping, and traceroute.

The following example shows the list of execute commands:

```
FI800B3913000018 # execute
backup backup
bypass-traffic bypass data traffic <enable|disable>
checklogdisk find and correct errors on the log disk
cleanup-db-transaction-log cleanup database transaction log files
date set date and time
dos-control dos-control <enable|disable>
factoryreset reset to factory default
formatlogdisk format log disk to enhance performance
log-rebuild rebuild log index
nslookup nslookup
ping ping <host name | host ip>
ping-option ping option settings
ping6 ping <host name | host ipv6>
ping6-option ping6 option settings
reboot reboot the system
reconfigure reconfigure
reload reload appliance
repair-database-tables repair database tables
restore restore
shutdown shutdown appliance
telnettest test if we can telnet to a server
traceroute traceroute
```

You can also use the `tree` command to display the list of commands:

```
FI800B3913000018 # tree execute
-- {reboot(0)}
|- {shutdown(0)}
|- backup -- config -- {tftp(0)} -- arg ...
|- {checklogdisk(0)}
|- {factoryreset(0)} -- arg ...
|- {formatlogdisk(0)} -- arg ...
|- nslookup -- {name(0)} -- arg ...
|- {ping(0)} -- arg ...
|- {ping-option(0)} -- data-size -- arg ...
|- pattern -- arg ...
|- repeat-count -- arg ...
|- source -- arg ...
|- timeout -- arg ...
|- tos -- arg ...
|- ttl -- arg ...
|- validate-reply -- arg ...
|- view-settings -- arg ...
+- df-bit -- arg ...
|- {ping6(0)} -- arg ...
|- {ping6-option(0)} -- data-size -- arg ...
|- pattern -- arg ...
|- repeat-count -- arg ...
|- source -- arg ...
|- timeout -- arg ...
|- tos -- arg ...
|- ttl -- arg ...
|- validate-reply -- arg ...
+- view-settings -- arg ...
|- {reload(0)} -- arg ...
|- reconfigure -- asic -- {tftp(0)} -- arg ...
|- {test-image(0)} -- arg ...
+- {production(0)} -- arg ...


|- restore -- config -- {tftp(0)} -- arg ...
+- syntax -- {tftp(0)} -- arg ...
|- {telnettest(0)} -- arg ...
|- {traceroute(0)} -- arg ...
|- {date(0)} -- arg ...
|- {log-rebuild(0)}
|- {cleanup-db-transaction-log(0)}
|- {dos-control(0)} -- arg ...
+- {repair-database-tables(0)}
```

diagnose commands

You can use the CLI diagnose commands to gather diagnostic information that can be useful to Fortinet Customer Care when diagnosing any issues with your system.

The following examples show the lists of diagnose commands:

```
FI800B3913000018 # diagnose
debug debug
hardware hardware
netlink netlink
sniffer sniffer
system system

FI-2KB3913000002 # diagnose debug
application set/get debug level for daemons
cli set/get debug level for CLI and CMDB
crashlog clear/get crashlog
disable disable debug output
enable enable debug output
kernel set/get debug level for kernel
mysql-log get mysql error log
nginx-log get nginx error log
rrd_cmd_check Perform RRD commands check
rrd_cmd_recreate Re-create RRD commands
```

```
rrd_tune Tune RRD database to eliminate dropcount limit.

FI800B3913000018 # diagnose hardware get
deviceinfo list device status and information
ioport read data from an I/O port
pciconfig list information on PCI buses and connected devices
sysinfo list system hardware information
tp2control read control registers from tp2
tp2reg read registers from tp2
tp2statistics read statistics registers from tp2


FI800B3913000018 # diagnose netlink
backlog set netlink backlog length
device display network devices statistic information
interface netlink interface
ip ip
ipv6 ipv6
neighbor netlink neighbor
neighbor6 netlink neighbor for ipv6
queue-len set netlink TX queue length
route netlink routing table
route6 netlink routing table
rtcache netlink realtime cache
tcp display tcp statistic information
udp display udp statistic information

FI800B3913000018 # diagnose sniffer packet mgmt2
interfaces=[mgmt2]
filters=[none]
pcap_lookupnet: mgmt2: no IPv4 address assigned
0.000000 172.30.144.11.62729 -> 172.30.153.113.22: ack 1556045916
0.000000 172.30.144.11.62729 -> 172.30.153.113.22: ack 1556046032
0.000000 172.30.153.17.68 -> 255.255.255.255.67: udp 300

FI800B3913000018 # diagnose system top
Run Time: 0 days, 23 hours and 12 minutes
0.6U, 1.7S, 119946561.8I; 7996T, 44F
tp2traffic 3501 R 9.4 1.3
mysqld 1351 S 0.0 6.4
flg_access 1354 S 0.0 1.6
tp2d 1358 S 0.0 1.6
flg_indexd 1352 S 0.0 1.7
tp2logd 1357 S 0.0 1.1
php-fpm 32341 S 0.0 2.1
php-fpm 20518 S 0.0 2.2
php-fpm 1600 S 0.0 2.2
cli 15253 S 0.0 1.6
php-fpm 1349 S 0.0 2.1
fddcli 14721 S 0.0 1.5
cmdbsvr 1297 S 0.0 1.5
cli 31603 R 0.0 1.6
sshd 15104 S 0.0 0.9
cli 1342 S 0.0 1.6
```

You can also use the `tree` command to display the list of commands:

```
FI-2KB3913000002 # tree diag
```

```
-- sniffer -- {packet(0)} -- arg ...
|- netlink -- {queue-len(0)} -- arg ...
|- {backlog(0)} -- arg ...
|- {route(0)} -- arg ...
|- {route6(0)} -- arg ...
|- {rtcache(0)} -- arg ...
|- {interface(0)} -- arg ...
|- {neighbor(0)} -- arg ...
|- {neighbor6(0)} -- arg ...
|- {tcp(0)}
|- {udp(0)}
|- {device(0)}
|- ip -- {flush(0)} -- arg ...
|- {list(0)} -- arg ...
|- {delete(0)} -- arg ...
+- {add(0)} -- arg ...
+- ipv6 -- {flush(0)} -- arg ...
|- {list(0)} -- arg ...
|- {delete(0)} -- arg ...
+- {add(0)} -- arg ...
|- debug -- {application(0)} -- info_centerd -- <level>
|- hasyncd -- arg ...
|- updated -- arg ...
|- miglogd -- <level>
|- cmdb_event -- <level>
|- sshd -- <level>
|- netd -- <level>
|- alertmaild -- <level>
|- tp2trafficd -- <level>
|- ntpd -- <level>
|- tp2rupd -- <level>
|- snmpd -- <level>
|- flg_indexd -- <level>
|- flg_reportd -- <level>
|- flg_accessd -- <level>
|- tp2oneclickd -- <level>
|- tp2d -- <level>
|- tp2init -- <level>
|- tp2logd -- <level>
|- tp2ird -- <level>
|- tp2threshd -- <level>
|- tp2portd -- <level>
|- tp2proxyd -- <level>
|- tp2stressd -- <level>
|- tp2diagd -- <level>
+- tp2dpcconfd -- <level>
|- {crashlog(0)} -- <level>
|- {cli(0)} -- <level>
|- {kernel(0)} -- <level>
|- {rrd_cmd_check(0)} -- arg ...
|- {rrd_cmd_recreate(0)} -- arg ...
|- {rrd_tune(0)} -- arg ...
|- {enable(0)}
|- {disable(0)}
|- {mysql-log(0)} -- <level>
|- {nginx-log(0)} -- <level>
+- {scramble(0)} -- <level>
```

```
|- hardware -- get -- {sysinfo(0)} -- arg ...
|- {deviceinfo(0)} -- arg ...
|- {biosinfo(0)}
|- {pciconfig(0)} -- arg ...
|- {ioport(0)} -- arg ...
|- {tp2reg(0)} -- arg ...
|- {tp2control(0)} -- arg ...
+- {tp2statistics(0)} -- arg ...
+- set -- {pciconfig(0)} -- arg ...
|- {ioport(0)} -- arg ...
|- {tp2reg(0)} -- arg ...
+- {tp2control(0)} -- arg ...
+- system -- {load(0)} -- arg ...
|- {top(0)} -- arg ...
|- {matrix(0)}
|- {scp(0)} -- arg ...
|- {slab(0)} -- arg ...
|- {shm(0)} -- arg ...
|- {sem(0)} -- arg ...
|- {cmdb(0)}
|- {ntp-status(0)}
|- {ha(0)} -- arg ...
|- {checkused(0)} -- arg ...
|- {fips(0)} -- arg ...
+- {file-system(0)} -- arg ...
```

Special Fortinet Support commands

The commands described in this section are useful when you are troubleshooting an issue with the help of Fortinet Technical Support. Your Fortinet contact might ask you to run these commands to gather data they need to troubleshoot system issues.

execute backup diag_info

This command exports diagnostic information to a remote TFTP server. The following information is exported:

- System status
- Current configuration
- Hardware register values
- Event and DDoS attack log database

Use the following command syntax:

```
# execute backup diag_info tftp <tftp_server_ipaddress>
```

The filename generated stems from the appliance serial number and date. For example, diag_info-FI-1KB0000000007-2015-03-07-16-57.tgz.

The archive includes four files with filenames similar to the following:

```
back_status-FI-1KB0000000007-2015-03-07-16-57
back_cfg-FI-1KB0000000007-2015-03-07-16-57
back_hw_reg-FI-1KB0000000007-2015-03-07-16-57
back_logs-FI-1KB0000000007-2015-03-07-16-57.tgz
```

The logs archive includes four files with filenames similar to the following:

```
elog@002e0000000001.MAI
```

```
elog@002e0000000001.MAD
dlog.MAI
dlog.MAD
```

execute backup hw_reg

This command exports a dump file of the TP2 register data to a remote TFTP server.

Use the following command syntax:

```
#execute backup hw_reg tftp <filename> <tftp_server_ipaddress>
NUM Boards = 2,
Connecting to tftp server <tftp_server_ipaddress> ...
Please wait...
#
Send TP2 Register dump to tftp server OK.
```

get system sensors

Use this command to display the status of system sensors. The following is an example:

```
FI200B3914000001 # get system sensors
Sensor ID | Reading | Units | State | LNR | LC | LNC | UNC | UC | UNR
------------------------------------------------------------------------------------------
    -------------
AD_+3.3V | 3.320 | Volts | ok | 2.976 | 3.044 | 3.148 | 3.474 | 3.578 | 3.646
AD_+5V | 5.047 | Volts | ok | 4.508 | 4.606 | 4.753 | 5.268 | 5.415 | 5.513
AD+12V | 12.272 | Volts | ok | 10.207 | 10.443 | 10.856 | 13.216 | 13.570 | 13.806
AD_VCORE_CPU | 0.921 | Volts | ok | 0.333 | 0.343 | 0.353 | 1.607 | 1.656 | 1.686
AD_VTT_CPU | 1.049 | Volts | ok | 0.951 | 0.970 | 1.000 | 1.107 | 1.137 | 1.156
AD_VSA_CPU | 0.931 | Volts | ok | 0.735 | 0.745 | 0.774 | 1.029 | 1.058 | 1.068
AD_DDR3_VDQ1 | 1.499 | Volts | ok | 1.215 | 1.245 | 1.284 | 1.597 | 1.627 | 1.656
NCT +1.05V | 1.040 | Volts | ok | 0.960 | 0.976 | 1.008 | 1.104 | 1.136 | 1.168
NCT 12V_FET | 12.096 | Volts | ok | 10.368 | 10.560 | 10.944 | 13.248 | 13.632 | 13.824
NCT MAIN_12V | 12.096 | Volts | ok | 10.368 | 10.560 | 10.944 | 13.248 | 13.632 | 13.824
NCT VCC1_8_BP | 1.792 | Volts | ok | 1.632 | 1.664 | 1.712 | 1.904 | 1.952 | 1.984
NCT +1.8V | 1.776 | Volts | ok | 1.632 | 1.664 | 1.712 | 1.904 | 1.952 | 1.984
NCT V1_2_56321| 1.216 | Volts | ok | 1.088 | 1.104 | 1.152 | 1.264 | 1.296 | 1.328
NCT +DDR3_VTT | 0.736 | Volts | ok | 0.688 | 0.704 | 0.720 | 0.800 | 0.816 | 0.832
NCT RPS_12V | 0.000 | Volts | nr | 10.368 | 10.560 | 10.944 | 13.248 | 13.632 | 13.824
NCT7904D 3VDD | 3.264 | Volts | ok | 2.976 | 3.072 | 3.168 | 3.504 | 3.600 | 3.648
NCT7904D 3VSB | 3.264 | Volts | ok | 2.976 | 3.072 | 3.168 | 3.504 | 3.600 | 3.648
NCT7904D VTT | 1.040 | Volts | ok | 0.960 | 0.976 | 1.008 | 1.104 | 1.136 | 1.168
NCT7904D VBAT | 3.072 | Volts | ok | 2.736 | 2.784 | 2.880 | 3.504 | 3.600 | 3.648
TD1 | 37.000 | degrees C | ok | na | na | na | 75.000 | 80.000 | 85.000
TD2 | 32.000 | degrees C | ok | na | na | na | 65.000 | 70.000 | 75.000
TD3 | 30.000 | degrees C | ok | na | na | na | 65.000 | 70.000 | 75.000
DTS CPU | 35.000 | degrees C | ok | na | na | na | 100.000 | 105.000 | 106.000
CPU Core 0 | 33.000 | degrees C | ok | na | na | na | 100.000 | 105.000 | 106.000
CPU Core 1 | 32.000 | degrees C | ok | na | na | na | 100.000 | 105.000 | 106.000
CPU Core 2 | 37.000 | degrees C | ok | na | na | na | 100.000 | 105.000 | 106.000
CPU Core 3 | 35.000 | degrees C | ok | na | na | na | 100.000 | 105.000 | 106.000
TS1 | 34.000 | degrees C | ok | na | na | na | 75.000 | 80.000 | 85.000
TS2 | 29.000 | degrees C | ok | na | na | na | 70.000 | 75.000 | 80.000
TS3 | 31.000 | degrees C | ok | na | na | na | 70.000 | 75.000 | 80.000
TS4 | 27.000 | degrees C | ok | na | na | na | 70.000 | 75.000 | 80.000
Fan 1 | 6400.000 | RPM | ok | 100.000 | 200.000 | 300.000 | 23000.00| 24000.00| 25000.000
Fan 2 | 6500.000 | RPM | ok | 100.000 | 200.000 | 300.000 | 23000.00| 24000.00| 25000.000
Fan 3 | 6500.000 | RPM | ok | 100.000 | 200.000 | 300.000 | 23000.00| 24000.00| 25000.000
Fan 4 | 6300.000 | RPM | ok | 100.000 | 200.000 | 300.000 | 23000.00| 24000.00| 25000.000
Fan 5 | 6400.000 | RPM | ok | 100.000 | 200.000 | 300.000 | 23000.00| 24000.00| 25000.000
PS Temp 1 | 22.000 | degrees C | ok | na | na | na | 70.000 | 75.000 | 80.000
PS Temp 2 | 26.000 | degrees C | ok | na | na | na | 80.000 | 85.000 | 90.000
PS Fan 1 | 7168.000 | RPM | ok | 128.000 | 256.000 | 384.000 | 15104.00| 18048.00|
    20096.000
PS VIN | 111.000 | Volts | ok | 34.000 | 36.000 | 38.000 | 250.000 | 254.000 | 255.000
PS VOUT_12V | 12.285 | Volts | ok | 10.836 | 11.088 | 11.466 | 12.600 | 12.915 | 13.167
PS IIN | 0.126 | Amps | ok | na | na | na | 13.230 | 13.608 | 13.923
PS IOUT_12V | 7.000 | Amps | ok | na | na | na | 30.500 | 31.500 | 32.000
PS PIN | 48.000 | Watts | ok | na | na | na | 472.000 | 488.000 | 500.000
PS POUT | 84.000 | Watts | ok | na | na | na | 380.000 | 392.000 | 400.000
PS Status | 0x0 | discrete | 0x0100| na | na | na | na | na | na
INA219 PS Vsht| 0.004 | Volts | ok | na | na | na | 0.030 | 0.031 | 0.032
INA219 PS Vbus| 12.160 | Volts | ok | 10.880 | 11.136 | 11.392 | 12.672 | 13.056 | 13.312
INA219 PS Pwr | 45.000 | Watts | ok | na | na | na | 380.000 | 392.500 | 400.000
INA219 PS Curr| 3.500 | Amps | ok | na | na | na | 30.500 | 31.500 | 32.000
Chassis FRU | 0x0 | discrete | 0x1080| na | na | na | na | na | na
Version change| na | discrete | na | na | na | na | na | na | na
TP2 1 Presence| 0x0 | discrete | 0x0100| na | na | na | na | na | na
TP2 2 Presence| 0x0 | discrete | 0x0200| na | na | na | na | na | na
TP2 1 Enable | 0x0 | discrete | 0x0100| na | na | na | na | na | na
TP2 2 Enable | 0x0 | discrete | 0x0100| na | na | na | na | na | na
TP2 1 TS | na | degrees C | na | na | na | na | 70.000 | 75.000 | 80.000
TP2 2 TS | 40.000 | degrees C | ok | na | na | na | 70.000 | 75.000 | 80.000
```

Solutions by issue type

This section includes the following topics:

- Browser compatibility
- Connectivity issues
- Resource issues

Browser compatibility

FortiDDoS 4.3.0 will not work with Internet Explorer version 9.0 and 10.0. In IE 11.0, the following setting must be adjusted for correct operation:

1. Click Settings > Internet options.
2. Click Settings under Browsing history.
3. Select **Every time I visit the webpage** under 'Check for newer versions of stored pages:'.

Connectivity issues

One of your first tests when configuring a new SPP should be to determine whether legitimate traffic is forwarded to protected resources.

Checking hardware connections

If packets are not forwarded by the FortiDDoS appliance, it might be a hardware problem.

To check hardware connections:

- Ensure the network cables are properly plugged into the interfaces on the FortiDDoS appliance.
- Ensure there are connection lights for the network cables on the appliance.
- Change the cable if the cable or its connector are damaged or you are unsure about the cable's type or quality.
- Connect the FortiDDoS appliance to different hardware to see if that makes a difference.
- In the web UI, select **System > Status > Dashboard**. In the **System Status** widget, ensure that the status indicators for the ports that are in use are green (indicating that physical connections are present) or flashing green (indicating that data is flowing). Hover over the indicator for further status information.

If any of these checks solve the problem, it was a hardware connection issue. You should still perform some basic software tests to ensure complete connectivity.

If the hardware connections are correct, and the appliance is powered on but you cannot connect using the CLI or web UI, you might be experiencing bootup problems. You might have to reimage the system. See Restoring firmware (“clean install”).

Data path connectivity

You can use `ping` and other methods to verify that FortiDDoS appliance forwards packets to protected servers.

Verifying the path between client and server

If you are testing connectivity using a client that is directly connected to the appliance, use `ping` to test the traffic flow. For indirect connections, use `traceroute`.

To verify routes between clients and your servers using ping

1. Try to communicate with the server from the client using the `ping` command. Use the following graphs to detect if the traffic has travelled through the FortiDDoS appliance:
   - Monitor > Port Statistics > Packets
   - Monitor > Port Statistics > Bits
   - Monitor > Specific Graphs > Protocols (protocol 1 or 58)

   If you do not see the expected count of bits or packets, continue to the next step.

1. Use the `ping` command on both the client and the server to verify that a route exists between the two. Test traffic movement in both directions: from the client to the server, and the server to the client. Servers do not need to be able to initiate a connection, but must be able to send reply traffic along a return path.

In networks that have asymmetric routes, routing success in one direction does *not* guarantee success in the other..

To verify routes between clients and your servers using traceroute

Use the `tracert` or `traceroute` command on both the client and the server (depending on their operating systems) to determine if there is a point of failure along the route.

If the route is broken when it reaches the FortiDDoS appliance, first examine its network interfaces and routes. To display network interface information, enter the CLI command:

```
show system interface
```

Testing data path routes and latency with traceroute

`traceroute` sends ICMP packets to test each hop along the route. It sends three packets to the destination, and then increases the time to live (TTL) setting by one, and sends another three packets to the destination. As the TTL increases, packets go one hop farther along the route until they reach the destination.

Most `traceroute` commands display their maximum hop count — that is, the maximum number of steps it will take before declaring the destination unreachable — before they start tracing the route. The TTL setting may result in routers or firewalls along the route timing out due to high latency.

Where `ping` only tells you if the signal reached its destination and returned successfully, `traceroute` shows each step of its journey to its destination and how long each step takes. If you specify the destination using a domain name, the `traceroute` output can also indicate DNS problems, such as an inability to connect to a DNS server.

By default, `traceroute` uses UDP with destination ports numbered from 33434 to 33534. The traceroute utility usually has an option to specify use of ICMP `ECHO_REQUEST` (type 8) instead, as used by the Windows `tracert` utility. If you have a firewall and you want `traceroute` to work from both machines (Unix-like systems and Windows) you will need to allow *both* protocols inbound through your firewall (UDP ports 33434 - 33534 and ICMP type 8).

To trace the route to a server from a Microsoft Windows computer:

1. Go to the Windows command line interface (cmd.exe).
2. Enter the command:

```
tracert {<destination_ipv4> | <destination_fqdn>}
```

If the appliance has a complete route to the destination, output similar to the following appears:

```
Tracing route to www.fortinet.com [66.171.121.34]
over a maximum of 30 hops:

1 <1 ms <1 ms <1 ms 172.16.1.2
2 2 ms 2 ms 2 ms static-209-87-254-221.storm.ca [209.87.254.221]

3 2 ms 2 ms 22 ms core-2-g0-1-1104.storm.ca [209.87.239.129]
4 3 ms 3 ms 2 ms 67.69.228.161
5 3 ms 2 ms 3 ms core2-ottawa23_POS13-1-0.net.bell.ca [64.230.164
.17]
(Output abbreviated.)
15 97 ms 97 ms 97 ms gar2.sj2ca.ip.att.net [12.122.110.105]
16 94 ms 94 ms 94 ms 12.116.52.42
17 87 ms 87 ms 87 ms 203.78.181.10
18 89 ms 89 ms 90 ms 203.78.181.130
19 89 ms 89 ms 90 ms fortinet.com [66.171.121.34]
20 90 ms 90 ms 91 ms fortinet.com [66.171.121.34]
```

```
Trace complete.
```

Each line lists the routing hop number, the 3 response times from that hop, and the IP address and FQDN (if any) of that hop. Typically a value of `<1ms` indicates a local router.

If the appliance does not have a complete route to the destination, output similar to the following appears:

```
Tracing route to 10.0.0.1 over a maximum of 30 hops

1 <1 ms <1 ms <1 ms 172.16.1.2
2 <1 ms <1 ms <1 ms 172.16.1.10
3 * * * Request timed out.
4 * * * Request timed out.
5 ^C
```

The asterisks (*) and “Request timed out.” indicate no response from that hop in the network routing.

To trace the route to a server from a Linux or Mac OS X computer

1. Open a command prompt, or on Mac OS X, use the Network Utility application.
2. Enter (the path to the executable varies by distribution):

```
traceroute {<destination_ipv4> | <destination_fqdn>}
```

If the appliance has a complete route to the destination, output similar to the following appears:

```
traceroute to www.fortinet.com (66.171.121.34), 30 hops max, 60 byte packets
1 172.16.1.2 (172.16.1.2) 0.189 ms 0.277 ms 0.226 ms
2 static-209-87-254-221.storm.ca (209.87.254.221) 2.554 ms 2.549 ms 2.503 ms
3 core-2-g0-1-1104.storm.ca (209.87.239.129) 2.461 ms 2.516 ms 2.417 ms
4 67.69.228.161 (67.69.228.161) 3.041 ms 3.007 ms 2.966 ms
5 core2-ottawa23_POS13-1-0.net.bell.ca (64.230.164.17) 3.004 ms 2.998 ms 2.963 ms
(Output abbreviated.)
16 12.116.52.42 (12.116.52.42) 94.379 ms 94.114 ms 94.162 ms
17 203.78.181.10 (203.78.181.10) 122.879 ms 120.690 ms 119.049 ms
18 203.78.181.130 (203.78.181.130) 89.705 ms 89.411 ms 89.591 ms
19 fortinet.com (66.171.121.34) 89.717 ms 89.584 ms 89.568 ms
```

Each line lists the routing hop number, the IP address and FQDN (if any) of that hop, and the 3 response times from that hop. Typically a value of `<1ms` indicates a local router.

If the appliance does not have a complete route to the destination, output similar to the following appears:

```
traceroute to 10.0.0.1 (10.0.0.1), 30 hops max, 60 byte packets
1 * * *
2 172.16.1.10 (172.16.1.10) 4.160 ms 4.169 ms 4.144 ms
3 * * *
4 * * *^C
```

The asterisks (*) indicate no response from that hop in the network routing.

Likewise, if the computer’s DNS query cannot resolve the host name, output similar to the following appears:

```
example.lab: Name or service not known
Cannot handle "host" cmdline arg `example.lab' on position 1 (argc 1)
```

Checking routing

`ping` and `traceroute` are useful tools in network connectivity and route troubleshooting for management network interfaces.

Since you typically use these tools for troubleshooting only, allow ICMP (the protocol used by these tools) on interfaces only when you need them. Otherwise, disable ICMP for improved security and performance.

By default, FortiDDoS appliances respond to `ping` and `traceroute`. However, if the appliance does not respond, and there are no firewall policies that block it, ICMP type 0 (`ECHO_REPSPONSE`) might be effectively disabled.

> Disabling ping only prevents FortiDDoS from *receiving* ICMP type 8 (`ECHO_REQUEST`) or type 30 and traceroute-related UDP.
>
> It does *not* disable FortiDDoS CLI commands such as `execute ping` or `execute traceroute` that *send* such traffic.

To enable ping and traceroute responses from FortiDDoS:

1. Go to System > Network > Interface.
   To access this part of the web UI, you must have Read-Write permission in your administrator's account access profile to items in the System category.
2. In the row for the management network interface which you want to respond to ICMP type 8 (`ECHO_REQUEST`) for `ping` and UDP for `traceroute`, click **Edit**.
   A dialog appears.
3. Enable ping.
4. Click **OK**.

The appliance should now respond when another device such as your management computer sends a `ping` or `traceroute` to that management interface.

Examining the routing table

When a route does not exist, or when hops have high latency, examine the routing table. The routing table is where the FortiDDoS appliance caches recently used routes.

If a route is cached in the routing table, it saves time and resources that would otherwise be required for a route lookup. If the routing table is full and a new route must be added, the oldest, least-used route is deleted to make room.

To check the routing table for the management network interface in the CLI, enter:

```
diagnose netlink route list
```

Resource issues

If the system resource usage appears to be abnormally high according to the System Resource widget or the CLI command `get system status`, you can view the current consumption by each process by entering this CLI command:

```
diagnose system top delay 10
```

The above command generates a list of processes every 10 seconds. It includes the process names, their process ID (pid), status, CPU usage, and memory usage.

The report continues to refresh and display in the CLI until you press `q` (quit).

If the issue recurs, and corresponds with a hardware or configuration change, you might need to change the configuration. Look especially into reducing frequent logging. If the issue persists, contact Fortinet Technical Support.

Resetting profile data or the system configuration

Table 88 summarizes “factory reset” options.

Table 88: “Factory reset” options

Task	Menu
Reset the threshold configuration for an SPP but do not clear traffic history. You might do this if you are conducting a demonstration or test, or you are troubleshooting an issue; or if you want to start over with a new learning period in Detection Mode and start with high thresholds that will not drop traffic.	See Restoring factory default threshold settings.
Reset the threshold configuration for an SPP and clear its traffic history. You might do this if characteristics of the traffic protected by an SPP change significantly (for example, you change which server or protocol that it protects).	See Performing a factory reset of SPP settings.
Reset the system to the factory state. All SPPs, statistics, and logs will be deleted. You might do this if you are selling your FortiDDoS appliance.	See Resetting the system.

Important: Before you perform a factory reset:

- Make a backup of the current configuration.
- Be ready to reconfigure the default gateway and IP address of the network interface that is used for connections to the web UI and CLI.
- Do not shut down the appliance while it is resetting.

Restoring firmware (“clean install”)

Restoring (also called re-imaging) the firmware can be useful in the following cases:

- You are unable to connect to the FortiDDoS appliance using the web UI or the CLI
- You want to install firmware *without* preserving any existing configuration (that is, perform a “clean install”)

Unlike updating firmware, restoring firmware re-images the boot device. Also, restoring firmware can only be done during a boot interrupt, before network connectivity is available, and therefore requires a local console connection to the CLI. It cannot be done through an SSH or Telnet connection.

Alternatively, if you cannot physically access the appliance’s local console connection, connect the appliance’s local console port to a terminal server to which you have network access. Once you have used a client to connect to the terminal server over the network, you will be able to use the appliance’s local console through it. However, be aware that from a remote location, you may not be able to power cycle the appliance if abnormalities occur.

Important: Back up the configuration before completing a clean install.

To restore the firmware

1. Download the firmware file from the Fortinet Technical Support website.
2. Connect your management computer to the FortiDDoS console port using a RJ-45-to-DB-9 serial cable or a null-modem cable.
3. Initiate a local console connection from your management computer to the CLI of the FortiDDoS appliance, and log in as the `admin` administrator.
4. Connect the MGMT1 port of the FortiDDoS appliance directly or to the same subnet as a TFTP server.
5. Copy the new firmware image file to the root directory of the TFTP server.
6. If necessary, start your TFTP server. (If you do not have one, you can temporarily install and run one such as `tftpd` on your management computer.)

TFTP is not secure, and it does not support authentication. You should run it only on trusted administrator-only networks, and never on computers directly connected to the Internet. Turn off tftpd off immediately after completing this procedure.

7. Verify that the TFTP server is currently running, and that the FortiDDoS appliance can reach the TFTP server.

   To use the FortiDDoS CLI to verify connectivity, enter the following command:

   ```
   execute ping 192.168.1.168
   ```

   where `192.168.1.168` is the IP address of the TFTP server.
8. Enter the following command to restart the FortiDDoS appliance:

   ```
   execute reboot
   ```

As the FortiDDoS appliances starts, a series of system startup messages appear.

```
Press any key to display configuration menu........
```

9. Immediately press a key to interrupt the system startup.

> You have only 3 seconds to press a key. If you do not press a key soon enough, the FortiDDoS appliance reboots and you must log in and repeat the `execute reboot` command.

   If you successfully interrupt the startup process, the following messages appears:

```
[G]: Get firmware image from TFTP server.
[F]: Format boot device.
[B]: Boot with backup firmware and set as default.
[Q]: Quit menu and continue to boot with default firmware.
[H]: Display this list of options.

Enter G,F,B,Q,or H:

Please connect TFTP server to Ethernet port "1".
```

10. If the firmware version requires that you first format the boot device before installing firmware, type `F`. Format the boot disk before continuing.
11. Type `G` to get the firmware image from the TFTP server.

    The following message appears:

```
Enter TFTP server address [192.168.1.168]:
```

12. Type the IP address of the TFTP server and press Enter.

    The following message appears:

```
Enter local address [192.168.1.188]:
```

13. Type a temporary IP address that can be used by the FortiDDoS appliance to connect to the TFTP server.

    The following message appears:

```
Enter firmware image file name [image.out]:
```

14. Type the file name of the firmware image and press Enter.

    The FortiDDoS appliance downloads the firmware image file from the TFTP server and displays a message similar to the following:

```
MAC:00219B8F0D94

###########################

Total 28385179 bytes data downloaded.

Verifying the integrity of the firmware image..

Save as Default firmware/Backup firmware/Run image without saving:[D/B/R]?
```

15. Type `D`.

The FortiDDoS appliance downloads the firmware image file from the TFTP server. The FortiDDoS appliance installs the firmware and restarts. The time required varies by the size of the file and the speed of your network connection.

The FortiDDoS appliance reverts the configuration to default values for that version of the firmware.

16. To verify that the firmware was successfully installed, log in to the CLI and type:

```
get system status
```

    The firmware version number is displayed.

17. Either reconfigure the FortiDDoS appliance or restore the configuration file.

If you are *downgrading* the firmware to a previous version, and the settings are not fully backwards compatible, the FortiDDoS appliance either removes incompatible settings, or uses the feature’s default values for that version of the firmware. You might need to reconfigure some settings.

Installing firmware overwrites any FortiGuard IP Reputation Service definitions and disables the service. After any firmware update, re-enable the IP Reputation feature. FortiDDoS downloads current definitions as part of the enabling process.

Additional resources

Fortinet also provides these resources:

- The Release Notes provided with your firmware
- Technical documentation (references, installation guides, and other documents)
- Knowledge base (technical support articles)
- Forums
- Online campus (tutorials and training materials)

Check within your organization. You can save time and effort during the troubleshooting process by checking if other FortiDDoS administrators experienced a similar problem before.

If you cannot resolve the issue on your own, contact Fortinet Technical Support.

Appendix A: DDoS Attack Log Reference

Event code (syslog): 1 - Layer 3, 2 - Layer 4, 4 - Layer 7

Subcode (syslog): a number of no particular significance

Name: Event Type in web UI, description field in syslog

Category: Filter category in web UI

Period: **Interrupt** - Rate Flood - means the first event is logged within 2 minutes after start of attack and reported every minute thereafter; **Periodic**- Events other than Rate Flood - means events are logged every 5 minutes. Source IP address is reported only for drops due to per-source thresholds (see Source tracking table).

Table 89: DDoS Attack Log Reference

Code	Sub	Name	Category	Period	Description	Parameter	Graph
1	0	Protocol flood	Rate flood	Interrupt	Effective rate limit for the protocol has been reached.	Protection Profiles > Thresholds > Protocols	Log & Report > Executive Summary > DDoS Attack Graphs -> Top Attacked Protocols, to identify Protocols of interest Then: Monitor > Layer 3 > Protocols and enter Protocol numbers to see rate and drop graphs.
1	1	Fragment flood	Rate flood	Interrupt	Effective rate limit for the fragment threshold has been reached.	Protection Profiles > Thresholds > Scalars: Fragment	Monitor > Layer 3 > Fragmented Packets

Code	Sub	Name	Category	Period	Description	Parameter	Graph
1	4	IP header anomaly	Header anomaly	Periodic	Drops due to predefined IP header rules: Invalid header length (less than 5 words) Total length less than 20 bytes End of Header before the data offset (while parsing options) Length field in the LSRR or SSRR IP option is other than (3+(n*4)) where n is a value greater than or equal to 1 Pointer in the LSRR or SSRR IP option is other than (n*4) where n is a value greater than or equal to 1	None. Dropped as Anomaly.	Monitor > Anomaly Drops > Layer 3 Monitor > Anomaly Drops > Layer 4 > Header -> Anomaly Detected
1	6	ST:Hash attack	-	Periodic	An issue with hash collisions in the source tracking (ST) table.	None. Internal Table issue. Report to Fortinet.	Monitor > Hash Attack Drops > Layer 3 > Source Table
1	7	ST:Out of memory	-	Periodic	An issue with the source tracking (ST) table internal logic or memory.	None. Internal Table issue. Report to Fortinet.	Monitor > Out of Memory Drops > Layer 3 > Source Table
1	8	Source flood	Rate flood	Interrupt	Effective rate limit for the most-active-source threshold has been reached. Source IP address is reported.	Protection Profiles > Thresholds > Scalars: Most Active Source	Monitor > Layer 3 > Most Active Source
1	10	DT:Hash attack	-	Periodic	An issue with hash collisions in the destination tracking (DT) table.	None. Internal Table issue. Report to Fortinet.	Monitor > Hash Attack Drops > Layer 3 > Destination Table

Code	Sub	Name	Category	Period	Description	Parameter	Graph
1	11	DT:Out of memory	-	Periodic	An issue with the destination tracking (DT) table internal logic or memory.	None. Internal Table issue. Report to Fortinet.	Monitor > Out of Memory Drops > Layer 3 > Destination Table
1	14	IP Header checksum error	Header anomaly	Periodic	Invalid IP header checksum.	None. Dropped as Anomaly.	Monitor > Anomaly Drops > Layer 3: IP Header Checksum
1	15	Source IP==dest IP	Header anomaly	Periodic	Identical source and protected IP addresses (LAND attack).	None. Dropped as Anomaly.	Monitor > Anomaly Drops > Layer 3: Source and Destination Address Match
1	16	Source/dest IP==localhost	Header anomaly	Periodic	Source/destination address is the local host (loopback address spoofing).	None. Dropped as Anomaly.	Monitor > Anomaly Drops > Layer 3: Source / Destination as localhost

Cod-e	Su-b	Name	Cat-egory	Period	Description	Parameter	Graph
1	17	L3 anomalies	Header anom-aly	Peri-odic	Drops due to predefined Layer 3 rules: - IP version other than IPv4 or IPv6. - EOP (End of Packet) before 20 bytes of IPv4 data. -EOP comes before the length specified by Total Length. -Reserved Flag set. -More Frag and Don't Frag Flags set. -Added Anomaly for DSCP and ECN.	None. Dropped as Anomaly.	Monitor > Anomaly Drops > Layer 3: Layer 3
1	50	Protocol denied	ACL	Peri-odic	Denied by an SPP ACL rule.	Protection Pro-files > Service > Service Con-fig: Protocols Protection Pro-files > Access Control List > Access Control List	Monitor > ACL Drops > Layer 3: Protocol Denied Drops

Code	Sub	Name	Category	Period	Description	Parameter	Graph
1	54	Fragment denied	ACL	Periodic	Denied by an SPP ACL rule.	Protection Profiles > Service > Service Config: Fragment Protection Profiles > Access Control List > Access Control List	Monitor > ACL Drops > Layer 3: Fragmented Packet Denied Drops
1	55	Source denied	ACL	Periodic	Denied by an SPP ACL rule.	Protection Profiles > Address Config (IPv4) Protection Profiles > Address Config IPv6 Protection Profiles > Access Control List > Access Control List	Monitor > ACL Drops > Layer 3: Address Denied Monitor > Layer 3 > Address Denied: Denied Address Drops
1	59	Denied: Geo-location	ACL	Periodic	Denied by the global geolocation ACL.	Global Settings > Address > Address Config (IPv4 Only) Global Settings > Access Control List > Access Control List: Select Geolocation and Country	Monitor > ACL Drops > Layer 3: Address Denied Monitor > Layer 3 > Address Denied > Geo Location Drop
1	60	Denied: IP address	ACL	Periodic	Denied by the global ACL.	Global Settings > Address > Address Config (IPv4) Global Settings > Address > Address Config IPv6 Global Settings > Access Control List > Access Control List: Select Deny	Monitor > ACL Drops > Layer 3: Address Denied Monitor > Layer 3 > Address Denied: Denied Address Drops

Code	Sub	Name	Category	Period	Description	Parameter	Graph
1	61	Denied: IP reputation	ACL	Periodic	Denied by the global IP Reputation ACL.	Global Settings > IP Reputation > IP Reputation: Enable Options	Monitor > ACL Drops > Layer 3: Address Denied Monitor > Layer 3 > Address Denied: IP Reputation Denied Drops
1	62	Denied: Local address antispoof	ACL	Periodic	Denied by the global local address antispoofing ACL.	Global Settings > Local Address Config (IPv4) Global Settings > Local Address Config IPv6 Global Settings > Settings > Settings > General Tab: Local Address Anti-Spoofing Options	Monitor > ACL Drops > Layer 3: Address Denied Monitor > Layer 3 > Address Denied: Local Address Antispoof Denied Drops
2	0	SYN flood	Rate flood	Interrupt	Effective rate limit for the syn threshold has been reached.	Protection Profiles > Thresholds > Scalars: SYN Protection Profiles > SPP Settings > SPP Settings > TCP Tab: TCP Session Feature control = SYN Validation	Monitor > Layer 4 > SYN Packets
2	6	State Anomalies: Foreign packet	State anomaly	Periodic	A foreign packet is a TCP packet that does not belong to any known connections. Tracked when the SPP setting foreign-packet is enabled.	Protection Profiles > SPP Settings > SPP Settings > TCP Tab: Options	Monitor > Anomaly Drops > Layer 4 > State: Foreign Packets

Code	Sub	Name	Category	Period	Description	Parameter	Graph
2	7	State Anomalies: Outside window	State anomaly	Periodic	Sequence number of a packet was outside the acceptable window. Tracked when the SPP setting seq-validation is enabled.	Protection Profiles > SPP Settings > SPP Settings > TCP Tab: Options	Monitor > Anomaly Drops > Layer 4 > State: Forward/Reverse Transmission Not Within Window
2	10	TCP SM:Hash attack	-	Periodic	An issue with hash collisions in TCP State Machine table.	None. Internal Table issue. Report to Fortinet.	Monitor >Hash Attack Drops > Layer 4 > TCP Connection Table: TCP Hash Attack Drops
2	11	TCP SM:Out of memory	-	Periodic	An issue with TCP State Machine table internal logic or memory.	None. Internal Table issue. Report to Fortinet.	Monitor >Out of Memory Drops > Layer 4 > TCP Connection Table: TCP Out of Memory Drops
2	12	State Anomalies: State transition error	State anomaly	Periodic	State of the TCP packet received was not consistent with the expected state. Tracked when the SPP setting state-transition-anomalies-validation is enabled.	Protection Profiles > SPP Settings > SPP Settings > TCP Tab: Options	Monitor > Anomaly Drops > Layer 4 > State: TCP State Transition

Code	Sub	Name	Category	Period	Description	Parameter	Graph
2	15	TCP zombie flood	Rate flood	Interrupt	Effective rate limit for the new-connections threshold has been reached. A spike in new connections from IP addresses formerly determined to be legitimate might be a sign of a zombie attack. “Zombies” are systems that are unwitting participants in an attack due to infection from a virus or a worm. Note, this Threshold is normally set to maximum by System Recommendations to avoid rate limiting new connections but New Connections are always shown on the graphs.	Protection Profiles > Thresholds > Scalars: New Connections	Monitor > Layer 4 > New Connections
2	17	TCP port flood	Rate flood	Periodic	Effective rate limit for the port has been reached.	Protection Profiles > Thresholds > TCP Ports	Log & Report > Report Browse > Executive Summary > Top Attacked TCP Ports, to identify specific ports Then: Monitor > Layer 4 > TCP Ports and enter TCP Port numbers of interest to see rate and drop information.

Code	Sub	Name	Category	Period	Description	Parameter	Graph
2	18	UDP port flood	Rate flood	Periodic	Effective rate limit for the port has been reached.	Protection Profiles > Thresholds > UDP Ports	Log & Report > Report Browse > Executive Summary > Top Attacked UDP Ports, to identify specific ports Then: Monitor > Layer 4 > UDP Ports and enter UDP Port numbers of interest to see rate and drop information.
2	19	ICMP flood	Rate flood	Periodic	Effective rate limit for the type/code has been reached.	Protection Profiles > Thresholds > ICMP Types and Codes	Log & Report > Report Browse > Executive Summary > Top Attacked ICMP / Type Code, to identify specific ICMP Types and Codes Then: Monitor > Layer 4 > ICMP Types/Codes and enter ICMP Types/Codes of interest to see rate and drop information.

Code	Sub	Name	Category	Period	Description	Parameter	Graph
2	20	Slow Connection: Aggressive Aging	State anomaly	Periodic	Slow Connection Aggressive Aging	Protection Profiles > SPP Settings > SPP Settings > TCP Tab: Aggressive aging set to Track Slow TCP Connections	Monitor > Anomaly Drops > Layer 4 > State
2	22	Slow Connection: Source flood	Rate flood	Interrupt	Slow connection attack detected and “Source blocking for slow connections” enabled. Source IP address is reported.	Protection Profiles > SPP Settings > SPP Settings > TCP Tab: Source Blocking for Slow Connections Options	Monitor > Layer 4 > Slow Connections
2	24	TCP checksum error	Header anomaly	Periodic	Invalid TCP checksum.	None. Dropped as Anomaly.	Anomaly Drops > Layer 4 > Header: TCP Checksum Error
2	25	UDP checksum error	Header anomaly	Periodic	Invalid UDP checksum.	None. Dropped as Anomaly.	Anomaly Drops > Layer 4 > Header: UDP Checksum Error
2	26	ICMP checksum error	Header anomaly	Periodic	Invalid ICMP checksum.	None. Dropped as Anomaly.	Anomaly Drops > Layer 4 > Header: ICMP Checksum Error
2	27	TCP invalid flag combination	Header anomaly	Periodic	Invalid TCP flag combination. If the urgent flag is set, then the urgent pointer must be non-zero. SYN, FIN or RST is set for fragmented packets.	None. Dropped as Anomaly.	Anomaly Drops > Layer 4 > Header: TCP Invalid Flag Combination

Code	Sub	Name	Category	Period	Description	Parameter	Graph
2	28	L4 anomalies	Header anomaly	Periodic	Drops due to predefined Layer 4 header rules: Data offset is less than 5 for a TCP packet; EOP (End of packet) is detected before the 20 bytes of TCP header; EOP before the data offset indicated data offset; Length field in TCP window scale option is a value other than 3; Length field in TCP window scale option is a value other than 3; Missing UDP payload; Missing ICMP payload ; TCP Option Anamoly based on Option Type	None. Dropped as Anomaly.	Anomaly Drops > Layer 4 > Header: Anomaly Detected
2	52	TCP port denied	ACL	Periodic	Denied by an SPP ACL rule.	Protection Profiles > Service > Service Config: TCP-Port Protection Profiles > Access Control List > Access Control List	Monitor > ACL Drops > Layer 4: TCP Port Denied Drops
2	53	UDP port denied	ACL	Periodic	Denied by an SPP ACL rule.	Protection Profiles > Service > Service Config: UDP-Port Protection Profiles > Access Control List > Access Control List	Monitor > ACL Drops > Layer 4: UDP Port Denied Drops

Code	Sub	Name	Category	Period	Description	Parameter	Graph
2	54	ICMP port denied	ACL	Periodic	Denied by an SPP ACL rule.	Protection Profiles > Service > Service Config: ICMP-Type-Code Protection Profiles > Access Control List > Access Control List	Monitor > ACL Drops > Layer 4: ICMP Type/Code Denied Drops
2	56	SYN flood from source	Rate flood	Interrupt	Effective rate limit for the syn-per-src threshold has been reached. Source IP address is reported.	Protection Profiles > Thresholds > Scalars: SYN-per-Source	Monitor > Layer 4 > SYN per Source
2	58	Excessive Concurrent Connections Per Source	Rate flood	Interrupt	Effective rate limit for the concurrent-connections-per-source threshold has been reached. Source IP address is reported.	Protection Profiles > Thresholds > Scalars: Concurrent-Connections-per-Source	Monitor > Layer 4 > Connection per Source
4	0	HTTP Method Flood	Rate flood	Interrupt	Effective rate limit for a particular HTTP method threshold has been reached.	Protection Profiles > Thresholds > HTTP Methods	Log & Report > Report Browse > Executive Summary > Top Attacked HTTP Methods, to identify specific HTTP Methods Then: Monitor > Layer 7 > HTTP > Methods and select specific Method from the list.

Code	Sub	Name	Category	Period	Description	Parameter	Graph
4	1	Undefined HTTP Method anomaly	Header anomaly	Periodic	Packets dropped due to the unknown-opcode-anomaly rule (Global Settings > Settings page).	None. Dropped as Anomaly.	Monitor > Anomaly Drops > Layer 7 > HTTP header: Unkown Method
4	2	HTTP version anomaly	Header anomaly	Periodic	Packets dropped due to the invalid-opcode-anomaly rule (Global Settings > Settings page).	None. Dropped as Anomaly.	Monitor > Anomaly Drops > Layer 7 > HTTP header: Invalid HTTP version
4	3	URL denied	ACL	Periodic	Denied by an SPP ACL rule.	Protection Profiles > Service > Service Config: URL Protection Profiles > Access Control List > Access Control List	Monitor > ACL Drops > Layer 7 > HTTP: URL Denied Drops
4	4	URL flood	Rate flood	Periodic	Effective rate limit for a particular URL threshold has been reached.	Protection Profiles > Thresholds > URLs NOTE: Learned URLs are hashed into 32,767 possible indexes per SPP. System will only display hash.	Log & Report > Report Browse > Executive Summary > Top Attacked URLs, to identify specific URL hashes. Then: Monitor > Layer 7 > HTTP > URLs and enter specific hash to see rates and drops.
4	5	Invalid HTTP Method anomaly	Header anomaly	Periodic		Global Settings > Settings > General-> HTTP Anamoly	Invalid Opcode Anomaly graph under HTTP Header anomaly graphs

Code	Sub	Name	Category	Period	Description	Parameter	Graph
4	6	HTTP L7 Host Flood	Rate flood	Interrupt	Effective rate limit for a particular Host header threshold has been reached.	Protection Profiles > Thresholds > Hosts Note: Learned Hosts are hashed into 512 possible indexes per SPP. System will only display hashes.	Log & Report > Report Browse > Executive Summary > Top Attacked Hosts, to identify specific Host hashes. Then: Monitor > Layer 7 > HTTP > Hosts and enter specific hash to see rates and drops.
4	7	HTTP L7 Host Deny	ACL	Periodic	Denied by an SPP ACL rule.	Protection Profiles > Service > Service Config: Host Protection Profiles > Access Control List > Access Control List	Monitor > ACL Drops > Layer 7 > HTTP: Host Denied Drops
4	8	HTTP L7 Referer Flood	Rate flood	Interrupt	Effective rate limit for a particular Referer header threshold has been reached.	Protection Profiles > Thresholds > Referer NOTE: Learned Referers are hashed into 512 possible indexes per SPP. System will only display hashes.	Log & Report > Report Browse > Executive Summary > Top Attacked Referers, to identify specific Host hashes. Then: Monitor > Layer 7 > HTTP > Referers and enter specific hash to see rates and drops.

Code	Sub	Name	Category	Period	Description	Parameter	Graph
4	9	HTTP L7 Referer Deny	ACL	Periodic	Denied by an SPP ACL rule.	Protection Profiles > Service > Service Config: Referer Protection Profiles > Access Control List > Access Control List	Monitor > ACL Drops > Layer 7 > HTTP: Referer Denied Drops
4	10	HTTP L7 Cookie Flood	Rate flood	Interrupt	Effective rate limit for a particular Cookie header threshold has been reached.	Protection Profiles > Thresholds > Cookie NOTE: Learned Cookies are hashed into 512 possible indexes per SPP. System will only display hashes.	Log & Report > Report Browse > Executive Summary > Top Attacked Cookies to identify specific Cookie hashes. Then: Monitor > Layer 7 > HTTP > Cookies and enter specific hash to see rates and drops.
4	11	HTTP L7 Cookie Deny	ACL	Periodic	Denied by an SPP ACL rule.	Protection Profiles > Service > Service Config: Cookie Protection Profiles > Access Control List > Access Control List	Monitor > ACL Drops > Layer 7 > HTTP: Cookie Denied Drops

Code	Sub	Name	Category	Period	Description	Parameter	Graph
4	12	HTTP L7 User Agent Flood	Rate flood	Interrupt	Effective rate limit for a particular User-Agent threshold has been reached.	Protection Profiles > Thresholds > User Agent NOTE: Learned User Agents are hashed into 512 possible indexes per SPP. System will only display hashes.	Log & Report > Report Browse > Executive Summary > Top Attacked User Agents to identify specific User Agent hashes. Then: Monitor > Layer 7 > HTTP > User Agents and enter specific hash to see rates and drops.
4	13	HTTP L7 User Agent Deny	ACL	Periodic	Denied by an SPP ACL rule.	Protection Profiles > Service > Service Config: User-Agent Protection Profiles > Access Control List > Access Control List	Monitor > ACL Drops > Layer 7 > HTTP: UserAgent Denied Drops
4	37	DNS Fragment Deny	ACL	Periodic	Denied by an SPP ACL rule.	Protection Profiles > Service > Service Config: DNS-Fragment Protection Profiles > Access Control List > Access Control List	Monitor > ACL Drops > Layer 7 > DNS: Frag Drops
4	38	DNS MX Deny	ACL	Periodic	Denied by an SPP ACL rule.	Protection Profiles > Service > Service Config: DNS-MX Protection Profiles > Access Control List > Access Control List	Monitor > ACL Drops > Layer 7 > DNS: MX Drops

Code	Sub	Name	Category	Period	Description	Parameter	Graph
4	39	Qtype ALL	ACL	Periodic	Denied by an SPP ACL rule.	Protection Profiles > Service > Service Config: DNS-All (ANY) Protection Profiles > Access Control List > Access Control List	Monitor > ACL Drops > Layer 7 > DNS: Qtype All Drops
4	40	Qtype Zone Transfer Deny	ACL	Periodic	Denied by an SPP ACL rule.	Protection Profiles > Service > Service Config: DNS-Zone-Transfer Protection Profiles > Access Control List > Access Control List	Monitor > ACL Drops > Layer 7 > DNS: Qtype Zone Transfer Drops
4	42	DNS Header Anomaly: Invalid Opcode	DNS Anomaly	Periodic	Invalid value in the OpCode field.	Protection Profiles > SPP Settings > Settings > DNS Anomaly Feature Controls tab	Monitor > Anomaly Drops > Layer 7 > DNS > Header
4	43	DNS Header Anomaly: Illegal Flag Combination	DNS Anomaly	Periodic	Invalid combination in the flags field.	Protection Profiles > SPP Settings > Settings > DNS Anomaly Feature Controls tab	Monitor > Anomaly Drops > Layer 7 > DNS > Header
4	44	DNS Header Anomaly: Same Source/Destination Port	DNS Anomaly	Periodic		Protection Profiles > SPP Settings > Settings > DNS Anomaly Feature Controls tab	Monitor > Anomaly Drops > Layer 7 > DNS > Header

Code	Sub	Name	Category	Period	Description	Parameter	Graph
4	45	DNS Query Anomaly: Query Bit Set	DNS Anomaly	Periodic	(QR) bit set to 1	Protection Profiles > SPP Settings > Settings > DNS Anomaly Feature Controls tab	Monitor > Anomaly Drops > Layer 7 > DNS > Query
4	46	DNS Query Anomaly: RA Bit Set	DNS Anomaly	Periodic	recursion allowed (RA) bit set	Protection Profiles > SPP Settings > Settings > DNS Anomaly Feature Controls tab	Monitor > Anomaly Drops > Layer 7 > DNS > Query
4	47	DNS Query Anomaly: Null Query	DNS Anomaly	Periodic	DNS query with count 0	Protection Profiles > SPP Settings > Settings > DNS Anomaly Feature Controls tab	Monitor > Anomaly Drops > Layer 7 > DNS > Query
4	48	DNS Query Anomaly: QD Count not One in query	DNS Anomaly	Periodic	Question count not 1	Protection Profiles > SPP Settings > Settings > DNS Anomaly Feature Controls tab	Monitor > Anomaly Drops > Layer 7 > DNS > Query
4	49	DNS Query Anomaly: RD Bit Set	DNS Anomaly	Periodic	recursion desired (RD) bit set	Protection Profiles > SPP Settings > Settings > DNS Anomaly Feature Controls tab	Monitor > Anomaly Drops > Layer 7 > DNS > Query
4	50	DNS Response Anomaly: QClass in reply	DNS Anomaly	Periodic	DNS response with QCLASS	Protection Profiles > SPP Settings > Settings > DNS Anomaly Feature Controls tab	Monitor > Anomaly Drops > Layer 7 > DNS > Response

Code	Sub	Name	Category	Period	Description	Parameter	Graph
4	51	DNS Response Anomaly: Qtype in reply	DNS Anomaly	Periodic	DNS response with a resource specifying a TYPE ID	Protection Profiles > SPP Settings > Settings > DNS Anomaly Feature Controls tab	Monitor > Anomaly Drops > Layer 7 > DNS > Response
4	52	DNS Response Anomaly: Query bit not set	DNS Anomaly	Periodic	(QR) bit set to 0	Protection Profiles > SPP Settings > Settings > DNS Anomaly Feature Controls tab	Monitor > Anomaly Drops > Layer 7 > DNS > Response
4	53	DNS Response Anomaly: QD count not 1 in response	DNS Anomaly	Periodic	Question count not 1	Protection Profiles > SPP Settings > Settings > DNS Anomaly Feature Controls tab	Monitor > Anomaly Drops > Layer 7 > DNS > Response
4	54	DNS Buffer Overflow Anomaly: Message too long	DNS Anomaly	Periodic	TCP/UDP query or response message that exceeds the maximum header length	Protection Profiles > SPP Settings > Settings > DNS Anomaly Feature Controls tab	Monitor > Anomaly Drops > Layer 7 > DNS > Buffer Overflow
4	55	DNS Buffer Overflow Anomaly: Name too long	DNS Anomaly	Periodic	DNS name that exceeds 255 characters	Protection Profiles > SPP Settings > Settings > DNS Anomaly Feature Controls tab	Monitor > Anomaly Drops > Layer 7 > DNS > Buffer Overflow
4	56	DNS Buffer Overflow Anomaly:Label length too large	DNS Anomaly	Periodic	Query or response with a label that exceeds the maximum length (63)	Protection Profiles > SPP Settings > Settings > DNS Anomaly Feature Controls tab	Monitor > Anomaly Drops > Layer 7 > DNS > Buffer Overflow

Code	Sub	Name	Category	Period	Description	Parameter	Graph
4	57	DNS Exploit Anomaly: Pointer loop	DNS Anomaly	Periodic	DNS message with a pointer that points beyond the end of data	Protection Profiles > SPP Settings > Settings > DNS Anomaly Feature Controls tab	Monitor > Anomaly Drops > Layer 7 > DNS > Exploit
4	58	DNS Exploit Anomaly: Zone Transfer	DNS Anomaly	Periodic	An asynchronous Transfer Full Range (AXFR) request (QTYPE=252)	Protection Profiles > SPP Settings > Settings > DNS Anomaly Feature Controls tab	Monitor > Anomaly Drops > Layer 7 > DNS > Exploit
4	59	DNS Exploit Anomaly: Class is not IN	DNS Anomaly	Periodic	A query/response in which the question/resource address class is not IN	Protection Profiles > SPP Settings > Settings > DNS Anomaly Feature Controls tab	Monitor > Anomaly Drops > Layer 7 > DNS > Exploit
4	60	DNS Exploit Anomaly: Empty UDP message	DNS Anomaly	Periodic		Protection Profiles > SPP Settings > Settings > DNS Anomaly Feature Controls tab	Monitor > Anomaly Drops > Layer 7 > DNS > Exploit
4	61	DNS Exploit Anomaly: Message ends prematurely	DNS Anomaly	Periodic		Protection Profiles > SPP Settings > Settings > DNS Anomaly Feature Controls tab	Monitor > Anomaly Drops > Layer 7 > DNS > Exploit
4	62	DNS Exploit Anomaly: TCP Buffer Underflow	DNS Anomaly	Periodic	A query/response with less than two bytes of data specified in the two-byte prefix field	Protection Profiles > SPP Settings > Settings > DNS Anomaly Feature Controls tab	Monitor > Anomaly Drops > Layer 7 > DNS > Exploit

Code	Sub	Name	Category	Period	Description	Parameter	Graph
4	63	DNS Info Anomaly:DNS type all used	DNS Anomaly	Periodic	DNS request with request type set to ALL (QTYPE=255)	Protection Profiles > SPP Settings > Settings > DNS Anomaly Feature Controls tab	Monitor > Anomaly Drops > Layer 7 > DNS > Info
4	64	DNS Data Anomaly: Invalid type class	DNS Anomaly	Periodic	A query/response with TYPE or CLASS reserved values	Protection Profiles > SPP Settings > Settings > DNS Anomaly Feature Controls tab	Monitor > Anomaly Drops > Layer 7 > DNS > Data
4	65	DNS Data Anomaly: Extraneous data	DNS Anomaly	Periodic	A query/response with excess data	Protection Profiles > SPP Settings > Settings > DNS Anomaly Feature Controls tab	Monitor > Anomaly Drops > Layer 7 > DNS > Data
4	66	DNS Data Anomaly: TTL too long	DNS Anomaly	Periodic	TTL value is greater than 7 days	Protection Profiles > SPP Settings > Settings > DNS Anomaly Feature Controls tab	Monitor > Anomaly Drops > Layer 7 > DNS > Data
4	67	DNS Data Anomaly: Name length too short	DNS Anomaly	Periodic	A query/response with a null DNS name	Protection Profiles > SPP Settings > Settings > DNS Anomaly Feature Controls tab	Monitor > Anomaly Drops > Layer 7 > DNS > Data
4	68	DNS UDP Unsolicited Response	Rate flood	Periodic	UDP Drops due to a response with no matching query.	Parameter	Monitor > Layer 7 > DNS > Unsolicited Response

Cod-e	Su-b	Name	Cat-egory	Period	Description	Parameter	Graph
4	69	DNS TCP Unsolicited Response	Rate flood	Periodic	TCP Drops due to a response with no matching query.		Monitor > Layer 7 > DNS > Unsolicited Response
4	70	DNS DQRM Horizontal Link Limit Crossed	-	Periodic		None. Internal Table issue. Report to Fortinet.	Monitor > Out of Memory Drops > Layer 7 > DNS Query Response Table
4	71	DNS DQRM Out of Memory	-	Periodic	An issue with DQRM table internal logic or memory.	None. Internal Table issue. Report to Fortinet.	Monitor > Out of Memory Drops > Layer 7 > DNS Query Response Table
4	72	DNS UDP Response same direction	Rate flood	Periodic	UDP drops due to response sent to port 53.	None.	Monitor >Layer 7 >DNS > Unsolicited Response: Unsolicited UDP Resoponse Same Port
4	73	DNS TCP Response same direction	Rate flood	Periodic	TCP drops due to response sent to port 53.	None.	Monitor >Layer 7 >DNS > Unsolicited Response: Unsolicited TCP Resoponse Same Port
4	74	DNS LQ: UDP Query flood	Rate flood	Periodic	UDP drops due to LQ check during flood.	Protection Profiles > SPP Settings > Settings > DNS Feature Controls tab: Allow Only Legitimate Queries under Flood	Monitor > Layer 7 > DNS > LQ Drops

Code	Sub	Name	Category	Period	Description	Parameter	Graph
4	75	DNS LQ: UDP Question flood	Rate flood	Periodic	UDP drops due to LQ check during flood.	Protection Profiles > SPP Settings > Settings > DNS Feature Controls tab: Allow Only Legitimate Queries under Flood	Monitor > Layer 7 > DNS > LQ Drops
4	76	DNS LQ: UDP Qtype All flood	Rate flood	Periodic	UDP drops due to LQ check during flood.	Protection Profiles > SPP Settings > Settings > DNS Feature Controls tab: Allow Only Legitimate Queries under Flood	Monitor > Layer 7 > DNS > LQ Drops
4	77	DNS LQ: UDP Qtype Zone Transfer flood	Rate flood	Periodic	UDP drops due to LQ check during flood.	Protection Profiles > SPP Settings > Settings > DNS Feature Controls tab: Allow Only Legitimate Queries under Flood	Monitor > Layer 7 > DNS > LQ Drops
4	78	DNS LQ: UDP Qtype MX flood	Rate flood	Periodic	UDP drops due to LQ check during flood.	Protection Profiles > SPP Settings > Settings > DNS Feature Controls tab: Allow Only Legitimate Queries under Flood	Monitor > Layer 7 > DNS > LQ Drops

Code	Sub	Name	Category	Period	Description	Parameter	Graph
4	79	DNS LQ: UDP Fragment flood	Rate flood	Periodic	UDP drops due to LQ check during flood.	Protection Profiles > SPP Settings > Settings > DNS Feature Controls tab: Allow Only Legitimate Queries under Flood	Monitor > Layer 7 > DNS > LQ Drops
4	80	DNS LQ: UDP Query flood due to Negative Response	Rate flood	Periodic	UDP drops due to LQ check during flood.	Protection Profiles > SPP Settings > Settings > DNS Feature Controls tab: Allow Only Legitimate Queries under Flood	Monitor > Layer 7 > DNS > LQ Drops
4	81	DNS TTL: UDP Query flood	Rate flood	Periodic	UDP drops due to TTL check during flood.	Protection Profiles > SPP Settings > Settings > DNS Feature Controls tab: Validate TTL For Queries From The Same IP	Monitor > Layer 7 > DNS > TTL
4	82	DNS TTL: UDP Question flood	Rate flood	Periodic	UDP drops due to TTL check during flood.	Protection Profiles > SPP Settings > Settings > DNS Feature Controls tab: Validate TTL For Queries From The Same IP	Monitor > Layer 7 > DNS > TTL

Code	Sub	Name	Category	Period	Description	Parameter	Graph
4	83	DNS TTL: UDP Qtype All flood	Rate flood	Periodic	UDP drops due to TTL check during flood.	Protection Profiles > SPP Settings > Settings > DNS Feature Controls tab: Validate TTL For Queries From The Same IP	Monitor > Layer 7 > DNS > TTL
4	84	DNS TTL: UDP Qtype Zone Transfer flood	Rate flood	Periodic	UDP drops due to TTL check during flood.	Protection Profiles > SPP Settings > Settings > DNS Feature Controls tab: Validate TTL For Queries From The Same IP	Monitor > Layer 7 > DNS > TTL
4	85	DNS TTL: UDP Qtype MX flood	Rate flood	Periodic	UDP drops due to TTL check during flood.	Protection Profiles > SPP Settings > Settings > DNS Feature Controls tab: Validate TTL For Queries From The Same IP	Monitor > Layer 7 > DNS > TTL
4	86	DNS TTL: UDP Fragment flood	Rate flood	Periodic	UDP drops due to TTL check during flood.	Protection Profiles > SPP Settings > Settings > DNS Feature Controls tab: Validate TTL For Queries From The Same IP	Monitor > Layer 7 > DNS > TTL
4	87	DNS Spoofed IP: UDP Query Flood drop during TC=1 check	Rate flood	Periodic	UDP drops due to TC=1 antispoofing check during flood	Same as Source flood above	Monitor > Layer 7 > DNS > Spoofed IP
4	88	DNS Spoofed IP: UDP Question Flood drop during TC=1 check	Rate flood	Periodic	UDP drops due to TC=1 antispoofing check during flood	Same as Destination flood above	Monitor > Layer 7 > DNS > Spoofed IP

Code	Sub	Name	Category	Period	Description	Parameter	Graph
4	89	DNS Spoofed IP: UDP Qtype All Flood drop during TC=1 check	Rate flood	Periodic	UDP drops due to TC=1 antispoofing check during flood	Global Settings > Address > Address Config (IPv4 Only) Global Settings > Access Control List > Access Control List: Select Geolocation and Country	Monitor > Layer 7 > DNS > Spoofed IP
4	90	DNS Spoofed IP: UDP Qtype Zone Transfer Flood drop during TC=1 check	Rate flood	Periodic	UDP drops due to TC=1 antispoofing check during flood	Global Settings > Address > Address Config (IPv4) Global Settings > Address > Address Config IPv6 Global Settings > Access Control List > Access Control List: Select Deny	Monitor > Layer 7 > DNS > Spoofed IP
4	91	DNS Spoofed IP: UDP Qtype MX Flood drop during TC=1 check	Rate flood	Periodic	UDP drops due to TC=1 antispoofing check during flood	Global Settings > IP Reputation > IP Reputation: Enable Options	Monitor > Layer 7 > DNS > Spoofed IP
4	92	DNS Spoofed IP: UDP Fragment Flood drop during TC=1 check	Rate flood	Periodic	UDP drops due to TC=1 antispoofing check during flood	Global Settings > Local Address Config (IPv4) Global Settings > Local Address Config IPv6 Global Settings > Settings > Settings > General Tab: Local Address Anti-Spoofing Options	Monitor > Layer 7 > DNS > Spoofed IP

Code	Sub	Name	Category	Period	Description	Parameter	Graph
4	93	DNS Spoofed IP: UDP Query Flood Drop during Retransmission Check	Rate flood	Periodic	UDP drops due to Retransmission antispoofing check during flood	Protection Profiles > Thresholds > Scalars: SYN Protection Profiles > SPP Settings > SPP Settings > TCP Tab: TCP Session Feature control = SYN Validation	Monitor > Layer 7 > DNS > Spoofed IP
4	94	DNS Spoofed IP: UDP Question Flood Drop during Retransmission Check	Rate flood	Periodic	UDP drops due to Retransmission antispoofing check during flood	Protection Profiles > SPP Settings > SPP Settings > TCP Tab: Options	Monitor > Layer 7 > DNS > Spoofed IP
4	95	DNS Spoofed IP: UDP Qtype All Flood Drop during Retransmission Check	Rate flood	Periodic	UDP drops due to Retransmission antispoofing check during flood	Protection Profiles > SPP Settings > SPP Settings > TCP Tab: Options	Monitor > Layer 7 > DNS > Spoofed IP
4	96	DNS Spoofed IP: UDP Qtype Zone Transfer Flood Drop during Retransmission Check	Rate flood	Periodic	UDP drops due to Retransmission antispoofing check during flood	None. Internal Table issue. Report to Fortinet.	Monitor > Layer 7 > DNS > Spoofed IP
4	97	DNS Spoofed IP: UDP Qtype MX Flood Drop during Retransmission Check	Rate flood	Periodic	UDP drops due to Retransmission antispoofing check during flood	None. Internal Table issue. Report to Fortinet.	Monitor > Layer 7 > DNS > Spoofed IP
4	98	DNS Spoofed IP: UDP Fragment Flood Drop during Retransmission Check	Rate flood	Periodic	UDP drops due to Retransmission antispoofing check during flood	Protection Profiles > SPP Settings > SPP Settings > TCP Tab: Options	Monitor > Layer 7 > DNS > Spoofed IP

Code	Sub	Name	Category	Period	Description	Parameter	Graph
4	99	DNS Cache: UDP Query Flood Drop Due To Response From Cache	Rate flood	Periodic	UDP drops because the response was served from the cache during a flood.	Protection Profiles > SPP Settings > Settings > DNS Feature Controls tab: Generate Response From Cache Under Flood	Monitor > Layer 7 > DNS > Cache
4	100	DNS Cache: UDP Question Flood Drop Due To Response From Cache	Rate flood	Periodic	UDP drops because the response was not served from the cache during a flood.	Protection Profiles > SPP Settings > Settings > DNS Feature Controls tab: Generate Response From Cache Under Flood	Monitor > Layer 7 > DNS > Cache
4	101	DNS Cache: UDP Qtype All Flood Drop Due To Response From Cache	Rate flood	Periodic	UDP drops because the response was not served from the cache during a flood.	Protection Profiles > SPP Settings > Settings > DNS Feature Controls tab: Generate Response From Cache Under Flood	Monitor > Layer 7 > DNS > Cache
4	102	DNS Cache: UDP Qtype Zone Transfer Flood Drop Due To Response From Cache	Rate flood	Periodic	UDP drops because the response was not served from the cache during a flood.	Protection Profiles > SPP Settings > Settings > DNS Feature Controls tab: Generate Response From Cache Under Flood	Monitor > Layer 7 > DNS > Cache

Code	Sub	Name	Category	Period	Description	Parameter	Graph
4	10-3	DNS Cache: UDP Qtype MX Flood Drop Due To Response From Cache	Rate flood	Periodic	UDP drops because the response was not served from the cache during a flood.	Protection Profiles > SPP Settings > Settings > DNS Feature Controls tab: Generate Response From Cache Under Flood	Monitor > Layer 7 > DNS > Cache
4	10-4	DNS Cache: UDP Fragment Flood Drop Due To Response From Cache	Rate flood	Periodic	UDP drops because the response was not served from the cache during a flood.	Protection Profiles > SPP Settings > Settings > DNS Feature Controls tab: Generate Response From Cache Under Flood	Monitor > Layer 7 > DNS > Cache
4	10-5	DNS Cache: UDP Query Flood Drop Due To No Response From Cache	Rate flood	Periodic	UDP drops because the response was not served from the cache during a flood.	Protection Profiles > SPP Settings > Settings > DNS Feature Controls tab: Generate Response From Cache Under Flood	Monitor > Layer 7 > DNS > Cache
4	10-6	DNS Cache: UDP Question Flood Drop Due To No Response From Cache	Rate flood	Periodic	UDP drops because the response was not served from the cache during a flood.	Protection Profiles > SPP Settings > Settings > DNS Feature Controls tab: Generate Response From Cache Under Flood	Monitor > Layer 7 > DNS > Cache

Code	Sub	Name	Category	Period	Description	Parameter	Graph
4	107	DNS Cache: UDP Qtype All Flood Drop Due To No Response From Cache	Rate flood	Periodic	UDP drops because the response was not served from the cache during a flood.	Protection Profiles > SPP Settings > Settings > DNS Feature Controls tab: Generate Response From Cache Under Flood	Monitor > Layer 7 > DNS > Cache
4	108	DNS Cache: UDP Qtype Zone Transfer Flood Drop Due To No Response From Cache	Rate flood	Periodic	UDP drops because the response was not served from the cache during a flood.	Protection Profiles > SPP Settings > Settings > DNS Feature Controls tab: Generate Response From Cache Under Flood	Monitor > Layer 7 > DNS > Cache
4	109	DNS Cache: UDP Qtype MX Flood Drop Due To No Response From Cache	Rate flood	Periodic	UDP drops because the response was not served from the cache during a flood.	Protection Profiles > SPP Settings > Settings > DNS Feature Controls tab: Generate Response From Cache Under Flood	Monitor > Layer 7 > DNS > Cache
4	110	DNS Cache: UDP Fragment Flood Drop Due To No Response From Cache	Rate flood	Periodic	UDP drops because the response was not served from the cache during a flood.	Protection Profiles > SPP Settings > Settings > DNS Feature Controls tab: Generate Response From Cache Under Flood	Monitor > Layer 7 > DNS > Cache

Code	Sub	Name	Category	Period	Description	Parameter	Graph
4	20-1	HTTP Header Range Present Anomaly	Header anomaly	Periodic	Drops due to HTTP Range Header rule (Global Settings > Settings page).	Global Settings > Settings > General tab: Drop HTTP Range Header Option	Monitor > Anomaly Drops > Layer 7 > HTTP Header: Rnge Present
4	11-1	DNS TCP Query Flood	Rate flood	Interrupt	Effective rate limit for the dns-query threshold has been reached.	Protection Profiles > Thresholds > Scalars: DNS-Query	Monitor > Layer 7 > DNS > Query: TCP Query Dropped
4	11-2	DNS TCP Question Flood	Rate flood	Interrupt	Effective rate limit for the dns-question-count threshold has been reached.	Protection Profiles > Thresholds > Scalars: DNS-Question-Count	Monitor > Layer 7 > DNS > Question Count: TCP Question Count Dropped
4	11-3	DNS TCP Fragment Flood	Rate flood	Interrupt	Effective rate limit for the dns-fragment threshold has been reached.	Protection Profiles > Thresholds > Scalars: DNS-Fragment	Monitor > Layer 7 > DNS > Fragment
4	11-4	DNS TCP Zone Transfer Flood	Rate flood	Interrupt	Effective rate limit for the dns-zone-xfer threshold has been reached.	Protection Profiles > Thresholds > Scalars: DNS-Zone-Transfer	Monitor > Layer 7 > DNS > Qtype Zone Transfer
4	11-5	DNS TCP MX Flood	Rate flood	Interrupt	Effective rate limit for the dns-mx threshold has been reached.	Protection Profiles > Thresholds > Scalars: DNS-MX-Count	Monitor > Layer 7 > DNS > Qtype MX
4	11-6	DNS TCP All Flood	Rate flood	Interrupt	Effective rate limit for the dns-all threshold has been reached.	Protection Profiles > Thresholds > Scalars: DNS-All	Monitor > Layer 7 > DNS > Qtype All

Code	Sub	Name	Category	Period	Description	Parameter	Graph
4	117	DNS UDP Duplicate Query before Response	Rate flood	Periodic	UDP Drops due to DQRM duplicate query check.	Protection Profiles > SPP Settings > Settings > DNS Feature Controls tab: Duplicate Query Check Before Response	Monitor > Layer 7 > DNS > Unexpected Query: UDP Duplicate Query before Response Drop
4	118	DNS TCP Duplicate Query before Response	Rate flood	Periodic	TCP Drops due to DQRM duplicate query check.		Monitor > Layer 7 > DNS > Unexpected Query: TCP Duplicate Query before Response Drop
4	119	DNS Query Restricted to Specific Subnet	ACL	Periodic	DNS Query ACL drops due to Query restricted to specific subnets	Protection Profiles > Address Config > IP Address/Netmask Protection Profiles > Address Config IPv6 > IP Address/Netmask Protection Profiles > Access Control List > Address/Address IPv6	Monitor > ACL Drops > Layer 7 > DNS
4	120	DNS Query Drop due to ACL	ACL	Periodic	DNS Query ACL Drops due to Blacklisted domains/ipv4 address	Global Settings > Blacklisted Domains Global Settings > Blacklisted IPv4 Addresses	Monitor > ACL Drops > Layer 7 > DNS
2	82	DNS Query flood from Source	Rate flood	Periodic		Protection Profiles > Thresholds > Scalars: DNS-Query-per-Source	Monitor > Layer 7 > DNS > Query per Source

Code	Sub	Name	Category	Period	Description	Parameter	Graph
2	83	DNS Packet Track Flood from Source	Rate flood	Periodic		Protection Profiles > Thresholds > Scalars: DNS-Packet-Track-per-Srce	Monitor > Layer 7 > DNS > Suspious Sources
2	86	Invalid ICMP Type/Code	Header Anamoly	Periodic		Global Settings > Settings > General tab	Monitor > Anamoly Drops > Layer 4 > Header
2	87	HTTP Method flood from source	Rate flood	Interrupt		Protection Profiles > Thresholds > Scalars	Monitor > Flood Drops > Layer 7 > HTTP

Table 90: DDoS Attack log Directionality for TCP

Setup	Traffic Direction	Source	Destination	Source Port	Destination Port	Attack Log Direction	Protected IP
SYN	Outbound	Inside	Outside	High	Low	Outbound	Inside
ACK	Inbound	Outside	Inside	Low	High	Outbound	Inside
SYN	Inbound	Outside	Inside	High	Low	Inbound	Inside
ACK	Outbound	Inside	Outside	Low	High	Inbound	Inside
SYN	Outbound	Inside	Outside	High	High	Outbound	Inside
ACK	Inbound	Outside	Inside	High	High	Outbound	Inside
SYN	Inbound	Outside	Inside	High	High	Inbound	Inside
ACK	Outbound	Inside	Outside	High	High	Inbound	Inside

Setup	Traffic Direction	Source	Destination	Source Port	Destination Port	Attack Log Direction	Protected IP
SYN	Outbound	Inside	Outside	Low	Low	Outbound	Inside
ACK	Inbound	Outside	Inside	Low	Low	Outbound	Inside
SYN	Inbound	Outside	Inside	Low	Low	Inbound	Inside
ACK	Outbound	Inside	Outside	Low	Low	Inbound	Inside

Table 91: DDoS Attack log Directionality for UDP

Traffic Direction	Source	Destination	Source Port	Destination Port	Attack Log Direction	Protected IP
Outbound	Inside	Outside	High	Low	Outbound	Inside
Inbound	Outside	Inside	Low	High	Inbound	Inside
Inbound	Outside	Inside	High	Low	Inbound	Inside
Outbound	Inside	Outside	Low	High	Outbound	Inside
Outbound	Inside	Outside	High	High	Outbound	Inside
Inbound	Outside	Inside	High	High	Inbound	Inside

Appendix B: Management Information Base (MIB)

The FortiDDoS SNMP agent supports a few management information blocks (MIBs).

Table 92: Supported MIBs

MIB or RFC	Description
Fortinet Core MIB	This Fortinet-proprietary MIB enables your SNMP manager to query for system information and to receive traps that are common to multiple Fortinet devices.
FortiDDoS MIB	This Fortinet-proprietary MIB enables your SNMP manager to query for FortiDDoS-specific information and to receive FortiDDoS-specific traps.
RFC 1213 (MIB II)	The FortiDDoS SNMP agent supports MIB II groups, except: • There is no support for the EGP group from MIB II (RFC 1213, section 3.11 and 6.10). • Protocol statistics returned for MIB II groups (IP, ICMP, TCP, UDP, and so on) do not accurately capture all FortiDDoS traffic activity. More accurate information can be obtained from the information reported by the FortiDDoS MIB.
RFC 2665 (Ethernet-like MIB)	The FortiDDoS SNMP agent supports “Ethernet-like MIB information,” except the dot3Tests and dot3Errors groups.
RFC 2863 (IF-MIB)	FortiDDoS SNMP uses the linkDown and linkUP traps from IF-MIB, RFC 2863, section 6.
RFC 3411	“An Architecture for Describing Simple Network Management Protocol (SNMP) Management Frameworks”
RFC 3414	Partial support for “User-based Security Model (USM).”

To communicate with the FortiDDoS SNMP agent, you must first compile these MIBs into your SNMP manager. If the standard MIBs used by the SNMP agent are already compiled into your SNMP manager, you do not have to compile them again. The FortiDDoS SNMP implementation is read-only.

To view a trap or query’s name, object identifier (OID), and description, open its MIB file in a plain text editor.

All traps sent include the message, the FortiDDoS appliance’s serial number, and hostname.

You can obtain the Fortinet MIB files from the Fortinet Service & Support website in the same section where you download firmware images. Go to https://support.fortinet.com/.

Figure 173: Download MIB files from the Fortinet Service & Support website

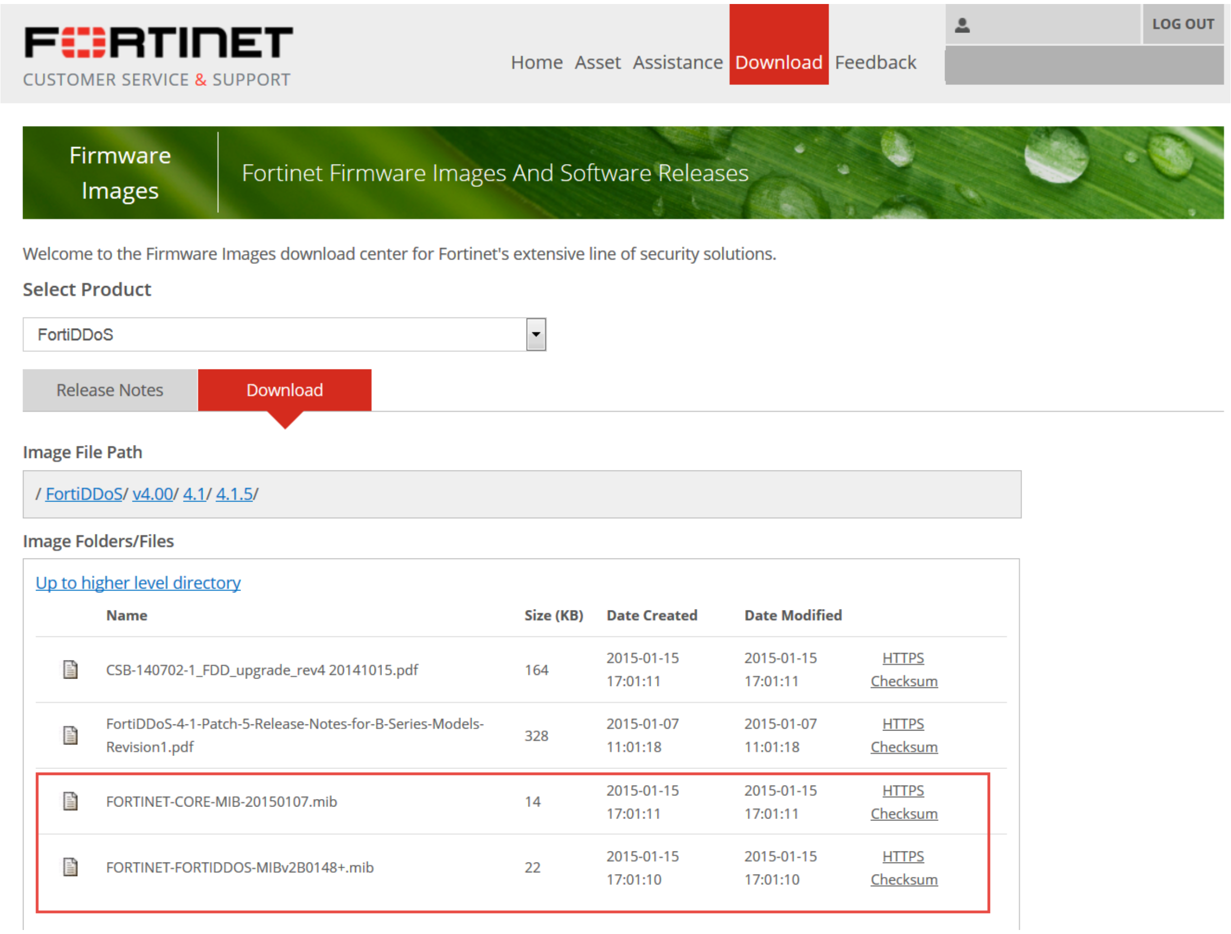

FORTINET
CUSTOMER SERVICE & SUPPORT
Home
Asset
Assistance
Download
Feedback
LOG OUT
Firmware Images
Fortinet Firmware Images And Software Releases

Welcome to the Firmware Images download center for Fortinet's extensive line of security solutions.

Select Product

FortiDDoS

Image File Path

/ FortiDDoS/ v4.00/ 4.1/ 4.1.5/

Image Folders/Files

Up to higher level directory

Name	Size (KB)	Date Created	Date Modified	
CSB-140702-1_FDD_upgrade_rev4 20141015.pdf	164	2015-01-15 17:01:11	2015-01-15 17:01:11	HTTPS Checksum
FortiDDoS-4-1-Patch-5-Release-Notes-for-B-Series-Models-Revision1.pdf	328	2015-01-07 11:01:18	2015-01-07 11:01:18	HTTPS Checksum
FORTINET-CORE-MIB-20150107.mib	14	2015-01-15 17:01:11	2015-01-15 17:01:11	HTTPS Checksum
FORTINET-FORTIDDOS-MIBv2B0148+.mib	22	2015-01-15 17:01:10	2015-01-15 17:01:10	HTTPS Checksum

Appendix C: Port Numbers

Communications between the FortiDDoS appliance, clients, servers, and FortiGuard Distribution Network (FDN) require that any routers and firewalls between them permit specific protocols and port numbers.

The following tables list the default port assignments used by FortiDDoS.

Table 93: Default ports used by FortiDDoS for incoming traffic (listening)

Port Number	Protocol / Service	Purpose
N/A	ICMP	`ping` and `traceroute` responses.
22	TCP	SSH administrative CLI access.
23	TCP	Telnet administrative CLI access.
80	TCP	HTTP administrative web UI access.
161	UDP	SNMP queries.
443	TCP	HTTPS administrative web UI access. FortiDDoS REST API. Cloud Signaling REST API.
3306	TCP	SQL queries.
6055	UDP	HA heartbeat. Multicast.
6056	UDP	HA configuration synchronization. Multicast.

Table 94: Default ports used by FortiDDoS for outgoing traffic

Port Number	Protocol / Service	Purpose
20, 21	TCP	FTP client.
25	TCP	SMTP for alert email.
53	UDP	DNS client.
69	UDP	TFTP client for backups, restoration, and firmware updates. See commands such as `execute backup` or `execute restore`.

Port Number	Protocol / Service	Purpose
123	UDP	NTP client.
162	UDP	SNMP traps.
389	TCP	LDAP authentication.
443	TCP	FortiGuard polling and update downloads. FortiDDoS REST API. Cloud Signaling REST API.
514	UDP	Syslog.
1812	TCP	RADIUS authentication.
6055	UDP	HA heartbeat. Multicast.
6056	UDP	HA configuration synchronization. Multicast.

Appendix D: Switch and Router Configuration

Switch configuration for load balancing

The following example load balancing configuration is for the FortiSwitch 248-B DPS Ethernet switch.

It configures two trunk groups with eight ports per trunk. Trunk 10 is used for Internet traffic and trunk 11 is used for server-side traffic.

You use the `load-balance-hash` command to specify `src-dst-ip-ipports` as the hash distribution algorithm (hash mode) to apply to all trunk groups. This mode uses a 4-tuple (source and destination IP address and source IP L4 port and destination IP L4 port) to ensure that all packets belonging to a session pass through the same port pair on FortiDDoS appliance in both directions.

```
(clientSide-84.82) #show run
!Current Configuration:
!
!System Description "FortiSwitch-248B-DPS 48x1G & 4x10G"
!System Software Version "5.2.0.2.4"

serviceport ip 192.168.22.98 255.255.255.0 0.0.0.0
vlan database
vlan name 10 "egress"
vlan name 11 "ingress"
exit

port-channel "egress" 1
interface 0/1
channel-group 1/1
exit
interface 0/3
channel-group 1/1
exit
interface 0/5
channel-group 1/1
exit
interface 0/7
channel-group 1/1
exit
interface 0/9
channel-group 1/1
exit
interface 0/11
channel-group 1/1
exit
interface 0/13
channel-group 1/1
exit
interface 0/15
channel-group 1/1
exit
port-channel "ingress" 2
```

```
interface 0/2
channel-group 1/2
exit
interface 0/4
channel-group 1/2
exit
interface 0/6
channel-group 1/2
exit
interface 0/8
channel-group 1/2
exit
interface 0/10
channel-group 1/2
exit
interface 0/12
channel-group 1/2
exit
interface 0/14
channel-group 1/2
exit
interface 0/16
channel-group 1/2
exit

mac-addr-table aging-time 60000

interface 0/1
no cdp run
switchport allowed vlan add 10
exit

interface 0/2
no cdp run
exit
interface 0/3
no cdp run
exit

interface 0/4
no cdp run
exit

interface 0/5
no cdp run
exit

interface 0/6
no cdp run
exit

interface 0/7
no cdp run
exit

interface 0/8
no cdp run
```

```
exit

interface 0/9
no cdp run
exit

interface 0/10
no cdp run
exit

interface 0/11
no cdp run
exit

interface 0/12
no cdp run
exit

interface 0/13
no cdp run
exit

interface 0/14
no cdp run
exit

interface 0/15
no cdp run
exit

interface 0/16
no cdp run
exit

interface 0/17
no cdp run
switchport allowed vlan add 10
switchport native vlan 10
exit

interface 0/18
no cdp run
switchport allowed vlan add 11
switchport native vlan 11
exit

interface 0/49
no cdp run
switchport allowed vlan add 10
switchport native vlan 10
exit

interface 0/50
no cdp run
switchport allowed vlan add 11
switchport native vlan 11
exit
```

```
interface 1/1
staticcapability
switchport allowed vlan add 10
switchport native vlan 10
lacp collector max-delay 0
exit

interface 1/2
staticcapability
switchport allowed vlan add 11
switchport native vlan 11
lacp collector max-delay 0
exit

interface 1/3
staticcapability
switchport allowed vlan add 10
switchport tagging 10
lacp collector max-delay 0
exit

interface 1/4
staticcapability
switchport allowed vlan add 11
switchport tagging 11
lacp collector max-delay 0
exit

router rip
exit
router ospf
exit
exit

(clientSide-84.82) #
(clientSide-84.82) #show load-balance
Hash Mode: src-dst-ip-ipport
```

Configuring the routers & switch for traffic diversion

The following example router and switch configuration is used for traffic diversion.

Router configuration

```
vlan internal allocation policy ascending
!
interface GigabitEthernet1/0/1
   no switchport
   no ip address
!
interface GigabitEthernet1/0/2
   switchport access vlan 3
   switchport trunk encapsulation dot1q
!
interface GigabitEthernet1/0/3
```

```
!
interface GigabitEthernet1/0/4
!
interface GigabitEthernet1/0/5
!
interface GigabitEthernet1/0/6
!
interface GigabitEthernet1/0/7
!
interface GigabitEthernet1/0/8
!
interface GigabitEthernet1/0/9
!
interface GigabitEthernet1/0/10
  switchport access vlan 2
!
interface GigabitEthernet1/0/11
ip address 10.100.0.250 255.255.255.0
no ip directed-broadcast
ip policy route-map FDD-X00A-PBR
!
interface GigabitEthernet1/0/12
!
interface Vlan2
  ip address 10.1.0.251 255.255.255.0
!
interface Vlan3
  ip address 192.168.100.51 255.255.255.0
!
!
ip classless
ip route 207.117.1.0 255.255.255.0 10.1.0.250
!
!
ip access-list extended zone-A
permit ip any 207.117.0.0 0.0.0.255
!
route-map FDD-X00A-PBR permit 100
match ip address zone-A
set ip next-hop 10.200.0.254
!
route-map FDD-X00A-PBR permit 101
description let thru all other packets without modifying next-hop
```

Switch configuration

```
interface GigabitEthernet1/0/1
  no switchport
  no ip address
  channel-group 1 mode on
!
interface GigabitEthernet1/0/2
  switchport access vlan 3
  switchport trunk encapsulation dot1q
!
interface GigabitEthernet1/0/3
  switchport access vlan 3
```

FORTINET®

High Performance Network Security