FORTINET®

FortiDDoS Web-based Manager v3.2

Reference Guide

March 27, 2013

28-320-183684-20130327

Copyright © 2013 Fortinet, Inc. All rights reserved. Fortinet®, FortiGate®, and FortiGuard®, are registered trademarks of Fortinet, Inc., and other Fortinet names herein may also be trademarks of Fortinet. All other product or company names may be trademarks of their respective owners. Performance metrics contained herein were attained in internal lab tests under ideal conditions, and performance may vary. Network variables, different network environments and other conditions may affect performance results. Nothing herein represents any binding commitment by Fortinet, and Fortinet disclaims all warranties, whether express or implied, except to the extent Fortinet enters a binding written contract, signed by Fortinet's General Counsel, with a purchaser that expressly warrants that the identified product will perform according to the performance metrics herein. For absolute clarity, any such warranty will be limited to performance in the same ideal conditions as in Fortinet's internal lab tests. Fortinet disclaims in full any guarantees. Fortinet reserves the right to change, modify, transfer, or otherwise revise this publication without notice, and the most current version of the publication shall be applicable.

Technical Documentation:	http://docs.fortinet.com
Knowledge Base:	http://kb.fortinet.com
Customer Service & Support:	https://support.fortinet.com
Training Services:	http://training.fortinet.com
Document feedback:	techdoc@fortinet.com

Table of Contents

Authentication and Management

User authentication & role-based management in FortiDDoS

FortiDDoS protects itself from unauthorized modifications and access to its functions and data, and allows authorized users to access only appropriate functions and data. For this purpose, FortiDDoS is able to identify and authenticate users prior to allowing access to its functions and data. FortiDDoS supports three predefined roles: a Super User role, an administrator role and an operator role, and includes a set of functions that allow effective management of the device's functions and data. You can add, modify and delete the roles and users depending on your specific needs.

To access the Web-based Manager of FortiDDoS you need to first authenticate yourself.

Authentication steps

1. The FortiDDoS User Management panel is password protected. Access the FortiDDoS using the IP address provided by your network administrator.
2. Use https://IP Address, e.g. https://192.168.3.210
3. You will see a security alert warning about the certificate, similar to shown in Figure 1. The warning will appear differently in different web browsers.
4. Please ignore the security alert by pressing *Yes* and continue. If you would like to avoid the security alert in future.

Figure 1: Security alert while authenticating

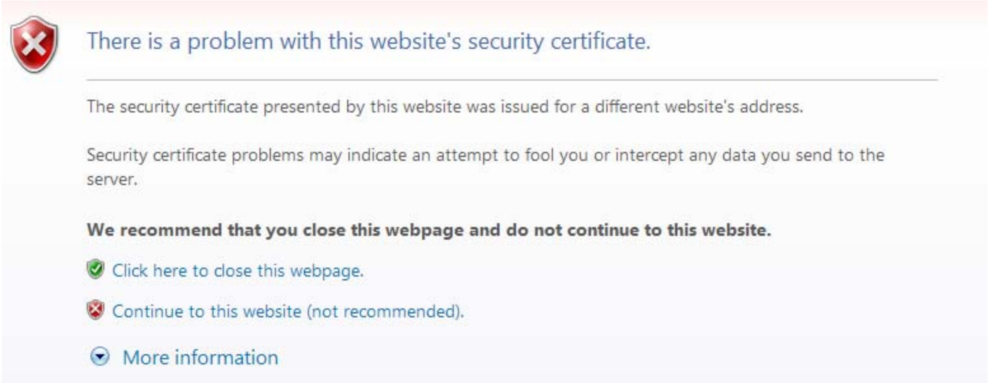

There is a problem with this website's security certificate.

The security certificate presented by this website was issued for a different website's address.

Security certificate problems may indicate an attempt to fool you or intercept any data you send to the server.

We recommend that you close this webpage and do not continue to this website.

Click here to close this webpage.
Continue to this website (not recommended).
More information

5. You will next see an authentication dialog shown in Figure 2. Enter your username and password and press *OK*.

Figure 2: User authentication

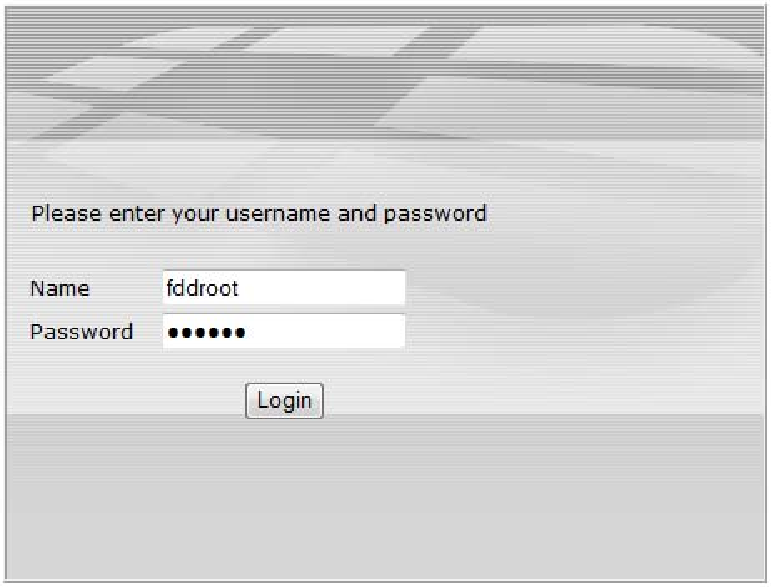

Please enter your username and password

Name fddroot

Password ••••••

6 You will then see the welcome screen of the Web-based Manager. At this stage you can access menu items based on your privilege decided by the administrator.

Note: Once you have logged in, your session will be valid until you are inactive for a 2 hour period. After this time, your session will timeout. You must authenticate again to gain access.

Some notes about role-based user management

- FortiDDoS is shipped with a known Super User *fddroot* with a default password.
- A Super User is one who has access to all menu items.
- The Super User is allowed to create new users with the super user's choice of password.
- The Super User can define roles. The roles include the list of function privileges for a given role. FortiDDoS ships the appliances with three predefined roles: the Super User, the Administrator and the Operator.
- Ensure that only one Super User account is created.
- The Super User can define a set of VID groups. The VID group includes a list of VIDs belonging to that group.
- A VID administrator is one who has access to the configuration menus for a particular VID group.
- All users can change their own passwords.
- Every user is associated with a VID group and has a specific role.
- It is the user's responsibility to avoid simultaneous configuration.

Recommended roles

The following 3 roles are recommended in a typical FortiDDoS installation:

- Super User

  This user should have access to all functions in every VID. Figure 3 shows the recommended role.

Figure 3: Super user’s recommended role

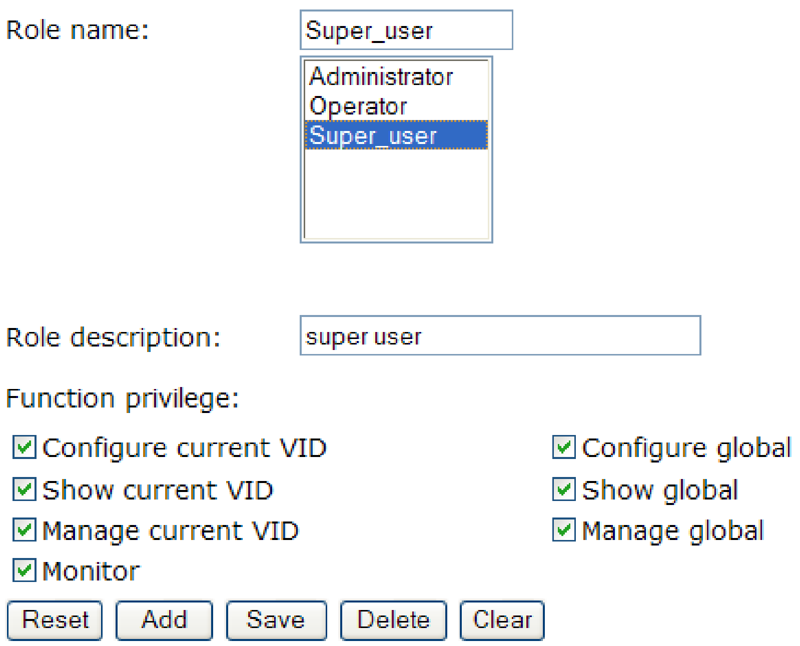

Role name: Super_user

Administrator
Operator
Super_user

Role description: super user

Function privilege:

☑ Configure current VID ☑ Configure global
☑ Show current VID ☑ Show global
☑ Manage current VID ☑ Manage global
☑ Monitor

- VID Administrator(s)

  You should define a VID administrator for each VID. You may define a single user who has access to multiple VIDs. Figure 4 shows the recommended role for a VID administrator.

Figure 4: VID administrator’s recommended role

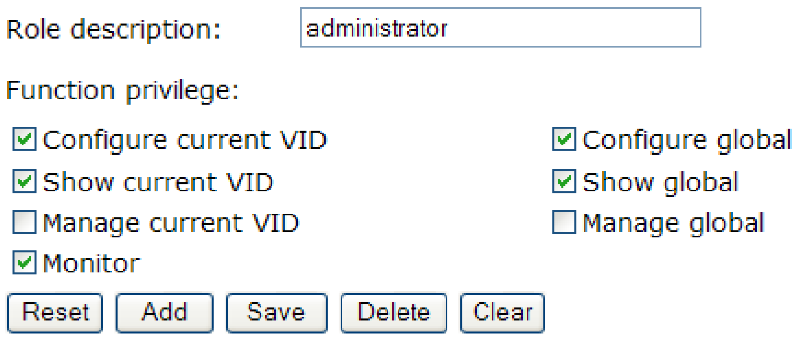

Role description: administrator

Function privilege:

☑ Configure current VID ☑ Configure global
☑ Show current VID ☑ Show global
☐ Manage current VID ☐ Manage global
☑ Monitor

- VID Operator

  You should define an operator for each VID who can only view and not modify any configurations. Figure 5 shows typical role for a VID operator. You may define a single operator who has access to multiple VIDs.

Figure 5: VID operator's recommended role

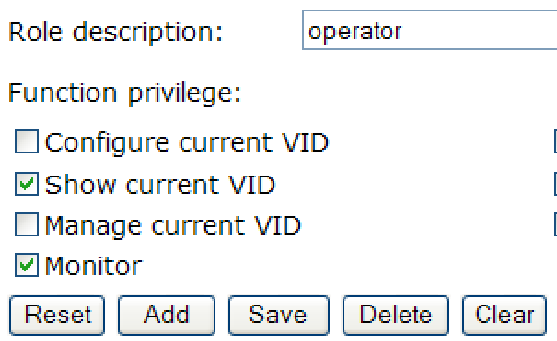

Role description: operator

Function privilege:

□ Configure current VID
☑ Show current VID
□ Manage current VID
☑ Monitor
□ Configure global
☑ Show global
□ Manage global

Managing roles

The user with the proper privilege, defined as Super User, can add, delete and modify roles. The *Add, Modify,* and *Delete* buttons save the information in the database.

To manage roles, click *Manage* > *Global* > *Roles.* You will see a screen as shown in Figure 6. The following paragraphs describe the functionality associated with this screen.

Figure 6: Managing roles

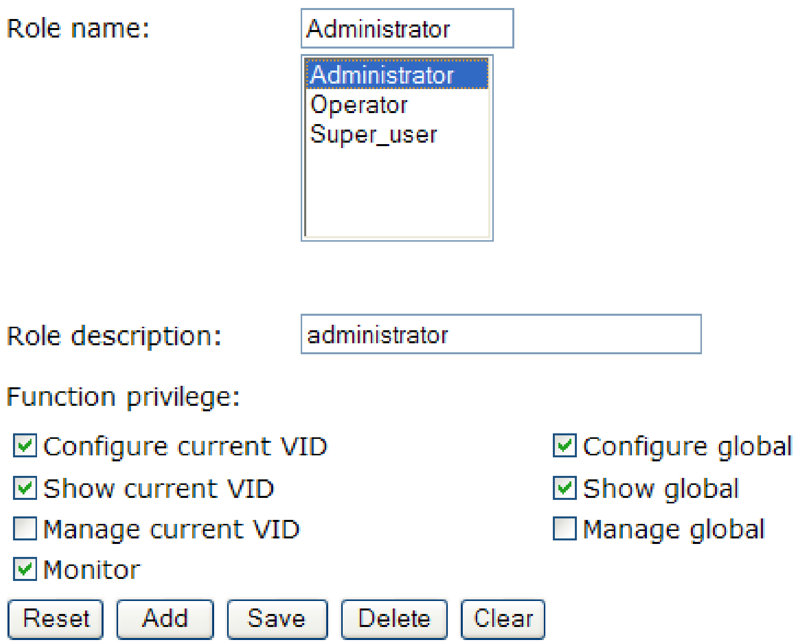

Role name: Administrator

Administrator
Operator
Super_user

Role description: administrator

Function privilege:

☑ Configure current VID
☑ Show current VID
□ Manage current VID
☑ Monitor
☑ Configure global
☑ Show global
□ Manage global

Adding a role

The *Add* button helps you add a new role to the database.

1. Start by pressing *Clear* button
2. Define a *Role name* with length between 2 and 30 alphanumeric characters
3. Define a *Role Description* not exceeding 100 alphanumeric characters
4. Ensure that the role does not already exist

5 Choose at least one function privilege

6 Click *Add* to add the role

Modifying a role

The *Modify* button helps you modify an existing role's properties.

1 Start by choosing a *Role name* from the list-box. This should retrieve the role's description and function privilege as defined in the database

2 Modify the *Role description* as desired. This should not exceed 100 alphanumeric characters

3 Adjust the *Function privilege* so that at least one remains chosen for the role

4 Ensure that you do not decrease the privilege of the last available role, or the last *super role*, which is the only one with the highest function privilege

5 Click *Modify* to modify the role

Deleting a role

The *Delete* button removes a role from the database.

1 Start by choosing a *Role name* from the list-box. This should retrieve the role's description and function privilege as defined in the database

2 Click *Delete* to delete the role

Note: You cannot delete a role that is currently in use by any user who is logged-on.

Note: You cannot delete the last role or the last super-role i.e. the role with the highest function privilege.

Note: *Reset* Button: This button resets the dialog to the values for the shown entity.

Note: *Clear* Button: This button clears the fields in the dialog so that the user can fill them from scratch.

Managing VID groups

The user with the proper privileges, normally defined as Super User, can add, delete and modify VID Groups. The *Add, Modify,* and *Delete* buttons save the information in the database.

To manage VID Groups, click *Manage* > *Global* > *VID Groups.* You will see a screen as shown in Figure 7. The following paragraphs describe the functionality related to this screen.

Figure 7: Managing VID groups

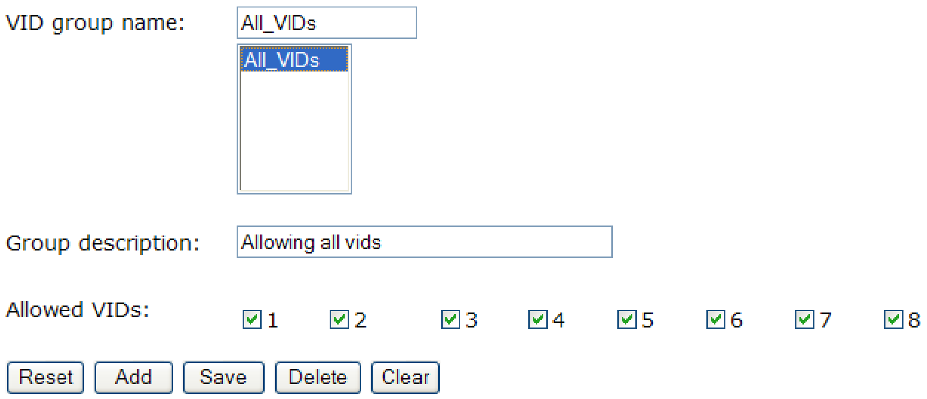

VID group name: All_VIDs

All_VIDs

Group description: Allowing all vids

Allowed VIDs: ☑ 1 ☑ 2 ☑ 3 ☑ 4 ☑ 5 ☑ 6 ☑ 7 ☑ 8

Adding a VID group

The *Add* button helps you add a VID Group in the database.

1. Start by pressing *Clear* button
2. Define a *VID Group name* with length between 2 and 30 alphanumeric characters
3. Define a *VID Group Description* not exceeding 100 alphanumeric characters
4. Ensure that the VID Group name does not already exist in the listbox
5. Choose at least one function VID from the list
6. Click *Add* to add the VID Group

Modifying a VID group

The *Modify* button helps you modify an existing VID Group's properties.

1. Start by choosing a *VID Group name* from the list-box. This should retrieve the group's description and VIDs associated with the group
2. Modify the *Group description* as desired. This should not exceed 100 alphanumeric characters
3. Adjust the *Allowed VIDs* so that at least one remains chosen for the group
4. Ensure that at least one VID Group has all VIDs associated with it
5. Click *Modify* to modify the VID Group

Deleting a VID group

The *Delete* button removes a VID Group from the database.

1. Start by choosing a *VID Group name* from the list-box. This should retrieve the group's description and VIDs associated with the group
2. Click *Delete* to delete the VID Group

Note: You cannot delete a VID Group that is currently associated with any user.

Note: You cannot delete the last VID Group that has all the VIDs associated with it.

Note: *Reset* Button: This button resets the dialog to the values for the shown entity.

Note: *Clear* Button: This button clears the fields in the dialog so that the user can fill them from scratch.

Managing users

The user with the proper privilege, normally defined as Super User, can add, delete and modify users. The *Add, Modify,* and *Delete* buttons save the information in the database.

To manage users, click *Manage* > *Global* > *Users.* You will see a screen as shown in Figure 8. The following paragraphs describe the functionality associated related to this screen.

Figure 8: Managing users

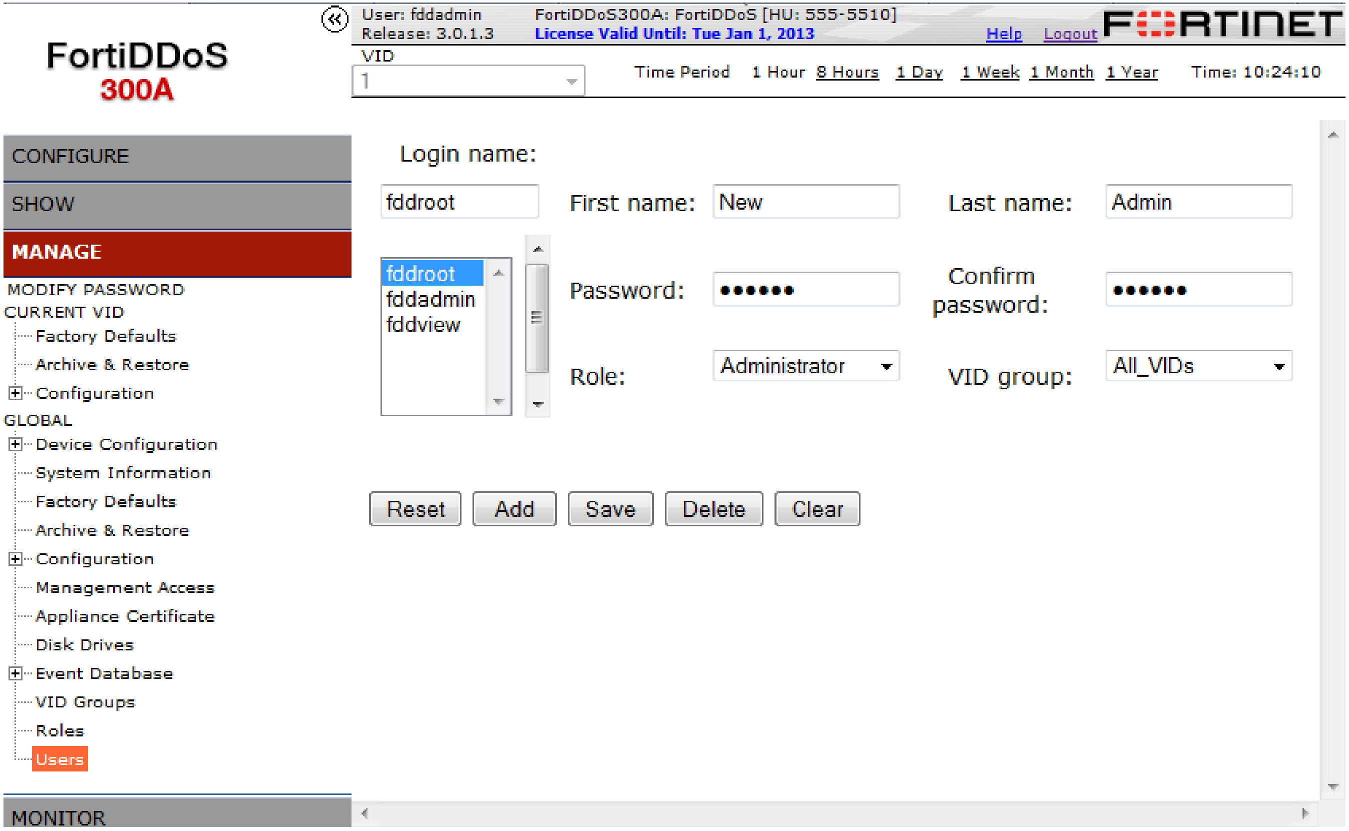

User: fddadmin
Release: 3.0.1.3
FortiDDoS300A: FortiDDoS [HU: 555-5510]
License Valid Until: Tue Jan 1, 2013

VID: 1

Time: 10:24:10

Login name: fddroot

First name: New

Last name: Admin

Password: ●●●●●●

Confirm password: ●●●●●●

Role: Administrator

VID group: All_VIDs

Adding a user

The *Add* button helps you add a user in the database.

1. Start by pressing *Clear* button
2. Enter the *Login name* with length between 3 and 30 alphanumeric characters
3. Enter the *First name* corresponding to the user with length between 3 and 30 alphanumeric characters
4. Enter the *Last name* corresponding to the user with length between 3 and 30 alphanumeric characters

5. Enter the *Password* corresponding to the user with length between 6 and 16 character. Only printable ASCII characters between '!' (Hex 21) and '~' (Hex 7d) are allowed
6. *Retype* the password
7. Select a *Role* for the user from the list box
8. Select a *VID Group* for the user from the list box
9. Click *Add* to add the user

Modifying a user

The *Modify* button helps you modify an existing User's properties.

1. Start by choosing a *Login name* from the list-box. This should retrieve the user's properties
2. Modify the *First name* corresponding to the user with length between 3 and 30 alphanumeric characters
3. Modify the *Last name* corresponding to the user with length between 3 and 30 alphanumeric characters
4. Modify the *Password* corresponding to the user with length between 6 and 16 character. Only printable ASCII characters between '!' (Hex 21) and '~' (Hex 7d) are allowed
5. *Retype* the password
6. If required, modify the *Role* for the user from the list box
7. Click *Modify* to complete the operation

Note: You should not decrease the role of the user if he is the last Super User.

Note: If required, modify the *VID Group* for the user from the list box.

Note: You should not modify the VID Group such that there is no Super User left with access to all VIDs.

Deleting a user

The *Delete* button removes a user from the database.

1. Start by choosing a *User* from the list-box. This should retrieve the user's properties
2. Click *Delete* to delete the user

Note: You cannot delete a user that is currently logged on.

Note: If the last user is a Super User, he cannot be deleted.

Note: *Reset* Button: This button resets the dialog to the values for the shown entity.

Note: *Clear* Button: This button clears the fields in the dialog so that the user can fill them from scratch.

Modifying your password

The *Modify Password* pane is in the form of top level and second level configuration panes that you are allowed to access in the Web-based Manager. This is a self-service screen. This screen is available to users who are already logged in. Refer to Figure 9. Click on *Manage > Modify Password* to access the screen.

Figure 9: Modifying your password

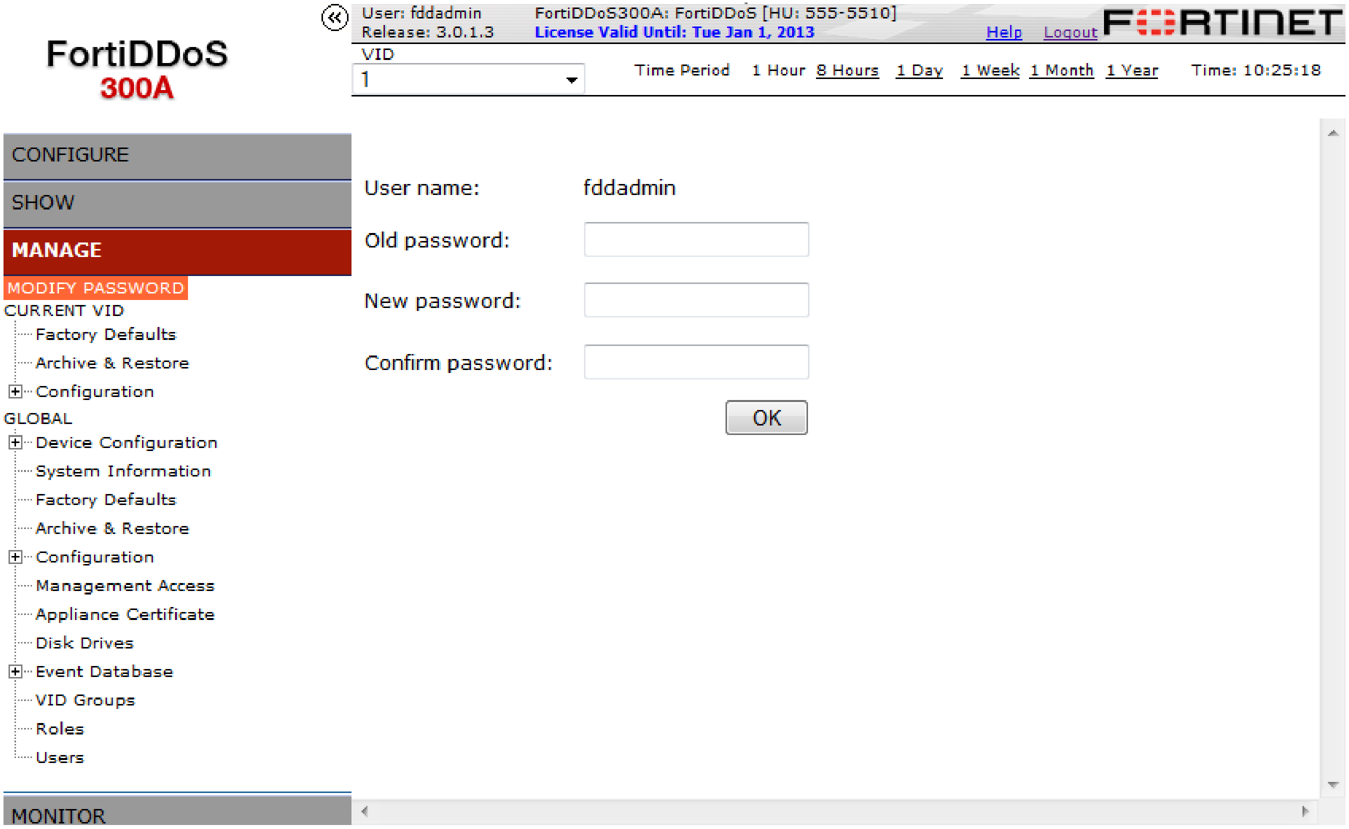

User: fddadmin
Release: 3.0.1.3
FortiDDoS300A: FortiDDoS [HU: 555-5510]
License Valid Until: Tue Jan 1, 2013
VID: 1
Time: 10:25:18

User name: fddadmin
Old password:
New password:
Confirm password:

1. Enter your *Old Password*
2. Enter a *New Password* with length between 6 and 16 character. Only printable ASCII characters between '!' (Hex 21) and '~' (Hex 7d) are allowed
3. Enter the same password in the *Retype Password* field
4. Click *OK* to Modify Password

Configuring a multi-VID deployment

This screen is available to the users with administrative privileges. The VIDs can be configured based on:

- IP addresses/netmasks.

To define your VID configuration, click *Configure* > *Global* > *VIDs.*

Note: Changes to the VID configuration MAY affect traffic patterns to that VID. All threshold estimation MAY therefore be affected. You must clear the traffic history for the affected VID if there is going to be significant change in the traffic pattern. After clearing the traffic history, you must re-train the appliance to learn the new traffic pattern for that VID.

IP address based VID configuration

Refer to Figure 10 for the IPv4 based VID configuration.

Figure 10: IPv4 based VID configuration

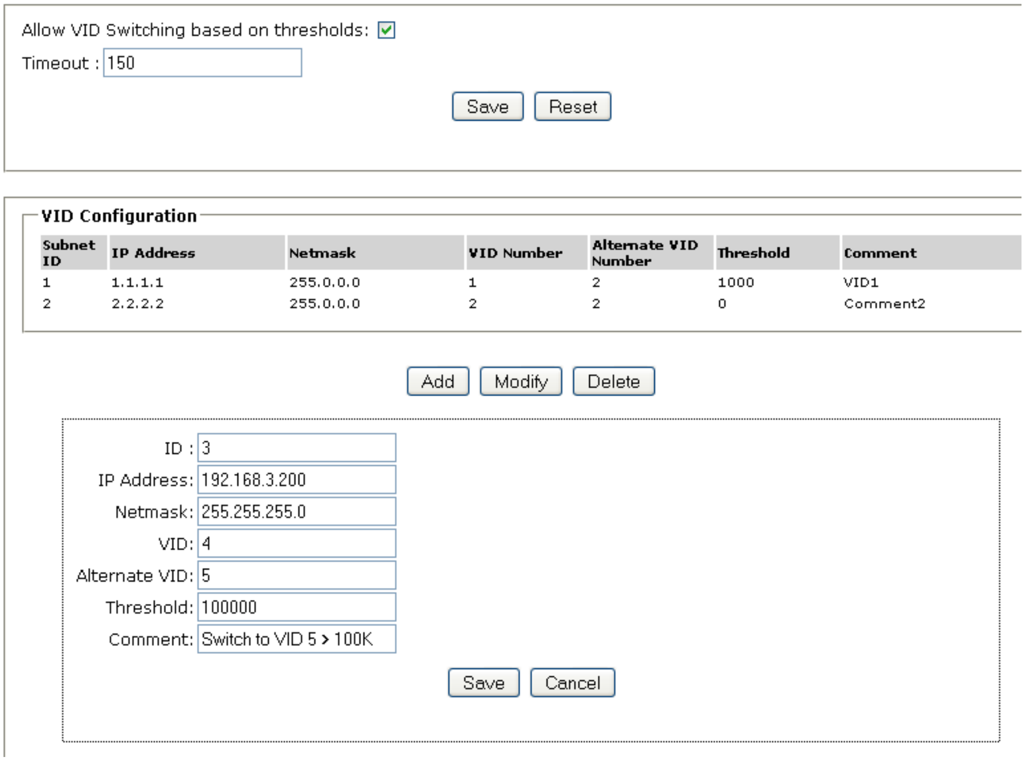

Allow VID Switching based on thresholds: ☑

Timeout : 150

VID Configuration

Subnet ID	IP Address	Netmask	VID Number	Alternate VID Number	Threshold	Comment
1	1.1.1.1	255.0.0.0	1	2	1000	VID1
2	2.2.2.2	255.0.0.0	2	2	0	Comment2

ID :	3
IP Address:	192.168.3.200
Netmask:	255.255.255.0
VID:	4
Alternate VID:	5
Threshold:	100000
Comment:	Switch to VID 5 > 100K

Click *Configure* > *Global* > *VIDs* to select IPv4 based configuration.

- For each subnet, a subnet ID (between 2 and 512), a VID number can be entered (between 1 - 8). All traffic matching a given network in either the source or the destination IP address will be assigned the given VID number.
- You can enable VID Switching based on threshold. This is a useful feature for scenarios where the administrator wants to keep different policies for low traffic and another policy once the traffic exceeds a threshold. Once the traffic is switched to another VID, you can enforce a different set of (e.g. more stringent) policies. The subnet remains in the alternate VID as long as the traffic remains above threshold. Once the traffic goes below the threshold and remains below for a timeout period, it is moved back to the original VID.
- If any packet matches multiple entries, the first match will be chosen since the table is searched from top to bottom and the first match succeeds.
- If no network matches, then the packet is assigned a VID that is entered in *Match all remaining traffic in this VID* entry under *Configure* > *Global* > *Default VID*. If this check-box is unchecked, and a packet does not match any other entry in the table, then the packet is deemed to be *Non-Tracked* and passes through the system unchecked.
- You must also enter a default subnet ID for such packets. This will be used for subnet-ID based reporting of events.
- Packets that have corrupted IP headers, need to be accounted for in some VIDs. By entering a value in the *Assign VID to packets which do not have valid IPs or have corrupted headers*, you associate a VID to such packets. You can do this configuration under *Configure* > *Global* > *Default VID.*
- You must also enter a default subnet ID for such packets. This will be used for subnet-ID based reporting of events.

Note: It is very important that the administrator is aware of the above feature, and always checks this box so that every packet passing through the device is checked for network attack.

Note: By assigning a VID by entering in the *Match all remaining traffic in this VID* box, you will affect the traffic statistics corresponding to that VID. Preferably, you must assign a different VID than those already assigned in the list.

Note: By assigning a VID by entering in the *Assign VID to packets which do not have valid IPs or have corrupted headers* box, you will affect the traffic statistics corresponding to that VID. Preferably, you must assign a different VID than the already assigned in the list.

- You can enter the same VID numbers for multiple networks if you desire. For example, you can use up five table slots but configure only two VIDs: three different networks assigned to VID 1 and two other networks assigned to VID 2.
- To assign a VID to a single IP address, you must enter a netmask of 255.255.255.255.
- You can associate each entry with a *Comment* for your easy reference later.
- The desired configuration only takes effect when the *Save* button is clicked. Choosing some other menu item in the left panel ignores the changes made in this interface and the VID configuration will remain unchanged.

Note: The order of defining networks is important in the above list. Items earlier in the network have a precedence. Therefore if you have overlapping ranges, ensure that specific networks are higher in the list.

Users and their associated VID groups

Once you have setup the VID in one of the above configurations, you can define VID groups and associated users who can manage the VIDs associated with the VID.

Assigning logical names to VIDs

You can assign a logical name to a VID. This will help you remember the VID from your own perspective. For example, you can assign VID 1 a logical name of web servers, VID 2 a logical name of DNS servers etc. Depending on your specific grouping, you can assign a name to the VID. These names are later shown in the Show and Configure Screens.

To assign a logical name to a VID, click *Configure* > *Global* > *VID Name.* In the dialog screen, enter the corresponding logical names against each VID and click *Save*.

Figure 11: Assigning logical names to VIDs

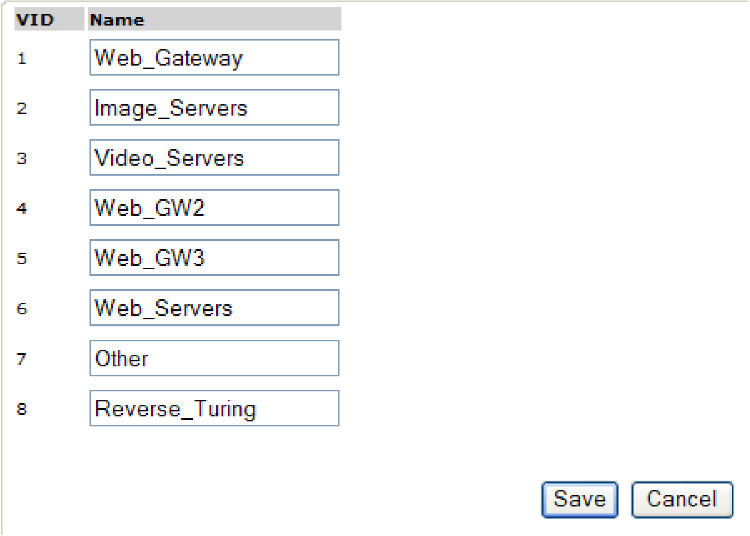

VID	Name
1	Web_Gateway
2	Image_Servers
3	Video_Servers
4	Web_GW2
5	Web_GW3
6	Web_Servers
7	Other
8	Reverse_Turing

Assigning name tags to servers

You can assign a logical name to a server that is being protected by the FortiDDoS. This will help you remember the names from your own perspective. For example, you can assign a name such as *Web Server Yosemite* to a specific IP address. These names are later shown in the Show Screens.

To assign a logical name to a Server, click *Configure* > *Global* > *Server Name Tags.* In the dialog screen, enter the corresponding logical Tag against each Server's IP Address and click *Save.*

Understanding the continuous learning graphs

FortiDDoS starts learning traffic patterns the moment it is introduced in the network and it never stops learning. It continuously records traffic statistics in a round robin fashion which you can view in graphical and tabular form. The traffic statistics are highly granular. Table 1 shows the traffic statistics learned by FortiDDoS for inbound and outbound directions for each VID.

Table 1: FortiDDoS Traffic Statistics

Layer	Type
3	Fragmented packets All 256 Protocols Count of unique sources (up to 1 million) Most active source rate (up to 1 million) Most active destination rate (up to 1 million) Rate of dark address scan
4	SYN packets SYN packets for most active SYN source SYN packets for most active SYN destination ACK packets for most active ACK destination RST packets for most active RST destination FIN packets for most active FIN destination ESTAB packets for most active ESTAB destination Maximum Packets per connection for most active connection across up to 1million connections Number of established TCP connections New TCP connection establishment rate Maximum concurrent connections per source Number of entries in the legitimate IP address table All 64K TCP Ports All 64K UDP Ports All 64K ICMP Type/Codes
7	HTTP Op codes (Methods) (Up to 8) URLs (up to 8192 in terms of Hash indexes) Hosts, Referers, Cookies, User-Agents (up to 512 in terms of Hash indexes) SIP Invite Per Source SIP Register Per Source SIP Concurrent Invite Per Source

Showing traffic statistics

You can show the traffic statistics mentioned in Table 12 granularity through the *Show* menu. All the traffic is recorded every 5 minutes and the rates are stored per second. They are the highest rates observed during those 5 minutes.

Some of these rates are for anomalous packets.

These rates give you an idea of traffic characteristics in your network.

Following graphs are available when you choose *Show* > *Current VID* > *Layer 3*:

- Fragmented Packets

  This graph illustrates the fragmented packets flowing through the VID in inbound and outbound directions.

- Unique Sources

  This graphs illustrates instantaneous count of unique sources flowing through the FortiDDoS. A spike in this graph shows a possibility of DDoS.

- Most Active Source

  This graph illustrates the traffic rate in inbound and outbound direction recorded every 5 minutes for the most active source among all the sources. A spike in this graph shows a possibility of attack from a single source.

- Most Active Destination

  This graph illustrates the traffic rate in inbound and outbound direction recorded every 5 minutes for the most active destination among all the destinations. A spike in this graph shows a possibility of attack to a single outbound destination.

- Protocols

  These graphs illustrate the traffic rate in inbound and outbound direction for every protocol. You can view the most common protocol using *My Graphs*.

- Dark Address Scan

  In case Dark Address Scan feature is enabled, these graphs illustrate the traffic rate in inbound and outbound direction for Dark Address Scan. If you have not enabled the Dark Address Scan prevention, you can enable it by clicking on *Configure* > *Current VID* > *Feature Control* > *Layer 3* -> *Source tracking* -> *Dark Address Scan.* Unlike other graphs, the rates are not per second, but are higher and are shown alongside the graph.

The following graphs are available when you choose *Show* > *Current VID* > *Layer 4:*

- SYN Packets

  This graph illustrates the SYN packets flowing through the VID in inbound and outbound directions.

- SYN Per Source Packets

  This graph illustrates the traffic rate in inbound and outbound direction recorded every 5 minutes for the most active source sending TCP SYN packets among all the sources. A spike in this graph shows a possibility of SYN attack from a single source or a few limited sources

- SYN Per Destination Packets

  This graph illustrates the traffic rate in inbound and outbound direction recorded every 5 minutes for the most active destination receiving TCP SYN packets among all the destinations in the VID. A spike in this graph shows a possibility of SYN attack to a single destination or a few limited destination.

- ACK Per Destination Packets

  This graph illustrates the traffic rate in inbound and outbound direction recorded every 5 minutes for the most active destination receiving TCP ACK packets among all the destinations in the VID. A spike in this graph shows a possibility of ACK attack to a single destination or a few limited destination.

- RST Per Destination Packets

  This graph illustrates the traffic rate in inbound and outbound direction recorded every 5 minutes for the most active destination receiving TCP RST packets among all the destinations in the VID. A spike in this graph shows a possibility of RST attack to a single destination or a few limited destination.

- FIN Per Destination Packets

  This graph illustrates the traffic rate in inbound and outbound direction recorded every 5 minutes for the most active destination receiving TCP FIN packets among all the destinations in the VID. A spike in this graph shows a possibility of FIN attack to a single destination or a few limited destination.

- ESTAB Per Destination Packets

  This graph illustrates the traffic rate in inbound and outbound direction recorded every 5 minutes for the most active destination receiving TCP established packets among all the destinations in the VID. A spike in this graph shows a possibility of connection establishment attack to a single destination or a few limited destination.

- Most Active Connection

  This graph illustrates the traffic rate in inbound and outbound direction recorded every 5 minutes for the most active TCP connection among all the connections. A spike in this graph shows a possibility of attack from a single connection or a limited set of connections.

- Number of TCP Connections

  A graph titled *Number of Established Connections* illustrates the count of maximum instantaneous TCP connections that have completed three way handshake. This is recorded every 5 minutes. A spike in this graph shows a possibility of a DoS or DDoS attack.

  A graph titled *Number of Entries in TCP Table* illustrates the count of maximum instantaneous TCP connections. This is recorded every 5 minutes. A spike in this graph shows a possibility of a DoS or DDoS attack. If the values in this graph are much higher than the first graph above, there is a possibility of SYN flood. For this graph and table, all entries in the connection table, including the half-open connections are used.

- TCP Connection Establishment Rate

  These graphs illustrate the traffic rate in inbound and outbound direction for new TCP connections established every second. A spike in the above two graphs show the possibility of a concerted DoS or DDoS attack.

- Entries in the Legitimate IP Address Table

  This graph illustrates the number of entries in the Legitimate IP address table.

- TCP Concurrent Connections Per Source

These graphs illustrate the count of concurrent connections for the busiest source. A spike in the above graph shows the possibility of a single source trying to establish too many connections.

- TCP Concurrent Connections Per Destination

  These graphs illustrate the count of concurrent connections for the busiest destination. A spike in the above graph shows the possibility of a single destination facing DDoS attack.

- Ports

  These graphs illustrate the traffic rate in inbound and outbound direction for every TCP and UDP port. You can view the most common ports using *My Graphs*.

- ICMP Types/Codes

  These graphs illustrate the traffic rate in inbound and outbound direction for every ICMP type and code combination. You can view the most common ports using *My Graphs.*

The following graphs are available when you choose *Show* > *Current VID* > *Layer 7* > *HTTP*:

- Op Codes

  This graph illustrates the traffic rate for HTTP op codes (methods). You can see these graphs as a group in My Graph or a specific one. The following eight methods are graphed:

  - GET
  - HEAD
  - OPTIONS
  - TRACE
  - POST
  - PUT
  - DELETE
  - CONNECT

- URLs, Hosts, Referers, Cookies, User-Agents

  These graphs illustrate the traffic rate in inbound and outbound direction for up to a certain number of hash indexes for URLs, Hosts, Referers, Cookies, User-Agents. You can view the most common ports using *My Graphs*.

  Since there can be infinite number of possible URLs, Hosts, Referers, Cookies, User-Agents, the hardware assigns the items to a certain number of buckets called hash-indexes. Due to this reason more than one item may fall into the same hash-index.

  To determine which hash index does a String falls into, click *Show* > *Current VID* > *Layer 7* > *HTTP* > *String to Hash Index* and enter the *String,* choose String Type and click *Submi***t** to determine the Hash Index.

  To determine the String that a given hash index corresponds to click *Show* > *Current VID* > *Layer 7* > *HTTP* > *Hash Index to String* and enter the *Hash Index* and click *Submit* to determine the possible String. Since there can be multiple Strings hashing to the same index, this result may not be accurate and will correspond to the last page that matched the Hash Index.

- SIP INVITEs Per Source

  This graph illustrates the traffic rate in inbound and outbound direction recorded every 5 minutes for the most active source sending SIP INVITE packets among all the sources. A spike in this graph shows a possibility of SIP INVITE attack from a single source or a few limited sources.
- SIP REGISTER Per Source

  This graph illustrates the traffic rate in inbound and outbound direction recorded every 5 minutes for the most active source sending SIP REGISTER packets among all the sources. A spike in this graph shows a possibility of SIP REGISTER attack from a single source or a few limited sources.
- SIP Concurrent INVITEs Per Source

  These graphs illustrate the count of concurrent SIP INVITE sessions for the busiest source. A spike in the above graph shows the possibility of a single source trying to establish too many sessions.

Showing dropped and blocked traffic statistics

You can show the dropped and blocked traffic statistics mentioned in Figure 12 granularity through the *Show* menu. All the traffic is recorded every 5 minutes. They are the actual counts of packets dropped and blocked during those 5 minutes.

Dropped packets are packets that have been dropped due to header, rate or state anomalies.

Blocked packets are packets that have been blocked due to access control lists (ACLs).

These graphs give you an idea of anomalous traffic characteristics in your network.

Understanding aggregate drops and drilling down

There is a certain hierarchy that you can follow while understanding dropped packets.

You must first view the *Aggregate Drops.* These are available when you choose *Show* > *Current VID* > *Aggregate Drops*. Any packet dropped by the system is accounted for here. You can drill down the drops from here onwards.

Following graphs are shown under aggregate drops:

- Layer 3 Drops

  Once you see non-zero drops in Layer 3 Drops under the Aggregate Drops graph or table, you can further drill down to Layer 3 Drops graph which is available in the same screen.

  Layer 3 Drops graph further break down the Layer 3 Drops into the following line-graphs:

  Protocol

  IPv4 options

  Fragmented packets

  L3 Anomalies

  Source Flood

  Misc. Source Flood

  Destination Flood

  Misc. Destination Flood

  Dark Address Scan

URL Source Tracking

Denied Countries

The interpretation of above graphs and tables are explained below. If you see any of the above line graphs having non-zero values, you can further drill down within Layer 3 graphs. You may also want to visit the Attack Reports. There is a similar hierarchy among Attack Reports starting from Top Attack Report.

- Layer 4 Drops

  Once you see non-zero drops in Layer 4 Drops under the Aggregate Drops graph or table, you can further drill down to Layer 4 Drops graph which is available in the same screen.

  Layer 4 Drops graph further break down the Layer 4 Drops into the following line-graphs:

  SYN Packets

  L4 Anomalies

  TCP Ports

  UDP Ports

  ICMP Ports

  Packets Per Connection

  Zombie Flood

  SYN Per Source Flood

  Excessive Concurrent Connections Per Source

  Excessive Concurrent Connections Per Destination

  TCP Packets Per Destination

  If you see any of the above line graphs having non-zero values, you can further drill down within Layer 4 graphs. You may also want to visit the Attack Reports. There is a similar hierarchy among Attack Reports starting from Top Attack Report.

- Layer 7 Drops

  Once you see non-zero drops in Layer 7 Drops under the Aggregate Drops graph or table, you can further drill down to Layer 7 Drops graph which is available in the same screen.

  Layer 7 Drops graph further break down the Layer 7 Drops into the following line-graphs:

  Op code Flood

  HTTP Anomalies

  URL, Host, Referer, Cookie, User-Agent Flood

  SIP Invite Per Source

  SIP Register Per Source

  SIP Concurrent Invite Per Source

  The interpretation of above graphs and tables are explained below.

Understanding layer 3 drop graphs and tables

The following dropped and blocked graphs are available when you choose *Show* > *Current VID* > *Layer 3*:

- Protocols

  These graphs illustrate the dropped and blocked packet count in inbound and outbound direction for every protocol. You can view the most common protocol using *My Graphs*.
- Fragmented Packets

  This graph illustrates the count of dropped and blocked fragmented packets through the VID in inbound and outbound directions.
- L3 Anomalies

  This graph illustrates the packets with anomalous layer 3 headers. These anomalies are:

  IP header checksum errors

  Land attacks

  Loopback address spoofing errors

  Other anomaly drops.
- Source Flood

  These graphs illustrate count of dropped and blocked packets from over-active sources during the period. Dropped sources correspond to sources that have been dropped due to floods while the blocked sources correspond to ACL entries in the layer 3 sources.
- Misc Source Flood

  Since the source table is a large dynamic table with up to 2 million entries, it can be attacked using hash-attacks or may run out of memory for pointers. This graph illustrates the drops due to hash attack and memory limitations.
- Destination Flood

  This graph illustrates count of dropped and blocked packets from over-active or blocked destinations. Dropped destinations correspond to sources that have been dropped due to floods while the blocked destinations correspond to ACL entries in the layer 3 destinations.
- Misc Destination Flood

  Since the destination table is a large dynamic table with up to 2 million entries, it can be attacked using hash-attacks or may run out of memory for pointers. This graph illustrates the drops due to hash attack and memory limitations.
- Dark Address Scan

  This graph illustrates dropped packets which correspond to dark address scan activity.

Understanding layer 4 drop graphs and tables

The following graphs are available when you choose *Show* > *Current VID* > *Layer 4*:

- SYN Packets

  This graph illustrates the count of dropped and blocked SYN packets through the VID in inbound and outbound directions.

- L4 Anomalies

  This graph illustrates the count of packets with anomalous layer 4 headers. These anomalies are:

  TCP, UDP and ICMP header checksum errors

  Invalid TCP flag combinations

  Other layer 4 header anomalies such as incomplete packet

- TCP Ports and UDP Ports

  These graphs illustrate the dropped and blocked packet count in inbound and outbound direction for every TCP and UDP port value. You can view the most common port values using *My Graphs*.

- ICMP Types/Codes

  These graphs illustrate the dropped and blocked packet count in inbound and outbound direction for every ICMP type and code. You can view the most common values using *My Graphs*.

- Packets Per Connection

  This graph illustrates the count of packets dropped due to connection flood in inbound and outbound direction.

- Misc. Connection Flood

  Since the TCP Connection Table is a large dynamic table with up to 1 million entries, it can be attacked using hash-attacks or may run out of memory for pointers. This graph illustrates the drops due to hash attack and memory limitations.

  In addition, a graph shows the count of packets with anomalous TCP packets which do not meet TCP state transition rules. These anomalies are:

  TCP Window size violations

  Foreign TCP packets

  TCP state transition anomalies

- Zombie Flood

  This graph illustrates the count of dropped traffic in inbound and outbound direction. Zombie flood is assumed when the number of allowed legitimate IP addresses during a SYN flood exceeds a set threshold. These packets show existing of non-spoofed IP addresses generating a large number SYN packets and thus creating a distributed DoS.

- SYN Per Source Packets

  This graph illustrates the count of dropped and blocked packets through the VID in inbound and outbound directions that were dropped due to a single source sending too many SYN packets.

- Excessive Concurrent Connections Per Source

  This graph illustrates the count of dropped and blocked packets through the VID in inbound and outbound directions that were dropped due to a single source building too many TCP concurrent connections.

- Excessive Concurrent Connections Per Destination

  This graph illustrates the count of dropped and blocked packets through the VID in inbound and outbound directions that were dropped due to a single destination being flooded with too many TCP concurrent connections.

Understanding layer 7 drop graphs and tables

The following graphs are available when you choose *Show* > *Current VID* > *Layer 7* > *HTTP*:

- Op Codes

  These graphs illustrate the dropped packet count in inbound and outbound direction for every HTTP Op Code (Method) value. You can also view them using *My Graphs*.

- URLs, Hosts, Referers, Cookies, User-Agents

  These graphs illustrate the dropped and blocked packet count in inbound and outbound direction for every Hash Index value. You can view the most common option values using *My Graphs*.

The following graphs are available when you choose *Show* > *Current VID* > *Layer 7* > *SIP*:

- INVITE Per Source Packets

  This graph illustrates the count of dropped and blocked packets through the VID in inbound and outbound directions that were dropped due to a single source sending too many INVITE packets.

- REGISTER Per Source Packets

  This graph illustrates the count of dropped and blocked packets through the VID in inbound and outbound directions that were dropped due to a single source sending too many REGISTER packets.

- SIP Concurrent INVITEs Per Source

  This graph illustrates the count of dropped and blocked packets through the VID in inbound and outbound directions that were dropped due to a single source building too many concurrent INVITE sessions.

Typical traffic graph

Figure 12 illustrates a typical graph. The following terms are useful in interpreting the graph.

Figure 12: Graph of inbound/outbound TCP traffic

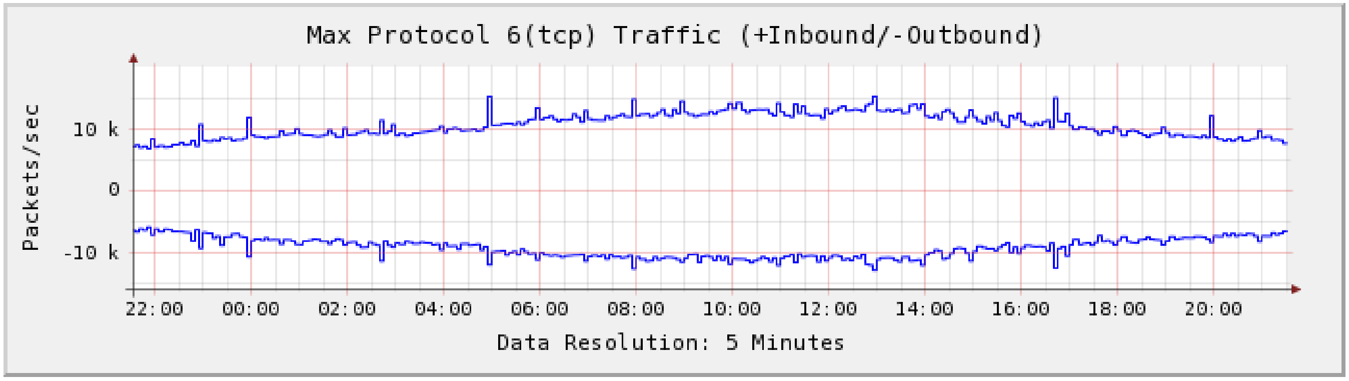

Title of the graph: In this case Max Protocol 6 (tcp0) Traffic (+Inbound/-Outbound)

- This graph illustrates daily graph of TCP protocol packets through the FortiDDoS.
- The positive side of the graph is showing the inbound traffic rate while the negative values correspond to the outbound traffic rate.
- The horizontal axis shows the time scale while the vertical scale shows the *Packets/sec*.

- You can view the daily seasonality in the graph as the traffic is highest around mid-morning and goes down at night.
- The *Data Resolution: 5 minutes* means that the samples have been taken every 5 minutes and are being shown that way. The value at any given time is the highest packet rate within 1 second for the TCP protocol during the 5 minute sampling period. If in a graph the data resolution is 1 hour, the value will be highest packet rate within 1 second for TCP protocol during the 1 hour sampling period. Same logic can be extended for data resolution of 3 hours and 45 hours which are used in graphs when you view periods of 1 month and 1 year respectively.

Note: Since different time periods are handled with a different resolution, each graph is correct only up to the last checkpoint period. For example the 1 hour graph with a 5 minute resolution is correct only up to the time before the last 5 minutes. The data in the last 5 minutes may not yet have been registered. So if you have a traffic peak in the last 5 minutes and you are trying to see it, it may not appear in the graphs immediately. Similarly, if you have a 1 month graph with a data resolution of 3 hours, the data for the last 3 hours may not appear immediately.

Other options while viewing traffic graphs

While viewing graphs you can change several parameters. These are explained below. Refer to Figure 13

Figure 13: Other controls in traffic graphs

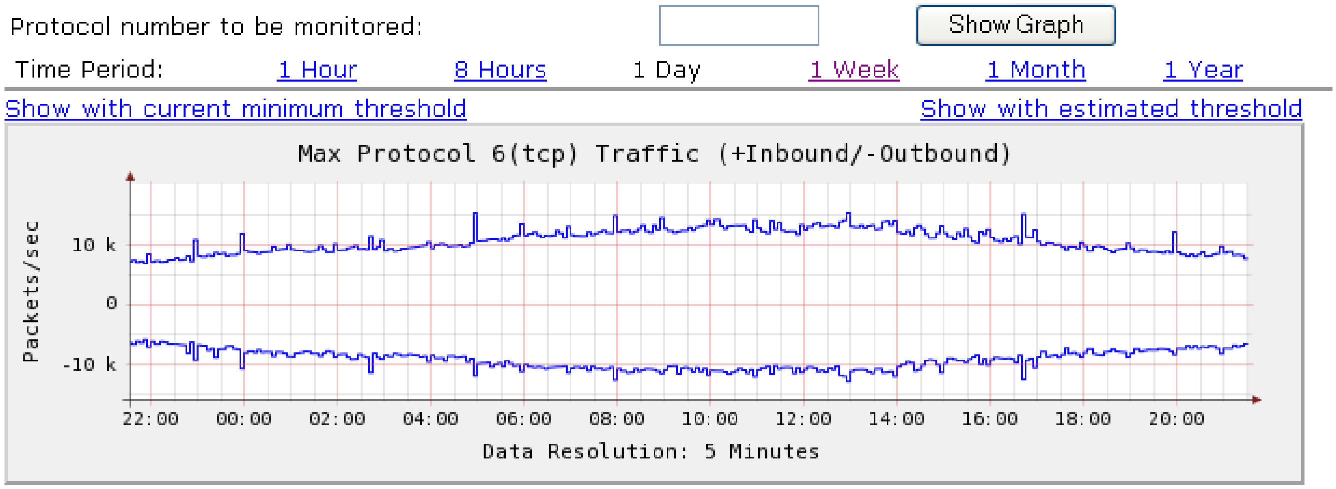
Protocol number to be monitored:
Show Graph
Time Period:
1 Hour
8 Hours
1 Day
1 Week
1 Month
1 Year
Show with current minimum threshold
Show with estimated threshold
Max Protocol 6(tcp) Traffic (+Inbound/-Outbound)
Packets/sec
10 k
0
-10 k
22:00
00:00
02:00
04:00
06:00
08:00
10:00
12:00
14:00
16:00
18:00
20:00
Data Resolution: 5 Minutes

Time period control: You can view the same graph in different resolution starting from last 1 hour, going all the way back to 1 year. This is a round-robin representation, i.e. it is always the last hour, last 8 hour and so on. According to the Time Period, the Data Resolution changes.

Show with current minimum threshold:

By clicking the hyperlink titled *Show with Current Minimum Threshold* you will start seeing the graph with the current minimum threshold that the administrator has set. The threshold is usually shown in black color. You can adjust the current minimum threshold through the *Configure* > *Current VID* > *Threshold* menu.

Refer to Figure 14 to see the effect of superimposing the current minimum threshold.

Figure 14: Effect of superimposing current minimum threshold

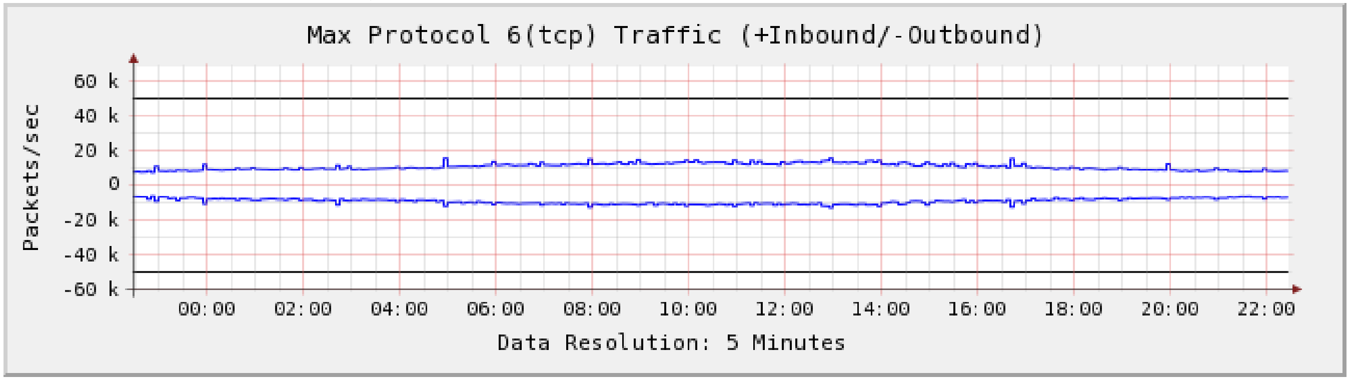

Show with estimated threshold:

By clicking hyperlink '*Show with Estimated Threshold*' you will start seeing the graph with the adaptive estimated threshold that the system calculates based on the past traffic. The threshold is usually shown in light blue color.

You cannot change the estimated thresholds as simply as you can change the current minimum threshold; it is learnt based on past traffic, its average, trend, and seasonality. However, you can adjust the behavior of adaptiveness by adjusting the Threat Level. That will have the effect of adjusting the estimated threshold up or down.

Refer to Figure 15 to see the effect of superimposing the estimated threshold.

Figure 15: Effect of superimposing the estimated threshold

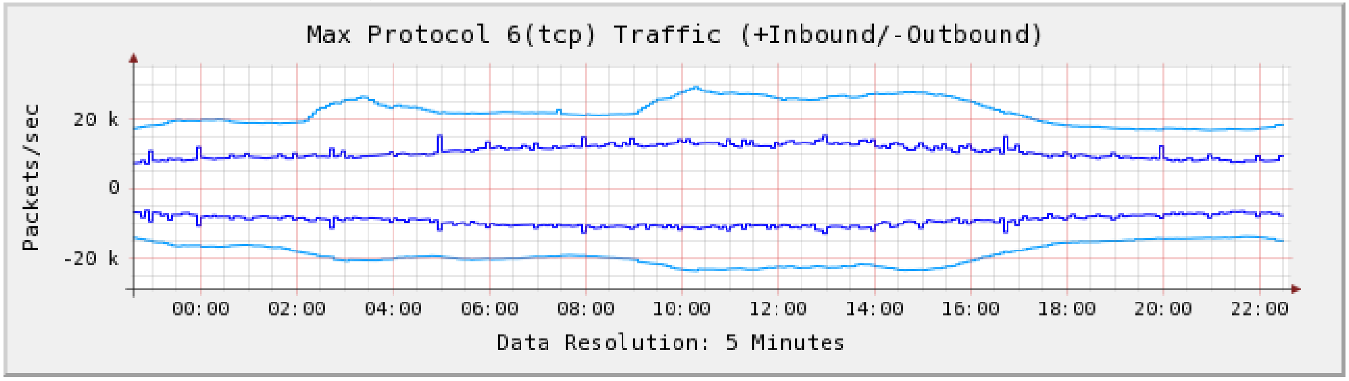

Show line or area graphs

Some graphs show multiple parameters. You can choose either line or area graphs for visualization. By clicking hyperlink *Show Area Graph* or *Show Line Graph* you can toggle the visualization.

Refer to Figure 16 to see the graphs plotted using line graphs.

Figure 16: Visualizing graphs using line graphs

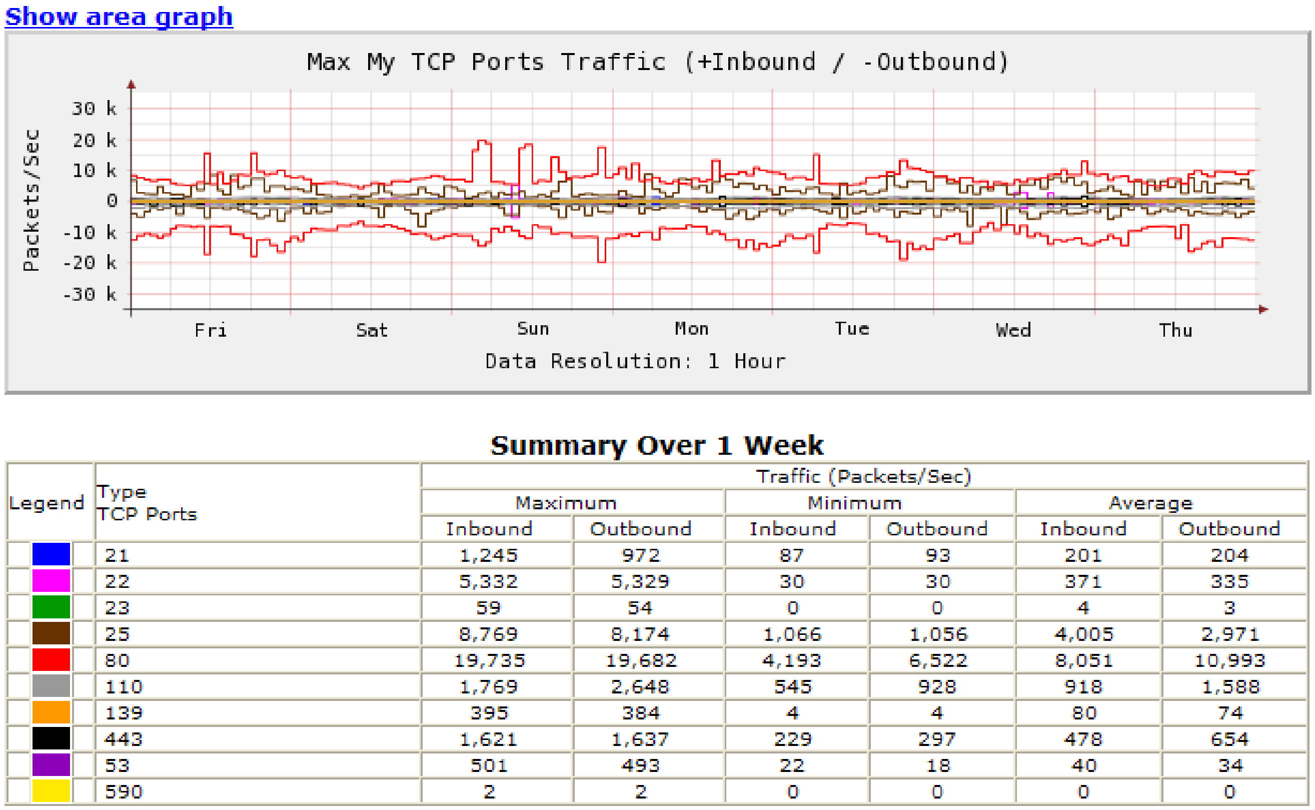

Summary Over 1 Week

Legend	Type TCP Ports	Traffic (Packets/Sec)					
		Maximum		Minimum		Average	
		Inbound	Outbound	Inbound	Outbound	Inbound	Outbound
🟦	21	1,245	972	87	93	201	204
🟪	22	5,332	5,329	30	30	371	335
🟩	23	59	54	0	0	4	3
🟫	25	8,769	8,174	1,066	1,056	4,005	2,971
🟥	80	19,735	19,682	4,193	6,522	8,051	10,993
⬜	110	1,769	2,648	545	928	918	1,588
🟧	139	395	384	4	4	80	74
⬛	443	1,621	1,637	229	297	478	654
🟪	53	501	493	22	18	40	34
🟨	590	2	2	0	0	0	0

Refer to Figure 17 to see the graphs plotted using area graphs.

Figure 17: Visualizing using area graphs

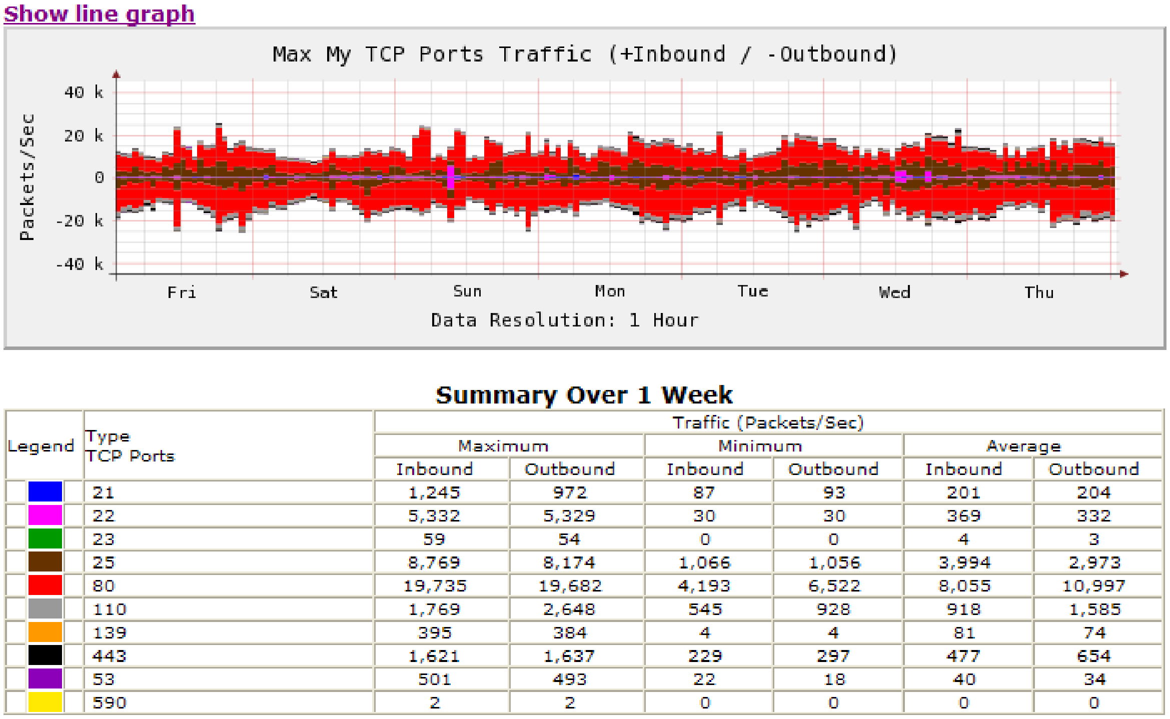

Summary Over 1 Week

Legend	Type TCP Ports	Traffic (Packets/Sec)					
		Maximum		Minimum		Average	
		Inbound	Outbound	Inbound	Outbound	Inbound	Outbound
🟦	21	1,245	972	87	93	201	204
🟪	22	5,332	5,329	30	30	369	332
🟩	23	59	54	0	0	4	3
🟫	25	8,769	8,174	1,066	1,056	3,994	2,973
🟥	80	19,735	19,682	4,193	6,522	8,055	10,997
▨	110	1,769	2,648	545	928	918	1,585
🟧	139	395	384	4	4	81	74
⬛	443	1,621	1,637	229	297	477	654
🟪	53	501	493	22	18	40	34
🟨	590	2	2	0	0	0	0

A typical graph showing dropped packets or blocked packets

Figure 18 illustrates a typical dropped packets graph. The following terms are useful in interpreting the graph.

Figure 18: A typical dropped packets graph

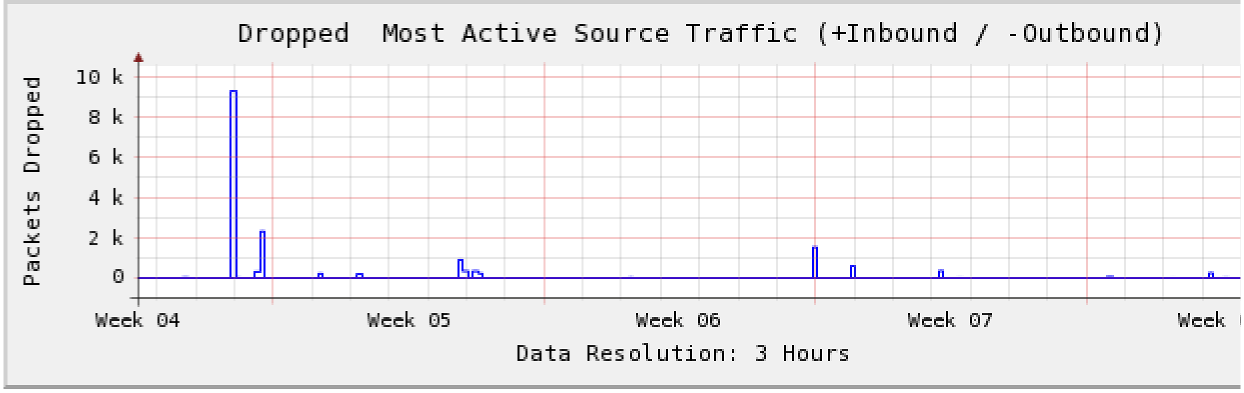

Title of the graph: In this case Dropped Most Active Source Traffic (+ Inbound/ -Outbound)

- The graph illustrates the number of packets dropped due to source flood.
- The positive side of the graph shows the packets dropped in inbound direction while the negative values correspond to the outbound dropped traffic.

- The horizontal axis shows the time scale In this case, the total period of 1 month has been divided in 4 weeks. The vertical scale shows the number of packets dropped during the data resolution period.
- The *Data Resolution: 3 hours* means that the samples have been taken every 3 hours minutes and are being shown that way. The value at any given time is the packets dropped during the 3 hours period. If, in a graph the data resolution is 1 hour, the value will be the count of the dropped packets due to source flood within 1 hour period. Same logic can be extended for data resolution of 3 hours and 45 hours which are used in graphs when you view periods of 1 month and 1 year respectively.

Note: Since different time periods are handled with a different resolution, each graph is correct only up to the last checkpoint period. For example the 1 hour graph with a 5 minute resolution is correct only up to the time before the last 5 minutes. The data in the last 5 minutes may not yet have been registered. So if you have dropped traffic in the last 5 minutes and you are trying to see it, it may not appear in the graphs immediately. Similarly, if you have a 1 month graph with a data resolution of 3 hours, the dropped packet data for the last 3 hours may not appear immediately.

A typical graph showing populations

Certain graphs are neither packets rates nor dropped or blocked packet counts. They show population of certain tables. Unlike other graphs, these graphs show the counts or population of certain tables at a given time.

These graphs include:

- Max count of unique sources, Figure 19
- Number of connections, Figure 20
- Number of entries in the TCP state table, Figure 21
- Entries in the legitimate IP address table, Figure 22

Note: Since different time periods are handled with a different resolution, each graph is correct only up to the last checkpoint period. For example the 1 hour graph with a 5 minute resolution is correct only up to the time before the last 5 minutes. The data in the last 5 minutes may not yet have been registered. So if you have a population peak in the last 5 minutes and you are trying to see it, it may not appear in the graphs immediately. Similarly, if you have a 1 month graph with a data resolution of 3 hours, the data for the last 3 hours may not appear immediately.

Figure 19: Max count of unique sources

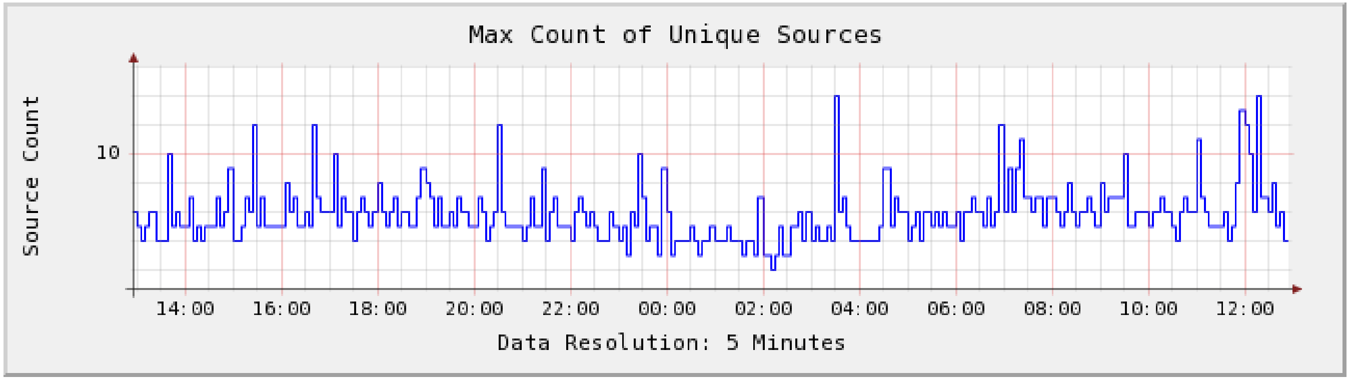

Figure 20: This graph illustrates the number of established TCP connections

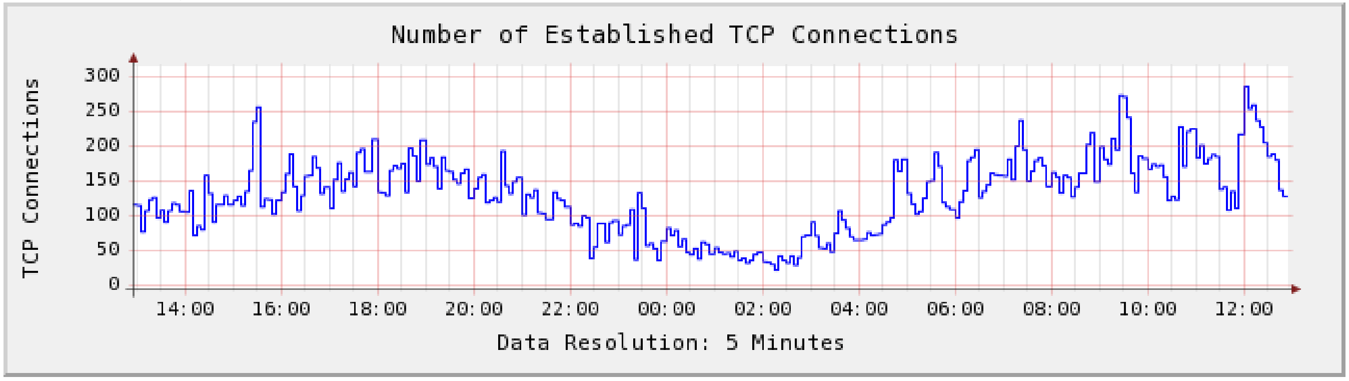

Figure 21: Graph showing the number of entries in the TCP state table

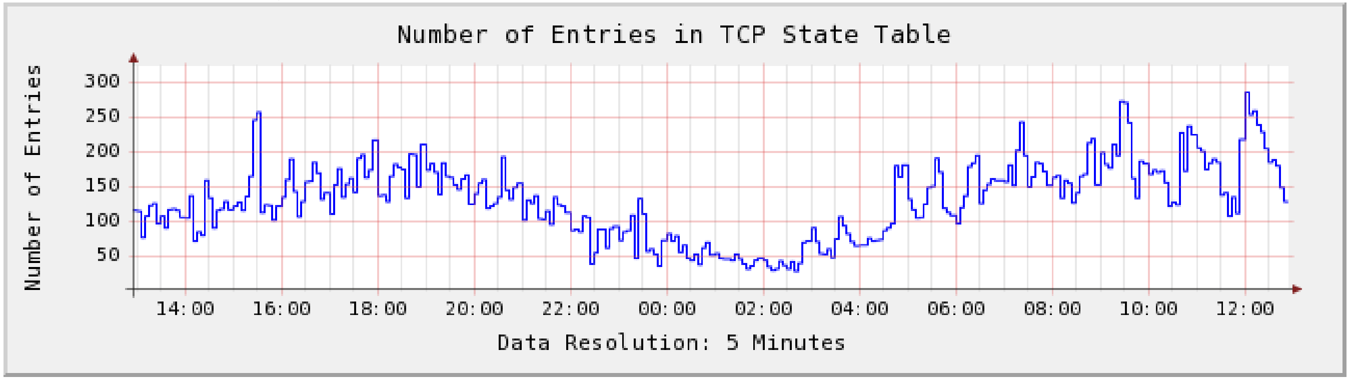

Figure 22: Graph showing entries in the legitimate IP address table

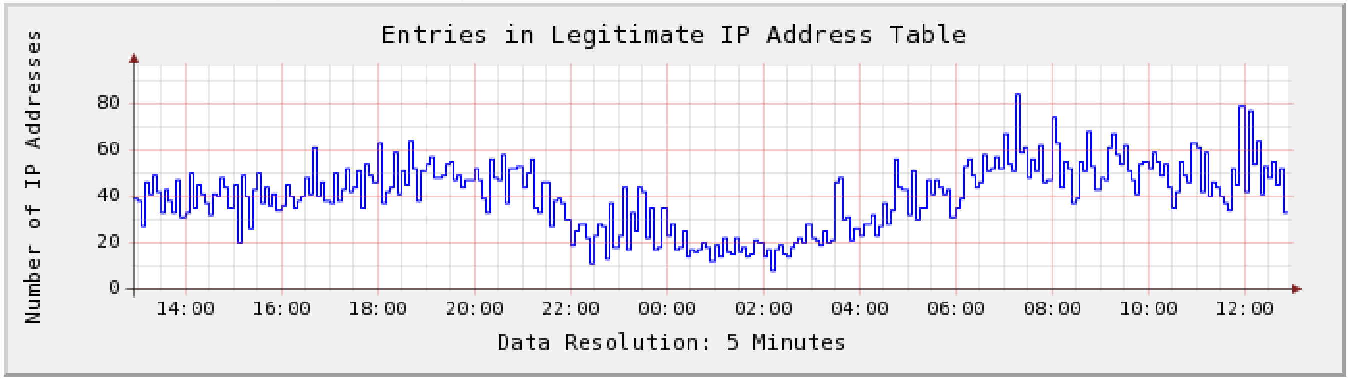

Interpreting the tabular traffic summaries

Figure 23 illustrates a typical traffic statistics summary.

Figure 23: Typical traffic summary

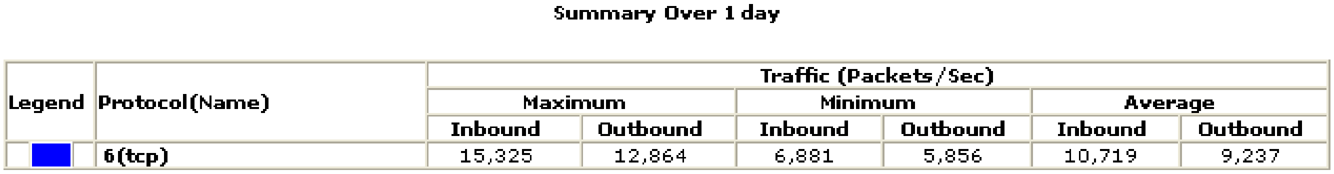

Summary Over 1 day

Legend	Protocol(Name)	Traffic (Packets/Sec)					
		Maximum		Minimum		Average	
		Inbound	Outbound	Inbound	Outbound	Inbound	Outbound
🟦	6(tcp)	15,325	12,864	6,881	5,856	10,719	9,237

- As the title states, the summary is for *1 day* period.
- During this 1 day period, the maximum inbound traffic rate was 15,325 packets/second while the outbound traffic clocked 12,864 packets/second during any second. Those were the highest recorded rates.
- During the same 1 day period, the minimum inbound and outbound traffic for protocol 6 were 6,881 and 5856 respectively. Those were the lowest recorded rates.
- During the same 1 day period, the average inbound and outbound traffic for protocol 6 were 10,719 and 9,237 packets/second respectively.

Note: Since different time periods are handled with a different resolution, each graph is correct only up to the last checkpoint period. For example the 1 hour graph with a 5 minute resolution is correct only up to the time before the last 5 minutes. The data in the last 5 minutes may not yet have been registered. So if you have a traffic peak in the last 5 minutes and you are trying to see it, it may not appear in the tables immediately. Similarly, if you have a 1 month table with a data resolution of 3 hours, the data for the last 3 hours may not appear immediately.

Viewing traffic on physical ports

While you can view traffic on many dimensions of layer 3, 4 and 7 protocols and also set thresholds on them and control the rate anomalies, you can also view the overall traffic being received on inbound and outbound directions on the physical ports. This information is useful from the perspective of understanding the overall traffic through the FortiDDoS.

There are four physical ports on the FortiDDoS 100A.

- LAN 1
- WAN 1
- LAN 2
- WAN 2

Note: The FortiDDoS 200A has additional ports that are marked LAN 3, WAN 3, LAN 4, WAN 4. The FortiDDoS 300A has additional ports that are marked LAN 5, WAN 5, LAN 6, WAN 6.

You can view traffic on each physical port by:

- Frame Type

  Click *Show* > *Global* > *Card [number]* > *Port Name* > *Frame Type*
- Frame Size

  Click *Show* > *Global* > *Card [number]* > *Port name* > *Frame Size*

- Transmit Errors

  Click *Show* > *Global* > *Card [number]* > *Port name* > *Transmit Errors*
- Receive Errors

  Click *Show* > *Global* > *Card [number]* > *Port name* > *Receive Errors*

Viewing traffic by frame type

Refer to Figure 24 and Figure 25. These show the graph and corresponding tabular form of traffic on LAN 1 arranged by frame types as they are received on the port.

The following frame types are shown:

- Unicast
- Flow Control
- Multicast
- Broadcast.

The table illustrates the same data in tabular text form. Depending on the period you have chosen, the graph and table resolution varies.

Inbound and outbound traffic minimum, maximum, average and total packets are shown in the table.

As an example, in the table shown, the highest unicast traffic ever during the last week over any 1 hour period was never more than 5,912,270 in the inbound direction. And it was never less than 1,930, 898 packets during any 1 hour period. A total of 6,011,563,850 packets were received in the inbound direction during the week.

Figure 24: Traffic graph by frame type

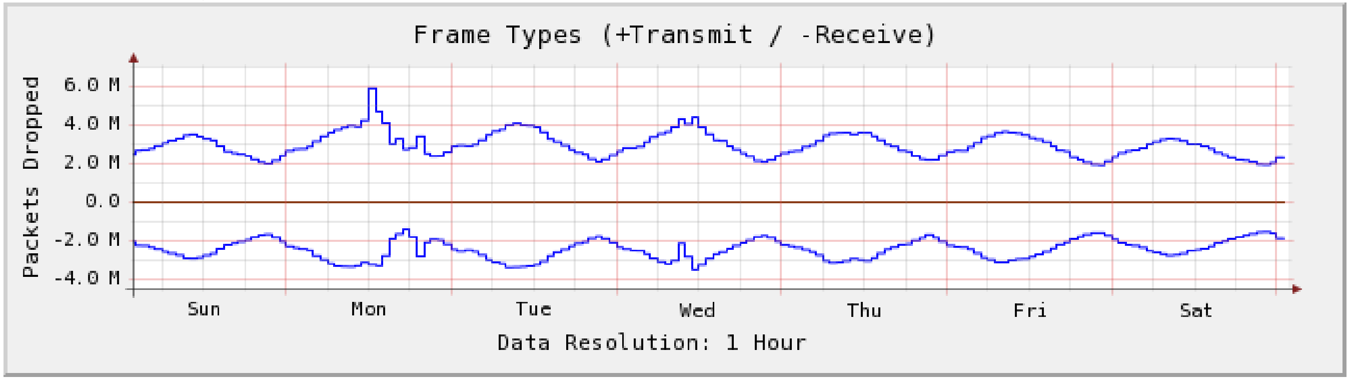

Figure 25: Traffic summary by frame type

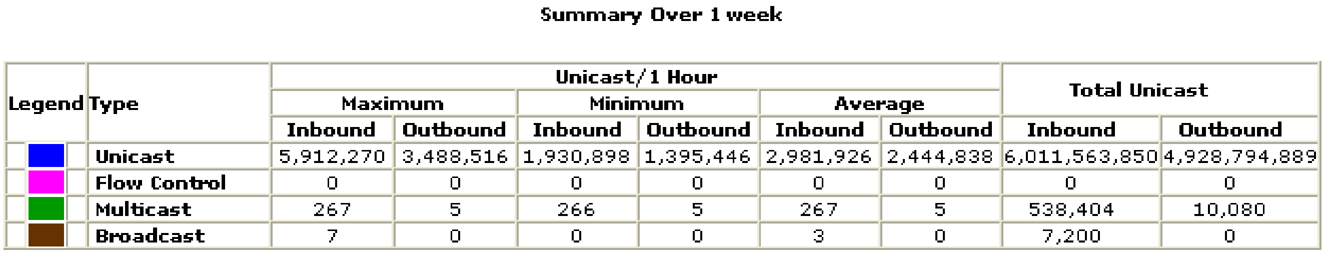

Summary Over 1 week

Legend	Type	Unicast/1 Hour						Total Unicast	
		Maximum		Minimum		Average			
		Inbound	Outbound	Inbound	Outbound	Inbound	Outbound	Inbound	Outbound
🟦	Unicast	5,912,270	3,488,516	1,930,898	1,395,446	2,981,926	2,444,838	6,011,563,850	4,928,794,889
🟪	Flow Control	0	0	0	0	0	0	0	0
🟩	Multicast	267	5	266	5	267	5	538,404	10,080
🟫	Broadcast	7	0	0	0	3	0	7,200	0

Viewing traffic by frame size

You can similarly see traffic graphs and a tabular summary arranged by frame sizes. Click on *Show* > *Global* > *Port name* > *Frame Size***.** Refer to Figure 26**.**

The following frame sizes are shown:

- 64
- 64-127
- 128-255
- 256-511
- 512-1023
- 1024-1518
- Jumbo

Figure 26: Example traffic graph and table by size

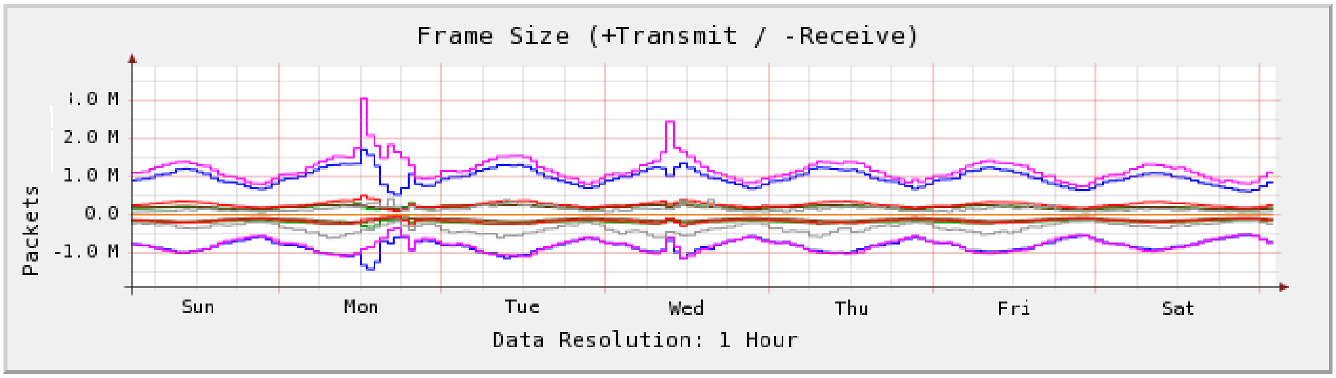

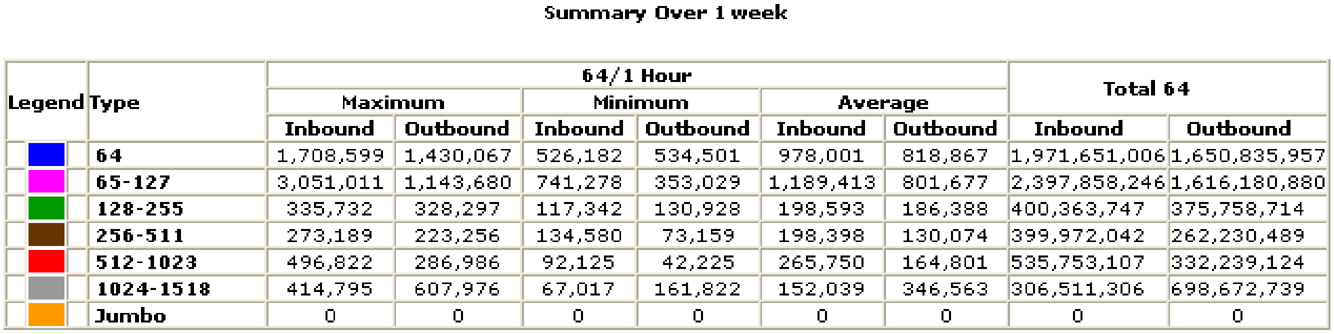

Summary Over 1 week

Legend	Type	64/1 Hour						Total 64	
		Maximum		Minimum		Average			
		Inbound	Outbound	Inbound	Outbound	Inbound	Outbound	Inbound	Outbound
🟦	64	1,708,599	1,430,067	526,182	534,501	978,001	818,867	1,971,651,006	1,650,835,957
🟪	65-127	3,051,011	1,143,680	741,278	353,029	1,189,413	801,677	2,397,858,246	1,616,180,880
🟩	128-255	335,732	328,297	117,342	130,928	198,593	186,388	400,363,747	375,758,714
🟫	256-511	273,189	223,256	134,580	73,159	198,398	130,074	399,972,042	262,230,489
🟥	512-1023	496,822	286,986	92,125	42,225	265,750	164,801	535,753,107	332,239,124
⬜	1024-1518	414,795	607,976	67,017	161,822	152,039	346,563	306,511,306	698,672,739
🟧	Jumbo	0	0	0	0	0	0	0	0

Viewing transmit errors on a port

You can see traffic graphs and a tabular summary for Transmit Errors. Click on *Show* > *Global* > *Port name* > *Transmit Errors*.

The following Transmit Errors are shown:

- Total Collisions
- Single Collisions
- Multiple Collisions
- Late Collisions
- Excessive Collisions
- Excessive Deferral
- Excessive Length Drop
- Underrun
- CRC Errors
- Pause Frames

For an understanding of these errors, you must refer to the IEEE Ethernet specification.

Viewing receive errors on a port

You can see traffic graphs and a tabular summary for Receive Errors. Click on *Show* > *Global* > *Port name* > *Receive Errors.*

The following Receive Errors are shown:

- FCS Errors
- Data Errors
- Align Errors
- Long Error
- Jabber Errors
- Very Long Errors
- Runt Errors
- Short Errors
- Carrier Extend Errors
- Sequence Errors

For an understanding of these errors, you must refer to the IEEE Ethernet specification.

Viewing benefit of FortiDDoS using port statistics

You can clearly see the benefit of FortiDDoS under attack by looking at the Port Statistics graph. Click on *Show* > *Global* > *Port Statistics.*

Following graph in Figure 27 illustrates Port Statistics under attack. The magenta line shows inbound and outbound traffic which is incoming. The blue area shows the inbound and outbound traffic that was allowed. The gap between the two is the benefit of the FortiDDoS solution. Under the normal circumstances, the blue area and magenta lines will hug each other. Under attack, since FortiDDoS will drop packets, the gap will increase. Larger an attack, larger the gap will be. In the graph shown below, you can see a periodic pattern from Friday until Wednesday. On Thursday, this periodicity is broken by an attack. The allowed traffic (which is shown as blue area graph) still follows the past pattern.The second graph illustrates the period of the same attack zoomed into 1 day from 1 week.

Figure 27: Example port statistics graph under attack

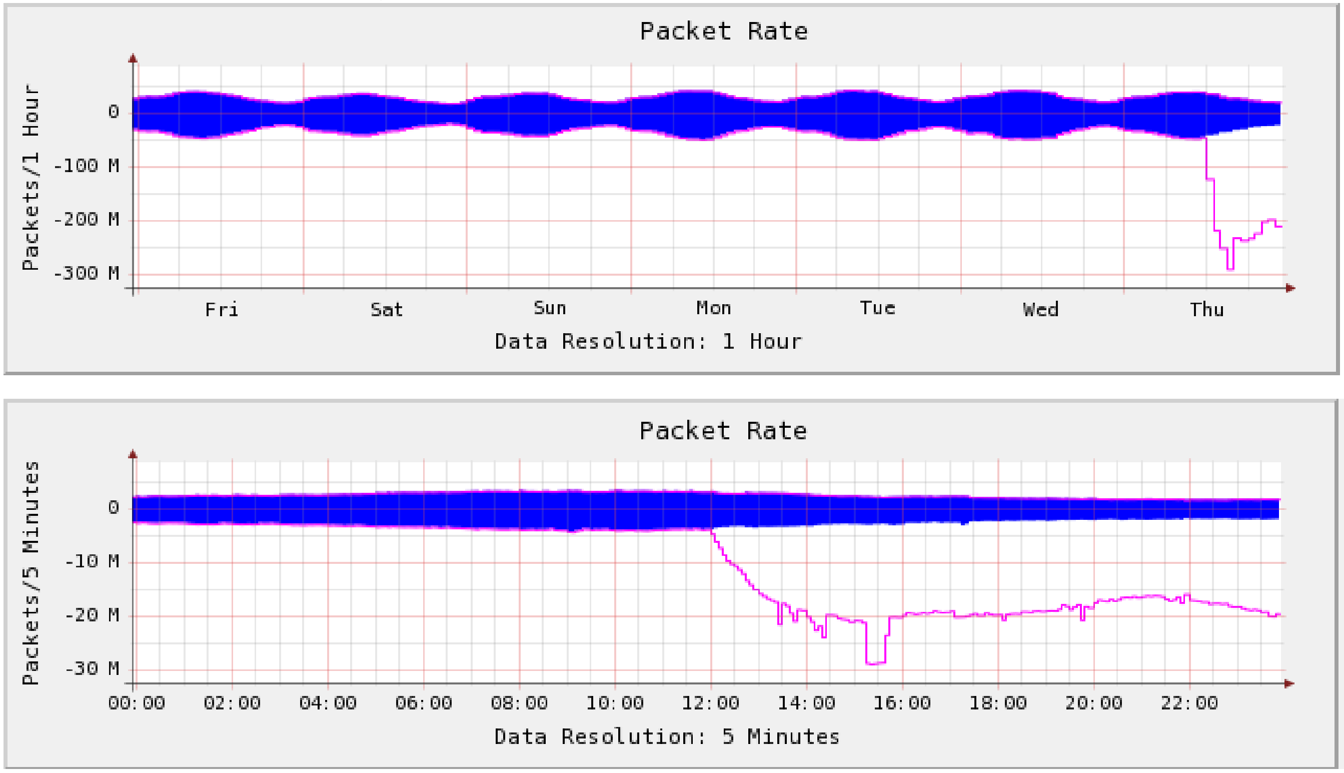

Viewing traffic in Mbps through FortiDDoS using port statistics

You can clearly see traffic passing through FortiDDoS in Mbps by looking at the Port Statistics graph. Click on *Show* > *Global* > *Port Statistics.*

Following graph in Figure 28 illustrates Port Statistics with traffic in Mbps. The lower half shows inbound traffic and the upper half shows the outbound traffic. The inbound side shows the ingress on the port as the line graph while the egress out of FortiDDoS as area graph. The outbound traffic is similarly graphed.

If the traffic is below Mbps (say in Kbps), it is shown with a letter in *m* in the graph and with a decimal in the tables. An example is shown in Figure 28.

If the traffic is above Mbps, it is shown normally as in Figure 29. Total Megabits passed through the appliance are shown in the table in the Total column. This can be used for billing purpose. Data for each port is available.

Figure 28: Example port statistics graph showing traffic in Mbps

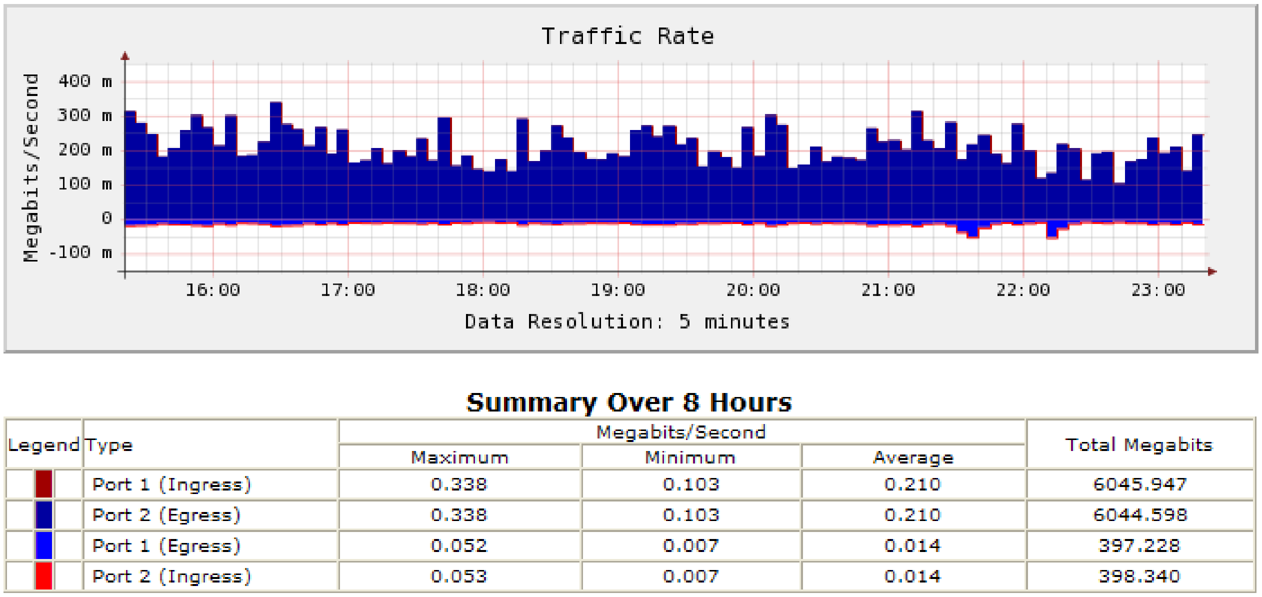

Summary Over 8 Hours

Legend	Type	Megabits/Second			Total Megabits
		Maximum	Minimum	Average	
🟥	Port 1 (Ingress)	0.338	0.103	0.210	6045.947
🟦	Port 2 (Egress)	0.338	0.103	0.210	6044.598
🟦	Port 1 (Egress)	0.052	0.007	0.014	397.228
🟥	Port 2 (Ingress)	0.053	0.007	0.014	398.340

Figure 29: Another example port statistics graph showing traffic in Mbps

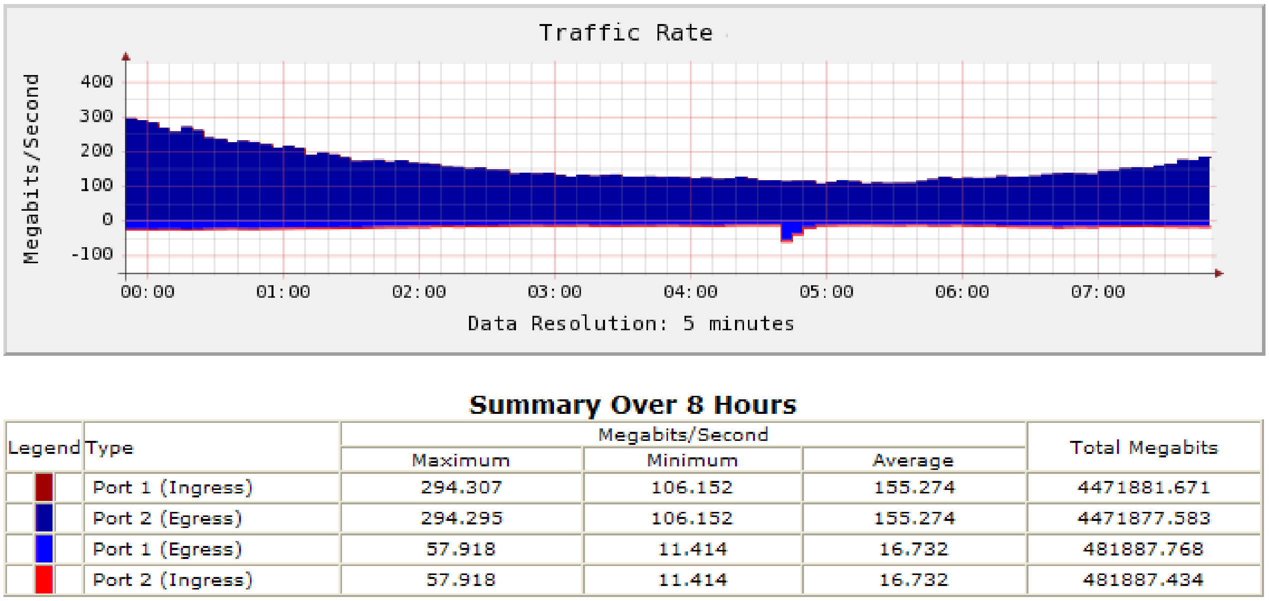

Summary Over 8 Hours

Legend	Type	Megabits/Second			Total Megabits
		Maximum	Minimum	Average	
🟥	Port 1 (Ingress)	294.307	106.152	155.274	4471881.671
🟦	Port 2 (Egress)	294.295	106.152	155.274	4471877.583
🟦	Port 1 (Egress)	57.918	11.414	16.732	481887.768
🟥	Port 2 (Ingress)	57.918	11.414	16.732	481887.434

Interpreting the traffic statistics reports

FortiDDoS starts learning traffic pattern the moment it is introduced in the network and it never stops learning. It continuously records traffic statistics in the form of round robin database which you can view in graphical and tabular form. The traffic statistics is highly granular.

Your network has both mission critical services and other sanctioned but low-priority applications. If you have a situation, where the low priority applications are taking over the mission critical services, you need to develop a traffic hierarchy to understand where the top bandwidth is being used.

Traffic Statistics Reports help you prepare such a traffic hierarchy by giving you a tabular view in a sorted form. This is a great tool for traffic discovery.

You can use these reports to set thresholds correctly so that your mission critical services and users get the right priority.

FortiDDoS collects maximum packets/second rate for all parameter during a 5 minute period. This rate is the rate of packets in either inbound or outbound direction for a given VID for a particular network parameter. The appliance collects the rate every second but stores the maximum for a 5 minute period. Thus the maximum packet rate/second is stored as a 5 minute maximum of these values. Traffic statistics provides following information for each VID, for each type of traffic parameter in inbound as well as outbound direction:

- Maximum packets/second

  This is the maximum value of the maximum packets/second over the observation period (1 hour, 8 hours, 1 day, 1 week, 1 month or 1 year). E.g. during 1 hour period, there are 12 5-minute observation periods. Each 5-minute interval has a maximum per second rate that has been observed. The maximum packets/second specifies the maximum value across these 12 periods of 5 minute intervals. The rate however is still in per second.
- Minimum packets/second

  This is the minimum value of the maximum packets/second over the observation period. E.g. during 1 hour period, there are 12 5-minute observation periods. Each 5-minute interval has a maximum per second rate that has been observed. The minimum packets/second specifies the minimum value across these 12 periods of 5 minute intervals. The rate however is still in per second.
- Average packets/second

  This is the average value of the maximum packets/second over the observation period. E.g. during 1 hour period, there are 12 5-minute observation periods. Each 5-minute interval has a maximum per second rate that has been observed. The average packets/second specifies the average value across these 12 periods of 5 minute intervals. The rate however is still in per second.

These rates correspond to the period you have chosen for the report.

E.g., if your weekly report has following information:

Table 2: Example Traffic Statistics Report Row

Type	Traffic (Packets/Sec)							
	Maximum		Minimum		Average		Threshold	
	Inbound	Outbound	Inbound	Outbound	Inbound	Outbound	Inbound	Outbound
SYN Packets	100	20	5	3	20	10	25	15

The above table indicates that Maximum SYN packets traffic in the inbound direction was 100 packets/second and never more during the observation period. Similarly the maximum SYN packets never went below 5 packets/second during the observation period. However, the average rate in the inbound direction during the observation period was 20 packets/second. The Inbound Threshold has been set to 25 and the Outbound Threshold has been set to 15 packets/second.

Traffic statistics reports are available per VID as well as on Global basis.

For some items that are vectors, such as protocols, options, ports, etc., a sorted list in the form of Top n values are available. Values are sorted in descending order for either Average Inbound/Outbound, or Maximum Inbound/Outbound traffic rate.

Note: Traffic Statistics Reports are a prerequisite for setting thresholds using One Click Adjustment menu to system recommended values.

Note: Generating Traffic Statistics Reports as a prerequisite for setting thresholds using One Click Adjustment menu to system recommended values, ensure that they are taken after a non-attack period. This can also be achieved by setting all thresholds to high or factory default values before the learning period begins.

Current VID traffic statistics

The following reports are available for current VID.

- Layer 3

  Fragmented Packets

  Most Active Source

  Most Active Destination

  Dark Address Scan

  Number of unique sources

  Unlike other values, this give population of the sources rather than their rate. Higher values are indication of DDoS

  Top Protocols

  The above report prepares a sorted list of protocols. The protocols in “My Protocols” list are highlighted. Those with zero traffic during the observation period are ignored.

- Layer 4

  SYN Packets

  SYN Packets/Source

  Most Active Connection

  Legitimate IPs

  SYN Packets/Destination

  ACK Packets/Destination

  RST Packets/Destination

  FIN Packets/Destination

  Established Connection Packets/Destination

  Number of Established TCP Connections

  Unlike other values, this give population of the connections rather than their rate. Higher values are indication of DDoS.

  TCP Ports (Top 100 ports)

  UDP Ports (Top 100 ports)

  ICMP Type/Codes (Top 100 items)

  The above three sections report a sorted list of top values. The option values in *MyLists* are highlighted. Those with zero traffic during the observation period are ignored.

  Ephemeral TCP Ports > 1023

  This set is treated as a range and the statistics is gathered for the whole set.

Ephemeral UDP Ports > 1023

This set is treated as a range and the statistics is gathered for the whole set.

- Layer 7

  HTTP Op codes

  URLs, Hosts, Referers, Cookies, User-Agents

Viewing current VID traffic statistics reports

To view the traffic statistics reports for the current VID:

- Click *Show* > *Current VID* > *Reports* > *Traffic Statistics.*
- The process of generating reports is a slow one.
- Different tables start appearing as they are generated. Until they are ready, you will see a comment: *Please wait while we retrieve the data...*
- Change the period for which you want to take the report by clicking on the appropriate time period hyperlink, 1 Hour, 8 Hours, 1 Day, 1 Week, 1 Month, 1 Year.
- Change the VID by clicking on the appropriate VID hyperlink in the top frame.

Global traffic statistics

Certain traffic parameters are independent of VIDs. Hence they are shown as global traffic statistics. These correspond to the four ports - LAN 1, WAN 1, LAN 2 and WAN 2.

Note: Please note that the time period of observation varies with the period you have selected. E.g. for 1 hour, 8 hours, 1 day periods, the traffic statistics shown is in terms of packets/5 minutes. For 1 week, the statistics is in terms of packets/hour, for 1 month report it is packets/3 hours, and for 1 year report is it packets/45 hours.

The following statistics are reported:

- Frame Type:

  Unicast

  Flow Control

  Multicast

  Broadcast

- Frame Size

  64

  65 - 127

  128 - 255

  256 - 511

  512 - 1023

  1024 - 1518

  Jumbo

- Concurrent Connections Per Source

  Since a single source can connect to multiple VID destinations, the concurrent connections per source statistics is maintained globally.

- Concurrent Connections Per Destination

  Since a single destination can connect to multiple VID sources or vice versa, the concurrent connections per source statistics is maintained globally.

Viewing global traffic statistics reports

To view the global traffic statistics reports:

- Click *Show* > *Global* > *Reports* > *Traffic Statistics*.
- Change the period for which you want to take the report by clicking on the appropriate time period hyperlink, 1 Hour, 8 Hours, 1 Day, 1 Week, 1 Month, 1 Year.

Configuration of layer 3, 4, and 7 thresholds

FortiDDoS is a network behavior analysis (NBA) system. For all continuously monitored traffic, you can set the thresholds for each VID and in each direction (inbound as well as outbound).

Effects of crossing threshold

Packets are forwarded until the traffic exceeds the threshold of a specific parameter. If the traffic is over the threshold, then that traffic is blocked for the configured blocking period. After blocking period, the threshold is checked again.

Example 1:
Too many packets of a certain ICMP Type/Code in a direction to network

- Packets of the type/code will be dropped for 15 seconds in that direction and destined to that network. All other packets will be forwarded
- Source tracking will be done to determine if this is a single source attack
- After 15 seconds the rate threshold will be checked again

Example 2:
Too many mail messages coming to an SMTP server

- TCP packets destined to port 25 to that mail server/network will be dropped for 15 seconds in that direction and destined to that network. All other packets will be forwarded.
- Source tracking will be done to determine if this is a single source attack. If there is a single source, all packets from that source will be blocked for 15 seconds (based on the Source blocking period).
- After 15 seconds the rate threshold will be checked again.
- The lost TCP packets will cause the mail clients to assume that the network is slowing down. They will back off and send the packets at a slower rate. No mail messages will be lost.

Example 3:
Too many SYN packets coming to a web server

- TCP SYN Packets destined to that web server/network will be checked in that direction and destined to that network if they come from an IP address in the legitimate IP(LIP) address table. If they do, they will be allowed as long as the rate of such SYN packets is lower than the zombie threshold. All other SYN packets will be forwarded without hurdles.
- If the IP address does not exist in the LIP table, and if the SYN cookie SYN flood mitigation method is enabled, a proxy three way handshake will be done to validate the IP address.
- After 15 seconds the rate threshold will be checked again.

Example 4:
Excessive Concurrent connections Per Source

- If there are too many concurrent TCP connections from a single source, new connections will be blocked until the concurrent connections goes below the threshold.
- Once the concurrent connection count goes down, new connection establishment from the same source is allowed.
- Source tracking will be done to determine if this is a single source attack. If there is a single source, all packets from that source will be blocked for 15 seconds (based on the Source blocking period).
- After 15 seconds the rate threshold will be checked again.

You can set thresholds for the following:

Table 3: FortiDDoS Configurable Traffic Thresholds

Layer	Type
3	Fragmented Packets Unique Sources Most Active Source Most Active Destination Dark Address Scan Protocols (256 values)
4	SYN Packets SYN Packets Per Source SYN Packets/Destination ACK Packets/Destination RST Packets/Destination FIN Packets/Destination Established Connection Packets/Destination TCP Connections Packets Per TCP Connection TCP Conn. Establishment Total Established TCP Connections TCP Ports UDP Ports ICMP Types/Codes Concurrent Connections Per Source (available as a Global Threshold) Concurrent Connections (available as a Global Threshold)

Table 3: FortiDDoS Configurable Traffic Thresholds (Continued)

Layer	Type
7	HTTP Op codes (Methods) URL, Hosts, Referers, Cookies, User-Agents Hash Indexes **Heuristic Thresholds for HTTP:** Associated Accesses Enforce Cookies Mandatory HTTP Headers Sequential Accesses Multiple Offending Accesses URLs/source **Thresholds for SIP:** SIP Invite Per Source SIP Register Per Source SIP Concurrent Invite Per Source

Adjusting the thresholds

You can adjust the thresholds at any given time. You can manipulate each individual threshold granularity in any direction independently for every VID.

FortiDDoS ships the appliance with high thresholds so that the appliance passes all traffic without blocking.

After a few days of training you can start adjusting the thresholds.

This process could be interactive based on the system's behavior. If you think that there are false positives, you can adjust the thresholds higher. If your attacks are not getting detected, you can adjust the thresholds lower.

Once you have adjusted the thresholds, there is usually no need to keep adjusting the thresholds, as the system keeps continuously learning the traffic characteristics and readjusts the thresholds upwards if there is a trend or seasonality in upward direction.

You need to adjust the thresholds if the traffic profile is thoroughly changing such as by addition of a new service or by replacement of servers etc.

Guessing the thresholds

Setting the right thresholds is key to performance of an NBA system. Since there are so many dimensions that you can control thresholds in FortiDDoS and it is not easy to guess what would be correct thresholds, FortiDDoS Web-based Manager helps you in setting the right thresholds.

You can use the estimated thresholds based on the past history of traffic and conveniently use them as the baseline if the history period was free of attacks. Alternatively, you can visually analyze the traffic graphs and guess a rate that you think will be right and keep the threshold at that level.

Avoiding disruptions while adjusting thresholds

One way to avoid disruptions while adjusting the thresholds is to keep the FortiDDoS in *Serial Detection Mode*. In this mode, while the system reports all attacks and drops, it does not actually drop the packets. You can use this mode to adjust the thresholds upward or downward until you stop seeing false positives. Use the *Monitor* or *Show* > *Current VID* > *Aggregate Drops or Show* > *Current VID* > *Reports* to see the drop rate.

Once you are satisfied, you can set the appliance in the *Serial Prevention Mode*.

Four Mechanisms for Adjusting Thresholds Using One Click Adjustment wizard

FortiDDoS allows you to set thresholds in four different ways:

1 Factory Defaults

- This option allows you to set the thresholds in a VID to factory defaults — which is the line rate value.

2 Adjust Minimum Thresholds

- This option allows you to adjust the minimum thresholds up or down by a certain percentage. This is useful when a flash flood is expected due to events such as advertisement campaigns.

3 Easy Setup

- This option is useful when the appliance has to be deployed in an unknown environment without much time left for training the appliance.
- Only certain key thresholds are adjusted here based on empirical knowledge and it is expected that these adjustments allow for protecting against common attacks.
- A later fine tuning/retraining is expected after the attack is over to set the thresholds properly via training.

4 System Recommended Thresholds

- This is the most common and recommended way to set the appliance threshold. The system recommended values are based on Traffic Statistics Report generated as part of the base-lining process.

Setting all thresholds to factory defaults (high values)

There are occasions when you may want to set all thresholds to factory defaults. These may be due to reasons such as:

- You have installed a new server or network behind a VID or behind FortiDDoS. Therefore you want the appliance to learn the traffic pattern without dropping the packets.
- You want to ensure that the devices is not dropping any packets due to rate thresholds.

To set all thresholds to factory defaults or high values follow the steps below:

- Click *Configure* > *Current VID* > *Blocking Thresholds* > *One Click Adjustment.*
- On the right screen, select the radio button entitled *Set minimum thresholds to factory defaults (high values).*
- Click *Save*.
- Similarly for Global Thresholds related to Concurrent Connections/Source and Concurrent Connections/Destination, Click *Configure* > *Global* > *Blocking Thresholds* > *One Click Adjustment.*
- On the right screen, select the radio button entitled *Set minimum thresholds to factory defaults (high values).*
- Click *Save*.

All thresholds will then be set to high values.

Note: By setting all thresholds to factory defaults, does not set the appliance to factory defaults. If you have used ACLs and feature control settings, you must revert them manually if you wish to do so.

Easy Setup

In the event you are already under attack and need to deploy the unit without an initial learning period, follow the instructions.

- In the main menu, go to *Configure* > *Current VID* > *Blocking Threshold* > *One Click Adjustment* > *Easy Setup*.
- Set the *SYN* threshold at 500.
- Set the *SYN Packets/Source* threshold at 200.
- Set the *Most Active Source* threshold at 10000.
- Click *Save*.

If you wish to set the outbound thresholds, check the *Outbound* checkbox under *Direction* and click *Save*. You can then enter outbound threshold values.

Wizard based configuration of rate thresholds based on past traffic

Setting the thresholds based on immediate past traffic is the easiest way to begin. FortiDDoS provides you with an easy wizard.

This wizard supports a way to set thresholds to system recommended values based on past traffic. In addition; in the advanced mode; it supports following features:

- Ability to set thresholds for each VID and in both the directions based on past traffic report taken for a period of user's choice.
- Ability to set certain thresholds to high values.
- Ability to set certain thresholds to a constant value.
- Ability to set certain thresholds to a constant value in case the traffic for that parameter is below a threshold.
- Ability to set constant value threshold for ports above 1023 based on either the maximum or average value across all ports.

Note: You need to generate Traffic Statistics Reports as a prerequisite for setting thresholds using One Click Adjustment menu to system recommended values. Please ensure that the reports are taken after a non-attack period. This can also be achieved by setting all thresholds to high or factory default values before the learning period begins.

System recommended option

This is the easiest option and can be used by users who are new to NBA concepts. If you are an experienced user, you can either use the *Advanced Option* or configure thresholds manually using *Configure* > *Current VID* > *Blocking Thresholds* > *Layer 3, Layer 4 and Layer 7* menus.

To set the thresholds for a VID to system recommended values based on past traffic, click *Configure* > *Current VID* > *Blocking Thresholds* > *One Click Adjustment.*

- If you have not taken any traffic reports previously, the screen will show a message

  Please select *Show* > *Current VID* > *Reports* > *Traffic Statistics* to generate a report.

  If you see the above message, please generate a traffic statistics report for the period desired and return to this menu.

  If you have already generated traffic statistics reports, you will see a list of reports which exist on the system along with the date and time they were taken. If any of them is relevant, select the corresponding radio button.

> If all of them are irrelevant, please select *Show* > *Current VID* > *Reports* > *Traffic Statistics* to generate a new report and restart the process.

- Select *Set minimum thresholds to system recommended values.*
- Click on *Save.*
- FortiDDoS will automatically configure the thresholds for the VID to system recommended values based on the traffic report and some heuristics.
- You can do a similar operation for Global Thresholds using *Show* > *Global* > *Reports* > *Traffic Statistics.*

Advanced option

This is another wizard based option which can be used with some experience. If you are an experienced user, you can configure thresholds manually using *Configure* > *Current VID* > *Blocking Thresholds* > *Layer 3, Layer 4 and Layer 7* menus.

To set the thresholds using the wizard based on past traffic, click *Configure* > *Current VID* > *Blocking Thresholds* > *One Click Adjustment.*

- If you have not taken any traffic reports previously, the screen will show a message as shown below:

  > Please select *Show* > *Current VID* > *Reports* > *Traffic Statistics* to generate a report.
  >
  > If you see the above message, please generate a traffic statistics report for the period desired and return to this menu.
  >
  > If you have already generated traffic statistics reports, you will see a list of reports which exist on the system along with the date and time they were taken. If any of them is relevant, select the corresponding radio button. If all of them are irrelevant, please select *Show* > *Current VID* > *Reports* > *Traffic Statistics* to generate a new report and restart the process.

- Select *Set minimum thresholds to system recommended values.*
- Click on *Advanced Options.*
- Empirical Exceptions
  - If you would leave the TCP, UDP protocol, TCP ports 22,25,80, and 443 to the currently configured values, please check *Leave Certain Key Thresholds to High Values*.
  - Else the above thresholds are adjusted based on the Traffic Statistics.
- You can adjust the settings for Layer 3, Layer 4 and Layer 7 thresholds.
- Go to Section titled Layer 3.
- If you would like to set all the Layer 3 rate thresholds to high values, select the radio button titled *Set minimum thresholds to factory defaults (high values)*.
- Alternatively, if you want to adjust the thresholds based on past traffic report, select the radio button titled *Set minimum threshold to.*
- Enter a percentage figure for adjustment over maximum traffic. If you enter 100%, the rate threshold will remain as the maximum traffic. If you enter 300%, the rate threshold will be set to 3 times the maximum traffic in the report.
- In addition, if you want to ensure that the rate threshold never falls below a number even if the traffic was low, check the box titled *If maximum traffic is below.* The same number will appear in the box titled *set it to.* This will ensure that the rate thresholds for layer 3 are never below the above number.

- Repeat the above procedure for Layer 4 and Layer 7 section as well.
- Layer 4 section has additional exceptions for TCP and UDP ports above 1023. These ports are the ephemeral ports used in protocols such as FTP, RealAudio, Windows media, other streaming protocols and voice over IP applications. These are usually negotiated ports and thus the whole range can be treated as a set - while you can make additional exceptions for fixed ports within this range. At the same time, you may have services running on ports above 1023. To exclude such ports, and thereby avoiding high values on ports above 1023, you must make those service ports as adaptive ports.
- Select one of the following options:
  - Maximum traffic rate of all TCP/UDP ports above 1023 - This is the recommended option. This will ensure minimum false positives but has potential for extra allowance if some ports have exceptionally high traffic during some time.
  - Maximum of the average traffic rate of all ports above 1023 - Use this option if you want to rate limit across all ports based on average behavior rather than maximum possibility.
  - Specific Value - Use this option if you an manually guess the value or you want to set it to a certain value of your choice.
- Click on *Save***.**
- FortiDDoS will configure the thresholds for the VID to system suggested values based on the traffic report and some heuristics.
- You can do a similar operation for Global Thresholds using *Configure* > *Global* > *Blocking Thresholds* > *Layer 4 menu*.

A typical threshold configuration screen

Various menu items under Layer 3, 4, and 7 provide a way to configure the Minimum Thresholds for different traffic rate parameters in the system.

These are grouped according to the network layers to which they belong i.e. Layer 3, and Layer 4.

The interfaces for all these menu-items are very similar to each other. Refer to Figure 30. Here we describe a typical screen shot taken from *Configure* > *Current VID* > *Blocking Threshold* > *Layer 3* > *Protocols.*

Figure 30: Typical threshold configuration screen

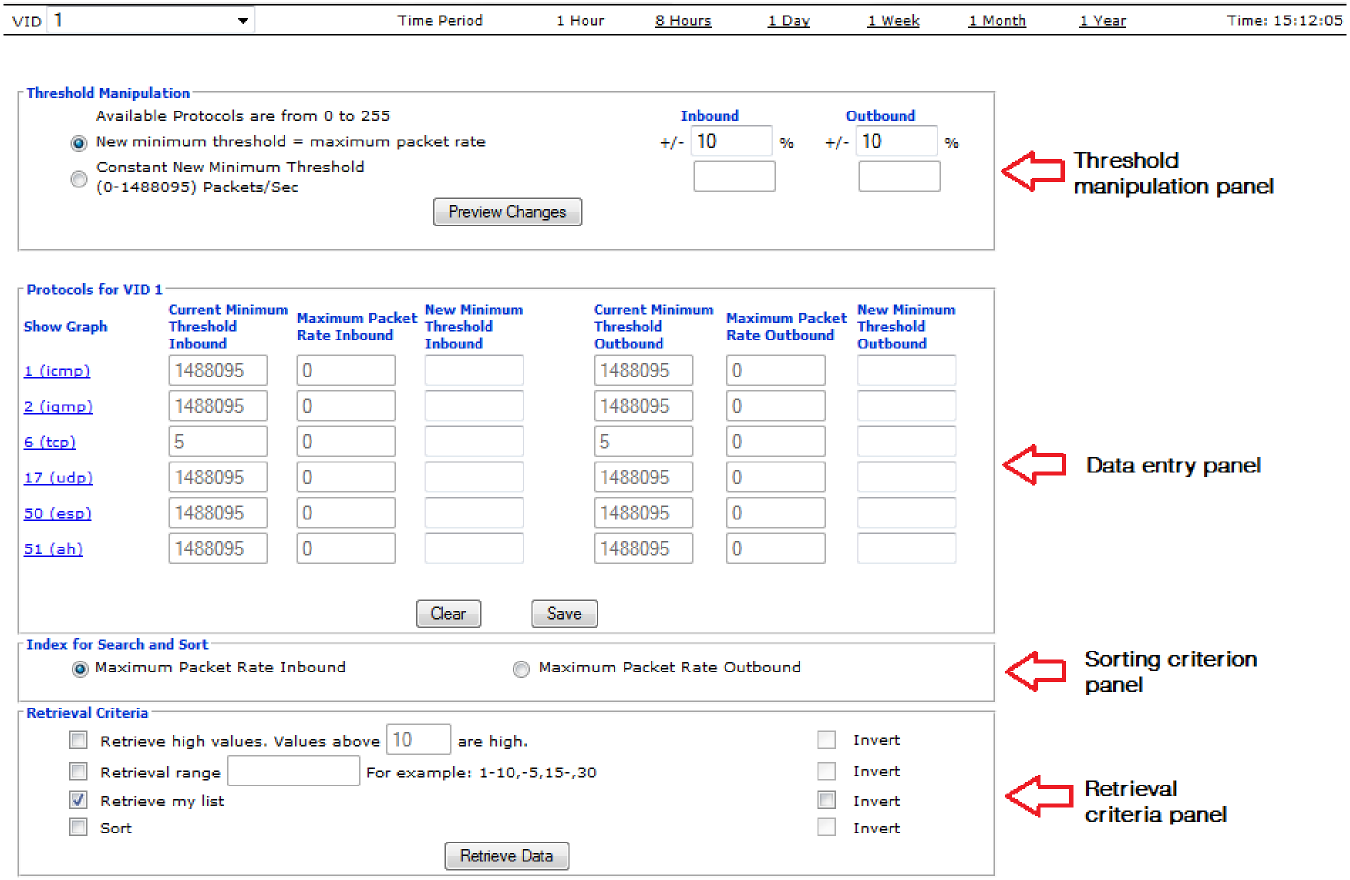

VID 1 ▾ Time Period 1 Hour 8 Hours 1 Day 1 Week 1 Month 1 Year Time: 15:12:05

Threshold Manipulation

Available Protocols are from 0 to 255

● New minimum threshold = maximum packet rate — Inbound +/- 10 % — Outbound +/- 10 %

○ Constant New Minimum Threshold (0-1488095) Packets/Sec

Protocols for VID 1

Show Graph	Current Minimum Threshold Inbound	Maximum Packet Rate Inbound	New Minimum Threshold Inbound	Current Minimum Threshold Outbound	Maximum Packet Rate Outbound	New Minimum Threshold Outbound
1 (icmp)	1488095	0		1488095	0	
2 (igmp)	1488095	0		1488095	0	
6 (tcp)	5	0		5	0	
17 (udp)	1488095	0		1488095	0	
50 (esp)	1488095	0		1488095	0	
51 (ah)	1488095	0		1488095	0	

Index for Search and Sort

● Maximum Packet Rate Inbound ○ Maximum Packet Rate Outbound

Retrieval Criteria

☐ Retrieve high values. Values above 10 are high. ☐ Invert

☐ Retrieval range For example: 1-10,-5,15-,30 ☐ Invert

☑ Retrieve my list ☐ Invert

☐ Sort ☐ Invert

This screen serves two purposes.

1 It provides a simple way to directly set thresholds to desired values.

2 It also acts as a wizard to set thresholds for selected items of interest to values based on arithmetic computations on observed traffic patterns.

Data entry panel

The Data Entry panel is in the center of the screen (labeled *Protocols for VID 1* in the figure). The first column shows the index and optionally a brief description of the items for which thresholds can be set. The desired threshold configuration values are entered in edit-boxes under *New Minimum Inbound/Outbound* columns. The *Current Minimum Inbound/Outbound* columns display the minimum threshold currently configured for that index in the given direction.

Setting thresholds to desired values

- Enter a value in the edit -boxes under *New Minimum Inbound/Outbound* columns.
- You are not required to enter all the values in case there are multiple edit-boxes. For example, you can enter just a single value in New Minimum Outbound. Only that new value will be applied and all blank edit boxes will retain their original values.
- If the Data Entry panel has more than a page-full of items, and if all the items are not displayed, pagination hyperlinks are provided for browsing through the items (not visible in above figure).
- You may use these pagination links to browse to any page and enter new threshold values in any of the edit-boxes.
- Click *Save* button.
- All values entered will be saved as the desired configuration and will be immediately visible under *Current Minimum Inbound/Outbound* columns.

Note: Any edit-boxes left unchanged will keep their existing values. For example, referring to above figure, there are 6 items each with Inbound and Outbound i.e. 12 possible data values to be entered.

Setting thresholds based on estimated thresholds

The *Maximum Estimated Inbound/Outbound* columns show the maximum of the estimated thresholds that the device has computed for that index item. This maximum value has been computed over a time-period that is defined by the currently selected *time-period* hyperlinks on the top of the screen.

- You have an option to display maximums in these columns computed over the past 1 hour, 8 hours, 1 day, etc.
- By clicking any of the *time-period* hyperlinks the values in the *Maximum Estimated Inbound/Outbound* columns are immediately updated to reflect the maximums over that time-period.
- By clicking the *Save* button, all the non-blank values entered in any page gets saved as the new configuration leaving other unchanged values intact.

Threshold manipulation panel

In addition to being able to enter desired threshold values by hand, this interface provides auto-filling of all the *New Minimum Inbound/Outbound* text-boxes with either a constant threshold value or with a value computed from the *Maximum Estimated Inbound/Outbound* value for that item.

Setting thresholds to a constant value

To fill all the text-boxes with a constant value:

- Select the *Constant New Minimum Threshold* radio-button and fill in the desired values in the corresponding Inbound and Outbound text boxes.
- Click the *Preview Changes* button. The screen refreshes and the values entered in the Inbound/Outbound text-boxes are copied into the *New Minimum Inbound and Outbound* columns respectively in the Data Entry panel. Note that clicking the *Preview Changes* button does not change any configuration. It just serves as a wizard to conveniently auto-fill the threshold text-boxes with constant values.
- You have an option to further edit the filled in columns.
- You can clear all the auto-filled values at anytime by clicking the *Clear* button.
- Click the *Save* button to save the desired values.

Setting thresholds based on estimated thresholds

To fill all the text-boxes with a computed value based on the Estimated Thresholds:

- Select the *New Minimum Threshold = Maximum Estimated Threshold* +/- radio-button and fills in desired values in the corresponding Inbound and Outbound edit-boxes.
- Click the *Preview Changes* button. The screen refreshes and the text-boxes for each row in *New Minimum Inbound and Outbound* columns are filled in with a value that is computed using the currently visible Estimated Threshold for that item.
- You can increment or decrement the estimated threshold by entering a percentage factor in the Inbound and Outbound edit-boxes. The percentage factor can be different in the Inbound and Outbound directions. If you enter a negative value, the Estimated Threshold is decremented by the given factor.
- Click the *Save* button the newly computed values will be applied.

Retrieval criteria panel

There are many layer 3 and 4 parameter which have a large dimension. E.g. in case of layer 3 protocols, there are 256 possible protocol values, each with Inbound and Outbound thresholds to be configured. For easy access to a subset of items that can be configured in one shot, you are provided with a *Retrieval Criteria* panel where you can define the filter parameters to narrow down on the list of items displayed in the *Data Entry* panel.

You are provided with three options to retrieve and filter the list:

1. Based on maximum estimated threshold cutoff
2. Based on a range of index values
3. Based on a preconfigured *My list*

Each of these filters can be optionally inverted. Also, any of these filters can be combined by checking the corresponding check-boxes. In this case only the items that match all the checked filters are displayed.

Note: It is important to note that there are only three filters and each filter can be independently inverted.

Limiting retrieval of only certain items to focus on

- Click on the checkbox *Retrieve High Values* and enter a value in the edit box. This will retrieve only those items whose Maximum Estimated Inbound or Outbound threshold is greater than the value configured in the *Values greater than... are high* edit-box. Check the *Invert* box if you want to retrieve values lower than the value in the edit-box. Select the appropriate radio button in the *Index for Search and Sort* panel. That determines whether *Maximum Estimated Inbound Threshold* or *Maximum Estimated Outbound Threshold* is used for this purpose.
- Click on the checkbox *Retrieval Range* to retrieve only those items whose index lies in range specified in the following edit-box. The range is defined as two numbers separated by a dash. The retrieved items include all indexes within and including the numbers. For example a range of 4-9 retrieves items 4,5,6,7,8 and 9. If first number is missing, it is assumed to be 0. If last number is missing it is assumed to be the maximum supported index (255 for Protocols). You can also enter a single digit in which case a single item of the specified index is retrieved. Thus the possible combinations are M-N, M-, -N, and M (note that M is equivalent to M-M). Click the *Invert* check box to retrieve items outside of the specified range.
- Click on checkbox *Retrieve my list* to retrieve only those items that have been previously configured in a user defined list. You can configure *My Lists* through *Configure* > *Current VID* > *My List.* When you enter this screen first time within a session, *Retrieve My List* is selected by default. Click the *Invert* check box in this row to retrieve items that are not in *My List*.
- Click on the *Sort* checkbox to sort items while displaying. By default the items that are retrieved are displayed in increasing index order. If you select *Sort*, items are sorted by either the *Maximum Estimated Inbound* or *Maximum Estimated Outbound threshold* values depending on what you have selected as the *Index for Search and Sort*. The currently selected radio button in the *Index for Search and Sort* panel determines whether Maximum Estimated Inbound or Maximum Estimated Outbound threshold is used for sorting. Click the *Invert* check box to reverses the sorting.

Sorting criterion panel

In addition to filtering the list of items that the user wishes to manipulate, you have a sort option to display the list in the desired sort order. The checkbox *Sort* does not affect the list of items that are retrieved, but only affects how they are displayed. Usually the items that are retrieved are displayed in increasing index order. If you select *Sort* checkbox, items are sorted by the Maximum Estimated Inbound or Outbound threshold values and displayed in decreasing order of the estimated threshold values. The currently selected radio button in the *Index for Search and Sort* panel determines whether Maximum Estimated Inbound or Maximum Estimated Outbound threshold is used for sorting. If you click the *Invert* checkbox, the sorting is done in reverse order.

Configuring layer 3 thresholds

The following types of configurable parameters are tracked at Layer 3.

- Fragmented packets
- Maximum packets per source
- Protocol
- TOS
- IP Options

To set the threshold and blocking periods at Layer 3 click on:

Configure > *Blocking Thresholds* > *Layer 3* and click on the relevant sub-category.

Configuring layer 4 thresholds

The following types of configurable parameters are tracked at Layer 4.

- SYN packets
- SYN Packets/Source
- SYN Packets/Destination
- ACK Packets/Destination
- RST Packets/Destination
- FIN Packets/Destination
- Established Connection Packets/Destination
- Maximum packets per connection
- Number of simultaneous TCP Connections
- Rate of establishment of new TCP connections
- TCP Ports
- UDP Ports
- ICMP Types/Codes

To set the threshold and blocking periods at Layer 4 click on:

Configure > *Current VID* > *Blocking Thresholds* > *Layer 4* and click on the relevant sub-category.

SYN flood and zombie flood prevention: some notes

You can prevent SYN floods using several built-in techniques within FortiDDoS. You must be aware of following SYN flood prevention schemes in place in the appliance:

- SYN flood thresholds are bi-directional and on a per VID basis as well as per destination (corresponding to the most active destination). You can control these individually.
- FortiDDoS store non-spoofed IP addresses that have done a three-way handshake successfully in a large table called Legitimate IP (LIP) Address table. This table retires entries every 5 minutes. Therefore this table has IP addresses which have recently connected successfully. Under SYN flood situation, i.e. when the SYN flood threshold is crossed, the LIP table is used to validate new connections. If the new connection request is from an address in this table it is allowed otherwise it is denied.
- Additionally, you can enable SYN Flood Mitigation methods through *Configure* > *Current VID* > *Advanced Options* > *Feature Controls* > *SYN Flood Mitigation*, new connection request can be validated for anti-spoofing on-the-fly. With this feature enabled, the FortiDDoS proxies the server and ensures that the client actually exists. If the client is able to responds to the FortiDDoS packets and they are valid, then the client IP is added to the LIP table.
- Additionally, you can set the right threshold for Legitimate IPs. This will ensure that too many legitimate IPs aka Zombies are not able to overwhelm the protected server. You can set this threshold using *Configure* > *Current VID* > *Layer 4* > *TCP Conns.* > *TCP Conn. Establishment*. This will ensure prevention of Zombie floods.
- Additionally, you can set the right threshold for Legitimate IPs. This will ensure that too many legitimate IPs aka Zombies are not able to overwhelm the protected server. You can set this threshold using *Configure* > *Current VID* > *Layer 4* > *TCP Conns* > *TCP Conn. Establishment threshold*. This will ensure prevention of Zombie floods.
- Next important thresholds related to distributed DoS attacks are SYN/Source rate, Concurrent Connections/Source and Concurrent Connections/Destinations. You must adjust these values correctly to thwart attacks which show abnormal behavior in these areas.

Configuring layer 7 thresholds

The following types of configurable parameters are tracked at Layer 7.

- Packets with specific HTTP Op Codes (Methods)
- Packets destined to specific URLs, Hosts, Referers, Cookies, User-Agents (hash indexes)
- SIP Invite Per Source
- SIP Register Per Source

The following thresholds are available as global. To configure thresholds, use *Configure* > *GLOBAL* > *Blocking Threshold* > *Layer 7* > *Concurrent Invite Per Source*.

- SIP Concurrent Invite Per Source

Configuring blocking periods

By default every attack event on FortiDDoS lasts 15 seconds. Thus if a flood is detected, the packets that belong to that flood are blocked for 15 seconds. You can adjust the blocking periods by clicking on *Configure* > *Global* > *Blocking Period.*

- To change the *Blocking Period* for most events, enter a number between 1 through 15 seconds.
- To change the *Blocking Period for Source Tracking*, enter a number between 1 through 65535 seconds (roughly 18 hours).
- Click *Save*.

Configuring adaptive limit

The *adaptive lim*it is an upper rate limit beyond which all traffic will be blocked. The *adaptive limit* is the upper limit for the *estimated threshold.*

The *adaptive limit* is defined as a percentage above the *configured threshold*. For example, an *adaptive limit* of 150% means that the FortiDDoS can use its Dynamic Threshold Estimation algorithm to raise the *calculated threshold* up to 150% of its original value, but no further. An *adaptive limit* of 100% means no Dynamic Threshold Estimation adjustment will take place once the *configured threshold* is reached.

The FortiDDoS device continuously monitors traffic patterns and estimates the applicable thresholds based on historical data. The threshold limits actually applied is never below the configured minimum thresholds (configured via the Layer 3, 4 and 7 menu items as explained above). In case the periodically estimated threshold is greater than the configured minimum threshold, the estimated threshold value is used. However there is also an upper limit to the threshold value applied. This upper limit is computed as a percentage of the current minimum threshold. The percentage value used is common amongst all the modules and is configured via this interface.

If a default value is set at 150 it means that the upper limit for the threshold for any module is equal to (“Current Minimum Threshold” * 150)/100. For example if the “Current Minimum Threshold” for say Layer3 > Protocols > 17 (UDP) in the Inbound direction is 10000, then the applicable threshold limit can never fall below 10000 and can never get above 15000.

- Use *Configure* > *Current VID* > *Blocking Thresholds* > *Adaptive Limit* to configure the Adaptive Limit.
- The valid range of values is from 101% to 200%.
- Click *Save* to save the new adaptive limit.

Disabling adaptiveness by adjusting adaptive limit

If you want to ensure that the thresholds you set are the thresholds you want to keep and you don’t want them to go above or below these values, an easy way to do that is to set the Adaptive Limit value to 101% - the minimum allowed value. By doing this the thresholds will vary between your set value and a 1% maximum deviation due to adaptiveness.

Adjusting all thresholds using one click

The FortiDDoS can be adjusted for one time anticipated changes in traffic.

This ensures that the legitimate traffic is not blocked.This increase could be due to a news flash or other important announcements.

All configured thresholds across all modules can be adjusted (incremented or decremented) by a certain factor via a screen interface. A user can enter a positive or negative integer (between –100 and +50) and all thresholds will be recomputed and applied. A negative value implies decrement and a positive value implies increment by the given factor. For example if the user enters +10 then all thresholds in the system will be recomputed as:

New Minimum Threshold = Current Minimum Threshold + (Current Minimum Threshold * 10)/100

- Use *Configure* > *Current VID* > *Blocking Thresholds* > *One Click Adjustment* to configure the thresholds with one click for the current VID.
- Click on the radio button titled *Adjust minimum threshold by +/.* Enter a percentage value.
- Click *Save* to save the new thresholds.

Note: To revert the thresholds back to original values, you will have enter a reverse value. E.g. the current minimum threshold was 100 and if you had adjusted the thresholds by +10%, the new value will become 110. To bring it back to 100, you will have to adjust the thresholds by -9.09%.

Note: Entering a value of -100% will lead to threshold values becoming zero. That means no traffic will pass through the appliance. Use the negative values carefully.

Setting Penalty Factors

The *Source Tracking Penalty* allows you to adjust the punishment given to a source being caught causing various floods. For inbound and outbound direction, there are independent penalty factors available for control. You may want to punish outsiders more than insiders depending on your situation. The Penalty is further specified through *Base Penalty Factor and Layer 7 Penalty Factor*. Layer 7 Penalty is calculated on top of the Base Penalty. E.g. if Base Penalty Factor is 2 and if Layer 7 Penalty Factor is 8, the total penalty used to track the source for Layer 7 drops is 2*8=16. Thus, e.g. under a User-Agent flood, if a source is sending a User-Agent that is overloaded, each time the source sends a layer 7 packet with the User-Agent, the packet will be considered equivalent to 16 packets. The Base Penalty Factor is used for all traffic whereas the Layer 7 Penalty Factor is used specifically for HTTP traffic.

This Penalty Factor is used to track Sources for floods associated with URL, Host, Referer, Cookie and User Agent

The *Destination Tracking Penalty Factor* allows you to adjust a factor that helps identify a destination being attacked. For inbound and outbound direction, there are independent penalty factors available for control. When multiple destinations are under attack, this tracking factor is used to multiply incoming packet counts. A threshold monitors the packets to individual destinations. The destination that reaches the threshold sooner is identified as an attacked destination. Too low a number will not correlate destinations faster and too high a number will falsely identify destinations that may not be under attack.

- Click *Save* to save the new penalty factors.

Figure 31: Setting Penalty Factors

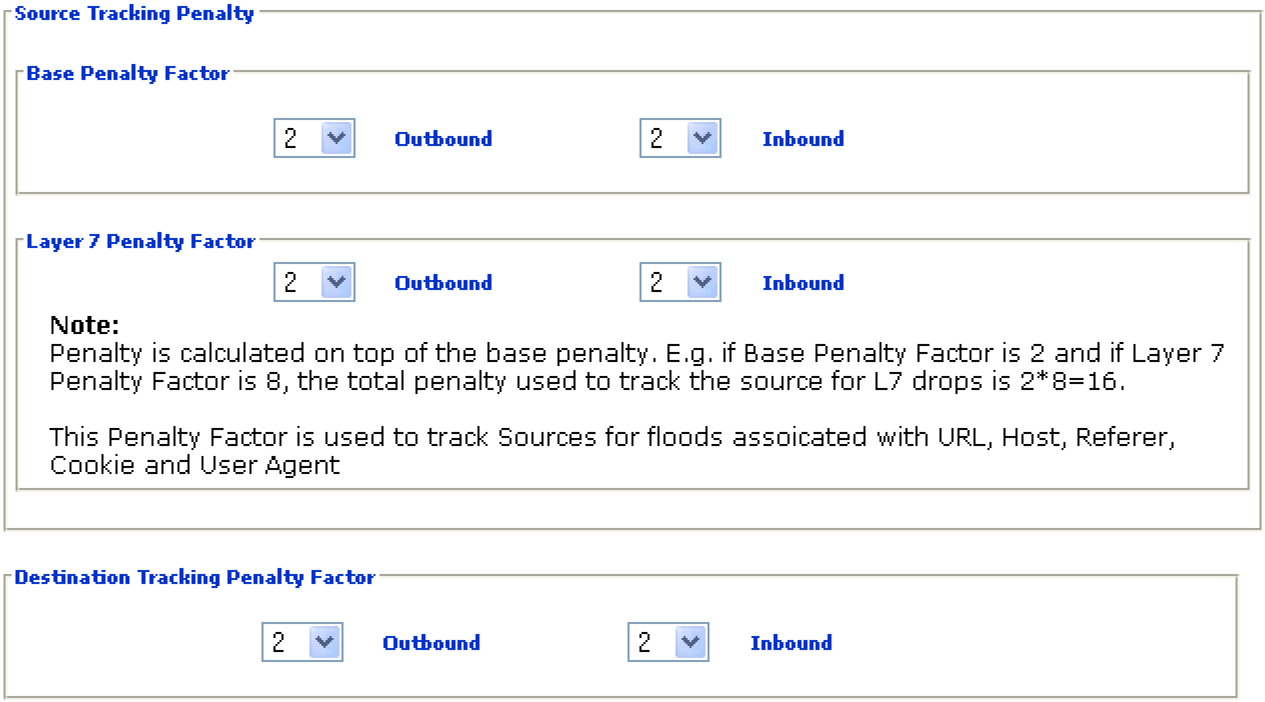

Source Tracking Penalty

Base Penalty Factor

2 Outbound 2 Inbound

Layer 7 Penalty Factor

2 Outbound 2 Inbound

Note:
Penalty is calculated on top of the base penalty. E.g. if Base Penalty Factor is 2 and if Layer 7 Penalty Factor is 8, the total penalty used to track the source for L7 drops is 2*8=16.

This Penalty Factor is used to track Sources for floods assoicated with URL, Host, Referer, Cookie and User Agent

Destination Tracking Penalty Factor

2 Outbound 2 Inbound

Configuring dark address scan prevention

FortiDDoS can monitor dark address scan activity on each VID independently.

We define dark scan as an activity in which a single source scans a dark destinations or a spoofed dark address is the source of a packet. Dark addresses are defined via a table.

Configuring dark address scan

Use *Configure* > *Current VID* > *Advanced Options* > *Feature Controls* > *Layer 3* > *Source Tracking* > *Source Tracking Feature Controls* > *Dark Address Scan* to access the screen. You will see a dialog-box shown in Figure 32.

Figure 32: Configuring dark address scan

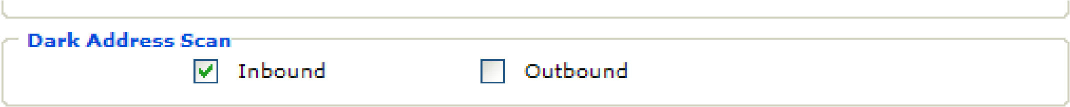

Dark Address Scan

☑ Inbound □ Outbound

- Enable Dark Address Scan by clicking *Inbound* and/or *Outbound* check boxes.
- Click *Save* to save the configuration.

Note: If you want to see the events corresponding to Dark Address Scan, remember to click on *Monitor* and check the *Event Categories* and *Scan Events* options.

Configuring dark address list

Use *Configure* > *Global* > *Dark Address List* to access the screen. You will see a dialog-box shown in Figure 33. This list needs to be filled with known dark or bogon addresses. A good place to start is at:

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

These addresses have been already entered as default values for your convenience.

Let us say, you want to deny inbound spoofing packets from the internet having source address as a private addresses 192.168.x.x, you need to define a dark address row as IP Address 192.168.0.0, Netmask as 255.255.0.0, Source checkbox as checked, Destination checkbox as unchecked, LAN 1 (corresponding to outbound port) unchecked and WAN 1 (corresponding to inbound port) should be checked.

As another example, let us say, you want to deny outbound spoofing, i.e. you want to ensure that addresses that are not in your inside LAN, you want to deny. Say you want to deny packets having source address as a private addresses 172.16.x.x, you need to define a dark address row as IP Address 172.16.0.0, Netmask as 255.255.0.0, Source checkbox as checked, Destination checkbox as unchecked, LAN 1 (corresponding to outbound port) checked and WAN 1 (corresponding to inbound port) as unchecked.

As another example, let us say, you want to an address range altogether as it is dark on inside as well as on outside. Say you want to deny packets having source address or destination address as a private addresses 10.x.x.x, you need to define a dark address row as IP Address 10.0.0.0, Netmask as 255.0.0.0, Source checkbox as checked, Destination checkbox as checked, LAN 1 (corresponding to outbound port) checked and WAN 1 (corresponding to inbound port) as checked.

To configure dark address list, use *Configure* > *Global* > *Dark Address List* to access the screen. You will see a dialog-box shown in Figure 34.

Figure 33: Configuring dark address scan

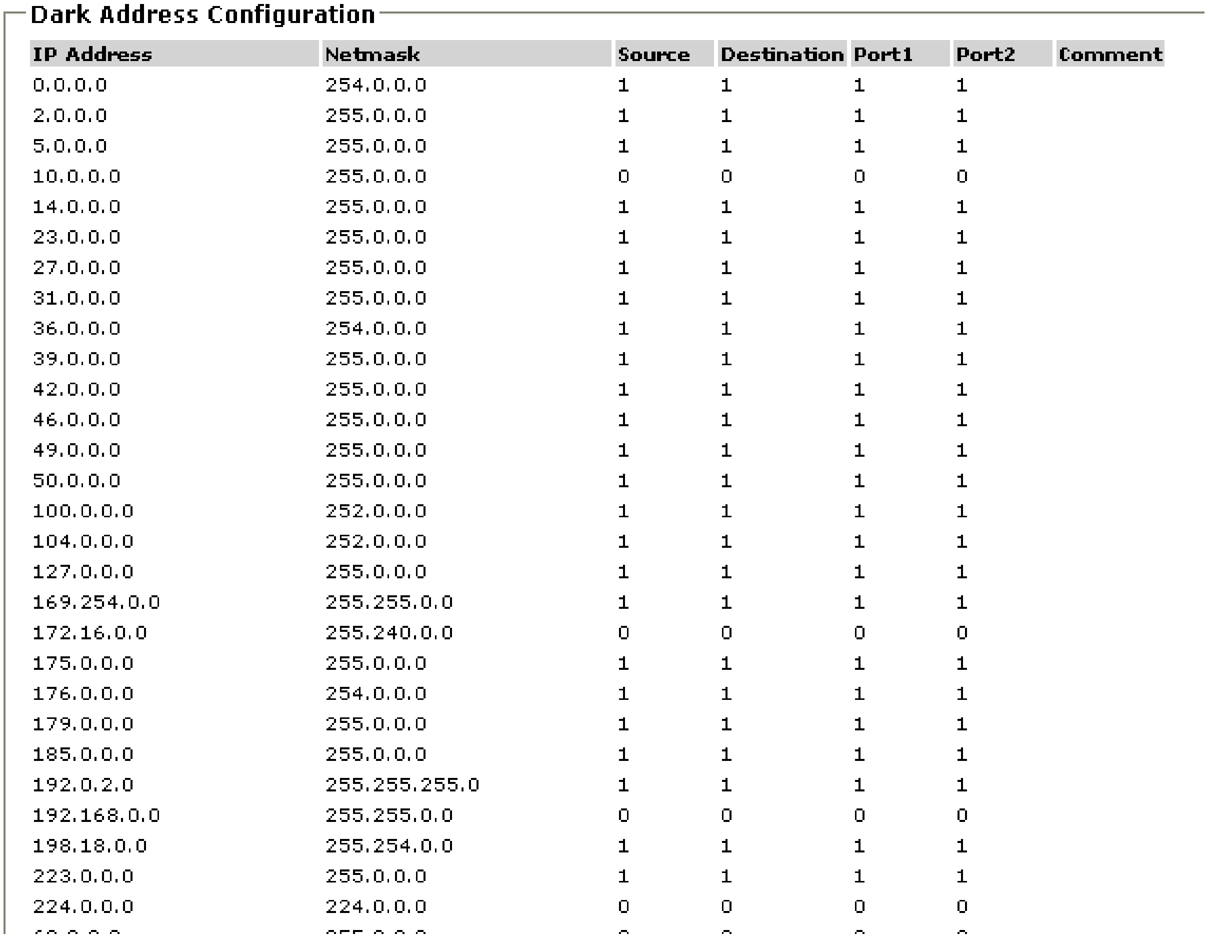

Dark Address Configuration

IP Address	Netmask	Source	Destination	Port1	Port2	Comment
0.0.0.0	254.0.0.0	1	1	1	1	
2.0.0.0	255.0.0.0	1	1	1	1	
5.0.0.0	255.0.0.0	1	1	1	1	
10.0.0.0	255.0.0.0	0	0	0	0	
14.0.0.0	255.0.0.0	1	1	1	1	
23.0.0.0	255.0.0.0	1	1	1	1	
27.0.0.0	255.0.0.0	1	1	1	1	
31.0.0.0	255.0.0.0	1	1	1	1	
36.0.0.0	254.0.0.0	1	1	1	1	
39.0.0.0	255.0.0.0	1	1	1	1	
42.0.0.0	255.0.0.0	1	1	1	1	
46.0.0.0	255.0.0.0	1	1	1	1	
49.0.0.0	255.0.0.0	1	1	1	1	
50.0.0.0	255.0.0.0	1	1	1	1	
100.0.0.0	252.0.0.0	1	1	1	1	
104.0.0.0	252.0.0.0	1	1	1	1	
127.0.0.0	255.0.0.0	1	1	1	1	
169.254.0.0	255.255.0.0	1	1	1	1	
172.16.0.0	255.240.0.0	0	0	0	0	
175.0.0.0	255.0.0.0	1	1	1	1	
176.0.0.0	254.0.0.0	1	1	1	1	
179.0.0.0	255.0.0.0	1	1	1	1	
185.0.0.0	255.0.0.0	1	1	1	1	
192.0.2.0	255.255.255.0	1	1	1	1	
192.168.0.0	255.255.0.0	0	0	0	0	
198.18.0.0	255.254.0.0	1	1	1	1	
223.0.0.0	255.0.0.0	1	1	1	1	
224.0.0.0	224.0.0.0	0	0	0	0	

- *Delete* addresses from the list which are being used as private addresses in your network.
- *Add* additional addresses that you think are dark addresses in your network on as sources or destinations.
- *Modify* addresses that you think are dark addresses in your network on as sources or destinations.

Figure 34: Delete, add, modify addresses

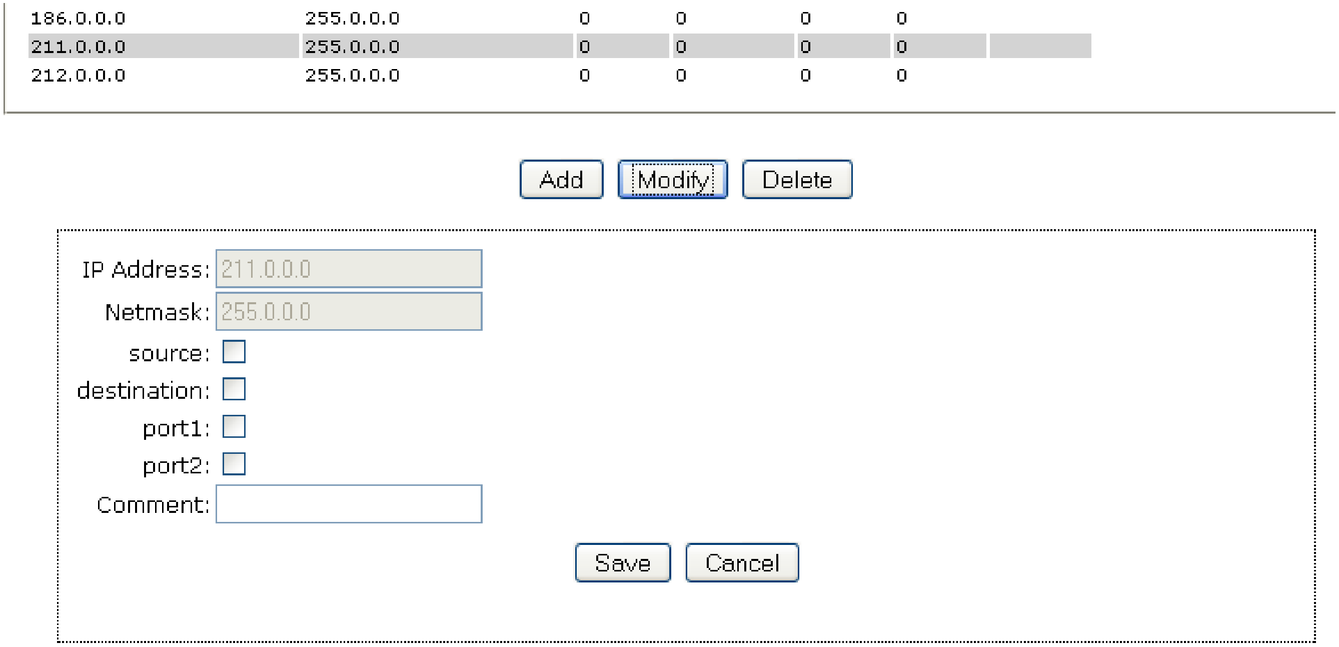

186.0.0.0	255.0.0.0	0	0	0	0
211.0.0.0	255.0.0.0	0	0	0	0
212.0.0.0	255.0.0.0	0	0	0	0

IP Address: 211.0.0.0

Netmask: 255.0.0.0

source: □

destination: □

port1: □

port2: □

Comment:

- If you want to enable or disable a particular column out of Source, Destination, LAN 1 or WAN 1, check/uncheck on the corresponding checkbox in the dialog box.
- Click *Save* to save the configuration.

Note: If you disable the feature control in Inbound direction and check the WAN 1 in dark address scan configuration, the dark address scan prevention will not work.

Note: If you disable the feature control in Outbound direction and check the LAN 1 in dark address scan configuration, the dark address scan prevention will not work.

Configuring dark address scan thresholds for source tracking

While every packet that contains either a dark address source or dark address destination is blocked by FortiDDoS, it is important to associate these packets with a source (if the source is not spoofed). Such an association is useful in isolating machines on the internal network which may have been infected due to a virus or worm.

FortiDDoS allows you to set a threshold to track the sources which are sending dark address packets. Use *Configure > Current VID > Blocking Threshold > Layer 3 > Dark Address Scan* to adjust the thresholds.

Unlike most thresholds, dark address scan thresholds (as well as rates) are calculated for a 60 seconds period. Thus the rates and thresholds shown on the screen are for 60 seconds period.

Configuration of My Lists for frequently used data

The menu items under *Configure* > *Current VID* > *My Lists* allow you to configure lists of items that you most often access in other parts of the user interface. This convenience feature is provided for all network parameters where the number of items to be configured (or monitored) is more than a handful. E.g., the Protocols module exposes 256 independent protocols for which thresholds need to be set or for which historical data needs to be viewed. But there are only a few protocols (e.g. TCP, UDP, ICMP) that you would usually be interested in since most of the internet traffic falls under these protocols.

A number of similar looking interfaces allow you to build lists of commonly used items for each module by adding or removing entries of your choice. Below we show a typical screenshot for these interfaces.

Figure 35: Managing My Lists

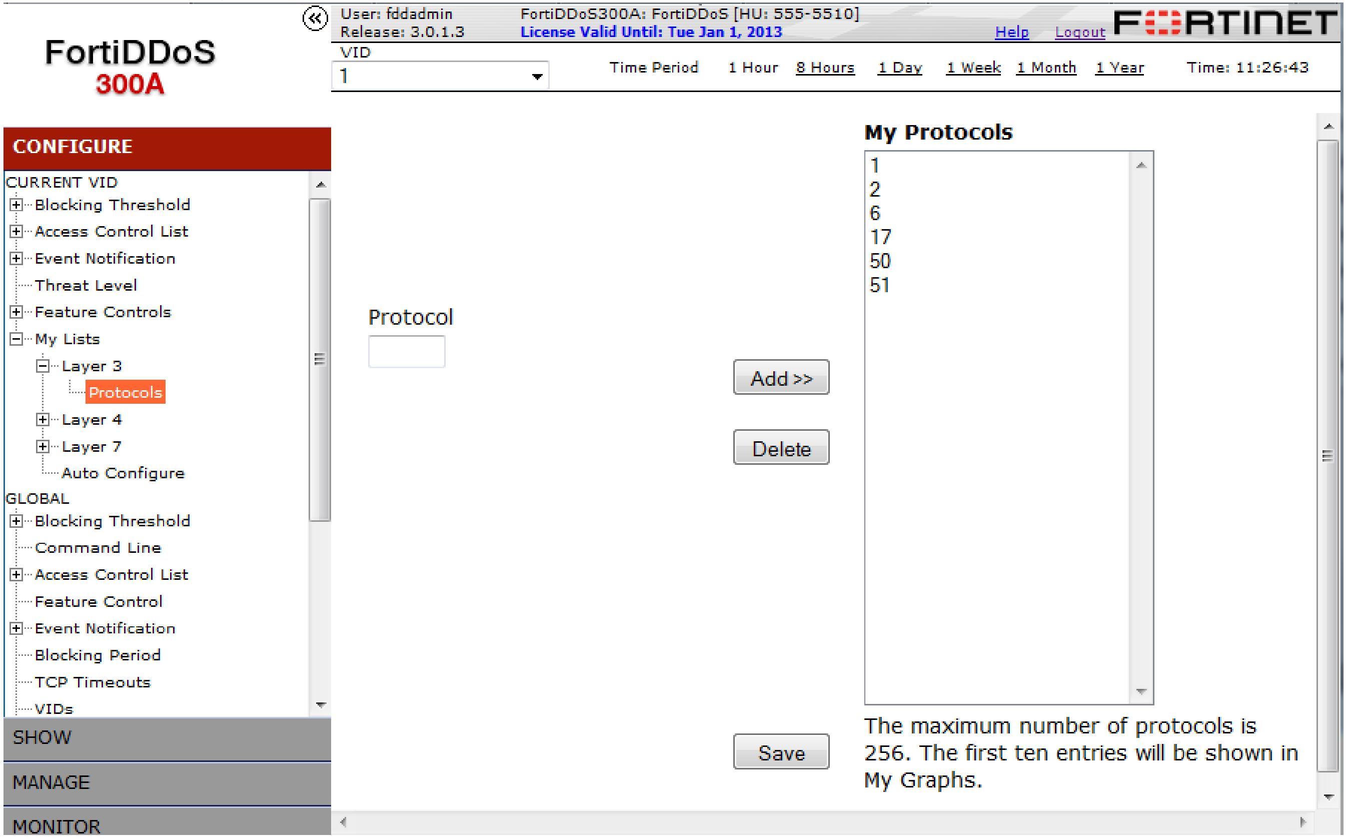

User: fddadmin
Release: 3.0.1.3
FortiDDoS300A: FortiDDoS [HU: 555-5510]
License Valid Until: Tue Jan 1, 2013
VID
1
Time: 11:26:43

My Protocols
1
2
6
17
50
51

The maximum number of protocols is 256. The first ten entries will be shown in My Graphs.

To add a item to *My List:*

- The list on the right includes items that are in the *My List.*
- You can add new values to the list by entering the values and then clicking *Add>>*
- You can delete existing values from *My List* by selecting them and then clicking *Delete.*

- Click *Save* to commit the changes to the FortiDDoS appliance.
- After the list is updated, the screen will be refreshed with the newly configured *My List* visible on the right side.

Automatically generating My Lists based on traffic statistics

Setting the My Lists based on immediate past traffic is the easiest way to begin. FortiDDoS provides you with an easy wizard.

This wizard supports a way to set My Lists to system recommended values based on past traffic. This wizard supports following features:

- Set My Protocols List to be top 10 protocols.
- Set Adaptive TCP Ports to be top 512 ports.
- Set Adaptive UDP Ports to be top 512 ports.
- Set Adaptive ICMP Type/Codes to be top 512 type/codes.

To set the above lists automatically based on past traffic:

- Click *Configure* > *Current VID* > *My Lists* > *Auto Configure*.
- If you have not taken any traffic reports previously, the screen will show a message:

  Please select *Show* > *Current VID* > *Reports* > *Traffic Statistics* to generate a report.
- If you see the above message, please generate a traffic statistics report for the period desired and return to this menu.
- If you have already generated traffic statistics reports, you will see a list of reports which exist on the system along with the date and time they were taken. If any of them is relevant, select the corresponding radio button. If all of them are irrelevant, please select *Show* > *Current VID* > *Reports* > *Traffic Statistics* to generate a new report and restart the process.
- Check the boxes corresponding to the lists you want to generate and corresponding to them select whether you want the lists to be generated based on inbound or outbound traffic.
- Click *Save* to generate the desired lists.

Utilizing My Lists

The configured lists are utilized in two places in the user interface:

- While configuring thresholds,
- While showing traffic graphs.

Utilizing My Lists to configure thresholds

While configuring thresholds, the number of items currently being configured can be narrowed down to just the items configured in these lists by checking “*My Lists*” in the “*Retrieval Criteria*” panel. This is the default filter used when the user configures these modules for the first time in a session.

Utilizing My Lists to show traffic graphs

You can use predefined *My Lists* to show commonly used traffic statistics.

These are available under:

- Show > Current VID > Layer 3 > My Graphs
- Show > Current VID > Layer 4 > My Graphs
- Show > Current VID > Layer 7 > My Graphs

Note: While you can define up to 512 items in each *My List***,** to avoid overcrowding of these graphs, if there are more than 10 items in these lists, only the first 10 items are shown in the graph.

Utilizing My Lists to define adaptive TCP/UDP ports, ICMP type/code

While FortiDDoS collects continuous traffic statistics for each and every TCP, UDP port and ICMP type/code, it limits the number of these items for adaptive threshold estimation to 512 each in each VID.

The periodic estimated thresholds that the FortiDDoS device computes are made only for the ports in My Lists and not for all the possible values.

Configuring SYN flood mitigation feature controls per VID

FortiDDoS provide powerful hardware based SYN flood mitigation. Three schemes for SYN flood mitigation are available one of which can be selected to work for each VID independently.

Figure 36: Configuring SYN flood mitigation control per VID

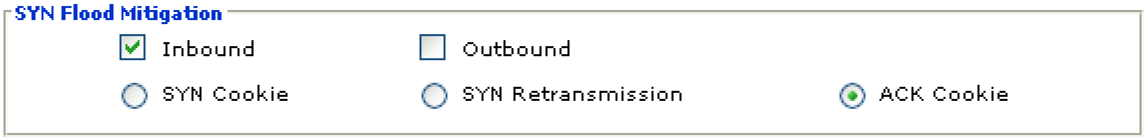

SYN Flood Mitigation

☑ Inbound □ Outbound

○ SYN Cookie ○ SYN Retransmission ● ACK Cookie

As shown in Figure 36 following three feature controls are available in Inbound and Outbound direction:

- SYN Cookie based SYN Flood Mitigation

  If this checkbox is checked anti-spoofing SYN proxy mechanism is enabled. SYN ACK is sent as part of proxy mechanism. If the IP which sent the SYN, responds back with an ACK, a RST/ACK is generated and the IP is added to the Legitimate IP address table for future use as a legitimate IP. If the client then retries, it will succeed in making a TCP connection.
- SYN Retransmission based SYN Flood Mitigation

  If this checkbox is checked a timeout based SYN retransmission scheme is enabled. With this scheme in place, the client is expected to send a SYN again to retry because the initial SYNs are dropped by the gateway hardware logic. If a preconfigured number of retransmitted SYNs arrive within a predefined time period, it ascertains existence of an non-spoofed source. The connection is then allowed to go through and is added into the LIP table.
- ACK Cookie based SYN Flood Mitigation

  If this checkbox is checked an ACK based anti-spoofing checking scheme is enabled. With this scheme in place, during SYN flood, the client is sent back an ACK packet with correct and another ACK packet with a wrong ACK number. Based on the behavior of the client through response, the hardware logic ascertains existence of an non-spoofed source. The connection is then allowed to go through and is added into the LIP table.

Configuring TCP state feature controls per VID

TCP is a stateful protocol. Depending on the deployment, you can select to allow or deny certain anomalies.

Following feature controls are available:

- Sequence Validation

  If this checkbox is checked, TCP state machine will ensure that TCP sequence numbers are validated across packets within a session.

- SYN Validation

  If this checkbox is checked, TCP state machine will ensure that only TCP SYNs from IP addresses in the legitimate IP addresses (that have done 3-way handshake in the past) are allowed.

- Track State Transition Anomalies

  If this checkbox is checked, TCP state machine will ensure that TCP state transitions follow the rules. E.g. if an ACK packet is received without a SYN/ACK packet being observed, it is a state transition anomaly.

- Foreign Packet Validation

  If this checkbox is checked, TCP state machine will ensure that foreign TCP packets without an existing TCP connection entry will be dropped.

  Here are some reasons for foreign packet occurrences even when you think they should not be present:

  1. Detection Mode:

  If the appliance is in Detection Mode, it thinks it dropped some packets or sessions, but actually they get forwarded. When the subsequent packets arrive for that connection, the appliance treats them as foreign packets - because the session may have been closed for those connections.

  2. Time Out Differences between Servers and IG appliance:

  The appliance TCP session timeout is kept lower than the typical Windows/Linux servers. This is done so that more new sessions can be handled per second. Sometimes this can create time differences and the sessions which have been purged in our appliance can be seen to send packets and the appliance will drop them. Most of the time, these packets are not data packets, they may be simple terminating packets.

  3. HTTP Browser Behavior

  Most people move on their browser from site to site without closing the browser. The browser eventually sends RST or FIN packets to close the connections to the site only when it is closed. The appliance sessions may have been purged by the time this arrives and there may not be any data associated with this session - just an RST packet.

  Therefore in most cases, the foreign packets are useful to filter out junk and you should not be unduly worried about them. Since the number is pretty high, we cannot afford to store the source/destination of each packet and therefore you may not be able to figure out who caused it.

- Allow Tuple Reuse

  If this checkbox is checked, TCP state machine will update the TCP entry if a tuple is reused. This is applicable only during the closed or close-wait, fin-wait, time-wait states when the connection is just about to retire.

- Allow Duplicate SYN in SYN_SENT

  If this checkbox is checked, TCP state machine will ensure that duplicate TCP SYN packets are allowed during the SYN-SENT state. They are allowed even if the sequence numbers are different.

- Allow Duplicate SYN in SYN_RECV

  If this checkbox is checked, TCP state machine will ensure that duplicate TCP SYN packets are allowed during the SYN-RECV state. They are allowed even if the sequence numbers are different.

- Allow SYN Anomaly , Allow SYN/ACK Anomaly, Allow ACK Anomaly , Allow Reset Anomaly, Allow FIN Anomaly

  If these checkboxes are checked, TCP state machine will ensure that duplicate TCP packets are allowed during any other state even if the sequence numbers are different from the existing connection entry. This is equivalent to allow the packet without updating an existing connection entry with the new information from the allowed packet.
- These features are available under: *Configure* > *Current VID* > *Advanced Options* > *Feature Controls* > *Layer 4* > *TCP State Machine* > *TCP State Feature Control* > *Session Feature Control*

Configuring aggressive aging feature controls per VID

Aggressive aging helps in reducing the connection overload on servers. When under attack from numerous slow connections, some servers get overloaded and need respite. They can be relieved by aggressively aging the connections based on some criteria such as aging those which are sitting idle.

FortiDDoS allow you to configure the aggressive aging features in following ways, as shown in Figure 37.

Figure 37: Configuring aggressive aging feature control per VID

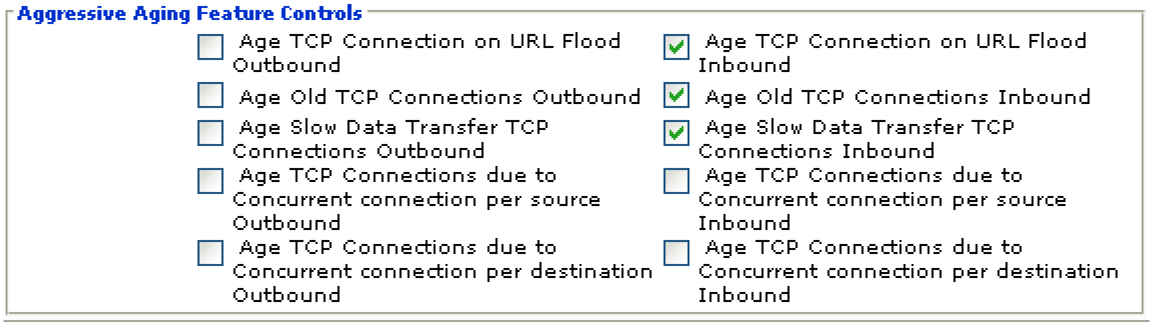

Aggressive Aging Feature Controls

☐ Age TCP Connection on URL Flood Outbound
☑ Age TCP Connection on URL Flood Inbound
☐ Age Old TCP Connections Outbound
☑ Age Old TCP Connections Inbound
☐ Age Slow Data Transfer TCP Connections Outbound
☑ Age Slow Data Transfer TCP Connections Inbound
☐ Age TCP Connections due to Concurrent connection per source Outbound
☐ Age TCP Connections due to Concurrent connection per source Inbound
☐ Age TCP Connections due to Concurrent connection per destination Outbound
☐ Age TCP Connections due to Concurrent connection per destination Inbound

- Age TCP Connection on URL Flood Inbound/Outbound

  These two feature controls allow the appliance to send RST packet to the servers to relieve them of excessive connection overload in case of URL flood (i.e. when a specific URL gets overloaded).
- Age Old TCP Connections Inbound/Outbound

  When the age idle connection feature control is on, the appliance will age out the connections from the device’s internal memory that are idle (no data transfer) for the programmed time and will also relieve the server by sending RST/ACK to it.

  Command to set the timeout to 120 seconds using *Command Line Interface* is:

  ```
  set tsm control OldConnTimeout 120
  ```

  Default timeout is 60 seconds, Max Value can be set is 36 hours (should program in seconds).

  You need to use Web-based Manager's *Command Line Interface* tab to make this command persistent.

- Age Slow Data Transfer TCP Connections Inbound/Outbound

  Identifies the TCP connections that has the slower data transfer rate based on the byte threshold and timeout settings.

  If any connection sends traffic less than the set threshold in slow connection byte count threshold (CLI command is: `set tsm control SlowConnByteCountThreshold 0x2` for 128 byte) within the time period specified by Slow connection timeout (CLI command is: `set tsm Control SlowConnTimeout 120` for 2 minutes) will be aged out by sending RST,ACK to the server for that connection.

  Byte count threshold is in terms of 64 bytes. A value of 1 means 64 bytes, 2 means 128 bytes and so on. Maximum value possible is: 4G * 64 bytes)

  Slow connection timeout is in seconds. Maximum value possible is 36 hours (programmed in seconds).

  Default value of byte count threshold is: 0 and timeout is: 60 seconds.

- Age TCP Connections due to Concurrent connection per source Inbound/Outbound
- Age TCP Connections due to Concurrent connection per destination Inbound/Outbound

  Ages the connections that are idle, i.e. have no data transfer for the programmed time specified by idle connections timeout. This is done by sending RST/ACK to the server.

  CLI command to set the timeout to 120 seconds is:

  ```
  set tsm Control ServerConnIdleTimeout 120
  ```

  Default value is 60 seconds, and Max value is 36 hours (program in seconds)

- These features are available under: *CConfigure > Current VID > Advanced Options > Feature Controls > Layer 4 > TCP State Machine > TCP State Feature Control > Aggressive Aging Feature Controls*

Configuring layer 7 HTTP heuristics feature controls (Per VID)

- These feature controls help you set features related to HTTP accesses based on heuristics.
- **Associate Resource Access** relates to the feature which ensures that a URL is accessed along with its associated resources. E.g. a typical index.html page is accessed along with other resources such as images and style sheets and never alone.
  - Normally a web page consists of multiple resource elements such as text, image files, css files etc.
  - Therefore when a webpage is accessed by a script, it may only access the text file without accessing the associated resources.
  - By ensuring a certain ratio between a main page and associated resource elements based on your site, you can ensure a heuristic access policy.
  - This ratio should be at least 1:1 or higher such as 1:2.
  - Only the denominator is shown in the listbox.

- **Enforce Cookies** relates to the feature which ensures that certain Cookie is presented to the server after an initial access. If your site does not set cookies, you may disable this feature.
  - You can force a cookie on subsequent accesses to your website after the first access.
  - Thus you can ensure that a typical session has more accesses with cookie than those without cookie.
  - You can disable this feature by disabling the check box if your site doesn't use cookies.
  - This ratio should be at least 1:1 or higher such as 1:2.
  - Only the denominator is shown in the listbox.
- **Mandatory HTTP Headers** relates to the feature which ensures that certain HTTP Headers are always present in a GET access to the URL. These headers are further defined in the HTTP Advanced menu.
  - In the Hypertext Transfer Protocol (HTTP), HTTP header fields contain the operating parameters of an HTTP request or response. With the request or response line (first line of message), they form the message header.The header fields define various characteristics of the data transfer that is requested or the data that is provided in the message body.
  - However, most scripted attacks may not have all the mandatory headers.
  - You can select a mandatory set of headers that you always expect in HTTP GET messages.
  - The threshold here is for a single IP to send a minimum number of packets without the selected mandatory headers before it is identified as an offender.
  - This period is defined under *Global* > *Advanced Options* > *Feature Control* > *URL Source Tracking Period for HTTP Accesses*.
- **URLs Per Source** relates to the feature which ensures that no single IP address retrieves more URLs/observation period than defined under HTTP Advanced menu.
  - The threshold here specifies how many HTTP accesses can any source within the VID perform per observation period.
  - This period is defined under *Global* > *Advanced Options* > *Feature Control* > *URL Source Tracking Period for HTTP Accesses*.
- **Sequential Accesses** relates to the feature which ensures that no single IP address retrieves same URL back to back without accessing any other URLs. This is a normal scripted access behavior and shows anomalous behavior.
  - The threshold here specifies how many HTTP accesses can any source within the VID do per observation period to the same URL back to back.
  - Such accesses would typically be done by scripted bots and are not done by humans.
  - This period is defined under *Global* > *Advanced Options* > *Feature Control* > *URL Source Tracking Period for HTTP Accesses*.

- **Multiple Offending Accesses** relates to the feature which ensures that no single IP address retrieves offends multiple anomalies sequentially. E.g. having less than mandatory headers combined with access without cookies. This is a normal scripted access behavior and shows anomalous behavior.
  - This feature combines 2 types of accesses which are considered offending viz. “mandatory HTTP headers” and “enforce cookies”.
  - If you enable this feature, any of the above is considered an offence.
  - During an observation period, if any source exceeds the threshold, it will be blocked.
  - This period is defined under *Global* > *Advanced Options* > *Feature Control* > *URL Source Tracking Period for HTTP Accesses*.
- **Rate Limit My URLs** relates to the feature which ensures that under attack a URL is rate limited per second rather than blocked for a duration. This ensures access to the URL while rate limiting it to a threshold set by the administrator.

Configuring TCP state feature controls (Global features)

- Track TCP Connections in Asymmetric Mode

  When the device is in asymmetric mode (i.e. traffic is seen only in inbound direction), device can still parse Layer 4 and Layer 7 headers for most floods and URL related features.

  If this feature is off, such floods won’t be detected if 2-way traffic is not completely seen by the appliance.
- Include all TCP states for Concurrent Connections Tracking

  During floods, it is possible that not all connections are established. Therefore the user can decide with this feature control whether to include only established connections or all connections (which may be also in the closing state).

  Using this feature is useful during some floods where the clients attack and close the connections quickly. If only established connections are counted, they appear to be lower number, but if closing connections are counted, one may get a true picture of the attack intensity.
- Drop Zero payload TCP Connections

During certain floods, the client botnets send empty packets without actually sending any data. This feature control allows you to remove such connections from the appliance’s memory table and send an RST/ACK packet to the server.

Using access control lists (ACLs)

ACLs on FortiDDoS can be used to allow or deny packets based on header values in layer 3, 4 and 7 headers. This is an easy-to-use function and allows simple rules to be set up. You can define rules based on a single header values. You cannot however combine multiple header values to make a rule. E.g. you can block packets from an IP address, or block packets coming to a destination port. But you cannot block packets coming from that a specified IP address to a specified port.

ACL functionality is available through header values for:

- Layer 3
- Layer 4
- Layer 7

Each VID administrator can define his/her own ACL rules.

ACL for layer 3

FortiDDoS allows granular ACLs for the following types of layer 3 packets in both inbound and outbound directions:

- Packets belonging to certain IP Protocols
- Fragmented packets
- Packets from specific sources or destinations
- Packets from specific geo-locations, i.e. countries
- Packets from specific IP addresses based on their reputation.

To configure ACLs for layer 3, click *Configure* > *Current VID* > *Access Control List* > *Layer 3.*

Configuring ACL for layer 3 protocols

Refer to Figure 38 for configuration of Layer 3 ACLs for protocols.

To configure ACLs for layer 3 protocols, click *Configure* > *Current VID* > *Access Control List* > *Layer 3* > *Protocols.*

- The list includes all 256 protocols with numbers starting from 0 through 256.
- Protocols that have a name assigned to them are shown with their names.

Figure 38: Configuring layer 3 ACLs for protocols

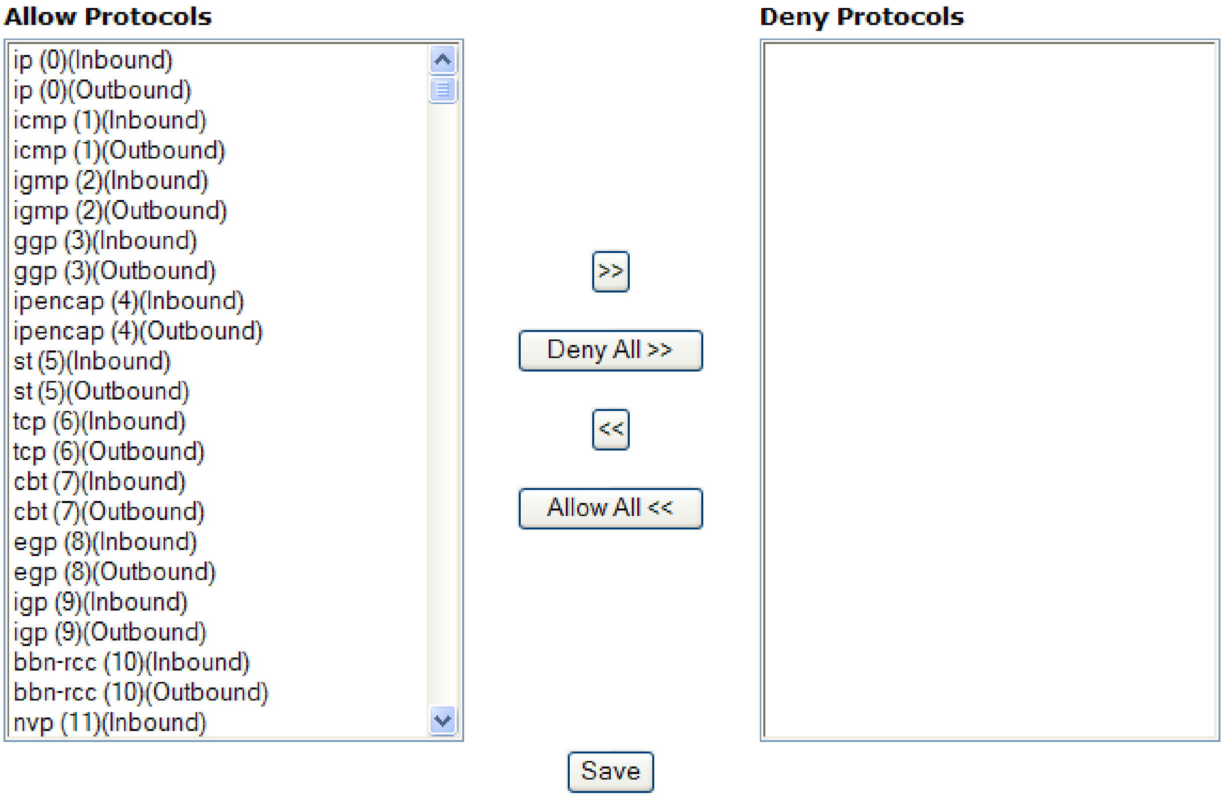

- User can select from the allow list and then click on **>>** button to move it to the deny list or click on **<<** button to move it to the allow list.
- The list on the left shows frames that are allowed for the VID.
- The list on the right shows frames that are to be denied for the VID.
- If you wish to deny all types of protocols, you can press *Deny All* >> button. Then you can add some protocols to be allowed.
- If you wish to allow all types of protocols, you can press *Allow All* << button. Then you can deny some protocols.
- You can specifically select certain frames and move them from allow list to deny list and back.
- Click *Save* to commit the changes to the FortiDDoS appliance.

Note: Some of the protocols such as icmp (1), tcp (6), udp (17) are ubiquitously used in the networks. You must take care that you understand your actions. You must understand your network and its packet behavior before you use the Layer 3 ACLs for protocols.

Configuring ACL for layer 3 fragmented packets

To configure ACL for layer 3 Option values, click *Configure* > *Current VID* > *Access Control List* > *Layer 3* > *Fragmented Packets.*

- The list includes inbound and outbound fragmented packets
- User can select from the allow list and then click on **>>** button to move it to the deny list or vice versa.
- The list on the left shows packet types that are allowed for the VID.
- The list on the right shows packet types that are to be denied for the VID.
- If you wish to deny all fragmented packets, you can press *Deny All* >> button.
- If you wish to allow all fragmented packets, you can press *Allow All* << button.
- You can specifically select certain values and move them from allow list to deny list and vice versa.
- Click *Save* to commit the changes to the FortiDDoS appliance.

Note: Some protocols such as multimedia streaming use fragmentation. You must take care that you understand your actions. You must understand your network and its packet behavior before you use the Layer 3 ACLs for fragmented packets.

Configuring ACL to deny specific source addresses

To configure ACL for denying layer 3 source addresses, click *Configure* > *Global* > *Access Control List* > *Layer 3* > *Deny Sources.*

Refer to Figure 39 for configuration of Layer 3 ACLs for denying specific source addresses.

Packets from these IP addresses will not affect the statistics for continuous learning for source addresses. However, other dimensions of the packets such as protocol, and ports will affect the corresponding statistics.

Figure 39: Configuring layer 3 ACLs for denying packets from specific source addresses

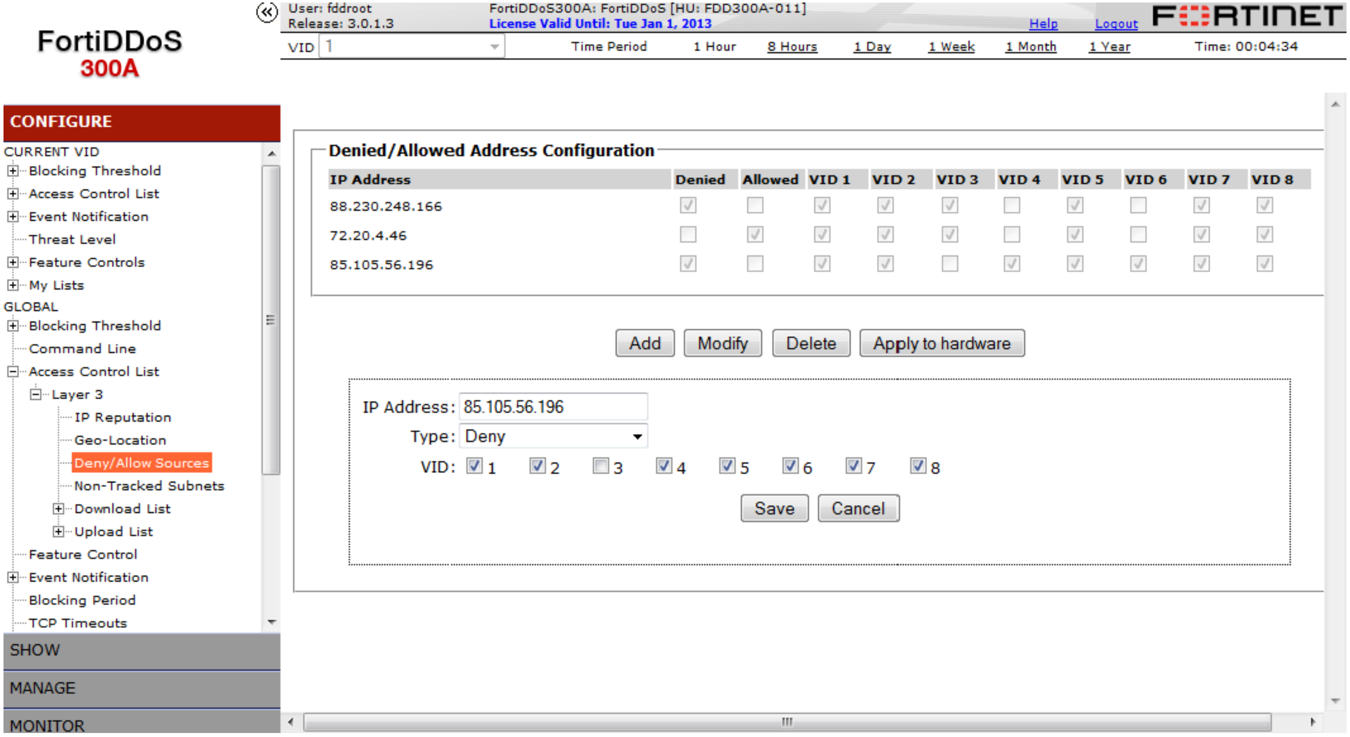

Denied/Allowed Address Configuration

IP Address	Denied	Allowed	VID 1	VID 2	VID 3	VID 4	VID 5	VID 6	VID 7	VID 8
88.230.248.166	☑	□	☑	☑	☑	□	☑	□	☑	☑
72.20.4.46	□	☑	☑	☑	☑	□	☑	□	☑	☑
85.105.56.196	☑	□	☑	☑	□	☑	☑	☑	☑	☑

IP Address: 85.105.56.196
Type: Deny
VID: ☑ 1 ☑ 2 □ 3 ☑ 4 ☑ 5 ☑ 6 ☑ 7 ☑ 8

- The list on the right includes sources that are denied.

To add new values to the list

1. Go to *Configure* > *Global* > *Access Control List* > *Layer 3* > *Deny/Allow Sources*.
2. Select the *Add* button.
3. Enter the IP address.
4. Select *Deny* from the *Type* drop down list.
5. Select VIDs for which the ACL should be applied.
6. Select *Save.*

To configure ACL to allow specific source addresses

1. Go to *Configure* > *Global* > *Access Control List* > *Layer 3* > *Deny/Allow Sources.*
2. Select the *Add* button.
3. Enter the IP address.
4. Select *Allow* from the *Type* drop down list.
5. Select VIDs for which the ACL should be applied.
6. Select *Save.*

- You can delete existing values in the list by selecting them and then clicking *Delete.*
- Once the list is created or updated, click *Apply to hardware* to commit the changes to the FortiDDoS appliance.

The user interface for these features is similar to denying certain sources except that this is an exception (or white) list. You can add IP addresses that are known to be good. Packets from those sources will always be allowed even if they exceed the thresholds.

- Packets from these IP addresses will not affect the statistics for continuous learning for source addresses. However, other dimensions of the packets such as protocol, and ports will affect the corresponding statistics.
- Add addresses for stations that perform backups and have high traffic profile.

Configuring ACL to enter non-tracked IP subnets

Certain IP addresses in network need to be left alone and should not affect the continuous learning and threshold estimation. FortiDDoS allow you configure such networks.

To configure ACL for entering non-tracked IP addresses, go to *Configure* > *Global* > *Access Control List* > *Layer 3* > *Non-Tracked Subnets.*

Refer to Figure 40 for configuration of Layer 3 ACLs for entering non-tracked source addresses.

Packets from these IP addresses will not affect the statistics for continuous learning for throughout the system. They go through the system without any restrictions or tracking as the name suggests.

Figure 40: Configuring layer 3 ACL for non-tracked sources

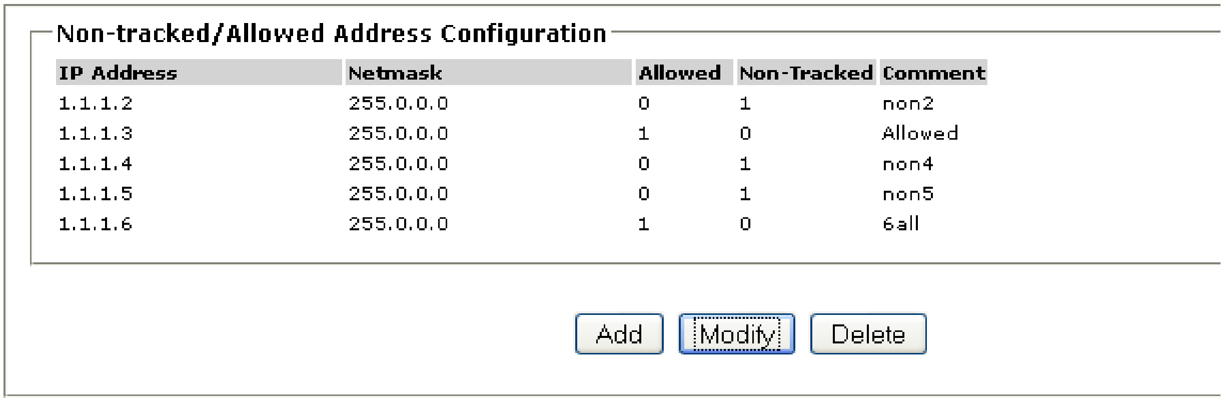

Non-tracked/Allowed Address Configuration

IP Address	Netmask	Allowed	Non-Tracked	Comment
1.1.1.2	255.0.0.0	0	1	non2
1.1.1.3	255.0.0.0	1	0	Allowed
1.1.1.4	255.0.0.0	0	1	non4
1.1.1.5	255.0.0.0	0	1	non5
1.1.1.6	255.0.0.0	1	0	6all

- The list includes sources that are not tracked by the system.
- You can add new values clicking *Add*.
- You can delete existing values in the list by selecting them and then clicking *Delete.*
- You can modify existing values in the list by selecting them and then clicking *Modify.*
- Click *Save* to commit the changes to the FortiDDoS appliance.

Configuring ACL to deny traffic from specific countries

Sometimes it is better to simply block traffic from geographic locations where you do not have any customers or expectation of traffic. Such a filtering can reduce the unnecessary traffic to servers by a significant percentage even during normal times.

To configure ACL for Geo-location, click *Configure* > *Global* > *Access Control List* > *Layer 3* > *Geo-location.*

Refer to Figure 41 for configuration of Layer 3 ACLs for entering Geo-locations.

Figure 41: Configuring layer 3 ACL for geo-locations

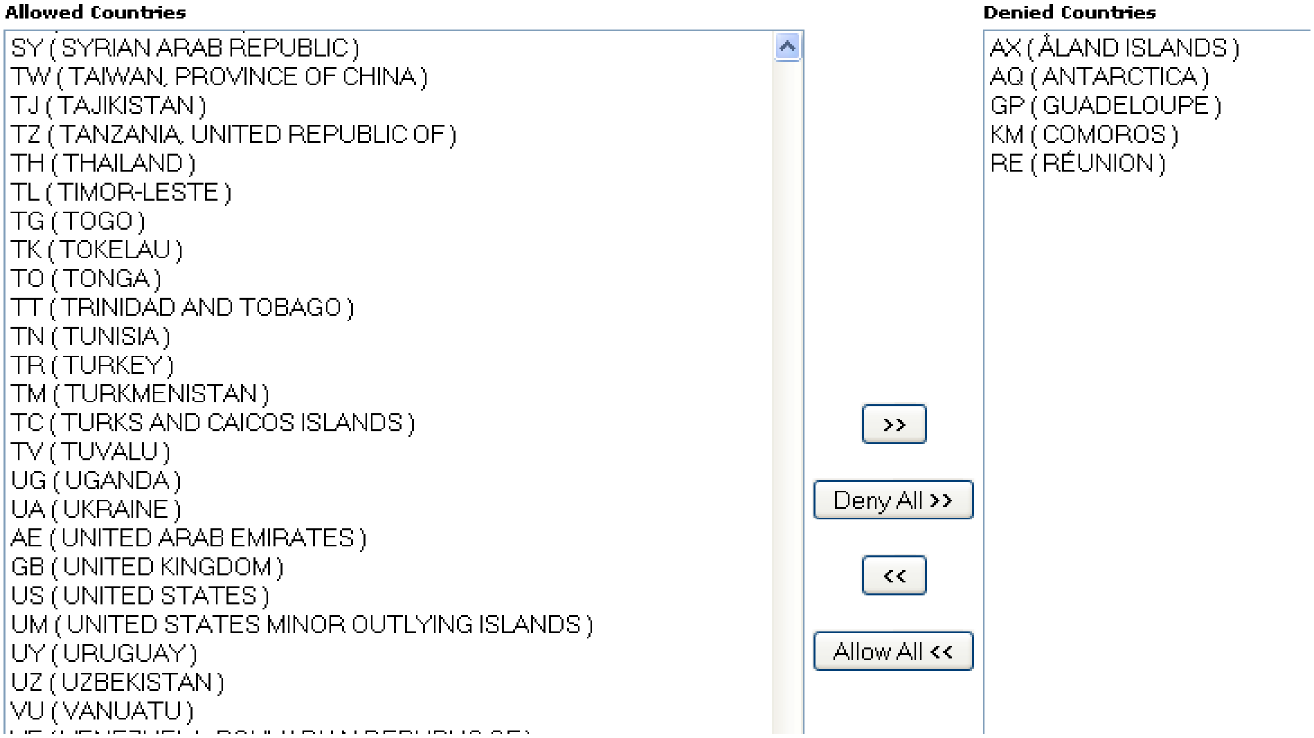

Allowed Countries

SY (SYRIAN ARAB REPUBLIC)
TW (TAIWAN, PROVINCE OF CHINA)
TJ (TAJIKISTAN)
TZ (TANZANIA, UNITED REPUBLIC OF)
TH (THAILAND)
TL (TIMOR-LESTE)
TG (TOGO)
TK (TOKELAU)
TO (TONGA)
TT (TRINIDAD AND TOBAGO)
TN (TUNISIA)
TR (TURKEY)
TM (TURKMENISTAN)
TC (TURKS AND CAICOS ISLANDS)
TV (TUVALU)
UG (UGANDA)
UA (UKRAINE)
AE (UNITED ARAB EMIRATES)
GB (UNITED KINGDOM)
US (UNITED STATES)
UM (UNITED STATES MINOR OUTLYING ISLANDS)
UY (URUGUAY)
UZ (UZBEKISTAN)
VU (VANUATU)

Denied Countries

AX (ÅLAND ISLANDS)
AQ (ANTARCTICA)
GP (GUADELOUPE)
KM (COMOROS)
RE (RÉUNION)

Such packets from denied countries will not be reported in any tables except in the aggregated drops. Subnet-based drops will also contain these drops.

- The Allow Countries List includes all countries currently registered with IANA.
- The Denied Country List includes all countries that you have denied so far. You can move countries from the Allowed Countries List to the Denied Countries List and vice versa by using >> and << buttons.
- You can move all to Allowed Countries List by clicking *Allow All >>.*
- You can move all to Denied Countries List by clicking *Deny All >>.*
- Click *Save* to commit the changes to the FortiDDoS appliance.

Configuring ACL to deny traffic from IP addresses based on their reputation

FortiDDoS allows you to create rules to block inbound and outbound access to thousands of IP addresses which have bad reputation. You can obtain lists of IP addresses and network ID ranges with bad reputation list by registering the device with FortiGuard IP Reputation Service. Once the appliance serial number is registered, user can schedule IP reputation list updates.

Your appliance's management port must have access to the Internet to download these lists. To configure ACL for IP Reputation, click *Configure* > *Global* > *Access Control List* > *Layer 3* > *IP Reputation*.

Refer to Figure 42 for configuration of Layer 3 ACLs for configuring IP Reputation based ACL. You can select which reputation category you want to download, can schedule list update frequency. These lists are automatically downloaded again by the appliance by the schedule.

IP reputation lists update status shows the status for the previous download attempt. If the download is successful and there is a new version of definitions are available, the lists are replaced, otherwise the previous versions are maintained.

Figure 42: Configuring layer 3 ACLs for IP reputation

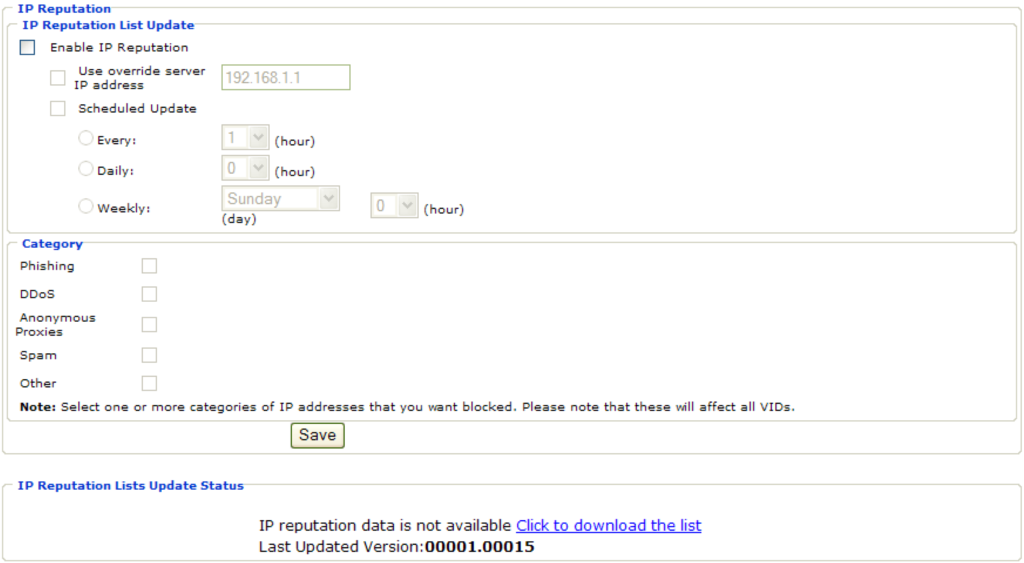

IP Reputation

IP Reputation List Update

☐ Enable IP Reputation

☐ Use override server IP address 192.168.1.1

☐ Scheduled Update

○ Every: 1 (hour)

○ Daily: 0 (hour)

○ Weekly: Sunday (day) 0 (hour)

Category

Phishing ☐

DDoS ☐

Anonymous Proxies ☐

Spam ☐

Other ☐

Note: Select one or more categories of IP addresses that you want blocked. Please note that these will affect all VIDs.

IP Reputation Lists Update Status

IP reputation data is not available Click to download the list

Last Updated Version:**00001.00015**

Configuring traffic from Proxy IP addresses based on their concurrent connections count

FortiDDoS detects Proxy IPs based on consistently higher number conncurrent connections from a single source IP over a specified period of time. If an IP meets this specification, then that IP is treated as a legitimate proxy behind which multiple IPs may be present. Proxy IPs will have higher thresholds than from a specific IP by a factor. This will ensure legitimate proxy IPs are treated different when compared to a non proxy source.

To configure Proxy IP, click *Configure* > *Global* > *Access Control List* > *Layer 3* > *Proxy IPs*.

Figure 43: Configuring layer 3 ACLs for Proxy IPs

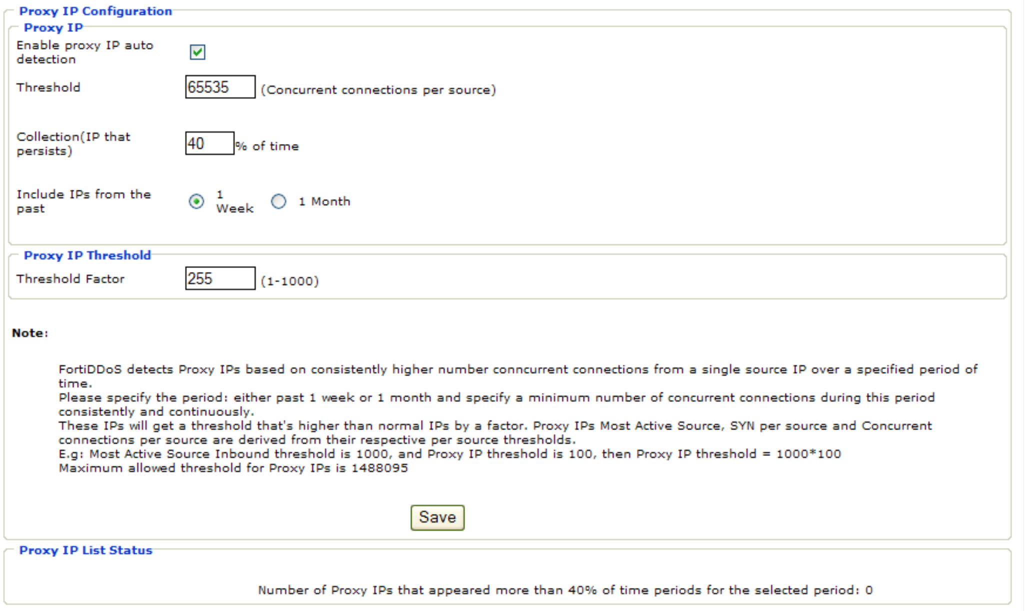

Proxy IP Configuration

Proxy IP

Enable proxy IP auto detection ☑

Threshold 65535 (Concurrent connections per source)

Collection(IP that persists) 40 % of time

Include IPs from the past ● 1 Week ○ 1 Month

Proxy IP Threshold

Threshold Factor 255 (1-1000)

Note:

FortiDDoS detects Proxy IPs based on consistently higher number conncurrent connections from a single source IP over a specified period of time.
Please specify the period: either past 1 week or 1 month and specify a minimum number of concurrent connections during this period consistently and continuously.
These IPs will get a threshold that's higher than normal IPs by a factor. Proxy IPs Most Active Source, SYN per source and Concurrent connections per source are derived from their respective per source thresholds.
E.g: Most Active Source Inbound threshold is 1000, and Proxy IP threshold is 100, then Proxy IP threshold = 1000*100
Maximum allowed threshold for Proxy IPs is 1488095

Proxy IP List Status

Number of Proxy IPs that appeared more than 40% of time periods for the selected period: 0

Configuring ACLs to deny traffic or not track traffic from specific sources using lists

Sometimes it is easier to simply upload a list of IP addresses and upload them to the FortiDDoS and block or not track such sources from sending traffic via the appliance. Such lists can be created using external sources such as web-server logs, firewall logs etc.

Another way is to download an existing list (if you have uploaded earlier) and append new addresses to the list or edit the list and upload again.

FortiDDoS allow you to download and upload lists of Non-tracked sources and denied sources.

To download these lists, click *Configure* > *Global* > *Access Control List* > *Layer 3* > *Download List.*

To upload these lists back to the appliance, click *Configure* > *Global* > *Access Control List* > *Layer 3* > *Upload List.*

Deny list file should contain only IP Addresses, no subnets, blank lines and non IP address characters.

Note: It is possible that one or more IPs will have a hash conflict. In such cases, only the latest IP of the hash conflicts will be programmed to the Firmware though the deny list has all the IPs that are uploaded.

Note: It is highly recommended that first download the existing deny list and append the new IPs to the list and upload back the modified list.

Note: Non-Tracked list file should contain only IP Addresses, subnets, 0/1 for Non-Track & Allow and a comment.

Note: List should not contain any blank lines and non IP/Mask address characters. It is highly recommended that first download the existing non-tracked list and append the new IPs to the list and upload back the modified list.
File Format:
IP Address SubnetMask Non-Track(0/1) Allow(1/0) Comment

Note: To Non-Track a Subnet column 3rd must be 1 and column 4th must be 0.

Note: To Allow a Subnet column 3rd must be 0 and column 4th must be 1.

Note: Non-Tracked Subnet example:
5.4.3.2 255.255.255.255 1 0 subnet5432

Note: Allowed Subnet example:
7.6.8.9 255.255.255.255 0 1 subnet7689

ACL for layer 4

FortiDDoS allows granular ACLs for the following types of layer 4 packets in both inbound and outbound directions:

- TCP, UDP Packets destined to some ports
- ICMP packets of certain type/code values.

To configure ACL for layer 4, click *Configure* > *Current VID* > *Access Control List* > *Layer 4.*

Configuring ACL to deny specific TCP and UDP ports

To configure ACL for denying layer 4 TCP ports, click *Configure* > *Current VID* > *Access Control List* > *Layer 4* > *TCP Ports.*

To configure ACL for denying layer 4 UDP ports, click *Configure* > *Current VID* > *Access Control List* > *Layer 4* > *UDP Ports.*

Refer to Figure 44 for configuration of Layer 4 ACLs for denying TCP ports.

Figure 44: Configuring layer 4 ACLs for TCP ports

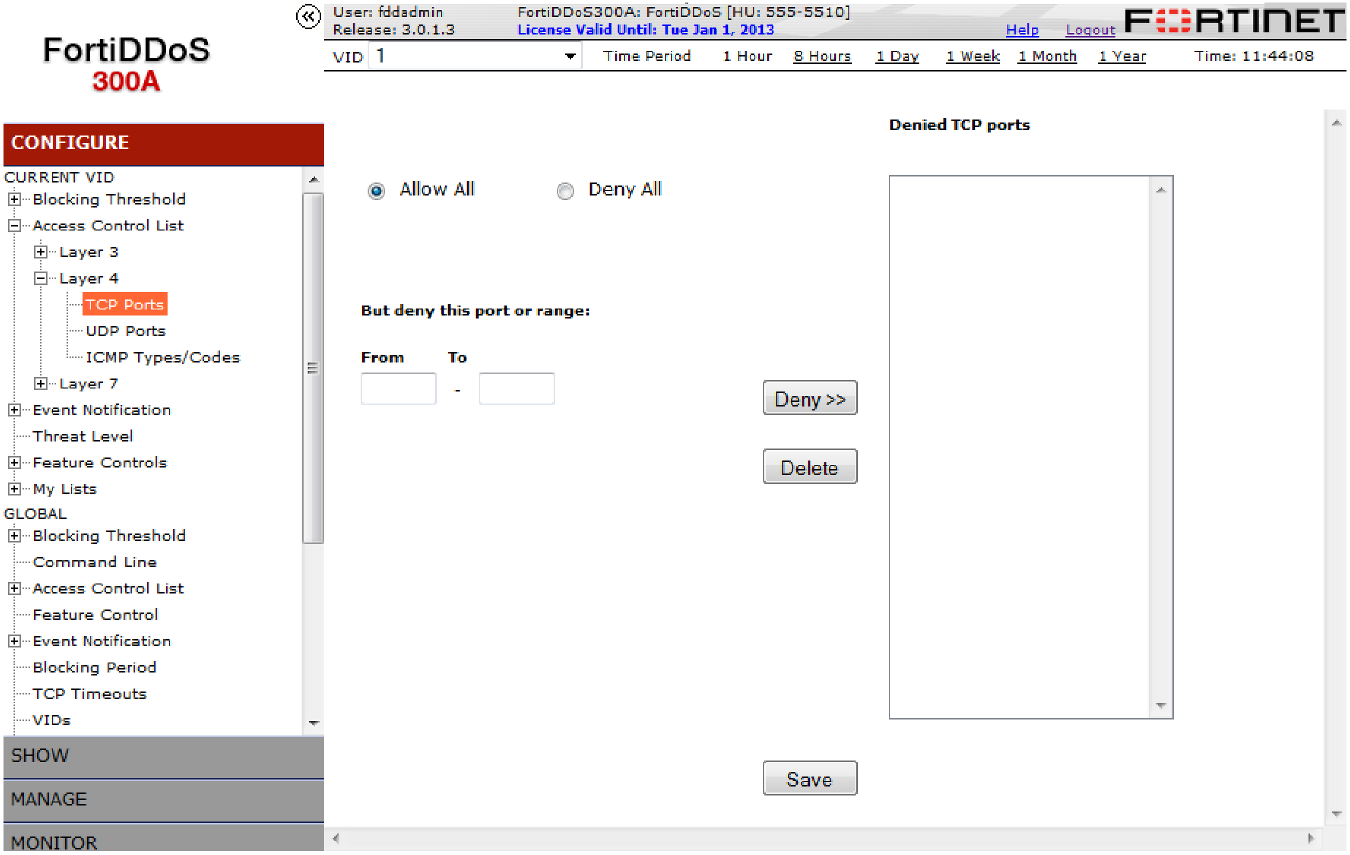

User: fddadmin
Release: 3.0.1.3
FortiDDoS300A: FortiDDoS [HU: 555-5510]
License Valid Until: Tue Jan 1, 2013
VID 1
Time: 11:44:08

● Allow All ○ Deny All

But deny this port or range:

From ____ - To ____

Denied TCP ports

- The list on the right includes ports that are denied if you have chosen the *Allow All* option.
- The list on the right includes ports that are allowed if you have chosen the *Deny All* option.
- You can add new values to the list by entering the Port or a Port Range by entering in the *From* and *To* fields then clicking *Move>>.*
- You can delete existing values in the list by selecting them and then clicking *Delete.*
- Click *Save* to commit the changes to the FortiDDoS appliance.

Configuring ACL to deny specific ICMP types/codes

To configure ACL for denying layer 4 ICMP type/codes, click *Configure* > *Current VID* > *Access Control List* > *Layer 4* > *ICMP Types and Codes.*

- The functionality of the user interface is similar to TCP and UDP ports.

ACL for layer 7

FortiDDoS allows granular ACLs for the following types of layer 7 packets in both inbound and outbound directions:

- Packets belonging to certain URL, Hosts, Referers, Cookies, User-Agents - Hash Indexes

To configure ACL for layer 7, click *Configure* > *Current VID* > *Access Control List* > *Layer 7.*

Configuring ACLs for denying specific URLs, hosts, referers, cookies, user-Agents

This feature allows you to deny specific URLs, Hosts, Referers, Cookies, and User-Agents. You can use this feature when a specific hash-index is under attack. The FortiDDoS will allow the TCP connection to establish between the source and the server. However, when it sees the specific hash-index, it will deny the packet and send a RST packet to server to aggressively age the connection. All subsequent packets from the source on that TCP connection will be denied.

E.g., to configure ACL for layer 7 HTTP URLs, click *Configure* > *Current VID* > *Access Control List* > *Layer 7* > *URL*

- Enter the *URL*.
- URL here means the text that follows the protocol and the web address. E.g. if you enter http://www.website.com/index.html in your browser to reach a specific URL, then you must enter /index.html.
- Click *Deny.*
- Denied URLs are listed in the *Denied URLs* list box.

Since there are infinite possible URLs, the hardware stores these in a hash based memory table. Up to 8192 such hash indexes are allowed in the current hardware. It is possible that there are duplicate hash-indexes. In that case, the last URL that corresponds to a hash-index takes overwrites any previous URLs in the URL field. All such URLs however affect the threshold and maximum packet rate calculations. 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.

Configuring MAC Addresses for Bypass Switch Heartbeat Packets

When a FortiDDoS appliance is used in conjunction with a bypass switch such as FortiBridge, you have to ensure that heartbeat packets from the bypass switch are allowed in Prevention Mode under all possible cases of packets being blocked by the FortiDDoS.

Typically all 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 passage of the heartbeat packets, FortiDDoS allows you to configure the MAC addresses of the bypass switch. These MAC addresses are used by the bypass switch for the heartbeat packets.

FortiBridge appliance allows you to view the MAC addresses in the status page.

Every FortiDDoS link pair can be connected via a FortiBridge link pair. E.g. LAN1, WAN1 can be bridged via a FortiBridge link and LAN2, WAN2 on a card can be similarly bridged via another FortiBridge link. Each of these link pairs will be associated with a pair of MAC addresses. Therefore if you are using two links you will need to configure 4 MAC addresses. If you are using one link then you need to specify just one pair of MAC addresses.

Figure 45: Configuring MAC Addresses for Bypass Switch Heartbeat Packets

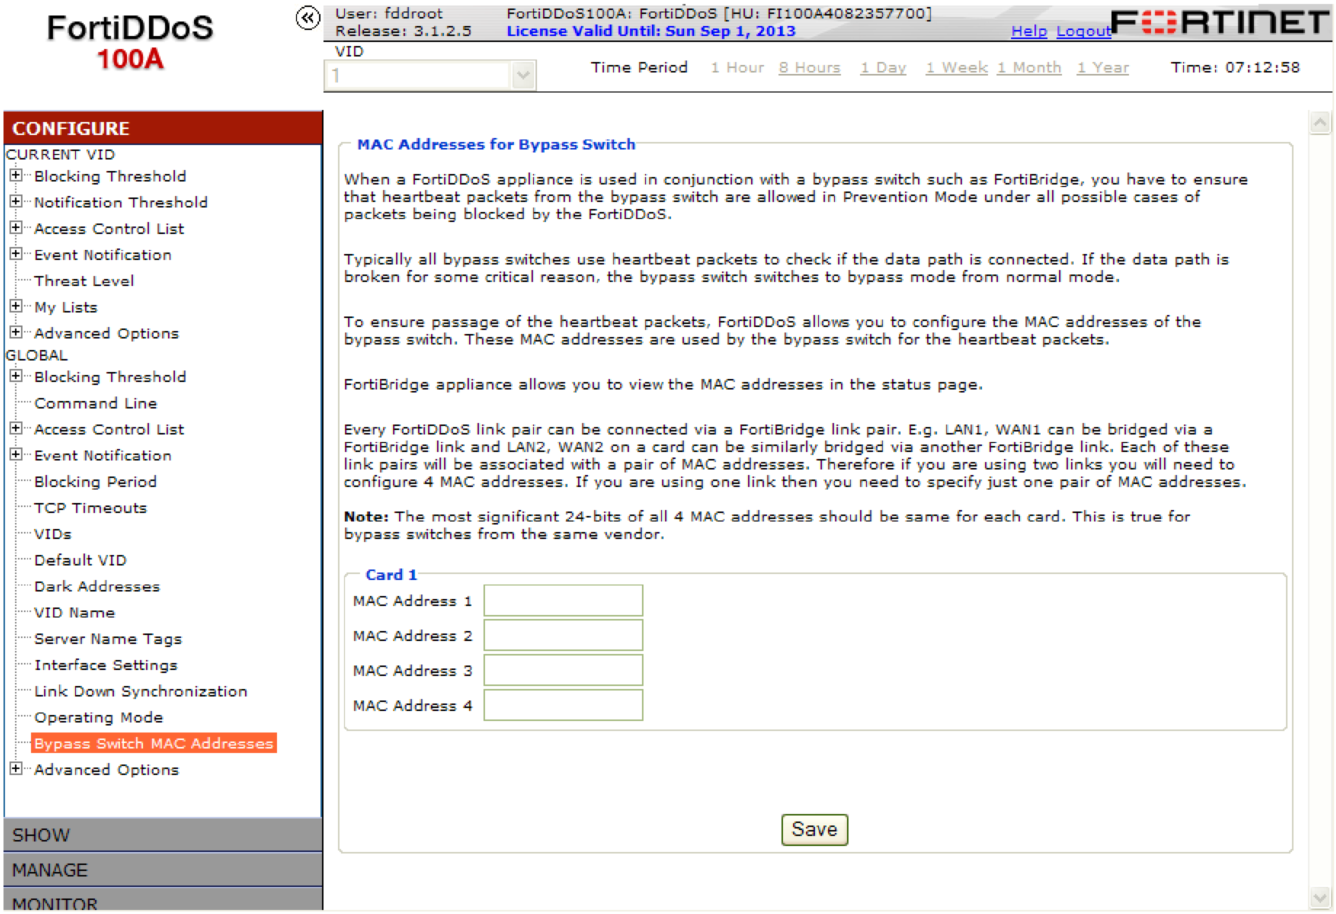

MAC Addresses for Bypass Switch

When a FortiDDoS appliance is used in conjunction with a bypass switch such as FortiBridge, you have to ensure that heartbeat packets from the bypass switch are allowed in Prevention Mode under all possible cases of packets being blocked by the FortiDDoS.

Typically all 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 passage of the heartbeat packets, FortiDDoS allows you to configure the MAC addresses of the bypass switch. These MAC addresses are used by the bypass switch for the heartbeat packets.

FortiBridge appliance allows you to view the MAC addresses in the status page.

Every FortiDDoS link pair can be connected via a FortiBridge link pair. E.g. LAN1, WAN1 can be bridged via a FortiBridge link and LAN2, WAN2 on a card can be similarly bridged via another FortiBridge link. Each of these link pairs will be associated with a pair of MAC addresses. Therefore if you are using two links you will need to configure 4 MAC addresses. If you are using one link then you need to specify just one pair of MAC addresses.

Note: The most significant 24-bits of all 4 MAC addresses should be same for each card. This is true for bypass switches from the same vendor.

Card 1

MAC Address 1

MAC Address 2

MAC Address 3

MAC Address 4

For FortiDDoS-200A, you can program up to 4 additional MAC addresses and for FortiDDoS-300A, can program up to 8 additional MAC addresses (depending on the number of links pairs are connected through the bypass switch).

Note: The most significant 24-bits of all 4 MAC addresses should be same for each card. This is true for bypass switches from the same vendor.

Interpreting events and understanding their implications

FortiDDoS provides four primary ways to view events:

- Event Monitor
- Event Notification through E-mail
- Event Notification to SNMP Managers through SNMP traps
- Consolidated Event Reporting for past events

Event Monitor provides you an interactive view of events as they are happening. This is available through the *Monitor* button on the Web-based Manager.

Event Notification is a mechanism to provide you emails on a periodic basis based on categories of your interest and in case the dropped packets have reached certain threshold. This is available through *Configure* > *Current VID* > *Event Notification.*

Event Notification to SNMP Managers is done through SNMP traps. This is available through *Configure* > *Global* > *SNMP.*

Event Reporting is an analysis tool available to get an overall picture of evens as they occurred in the past. This is available through *Show* > *Current VID* > *Reports* as well as *Show* > *Global* > *Reports.*

Event monitor

Event Monitor provides a comprehensive way of displaying network attacks so that you can investigate them interactively. You can choose a particular date range or number of events to be displayed. In addition, FortiDDoS provides categorized event entries as well as VID and database choices so that users can see only the events of their interest.

If you are a VID administrator, you can also see events for your own VID(s) through the menu *Show* > *Current VID* > *Reports.* However, if you are a Super User you can access *Show* > *Global* > *Reports* to view reports for across all VIDs**.**

The user interface of the event monitor contains two sections as shown in Figure 46:

- Event Parameters form
- Event table

Figure 46: Monitoring events

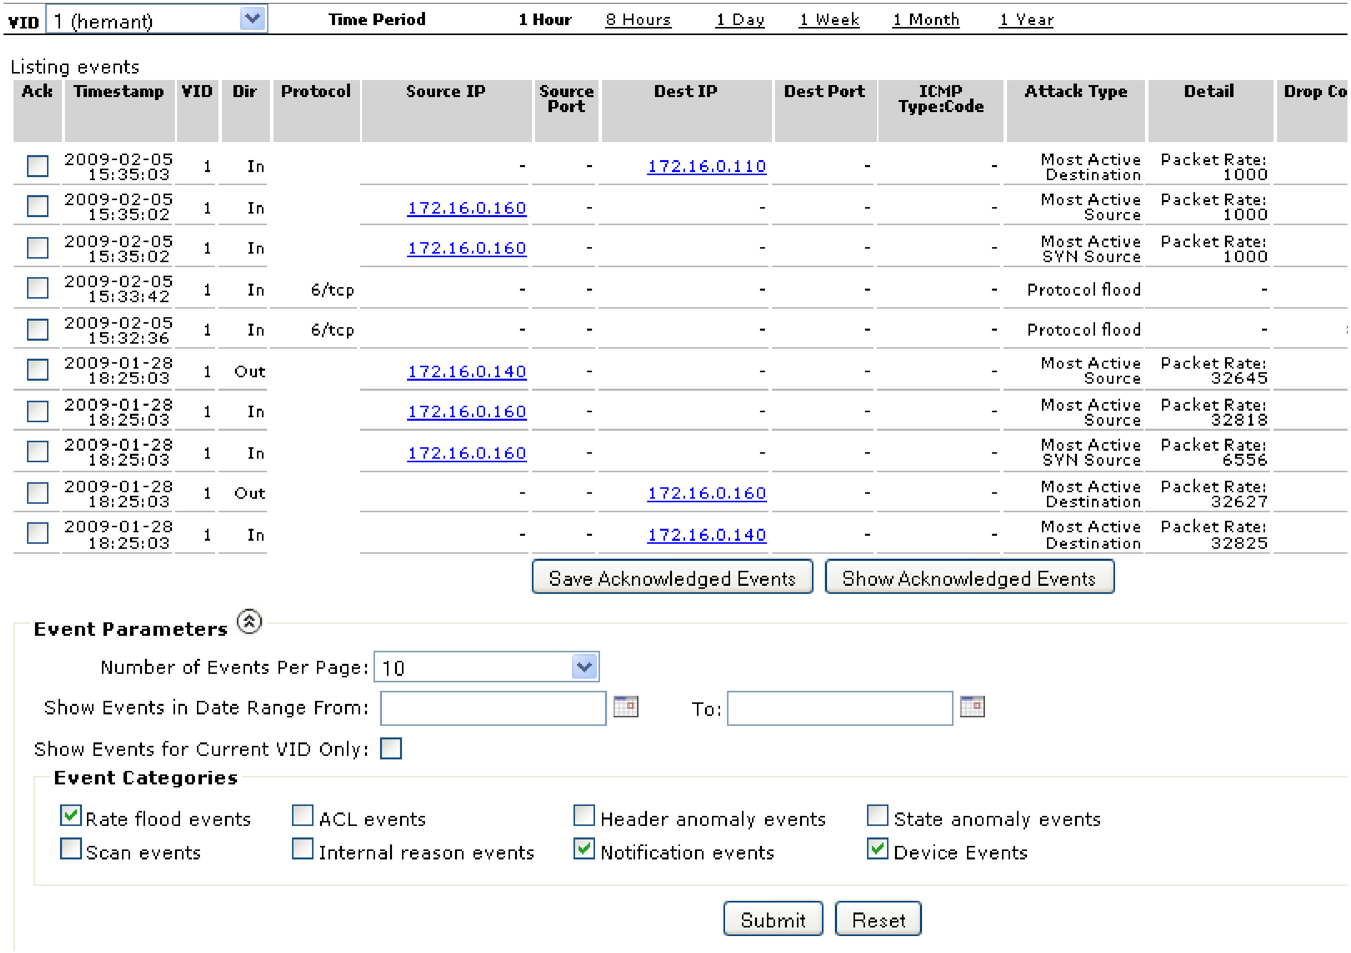

VID 1 (hemant) Time Period 1 Hour 8 Hours 1 Day 1 Week 1 Month 1 Year

Listing events

Ack	Timestamp	VID	Dir	Protocol	Source IP	Source Port	Dest IP	Dest Port	ICMP Type:Code	Attack Type	Detail	Drop Co
□	2009-02-05 15:35:03	1	In		-	-	172.16.0.110	-	-	Most Active Destination	Packet Rate: 1000	
□	2009-02-05 15:35:02	1	In		172.16.0.160	-	-	-	-	Most Active Source	Packet Rate: 1000	
□	2009-02-05 15:35:02	1	In		172.16.0.160	-	-	-	-	Most Active SYN Source	Packet Rate: 1000	
□	2009-02-05 15:33:42	1	In	6/tcp	-	-	-	-	-	Protocol flood	-	
□	2009-02-05 15:32:36	1	In	6/tcp	-	-	-	-	-	Protocol flood	-	
□	2009-01-28 18:25:03	1	Out		172.16.0.140	-	-	-	-	Most Active Source	Packet Rate: 32645	
□	2009-01-28 18:25:03	1	In		172.16.0.160	-	-	-	-	Most Active Source	Packet Rate: 32818	
□	2009-01-28 18:25:03	1	In		172.16.0.160	-	-	-	-	Most Active SYN Source	Packet Rate: 6556	
□	2009-01-28 18:25:03	1	Out		-	-	172.16.0.160	-	-	Most Active Destination	Packet Rate: 32627	
□	2009-01-28 18:25:03	1	In		-	-	172.16.0.140	-	-	Most Active Destination	Packet Rate: 32825	

Event Parameters

Number of Events Per Page: 10

Show Events in Date Range From: ______ To: ______

Show Events for Current VID Only: □

Event Categories

☑ Rate flood events	□ ACL events	□ Header anomaly events	□ State anomaly events
□ Scan events	□ Internal reason events	☑ Notification events	☑ Device Events

Monitoring events of your interest

The query criteria for events are displayed as a form, which allows you to control the display using following criteria:

- Number of Events per page,
- Date range of events,
- Show events for Current VID only,
- Event Categories.

The events that meet the criteria are displayed in a table which is scrollable.

Number of events per page

- Select a number from the list. The range varies from 10 to 50. At a time the number of events displayed will be limited to the number you have selected.
- This number is remembered as a session parameter and is used as long as you are logged in.

Date range

- Enter start and end dates. You can use the date controls to enter the date in the right format.
- This number is remembered as a session parameter and is used as long as you are logged in.

Show events for current VID only

- Select this checkbox if you have access to multiple VIDs, but if you want to see only this VID.
- This selection is remembered as a session parameter and is used as long as you are logged in.

Event categories

- Select at least one of the *Event Categories*. Events of these categories will be displayed after you press submit button. The states of check boxes in *Event Categories* are remembered as long as you are logged in.
- The categories and their corresponding events are summarized in Table 4.

Table 4: Event Categories

Layer	Event	Level
3	Rate Flood	Protocol
		Fragment
		Source
		Tracked Source
		Destination
	ACL Events	Protocol
		Fragment
		Source
		Destination
	Scan Events	Dark Address Scan
	Header Anomaly Events	IP header anomaly
		IP Header checksum error
		Source IP==dest IP
		Source/dest IP==localhost
		L3 anomalies
		Excessive hash collisions
	Internal Reason Events	Out of memory
		ST:Hash attack
		ST:Out of memory
		DT:Hash attack
		DT:Out of memory
		Distributed Source Attack
	Notification Events	Most Active Source
		Most Active Destination

Layer	Event	Level
4	Rate Flood Events	SYN
		SYN Flood From Source
		SYN Flood to Destination
		ACK Flood to Destination
		RST Flood to Destination
		FIN Flood to Destination
		Established Flood to Destination
		TCP connection
		Legitimate IP
		TCP port
		UDP port
		ICMP type/code
		TCP zombie
		Excessive Concurrent Connections Per Source
		Excessive Concurrent Connections Per Destination
	ACL Events	TCP port
		UDP port
		ICMP type/code
	Header Anomaly Events	TCP checksum error
		UDP checksum error
		ICMP checksum error
		TCP invalid flag combination
		L4 anomalies
		Outside TCP/UDP window
	State Anomaly Events	Foreign packet
		State transition anomalies
		Dark Address Scan
	Scan Events	Excessive hash collisions
	Internal Reason Event	Out of memory
		Excessive instantaneous TCP connections (DDoS)
	Notification Events	Most Active SYN Source
7	Rate Flood Events	Op Code Flood
		URL, Hosts, Referers, Cookies, User-Agents Flood,
		SIP Invite Per Source
		SIP Register Per Source
		SIP Concurrent Invite Per Source
	ACL Events	Op Code
		URL, Host, Referer, Cookie, User-Agent
	Header Anomaly Events	Undefined HTTP Op code Anomaly
		Unknown HTTP Op code Anomaly
		Invalid HTTP Version

Layer	Event	Level
NA	Device Events	Disk not OK Hardware Failure Software Failure Disk OK RAID Array Failure RAID Array OK Excessive Events

Acknowledging events and showing only unacknowledged events

Once you have seen an event, you can choose to acknowledge the events and thus hiding them from the display.

To acknowledge multiple events, click on each of them and they will get highlighted.

To select an event, click on it once. To unselect a given event, just click the event again and it will be unselected. To unselect all events, click on the *Unselect All button*.

If you want to display acknowledged events, click on the *Show Acknowledged Events* button. If you want to hide the acknowledged events, toggle the S*how Acknowledged Events* button.

VID and event database

The top line VID numbers are corresponding to the VIDs to which you have access to. The number of VIDs is limited to the maximum number of VIDs supported by the system. Within a session, the currently selected VID is remembered across screens of display. When you select any of the VID numbers, *Event Monitor* shows events corresponding to the VIDs to which you have access to unless you have also checked the *Show Events for Current VID Only* check box. If you check *Show events for current VID only* box, only the events for the VID selected in the top bar are shown.

Submit button

Click on *Submit* button to see events corresponding to your choice.

Event table

Event table displays the events with fields of timestamp, VID, direction, protocol, source IP, source port, destination IP, destination port, ICMP type:code, attack type, detail and drop count. Some fields may display '-' in the case of no valid data for those fields.

For *Protocol* field, a number/name format is used. If the name cannot be determined, a '-' is shown.

For *Destination Port* field, a number/name format is used. If the name cannot be determined, a '-' is shown.

For *ICMP Type:Code* field, a number:number/name format is used. If the name cannot be determined, a '-' is shown.

Viewing the who-is information for source IP addresses

If you see an event which contains an IP address which has a hyperlink available, you can click on the hyperlink to get the who-is information for the IP address.

Maximizing the display area for event monitor

Event table takes substantial area on the screen. If you want to display only events on the screen, you can hide the unwanted parts of the screen. You can do this in two steps:

- You can first hide the left side menu-bar by clicking on the Left arrow button shown in Figure 47. You can unhide the menu-bar by clicking on the right Right arrow that appears in its place.

Figure 47: Hiding the menu bar

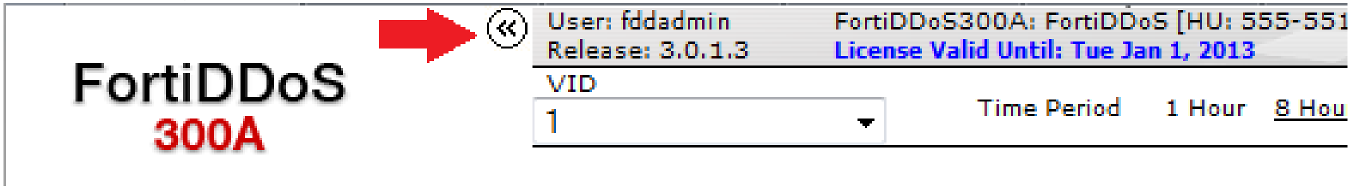

User: fddadmin
Release: 3.0.1.3

FortiDDoS300A: FortiDDoS [HU: 555-551
License Valid Until: Tue Jan 1, 2013

VID: 1

- You can then hide the Event Parameters Section by clicking on the Up arrow button shown in Figure 48. You can unhide the Event Parameters Section by clicking on the Down arrow that appears in its place.

Figure 48: Hiding the event parameter section

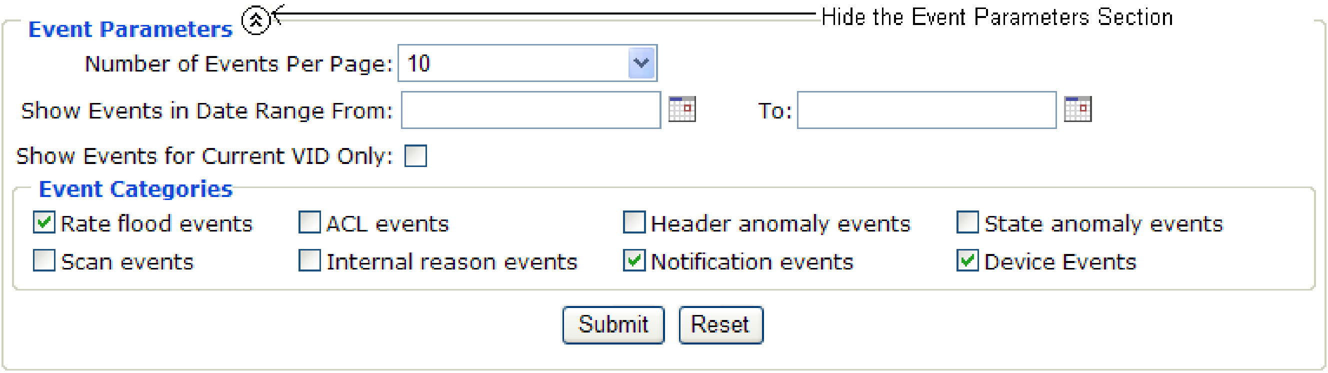

Number of Events Per Page: 10

Show Events in Date Range From: To:

Show Events for Current VID Only: □

☑ Rate flood events □ ACL events □ Header anomaly events □ State anomaly events

□ Scan events □ Internal reason events ☑ Notification events ☑ Device Events

Configuring event notification through email

Event Notification allows user to receive email messages when Network attacks happen. This feature can be used to send emails to pagers or PDAs as well. FortiDDoS provides graphic interface for users to set up notification properties. You can choose the type of events for which you should be notified and only above certain thresholds.

- Use *Configure* > *Current VID* > *Event Notification* to configure the event notification for the current VID.
- Depending on your role and corresponding privileges, you can add, delete or modify their notification event by clicking “Add”, “Delete” or “Modify” button on Configure/Event Notification screen.
- Refer to Figure 49.

Viewing configured event notifications

- Click on one of the items on the *Notification* list, the selected notification's name will appear in the *Notification Event Name* text box and its properties are displayed in Selected Notification Properties area.
- "*Notify event name*", located at the top right of the pane, together with the VID value, identifies a particular notification.

Temporarily disabling an event notification

- A notification event can be temporarily disabled by un-checking the "*Enable selected event notification*" check box.

Figure 49: Event notification

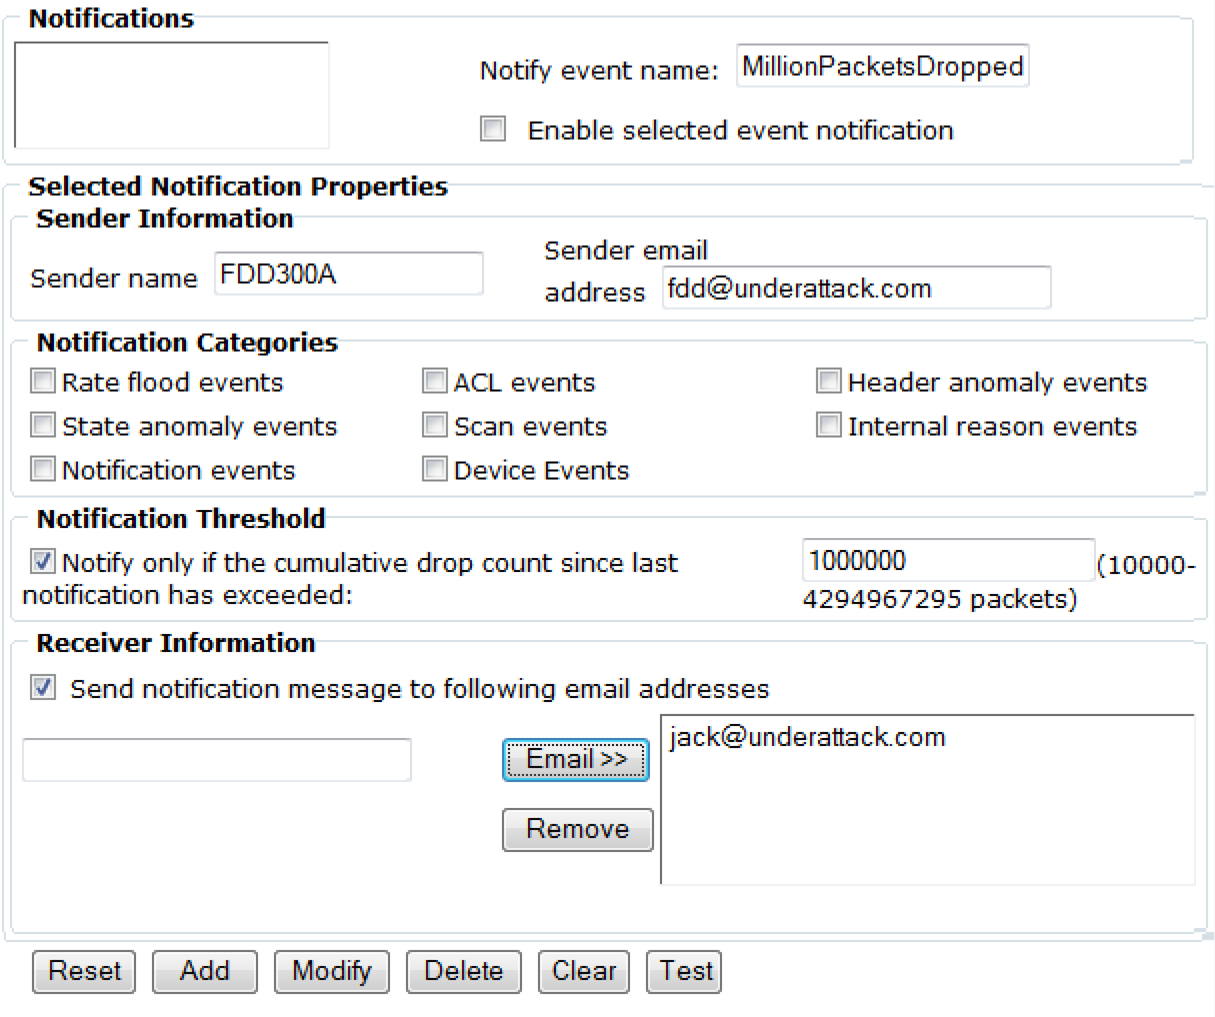

Notifications

Notify event name: MillionPacketsDropped

□ Enable selected event notification

Selected Notification Properties

Sender Information

Sender name FDD300A

Sender email address fdd@underattack.com

Notification Categories

□ Rate flood events □ ACL events □ Header anomaly events

□ State anomaly events □ Scan events □ Internal reason events

□ Notification events □ Device Events

Notification Threshold

☑ Notify only if the cumulative drop count since last notification has exceeded: 1000000 (10000-4294967295 packets)

Receiver Information

☑ Send notification message to following email addresses

jack@underattack.com

Adding an event notification

1. Click on *Clear* button to clear the fields
2. Enter a *Notify event name* that does not already exist in the database. The length should be between 3 and 30 alphanumeric characters
3. Enter the *Sender name* with length between 3 and 30 alpha characters

4 Enter a valid *Sender email address*. The notifications will appear to come from Sender name and the corresponding email address

5 Select at least one category from *Notification Categories* based on items for which you require event notification. The notification will be sent only if the events in the selected categories are present

> 1 If you want to limit the notification based on thresholds to avoid frequent notifications, check the *Notification Threshold* checkbox and enter a *Threshold* value. Event notification will happen only after the number of dropped packets exceed this threshold. If you have set the *Notification Threshold*, the program counts the cumulative drop-count for all events in the selected categories between the last notification time and the current time, and checks if the count has exceeded the threshold. If it has exceeded, a summary report is sent to notify the events after due filtering.
>
> 2 Multiple email addresses can receive the email notifications. Enter the email address of individuals or lists after checking the box entitled *Send notification message to following email addresses*. Use the Email >> or *Remove* buttons to add or remove email addresses.

3 Click on *Add* button to add the new notification event in the database

Deleting an event notification

The *Delete* button removes a configured event notification setup.

1 Start by choosing an event from the list-box. This should retrieve the event notification's properties including its *Notify event name*

2 Click *Delete* to delete the event notification. *Delete* button removes a notification event from the database permanently

Modifying an event notification

The *Modify* button changes a notification's properties in the database.

1 Start by choosing an *Notify Event Name* from the list-box. This should retrieve the notification's properties

2 Modify the properties per your requirements while following the same rules as given above

3 Select at least one item from Notification Categories

4 Click *Modify* button to modify the properties

Note: *Reset* Button: This button resets the dialog to the values for the shown entity.

Note: *Clear* Button: This button clears the fields in the dialog so that the user can fill them from scratch.

Event notification to a SNMP managers through SNMP traps

SNMP is an application-layer protocol that facilitates the exchange of management information between network devices. SNMP enables network administrators to manage network performance, find and solve network problems, and plan for network growth.

There are three versions of SNMP:

Version 2 (SNMPv2c):The second release of SNMP, described in RFC 1902, has additions and enhancements to data types, counter size, and protocol operations.

Version 3 (SNMPv3):This is the most recent version of SNMP and is fully described in RFC 2571, RFC 2572, RFC 2573, RFC 2574, and RFC 2575. SNMPv3 has significant enhancements to administration and security.

SNMPv3 provides secure access to devices by authenticating and encrypting packets over the network. The security features provided in SNMPv3 are:

- Message integrity—Ensuring that a packet has not been tampered with in transit
- Authentication—Determining the message is from a valid source
- Encryption—Scrambling contents of packet to prevent it from being seen by an unauthorized source

The following table describes the security levels of SNMP versions:

Table 5: Security Levels of SNMP versions

SNMP Version	Access Type	Authentication	Encryption
V2C	No Authentication No Privacy	Community String	None
V3	No Authentication No Privacy	User name	None
V3	Authentication No Privacy	MD5/SHA	None
V3	Authentication Privacy	MD5/SHA	DES

FortiDDoS allows event notifications to be sent to configured SNMP Managers through SNMP traps. You can configure up to 5 trap receivers. FortiDDoS supports Trap Receivers that are SNMP V2C or V3 compatible. FortiDDoS provides a graphical user interface for users to set up trap receiver properties. All events which are generated by FortiDDoS are sent as traps. When an NMS receives a trap, it needs to know how to interpret it. FortiDDoS sends the trap in a published format. Please contact Fortinet, Inc. for the MIB which defines the trap if you plan to use the SNMP features.

To allow your FortiDDoS to send SNMP traps to your NMS, you must first enable SNMP support. You must then add trap receiver properties. You must then add communities. SNMP configuration is available through *Configure* > *Global* > *Event Notification* > *SNMP.*

Configuring FortiDDoS SNMP agent

To configure SNMP agent, in the *Configuration* section of the screen, check the *Enable SNMP Access* box. Then enter following information.

1. *System Description:* You should set this to fully qualified domain name of the FortiDDoS or its name. You can enter up to 255 characters.
2. *System Location:* You should set this field to the physical location of the FortiDDoS. You can enter up to 255 characters. This information is useful for determining where the FortiDDoS is located, particularly if you have several FortiDDoS appliances spread over a wide area.
3. *System Contact:* You should set this field to primary contact for this FortiDDoS. You can enter up to 255 characters. This information is useful to the operations staff for determining who needs to be contacted in case of some failures.
4. Click *Save* to save the configuration information.
5. If you want to restart filling the information, click *Reset* and start over.
6. If you want to disable SNMP access, you must uncheck the *Enable SNMP Access* box and click *Submit.*

Configuring SNMP trap receivers

After you have enabled SNMP access, you must configure SNMP Trap Receivers if you want external NMS to receive traps. You can configure up to 1 Trap Receiver per VID. This configuration is available within the *Trap Receivers* section of the screen.

To add a new SNMP Trap Receiver, you must click *Add* and enter following information:

1. *Trap Receiver Host:* You should set this to the IP address of the SNMP manager.
2. *Trap Receiver Port:* You should set this field to the UDP port on which the manager is listening to the traps.
3. *Community/User Name:* If you use SNMP V2, you should set this field to the community. If you use SNMP V3, you must set this field to the user name. You can enter up to 255 characters. This is the Community/User Name that is used in the traps sent by the SNMP agent in the appliance.
4. *SNMP Version:* You should set this field to V2c or V3 depending on the Trap Receiver's SNMP support.
5. *Trap Receiver Engine ID:* This field is valid only for SNMP V3 trap receivers. You should set this field to Engine ID of the Trap Receiver. This is an administratively unique identifier for the SNMP engine. This is used for identification and not for addressing. Please refer to RFC 2571 for detailed discussions on Engine IDs.
6. *V3 Access Type*: This field is valid only for SNMP V3 trap receivers. You should set this field to one of the following:
    - No Authentication
    - Authentication
    - Privacy

      Depending on the desired security level as described above, you must choose one of the access types.

7 *Authentication Passphrase and Reenter Authentication Passphrase:* This field is valid only for SNMP V3 trap receivers and if you have chosen Authentication or Privacy Access Type. You must enter a string between 8 and 128 characters long that should not contain spaces in the beginning or at the end and must not contain double quotes.

8 *Privacy Passphrase and Reenter Privacy Passphrase:* This field is valid only for SNMP V3 trap receivers and if you have chosen Privacy Access Type. You must enter a string between 8 and 128 characters long that should not contain spaces in the beginning or at the end and must not contain double quotes.

9 Click *Save* to save the Trap Receiver information.

10 If you want to modify information related to an existing Trap Receiver, select the Trap Receiver from the list box, click *Modify* and then update the information as you desire and then click *Save*.

11 If you want to delete an existing Trap Receiver, select the Trap Receiver from the list box, and click *Delete*.

Configuring SNMP MIB-II support

After you have enabled SNMP access, you can configure FortiDDoS to support MIB-II management group. RFC 1213 defines MIB-II. FortiDDoS supports read-only access to MIB-II groups and objects. Only System and Interface groups are supported. The Interface object corresponds to the management interface of the FortiDDoS. You can configure up to 5 communities for MIB-II read-only access. This configuration is available within the *Communities* section of the screen.

To add a new Community, you must click *Add* and enter following information:

1 *Community/User Name*: If you use SNMP V2, you should set this field to the community. If you use SNMP V3, you must set this field to the user name. You can enter up to 255 characters. This is the Community/User Name that the external SNMP manager has to use to query the SNMP agent in the appliance.

2 *SNMP Version*: You should set this field to V2c or V3 depending on the SNMP Manager’s version support.

3 *V3 Access Type*: This field is valid only for SNMP V3 trap receivers. You should set this field to one of the following:

- No Authentication
- Authentication
- Privacy

  Depending on the desired security level as described above, you must choose one of the access types.

4 *Authentication Passphrase and Reenter Authentication Passphrase:* This field is valid only for SNMP V3 trap receivers and if you have chosen Authentication Access Type. You must enter a string between 8 and 128 characters long that should not contain spaces in the beginning or at the end and must not contain double quotes.

5 *Privacy Passphrase and Reenter Privacy Passphrase*: This field is valid only for SNMP V3 trap receivers and if you have chosen Privacy Access Type. You must enter a string between 8 and 128 characters long that should not contain spaces in the beginning or at the end and must not contain double quotes.

6 Click *Save* to save the Community information.

7 If you want to modify information related to an existing Community, select the Community from the list box, click *Modify* and then update the information as you desire and then click *Save.*

8 If you want to delete an existing Community, select the Community from the list box, and click *Delete.*

Configuring centralized monitoring of connection statistics and drop statistics using FortiDDoS's SNMP MIB

After you have enabled SNMP access, you can configure FortiDDoS to supply connection and drop statistics to a centralized SNMP monitor such as MRTG viewer.

You can get the FortiDDoS SNMP MIB from support@fortinet.com.

You can request support@fortinet.com for an application note describing usage with *MRTG.*

You can request support@fortinet.com for an application note along with the necessary xml files and describing usage with *Cacti*.

Preparing and understanding the event reports

Event Reporting is an analysis tool available to get an overall picture of events as they occurred in the past. If you are a VID administrator, you can see reports for events for your own VID(s) through the menu *Show* > *Current VID* > *Reports*. However, if you are a Super User you can access *Show* > *Global* > *Reports* to view reports for across all VIDs.

The following reports are available:

- Top Attacked TCP Services

  This report depicts network attacks that occurred on particular TCP ports in descending order of packets dropped.

- Top Attacked UDP Services

  This report depicts network attacks that occurred on particular UDP ports in descending order of packets dropped.

- Top Attacked ICMP Type and Code

  This report depicts network attacks that occurred on particular ICMP Types and Codes in descending order of packets dropped.

- Top Attacked URLs

  This report depicts network attacks that occurred on particular URLs in descending order of packets dropped.

- Top Attacked Protocols

  This report depicts network attacks that occurred on particular IP protocols in descending order of packets dropped.

- Top Attacks

  This report depicts network attacks that occurred across all types in descending order of packets dropped.

- Top Attackers

  This report depicts network attacks where the source IP address has been tracked. It is shown in descending order of packets dropped.

- Top Attack Destinations

  This report depicts network attacks where the destination IP address has been identified. It is shown in descending order of packets dropped.
- Top Attacked Connections

  This report depicts network attacks where the TCP connection tuple (Source and destination IP address and source and destination ports) has been identified. It is shown in descending order of packets dropped.
- Top Dark Address Scan

  If you have enabled Dark Address Scan through *Configure* > *Current VID* > *Feature Control* > *Layer 3* > *Source tracking* > *Dark Address Scan,* this report depicts network attacks where the source has been identified for packets destined to dark addresses. The report is shown in descending order of packets dropped.
- All

  This report shows all the above reports in one single page and can be used for management reporting purpose.

Generating reports

- The top line VID numbers correspond to the VIDs to which you have access to. The number of VIDs is limited to the maximum number of VIDs supported by the system. The currently selected VID is remembered as a session parameter across all screens of display.
- For all the reports, you can see the data for the following periods:
  - 1 hour
  - 8 hours
  - 1 day
  - 1 week
  - 1 month
  - 1 year.
- The time period is remembered as a session parameter across all screens of display.
- Each of the table entries has a hyperlink which when clicked shows the detailed events corresponding to the hyperlink.

Figure 50 shows a typical report.

Figure 50: An event report sample

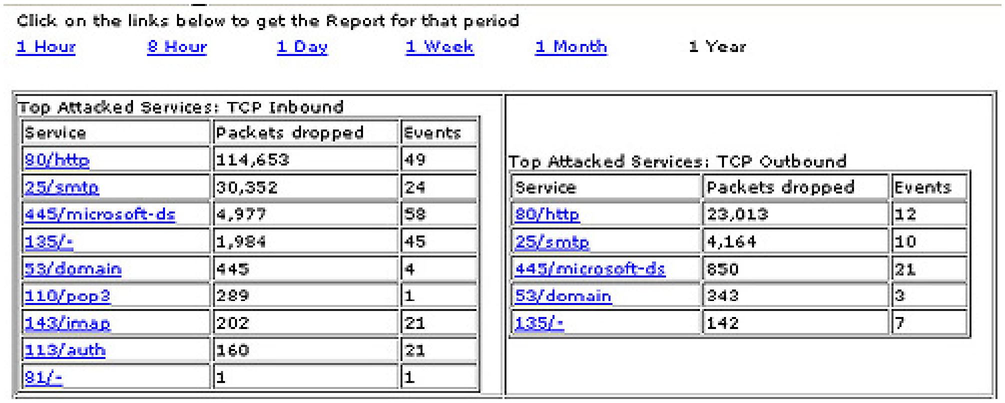

Click on the links below to get the Report for that period

1 Hour 8 Hour 1 Day 1 Week 1 Month 1 Year

Top Attacked Services: TCP Inbound

Service	Packets dropped	Events
80/http	114,653	49
25/smtp	30,352	24
445/microsoft-ds	4,977	58
135/-	1,984	45
53/domain	445	4
110/pop3	289	1
143/imap	202	21
113/auth	160	21
81/-	1	1

Top Attacked Services: TCP Outbound

Service	Packets dropped	Events
80/http	23,013	12
25/smtp	4,164	10
445/microsoft-ds	850	21
53/domain	343	3
135/-	142	7

The following is the event details of Top Attacked inbound TCP Services at port 143/IMAP within one year:

Figure 51: Top attacked events report

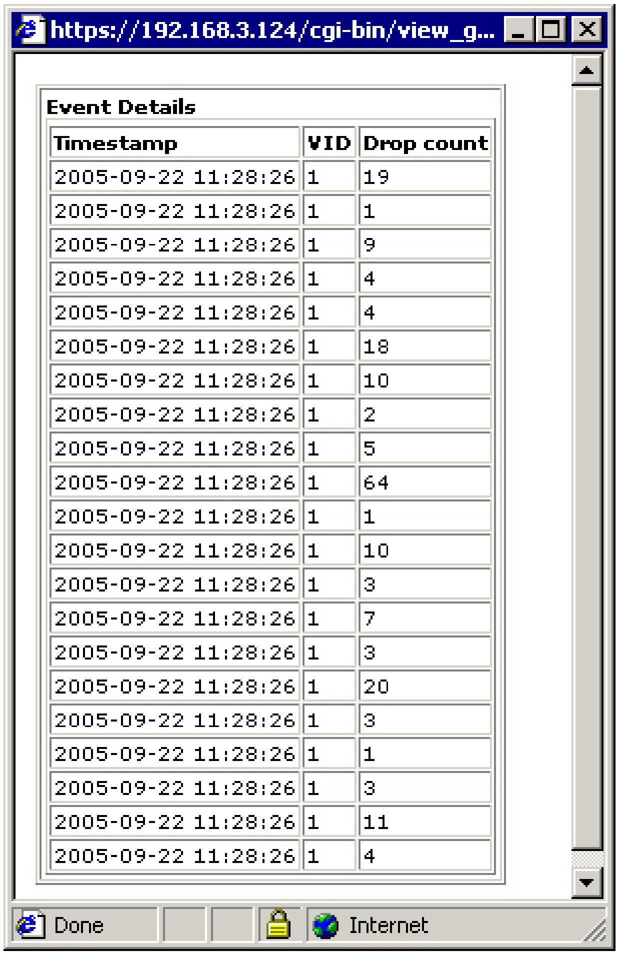

Event Details

Timestamp	VID	Drop count
2005-09-22 11:28:26	1	19
2005-09-22 11:28:26	1	1
2005-09-22 11:28:26	1	9
2005-09-22 11:28:26	1	4
2005-09-22 11:28:26	1	4
2005-09-22 11:28:26	1	18
2005-09-22 11:28:26	1	10
2005-09-22 11:28:26	1	2
2005-09-22 11:28:26	1	5
2005-09-22 11:28:26	1	64
2005-09-22 11:28:26	1	1
2005-09-22 11:28:26	1	10
2005-09-22 11:28:26	1	3
2005-09-22 11:28:26	1	7
2005-09-22 11:28:26	1	3
2005-09-22 11:28:26	1	20
2005-09-22 11:28:26	1	3
2005-09-22 11:28:26	1	1
2005-09-22 11:28:26	1	3
2005-09-22 11:28:26	1	11
2005-09-22 11:28:26	1	4

Preparing and understanding the subnet-ID based event reports

When a web host or a service provider protects multiple subnets using a single appliance, each of those subnets may belong to a customer that may demand a separate event report.

Subnet-ID based Reports help you achieve this goal.

If you are a Global administrator, you can see reports for events for your all subnets through the menu *Show* > *Global* > *Reports* > *Subnet Drops.*

To get such a report, you must have first configured the VIDs with the subnet IDs identified individually.

Figure 52: Subnet ID based report

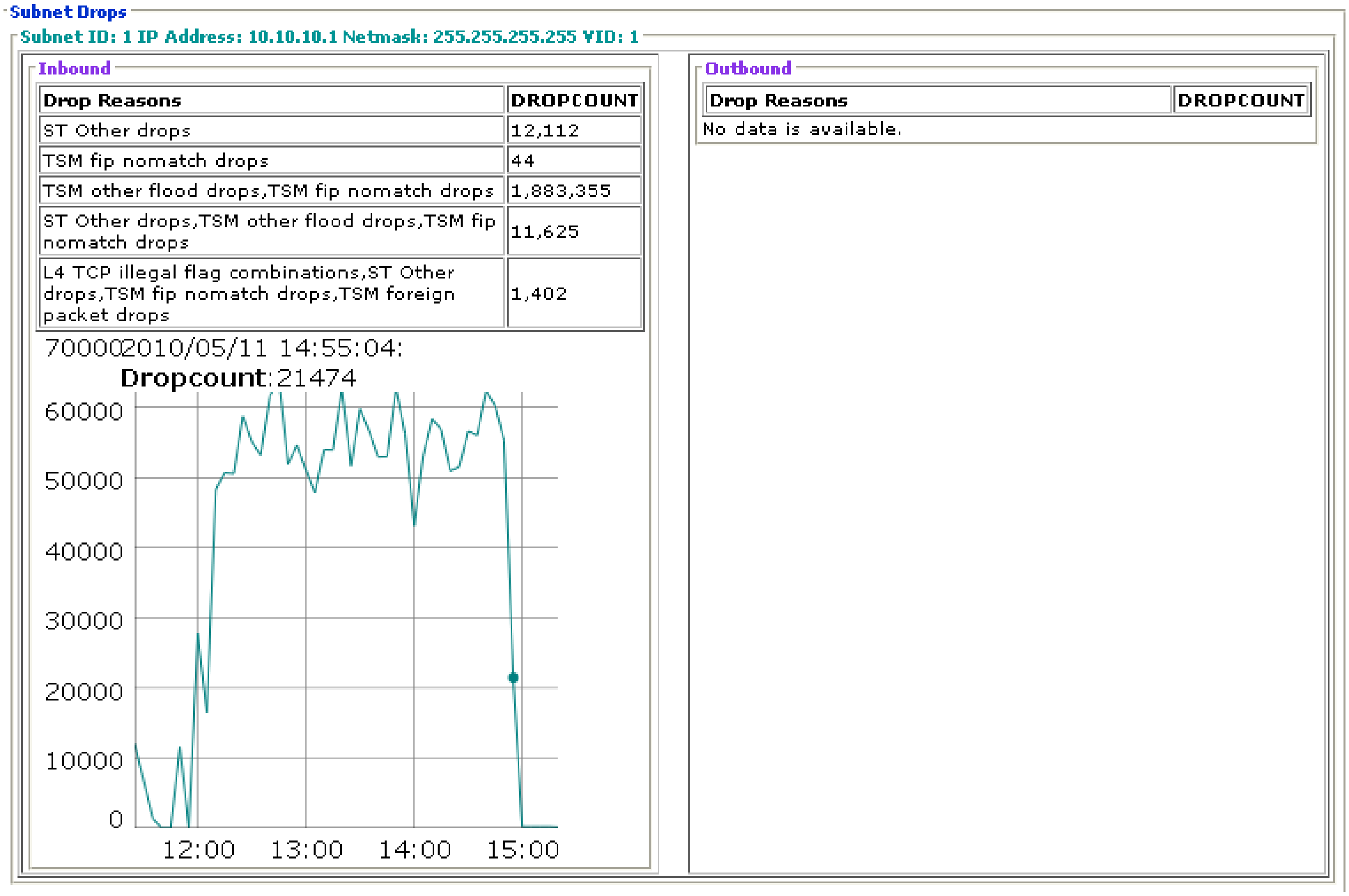

Subnet Drops

Subnet ID: 1 IP Address: 10.10.10.1 Netmask: 255.255.255.255 VID: 1

Inbound

Drop Reasons	DROPCOUNT
ST Other drops	12,112
TSM fip nomatch drops	44
TSM other flood drops,TSM fip nomatch drops	1,883,355
ST Other drops,TSM other flood drops,TSM fip nomatch drops	11,625
L4 TCP illegal flag combinations,ST Other drops,TSM fip nomatch drops,TSM foreign packet drops	1,402

Outbound

Drop Reasons	DROPCOUNT
No data is available.	

Running diagnostics to view internal connection tables

Under attacks, to further diagnose, you may want to dump the internal tables and view their contents in a sorted way. FortiDDoS provides following ways to view the tables:

- Top Servers
- Port
- Server
- Client
- Top URLs
- Top Hosts
- Top Referers
- Top Cookies
- Top User-Agents
- Top HTTP Accesses

Viewing top servers

At any given time, you may have several servers behind a given VID. Each server may have a busy service port. You can view the top servers (having the most number of simultaneous connections), top destinations, and top destination ports at any given second. Associated with each Source and Destination IP address pair is the destination port, direction (whether inbound or outbound), and the number of connections. Associated with the top destination is the TCP state (whether the connections established, closing, etc.) and the number of connections. Associated with the top destination ports for the top destination IPs are the number of connections. This information is transient in nature and is updated every second. This information can be useful to determine the servers which are getting the most traffic and within those, the ports which are getting affected.

If you are a VID administrator, you can see Top Servers for your VID through the menu *Show* > *Current VID* > *Session Diagnostics* > *Top Servers*.

Denying sources from Top 10 TCP Pairs table

During attacks, it is sometimes useful to directly deny certain sources that you see in the diagnostics tables.

If you feel that the sources being shown in the Top Servers diagnostics could be the culprits, you can simply press the *Deny Sources from Top TCP Pairs Table*. Refer to Figure 53.

Figure 53: Denying sources from Top TCP Pairs table (the IP addresses have been intentionally hidden)

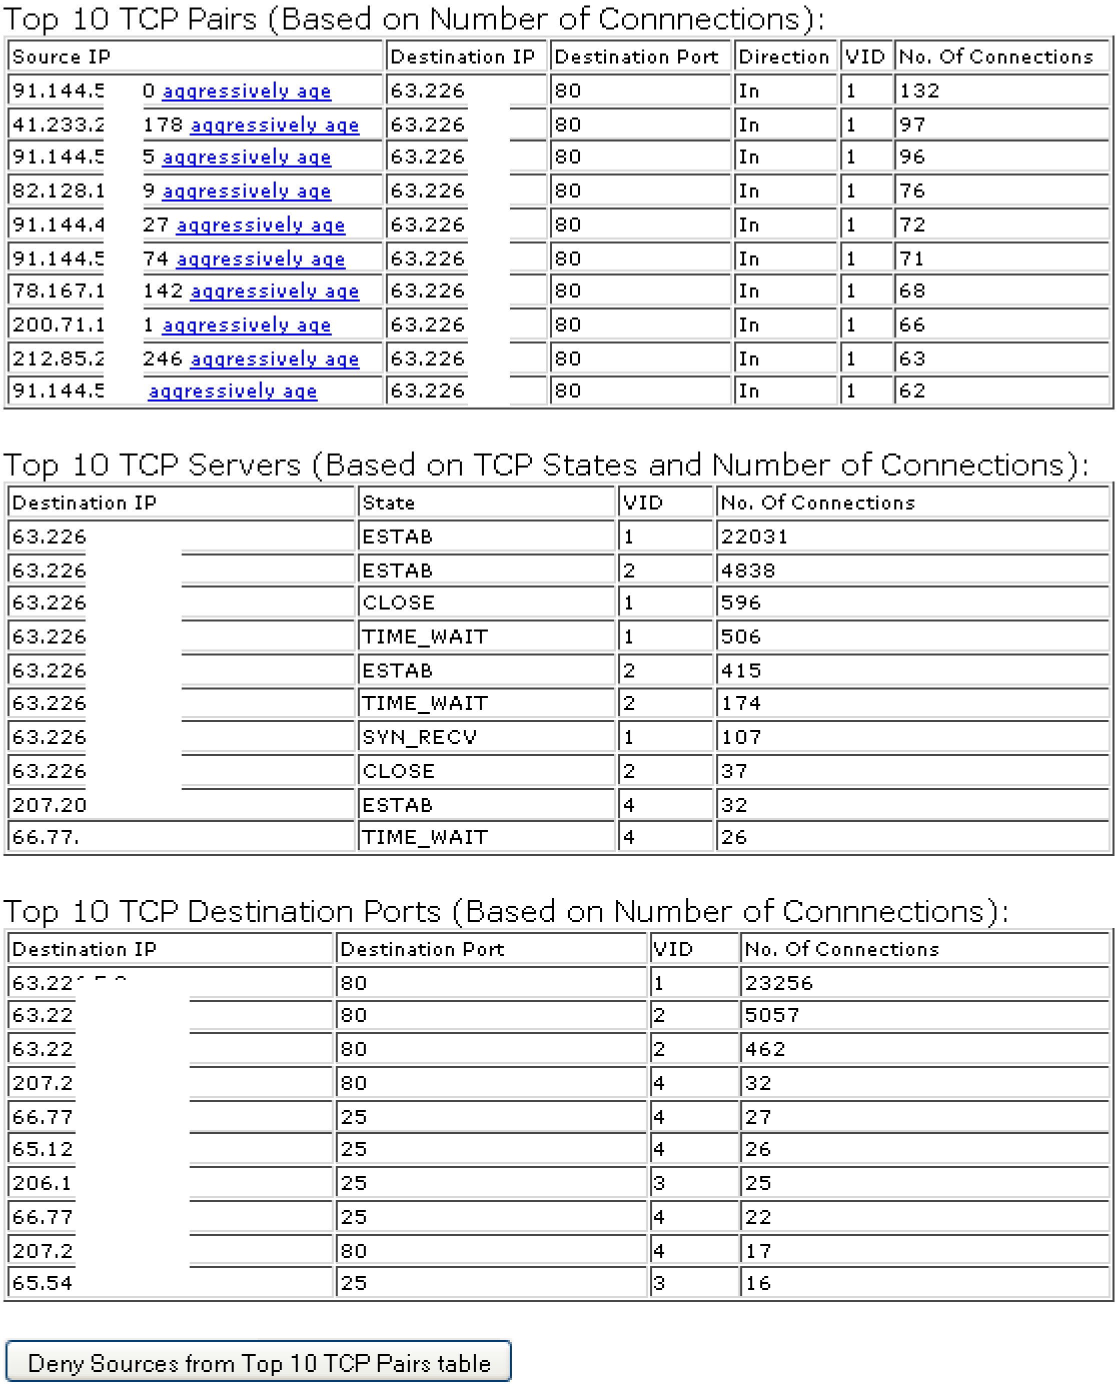

Top 10 TCP Pairs (Based on Number of Connnections):

Source IP	Destination IP	Destination Port	Direction	VID	No. Of Connections
91.144.5 0 aggressively age	63.226	80	In	1	132
41.233.2 178 aggressively age	63.226	80	In	1	97
91.144.5 5 aggressively age	63.226	80	In	1	96
82.128.1 9 aggressively age	63.226	80	In	1	76
91.144.4 27 aggressively age	63.226	80	In	1	72
91.144.5 74 aggressively age	63.226	80	In	1	71
78.167.1 142 aggressively age	63.226	80	In	1	68
200.71.1 1 aggressively age	63.226	80	In	1	66
212.85.2 246 aggressively age	63.226	80	In	1	63
91.144.5 aggressively age	63.226	80	In	1	62

Top 10 TCP Servers (Based on TCP States and Number of Connections):

Destination IP	State	VID	No. Of Connections
63.226	ESTAB	1	22031
63.226	ESTAB	2	4838
63.226	CLOSE	1	596
63.226	TIME_WAIT	1	506
63.226	ESTAB	2	415
63.226	TIME_WAIT	2	174
63.226	SYN_RECV	1	107
63.226	CLOSE	2	37
207.20	ESTAB	4	32
66.77.	TIME_WAIT	4	26

Top 10 TCP Destination Ports (Based on Number of Connnections):

Destination IP	Destination Port	VID	No. Of Connections
63.22	80	1	23256
63.22	80	2	5057
63.22	80	2	462
207.2	80	4	32
66.77	25	4	27
65.12	25	4	26
206.1	25	3	25
66.77	25	4	22
207.2	80	4	17
65.54	25	3	16

Viewing top URLs, hosts, referers, cookies, user-agents, HTTP accesses

At any given time, you may have several servers behind a given VID. Each server may have many busy URLs, Hosts, Referers, Cookies, User-Agents, HTTP Accesses. You can view the top items (having the most number of simultaneous connections) at any given second. This information is transient in nature and is updated every second. This information can be useful to determine the URLs which are getting the most traffic and within those, the URLs which are getting affected.

E.g., if you are a VID administrator, you can see Top URLs for your VID through the menu *Show* > *Current VID* > *Session Diagnostics* > *Top URLs*.

Denying sources from Top URLs table

During attacks, it is sometimes useful to directly deny certain URLs that you see in the diagnostics tables.

If you feel that the URLs being shown in the Top URLs diagnostics could be the affected, you can simply press the "*Update the URL ACL list with the URLs from the above table*".

Aggressively aging sources from the protected servers

During attacks, sometimes, the servers get overloaded from connections. Even though FortiDDoS may drop some packets and stop them from going further to the server, the connections still stay in the server's memory. It is sometimes useful to aggressively age connections from certain sources that you see in the diagnostics tables.

If you feel that the sources being shown in the Top Servers diagnostics could be the culprits, you can simply click on the "*Aggressively Age*" hyperlink against the IP address. This will age all existing connections from that source on the sever by sending a TCP RST packet to the server.

Viewing top TCP destinations and top TCP client/server pairs of a given port (based on number of connections)

FortiDDoS maintains a table of TCP connection tuples. A TCP connection tuple consists of:

- Source IP address
- Source Port
- Destination IP address
- Destination Port

Associated with each tuple is the current packet rate, direction of connection initiation and whether it is currently blocked due to an associated flood. This information is transient in nature and is updated every second. FortiDDoS maintains a table of up to 1 million entries in memory.

As a diagnostic tool, you can view top destinations with the highest number of connections within a given VID given a destination port.

As a diagnostic tool, you can also view top TCP client/server pairs with the highest number of connections within a given VID given a destination port.

If you are a VID administrator, you can see this information for your own VID through the menu *Show* > *Current VID* > *Session Diagnostics* > *Port*. Enter the *Port* you are interested in and enter the *number of entries* and click *OK*. You will see the results in a tabular form.

Viewing top TCP destination ports and top TCP client/server pairs of a given host (based on number of connections)

FortiDDoS maintains a table of TCP connection tuples. A TCP connection tuple consists of:

- Source IP address
- Source Port
- Destination IP address
- Destination Port

Associated with each tuple is the current packet rate, direction of connection initiation and whether it is currently blocked due to an associated flood. This information is transient in nature and is updated every second. FortiDDoS maintains a table of up to 1 million entries in memory.

As a diagnostic tool, you can view top destination ports with the highest number of connections within a given VID given a server.

As a diagnostic tool, you can also view top TCP client/server pairs with the highest number of connections within a given VID given a server.

If you are a VID administrator, you can see this information for your own VID through the menu *Show* > *Current VID* > *Session Diagnostics* > *Server*. Enter the *valid server address* and the *number of entries* and click *OK***.** You will see the results in a tabular form.

Viewing top TCP client/server pairs of a given host (based on number of connections)

FortiDDoS maintains a table of TCP connection tuples. A TCP connection tuple consists of:

- Source IP address
- Source Port
- Destination IP address
- Destination Port

Associated with each tuple is the current packet rate, direction of connection initiation and whether it is currently blocked due to an associated flood. This information is transient in nature and is updated every second. FortiDDoS maintains a table of up to 1 million entries in memory.

As a diagnostic tool, you can view top client and server pairs with the highest number of connections within a given VID given a client IP address.

If you are a VID administrator, you can see this information for your own VID through the menu *Show* > *Current VID* > *Session Diagnostics* > *Client*. Enter the *valid client address* and the *number of entries* and click *OK***.** You will see the results in a tabular form.

Archive

Archiving, and restoring data

FortiDDoS maintains 3 primary types of data on its internal storage for continued operation:

- Traffic History

  FortiDDoS stores traffic history for various network parameters in each direction for each VID. This helps in forecasting the traffic pattern in future and in estimating the thresholds adaptively. The traffic history files do not grow in size as time passes as they are in a round-robin form.

- Configuration

  FortiDDoS stores the configuration related to thresholds, ACLs, My Lists, Notification Parameters, etc. in the configuration database. The configuration database does not significantly grow as time passes.

- Events

  FortiDDoS stores the events associated with a VID in an event database. This database grows as time passes.

Note: The default maximum size of the event database is 1 million entries. Once this limit is reached, the system does an auto-truncation of the past entries so that the entries are once again restricted to the 1 million mark. If past events are important to you, you must archive them using the archive function.

Archiving operation allows you to upload any or all of the above to an FTP server.

Backing up events

Backing up operation allows you to download the events from the event database. The database is downloaded in a compressed form.

Click on *Manage* > *Global* > *Event Database* > *Backup.* You will see a screen shown in Figure 54.

Figure 54: Backing up events

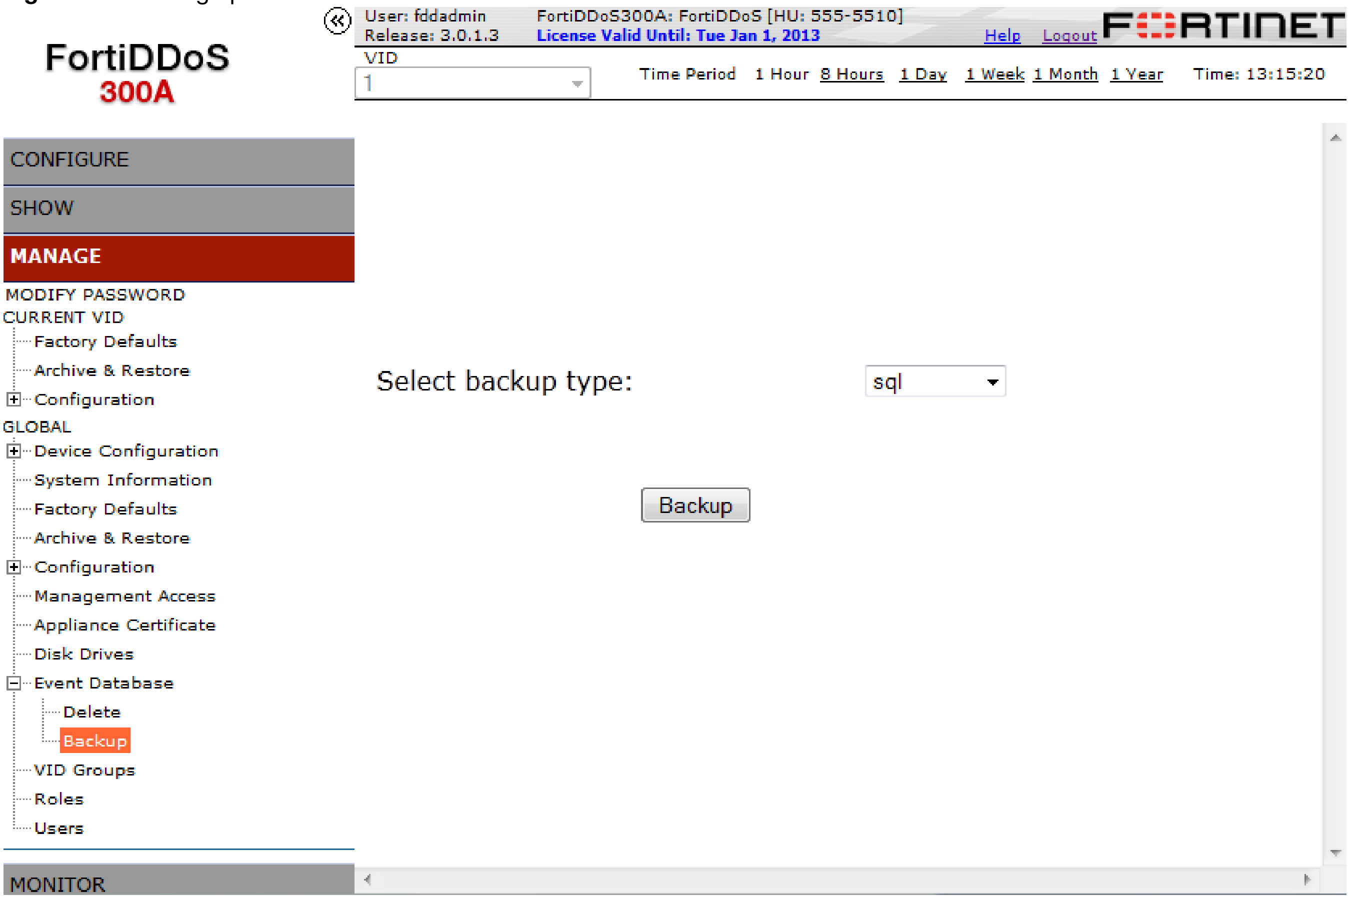

Backing up the event database

The event database contains currently stored events. To back up the event database, click on *Manage* > *Global* > *Event Database* > *Backup*.

- Select backup type *sql* if you want to get the SQL file.
- Select backup type *sql.gz* if you want to get the compressed SQL file that you can uncompress on your workstation.
- Click *Backup* to download the file to your workstation. You can save and use this file with applications that understand the SQL file syntax.

Backing up configuration

The configuration consists of all changes you have done to various policies and setup related to thresholds and features controls on the data path.

To back up the global configuration, click on *Manage* > *Global* > *Configuration* > *Backup*.

To back up the current VID configuration, click on *Manage* > *Current VID* > *Configuration* > *Backup*.

- Click *Backup* to initiate backup.
- Once the backup process is over, it shows a link called *Click here to download the last backup file*.
- Click the above hyperlink and download the file on your workstation.

Restoring configuration

Once you have backed up the configuration using the above process, you can restore it at a subsequent time.

To restore the global configuration, click on *Manage* > *Global* > *Configuration* > *Restore*.

To restore the current VID configuration, click on *Manage* > *Current VID* > *Configuration* > *Restore*.

- Click the *Browse* Button to locate the file.
- Click the *Restore* button to initiate the restore process.

Deleting Events

After you have backed up the events, you can delete them, thus freeing up space on the hard disk.

Click on *Manage* > *Global* > *Event Database* > *Delete.* You will see a screen shown in Figure 55.

Figure 55: Deleting events

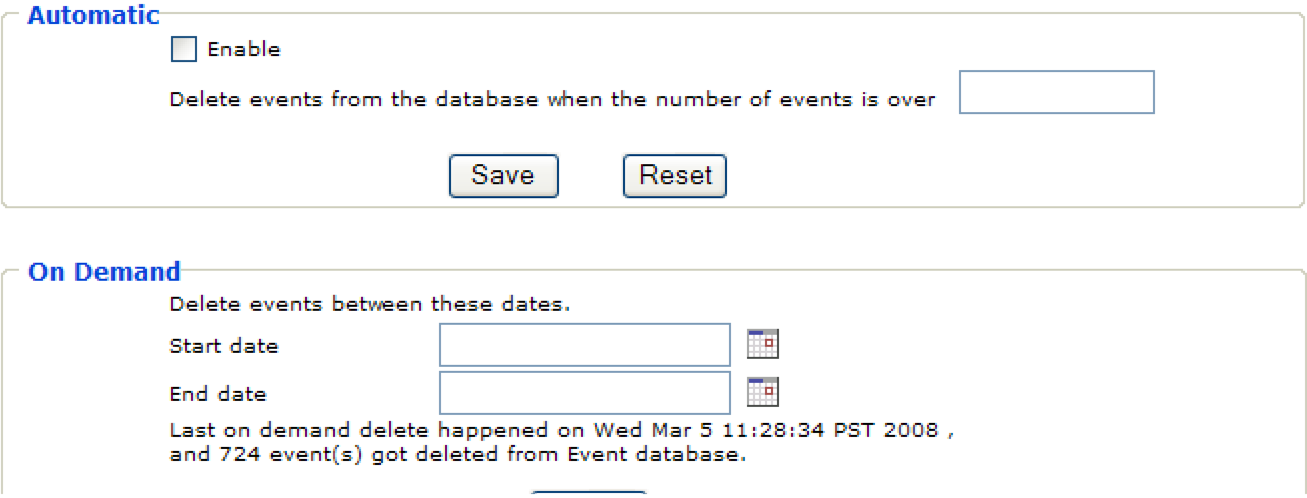

Automatic

□ Enable

Delete events from the database when the number of events is over

Save Reset

On Demand

Delete events between these dates.

Start date

End date

Last on demand delete happened on Wed Mar 5 11:28:34 PST 2008 , and 724 event(s) got deleted from Event database.

Deleting operation can be *automatic* or *on-demand*.

If you choose the automatic option, the system will delete the events anytime the size of the number of the events in the database exceeds certain thresholds. This is useful as a safety mechanism.

If you choose manual option, you can delete on demand.

Automatic deletion

1 Click on *Enable* check-box. Enter a number in the **"***Delete events from the database when the number of events is over***"** edit-box

2 Click *Delete* to setup the deletion of the selected events

Manual or on demand deletion

1 Enter a *Start date* and an *End date*. You can use the calendar controls to select dates. You can enter future dates for *End date.*

2 Click *Delete* to delete the events

Once you have deleted events from the database, you cannot use the data for analysis on **Monitor** or **Reports**.

Other administrative tasks

As an administrator, you can perform a few other tasks on FortiDDoS. These including:

- Setting system date
- Changing the management IP address
- Changing the email delivery method
- Rebooting
- Restarting services
- Upgrading the firmware
- Restricting management access to specific IP addresses
- Setting a VID to factory defaults

The following sections describe these in detail.

Setting system date

The FortiDDoS appliance ships with Pacific Standard Time (PST) timezone and the corresponding date/time.

If you want to change the date and/or time and the time zone, click on *Manage* > *Global* > *Device Configuration* > *System Date.*

- Set the *month, date and the year*
- Set the *hour and minute*
- Set the *Time Zone*
- Click *Save*
- Reboot the appliance using *Manage>Device Configure>Reboot*

Changing the management IP address

If you wish to change the IP address for the management port, click on *Manage* > *Global* > *Device Configuration* > *IP Address.*

- Set the *IP Address.*
- Set the *Netmask.*
- Set the *Default Gateway.*
- Set the *DNS Server1 and DNS Server 2.*
- Click *Save.*

Changing the email delivery method

FortiDDoS can be used to send email notifications to external email addresses. In most cases, you will be able to use DNS to deliver email notifications based on fully qualified domain names.

However, some ISPs, in an effort to reduce e-mail spam originating at their customer's IP addresses, will not allow their customers to communicate directly with the recipient's mail server via the default SMTP port number 25. In this case, the customer has to use the smart host provided by the ISP. To re-route e-mail alerts via your ISP, select Forward all emails to your ISP and enter the e-mail server address in the Smart Host Server Name field.

If you wish to change the Email Delivery Method, click on *Manage* > *Global* > *Device Configuration* > *IP Address.*

- Choose either *Use DNS to route emails*
- Or *Forward all emails to your ISP (Smart host).*
- If you choose Smart host option, enter the *Smart Host Server Name*
- Click *Save*.

Shutting down

If you wish to gracefully shutdown the appliance through the Web-based Manager, click on *Manage* > *Global* > *Device Configuration* > *Shutdown.*

- Click *Shutdown*.
- Wait until the LEDs on the front panel are off. The process normally takes about 2 minutes to shutdown.

Rebooting

If you wish to reboot the appliance through the Web-based Manager, click on *Manage* > *Global* > *Device Configuration* > *Reboot.*

- Click *Reboot*.
- Wait until you can ping the management IP address or connect to the Web-based Manager.

Restarting services

This is a diagnostic option. Use this under supervision from the FortiDDoS Support Team.

If you wish to restart services on the appliance through the Web-based Manager, click on *Manage* > *Global* > *Device Configuration* > *Restart Services.*

- Click *Restart Services*.
- Wait until you can connect to the Web-based Manager.

Upgrading the firmware

FortiDDoS allow you to upgrade the firmware remotely. You will receive firmware upgrades from Fortinet, Inc. These are encrypted and encoded files. You can receive them either via email or from a File Transfer Protocol (FTP) server.

Once you have this file on your workstation, you can upgrade the FortiDDoS by uploading the file to the appliance.

Only administrator with the right privilege can perform this operation.

Please remember that upgrades can only be reverted to one prior version.

If the system requires a reboot after upgrade, you will be informed. Before upgrades, ensure that you have read the release notes and schedule the operation with your users.

To perform the upgrade, click on *Manage* > *Global* > *Device Configuration* > *Upgrade Device.*

- You will see Upgrade History as a list. This shows you past upgrades that were performed on this appliance.
- In the *Upgrade* section of the screen, click on *Browse* button and point to the file that you have received from Fortinet, Inc.
- Click on *Upgrade* button.
- If for some reason, the upgrade cannot be performed, the appliance reverts back to the prior version.
- If for some reason, you want to revert to the prior version click the *Revert to Previous Version* button. System determines if such an operation is permissible. If the operation is not permissible, this button is disabled.

Logging to external syslog servers

FortiDDoS use the syslog protocol to manage system logs and alerts. Since the space on an appliance is limited, you can use an external syslog server. The storage size then does not depend on the FortiDDoS's resources and is limited only by the available disk space on the external syslog server. This option is not enabled by default.

Before configuring an FortiDDoS device to send syslog messages, make sure that it is configured with the right date, time, and time zone. Syslog data would be useless for troubleshooting if it shows the wrong date and time. Setting the devices with the accurate time is helpful for event correlation.

Because this is external to the FortiDDoS, there is an added benefit: a syslog server also provides for centralized logging for all network devices.

You can select which events need to be logged. Depending on what you select, the amount of data transferred will vary.

There are three kinds of events that are sent to the syslog server:

- VID based events like Flood events, ACL events, Scan events etc. You can enable logging of these events as described in "Enabling event syslog feature for the VID" on page 106.
- Non-VID based events like hardware and software malfunctions and RAID events. You can enable logging of these events as described in "Enabling event syslog feature for the appliance" on page 106
- General information about software processes running on the system. You can enable logging of these events as described in "Enabling syslog feature for the appliance" on page 105.

Enabling syslog feature for the appliance

FortiDDoS devices use a severity level of debug through fatal to generate error messages about software or hardware malfunctions. The debugging level displays the output of debug commands. The Notice level displays interface up or down transitions and system restart messages. The informational level reloads requests and low-process stack messages.

Specifies the kind of messages, by severity level, to be sent to the syslog server. The default is informational and lower. The possible values for level are as follows:

- Fatal
- Alert
- Critical
- Error
- Warning
- Notice
- Informational
- Debug

When you specify a logging level, the appliance is configured to send messages with lower severity levels as well. For example, if you configure the Warning level, the device will send all messages with the severity warning, error, critical, and fatal. Similarly, if you specify debug level, the device will send all messages to the syslog server. Exercise caution while enabling the debug level.

If you wish to enable logging of the events to an external syslog server, click on *Manage* > *Global* > *Device Configuration* > *Logging.*

- In *Syslog Redirection* box, click *Enable Syslog Redirection.*
- Enter *Remote IP Address.*
- Select the *Priority.*
- Click *Save*.

Enabling event syslog feature for the appliance

If you wish to enable logging of Non-VID based events like hardware and software malfunctions and RAID events to an external syslog server, click on *Configure* > *Global* > *Event Notification* > *Syslog.*

- In the *Events Syslog Redirection* box, click *Enable Event Syslog***.**
- Enter *Remote IP Address.*
- Click *Save.*

Enabling event syslog feature for the VID

If you wish to enable logging of VID based events like Flood events, ACL events, Scan events etc. to an external syslog server, click on *Configure* > *Current VID* > *Event Notification* > *Syslog.*

- In the *Events Syslog Redirection* box, click *Enable Event Syslog.*
- Enter *Remote IP Address.*
- Click *Save.*

Restricting management access to specific IP addresses

FortiDDoS allow you to restrict management access to certain specified IP addresses. This helps in ensuring that only certain IP addresses can manage the appliance.

You can enable/disable this feature using a single checkbox. You can specify a list of IP addresses and subnets from which management access is allowed. The management protocols associated with FortiDDoS are HTTP, CLI over SSH (for diagnostics), SNMP, and ICMP for connectivity checks. By default management access is open to everyone. After initial installation, you must restrict the access to conform to your organization’s policies. Once you configure restricted management access, the appliance acts creates a firewall which only allows specified IP addresses to access the specified protocols, all other IP addresses are blocked. To ensure that you don’t block your current IP address, your own IP address is shown to you on the screen. You must ensure that this IP address is in the list of allowed IP addresses. You can specify up to 10 networks from where the management access is allowed.

Figure 56: Management access

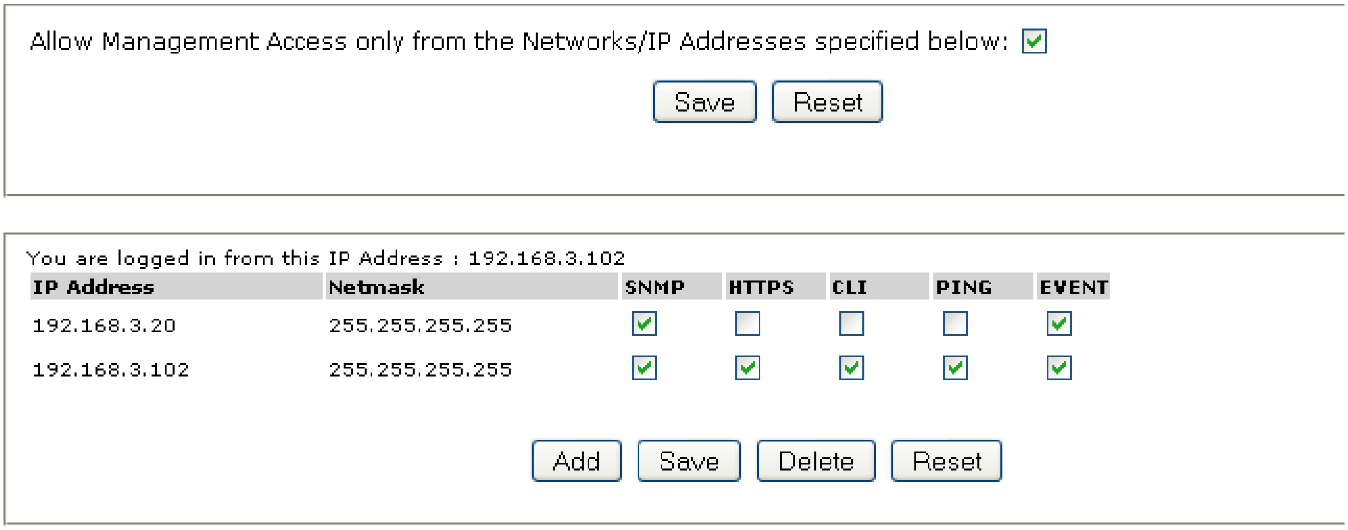

Allow Management Access only from the Networks/IP Addresses specified below: ☑

You are logged in from this IP Address : 192.168.3.102

IP Address	Netmask	SNMP	HTTPS	CLI	PING	EVENT
192.168.3.20	255.255.255.255	☑	□	□	□	☑
192.168.3.102	255.255.255.255	☑	☑	☑	☑	☑

To restrict management access, click on *Manage* > *Global* > *Management Access.*

- If you want to restrict management access, check the box titled *Allow Management Access From the Following Networks.*
- If the above checkbox is checked, the remaining items below are enabled for entry.
- Your current IP address is shown in the screen.
- Enter up to 10 IP addresses the corresponding netmasks.
- For each allowed network, check the protocols viz. SNMP, Web, CLI, PING, Event database access that you want allowed.
- Click *Save* to save the configuration.
- If you want to restart from the last saved configuration, click *Reset* and start all over again.

Management of the appliance on a port other than the default SSL TCP port 443

FortiDDoS allow you to change the SSL access on a TCP port other than default SSL port 443. Associated with this is the restriction of management access which automatically enables access to that new TCP port.

E.g. you can run the management access on TCP port 1443.

If you change the port from the default value of 443, the management webserver will be restarted.

You then have to logout and access the Web-based Manager using a different URL.

E.g. if the new value of the port is 1443, you have to access using: https://ipaddress:1443

To change management access SSL Port, click on *Manage* > *Global* > *Management Access.*

- Enter *New SSL Port*.
- Click *Save* to save the configuration.

Setting a VID to factory defaults

If the traffic pattern associated to a VID changes drastically, you must set the VID to factory defaults.

This involves setting following 3 key data to default values:

- Traffic History

  FortiDDoS stores traffic history for various network parameters in each direction for each VID. This helps in forecasting the traffic pattern in future and in estimating the thresholds adaptively.

  Since the traffic history is very much tied to the network, any change in network configuration or drastic changes in traffic patterns affect the threshold estimation.

  In case, you see changes in traffic pattern is drastically changing due to addition of a new server or new service or something similar, you must clear the traffic history for that VID.

  This is a slow process as a whole year worth of round-robin history has to be initialized - even though you may not have a year's worth of uptime. Once you erase the history for a VID, all traffic graphs will be initialized and will not show past data. The system has to be retrained at least for 2 days.
- Configuration

  FortiDDoS stores the configuration related to thresholds, ACLs, My Lists, Notification Parameters, etc. in the configuration database. If you are changing the traffic pattern to a VID, you must consider setting the configuration to factory defaults as well.
- Events

  FortiDDoS stores the events associated with a VID for each VID separately. If you are changing the traffic pattern to a VID, and the past events are of no significance, you must consider setting the events to factory defaults as well - i.e. removing the past events altogether.

You must ensure that the system is not shut down during this process.

To set a VID to factory defaults, plan a shutdown, click on *Manage* > *Current VID* > *Factory Defaults*

- Read the warning.
- Click *Traffic History*, *Config*, *Event* depending on your needs.
- Enter *Yes* in the text box.
- Click *OK* to set the VID to factory defaults and wait until the process is complete. This process takes about 15 minutes per VID. During this time services to all other VIDs are affected.

Setting the appliance to factory defaults

If the traffic pattern associated with the appliance, i.e. all VIDs changes drastically, you must set the appliance to factory defaults. This will set all the VIDs and global configurations to the factory defaults.

This involves setting following 3 key data to default values:

- Traffic History

  FortiDDoS stores traffic history for various network parameters in each direction for each VID. This helps in forecasting the traffic pattern in future and in estimating the thresholds adaptively.

  Since the traffic history is very much tied to the network, any change in network configuration or drastic changes in traffic patterns affect the threshold estimation.

  This is a slow process as a whole year worth of round-robin history has to be initialized - even though you may not have a year's worth of uptime. Once you erase the history for a VID, all traffic graphs will be initialized and will not show past data. The system has to be retrained at least for 2 days.

- Configuration

  FortiDDoS stores the configuration related to thresholds, ACLs, My Lists, Notification Parameters, etc. in the configuration database. If you are changing the traffic pattern to a VID, you must consider setting the configuration to factory defaults as well.

- Events

  FortiDDoS stores the events associated with a VID for each VID separately. If you are changing the traffic pattern to a VID, and the past events are of no significance, you must consider setting the events to factory defaults as well - i.e. removing the past events altogether.

You must ensure that the system is not shut down during this process.

To set an FortiDDoS to factory defaults, plan a shutdown, click on *Manage* > *Global* > *Factory Defaults.*

- Read the warning.
- Click Traffic History, Config, Event depending on your needs.
- Enter *Yes* in the text box.
- Click *OK* to set the appliance to factory defaults and wait until the process is complete. This process takes about 15 minutes per VID. During this time services to all other VIDs are affected.

Managing the appliance's SSL certificate

Introduction to SSL

The Secure Socket Layer protocol was created to ensure secure transactions between web server interfaces such as that in FortiDDoS and browsers. The protocol uses a third party, a Certificate Authority (CA), to identify one end or both end of the transactions. This is in short how it works.

- A browser requests a secure page (usually https://192.168.1.1).
- The FortiDDoS webserver sends its public key with its certificate.
- The browser checks that the certificate was issued by a trusted party (usually a trusted root CA), that the certificate is still valid and that the certificate is related to the site contacted.
- The browser then uses the public key, to encrypt a random symmetric encryption key and sends it to the server with the encrypted URL required as well as other encrypted http data.

- The FortiDDoS web server decrypts the symmetric encryption key using its private key and uses the symmetric key to decrypt the URL and http data.
- The web server sends back the requested html document and http data encrypted with the symmetric key.
- The browser decrypts the http data and html document using the symmetric key and displays the information.

Avoiding common browser warning messages related to certificates

When you access the FortiDDoS through HTTPS, your browser may display a Client Authentication message followed by a Security Alert message. Messages dialogs display when you first establish an HTTPS session with the FortiDDoS. You may get any one of the following messages:

- **This website’s address does not match the address in the security certificate:**

  This error indicates the website is using a digital certificate that was issued to a different web address. This error occurs because when the appliance is shipped, the factory is not aware of the address that will be used for the appliance. The address on the certificate by default is FortiDDoS.

  You can get rid of this error by:

  - Creating a fresh self-signed certificate with the *Common Name* corresponding to the management IP address of the appliance. Refer to “Creating self-signed certificates” on page 111.
  - Generating a certificate request followed by uploading a new signed certificate. Refer to “Installing an externally signed certificate” on page 112.
- **This website’s security certificate is out of date**:

  This error occurs when the current date is either before or after the time period during which the certificate is valid. You may see this error after a prolonged use of the appliance. The factory created certificate is valid for a certain period. Once that period is over, you will see this error.

  You can get rid of this error by:

  - Creating a fresh self-signed certificate with the correct *Valid Until* date. Refer to “Creating self-signed certificates” on page 111.
  - Generating a certificate request followed by uploading a new signed certificate. Refer to “Installing an externally signed certificate” on page 112.
- **This website’s security certificate is not from a trusted source**:

  This error occurs when the certificate has been issued by a certification authority that is not recognized by your browser. The default certificate shipped with the FortiDDoS is signed by FortiDDoS. Since FortiDDoS may not be a trusted certification authority in your browser’s database, you will see this error.

  You can get rid of this error by one of the following:

  - Adding the certificate authority for existing certificate in your browser as a trusted root. Refer to “Managing your browser’s trusted roots” on page 115
  - Generating a certificate request followed by uploading a new signed certificate from a trusted authority. Refer to “Installing an externally signed certificate” on page 112.

Creating self-signed certificates

To create a self-signed certificate for an FortiDDoS, click on *Manage* > *Global* > *Appliance Certificates.*

- You will see the details of current certificate in use.

Figure 57: Current certificate details

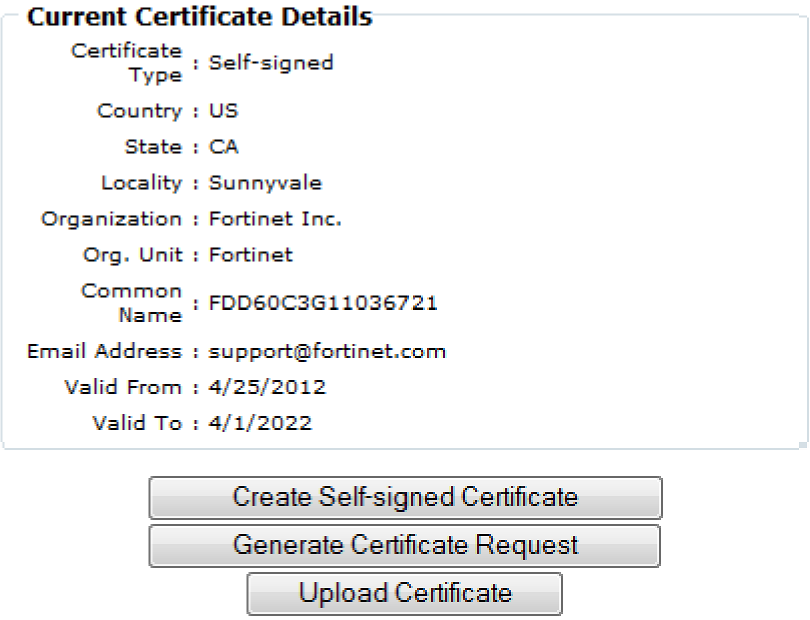

Current Certificate Details

Certificate Type :	Self-signed
Country :	US
State :	CA
Locality :	Sunnyvale
Organization :	Fortinet Inc.
Org. Unit :	Fortinet
Common Name :	FDD60C3G11036721
Email Address :	support@fortinet.com
Valid From :	4/25/2012
Valid To :	4/1/2022

- If you would like to create a self-signed certificate (which is the easier option of the two.), click *Create Self-signed Certificate.*
- In the dialog requesting details of the self-signed certificate, enter the details of the self-signed certificate to be created.
- Enter you *Country, State, and Locality*. It is very important that you do not abbreviate the names of the state or city. Please refer to ISO 3166 for 2 character country code names. These are described here:

  http://www.iso.ch/iso/en/prods-services/iso3166ma/02iso-3166-code-lists/list-en1.html.
- Enter your *Organization* (O) and your *Organizational Unit* (OU). For example, if your company is called Widgets and you are setting up the FortiDDoS for the Sales department, you would enter Widgets for the *Organization* and Sales for your *Organizational Unit*.
- Input the *Common Name* (CN) for your FortiDDoS. This should be the same name that the user will input when connecting to the management interface. For example, if a user inputs http://FortiDDoS.widgets.com to access your FortiDDoS, then your Common Name would be FortiDDoS.widgets.com. Alternatively use the IP address of the management interface.
- Enter the contact information for the person responsible for this FortiDDoS in the *Email Address* field. This is usually how the Certificate Authority contacts you, and then click Next.
- Enter the *Valid Until* date. This is the date until you want the certificate to be valid.

- Click *Create* to create the self-signed certificate. Your browser will then be redirected to the login screen.

Figure 58: Creating a self-signed certificate

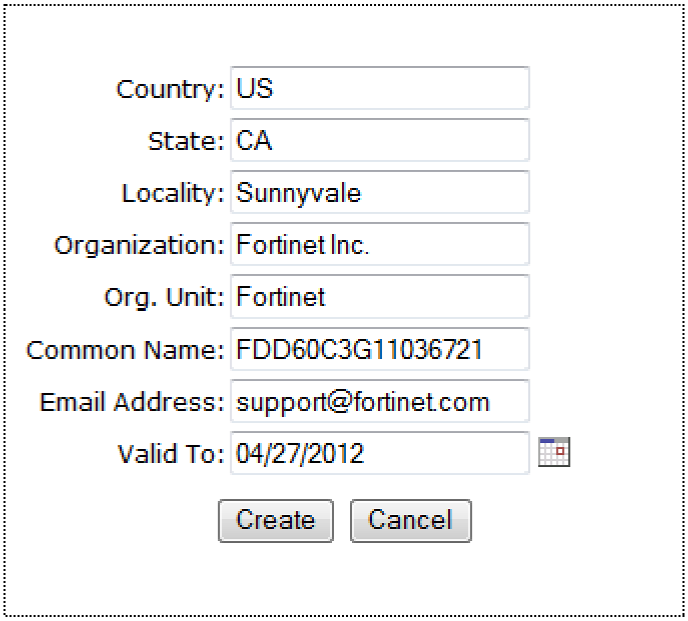

Country:	US
State:	CA
Locality:	Sunnyvale
Organization:	Fortinet Inc.
Org. Unit:	Fortinet
Common Name:	FDD60C3G11036721
Email Address:	support@fortinet.com
Valid To:	04/27/2012

Installing an externally signed certificate

To create an externally signed certificate and upload the certificate, you need to follow two separate steps:

- Generate Certificate Request
    - The first step you will need to perform to get a server certificate is creating the Certificate Signing Request (CSR) file. The CSR is a string of text generated by your FortiDDoS. You provide this string of text to your Certificate Authority (CA) during the enrollment process.

Figure 59: Generating the certificate request

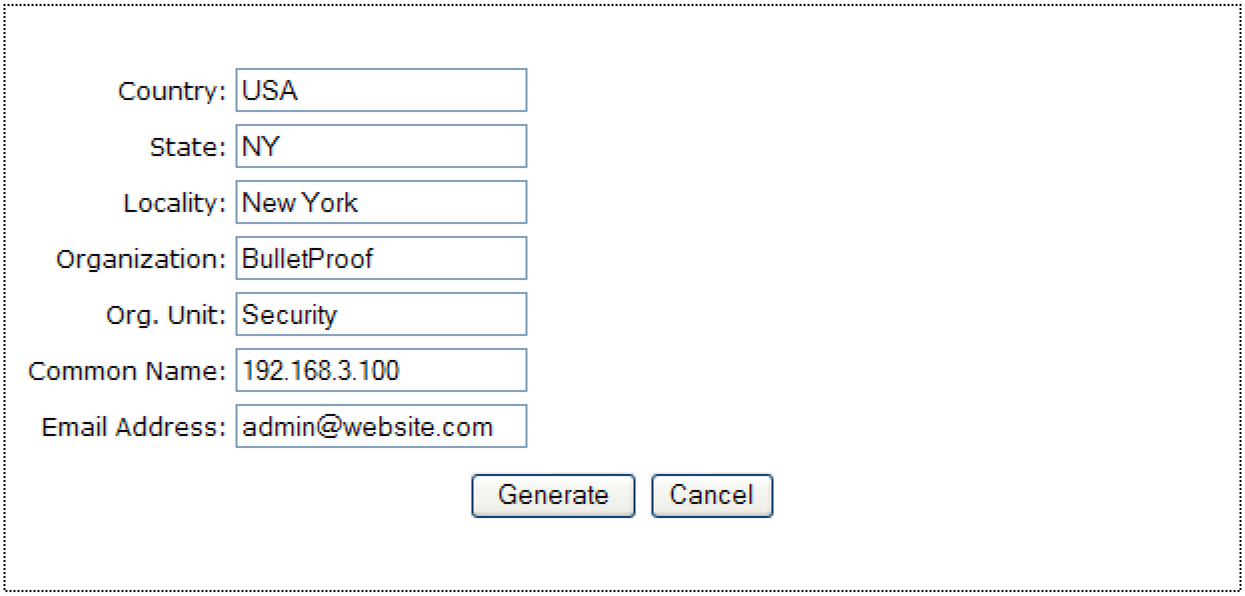

Country: USA

State: NY

Locality: New York

Organization: BulletProof

Org. Unit: Security

Common Name: 192.168.3.100

Email Address: admin@website.com

- Enter you *Country, State, and Locality*. It is very important that you do not abbreviate the names of the state or city. Please refer to ISO 3166 for 2 character country code names. These are described here:

  http://www.iso.ch/iso/en/prods-services/iso3166ma/02iso-3166-code-lists/list-en1.html.
- Enter your *Organization* (O) and your *Organizational Unit* (OU). For example, if your company is called Widgets and you are setting up the FortiDDoS for the Sales department, you would enter Widgets for the *Organization* and Sales for your *Organizational Unit*.
- Input the *Common Name* (CN) for your FortiDDoS. This should be the same name that the user will input when connecting to the management interface. For example, if a user inputs http://FortiDDoS.widgets.com to access your FortiDDoS, then your Common Name would be FortiDDoS.widgets.com. Alternatively use the IP address of the management interface.
- Enter the contact information for the person responsible for this FortiDDoS in the *Email Address* field. This is usually how the Certificate Authority contacts you, and then click Next.
- Enter the *Valid Until* date. This is the date until you want the certificate to be valid.
- Click *Generate***.** You will see a message saying *Generating Certificate Request*. Once the process is completed, you will see a text area under *Current Certificate Request*. This text will contain all the information you entered here in an encoded form, as well as your public key for your FortiDDoS.

Figure 60: Generated certificate request

```
-----BEGIN CERTIFICATE REQUEST-----
MIIB8jCCAVsCAQAwgYwxCzAJBgNVBAYTAlVTMQswCQYDVQQIEwJDQTESMBAGA1UE
BxMJU3Vubnl2YWxlMQowCAYDVQQKEwFPMQswCQYDVQQLEwJPVTEYMBYGA1UEAxMP
MTkyLjE2OC4xMDAuMjA4MSkwJwYJKoZIhvcNAQkBFhpncWluQGludHJ1Z3VhcmRk
ZXZpY2VzLmNvbTCBnzANBgkqhkiG9w0BAQEFAAOBjQAwgYkCgYEAzxzoladVeGgN
imoj2vLUIy7Hi2pdwPGVdVKaLhVWEslObt13FLIX1yR1VvmyyLzd7+mtPGOTmutF
1l7EcXrhAPIkkGYJR6XVGmtVMUVlW8XEZuHilpmb668pWm0nAd3W8E4GIwmHUXxb
WIXJSg9u9Ra77jUCMV0JFYSAn3ONp8ECAwEAAaAlMCMGCSqGSIb3DQEJBzEWExRB
IGNoYWxsZW5nZSBwYXNzd29yZDANBgkqhkiG9w0BAQUFAAOBgQBI7aMFQ2ffzmmg
cLYOeRZQ4spxDxH2Fue202KY4TaFeQ+/8EO3iLsO21mPSULR55XoFzZEwTVl7kgU
HqEGbspdr0iofj+YimhmxI2p955KNmeaZ1b6QRKN2FhABpKcrXagBSsky5JceUHO
8DeKOWtQZ6IWSdtsq2yLrhkzQmwtHA==
-----END CERTIFICATE REQUEST-----
```

- Copy and paste the *Generated Certificated Request* into a file and you have created your certificate request file.

Note: Between creating the request file and uploading the certificate, do NOT perform any of the following actions:
Change the management IP address of the FortiDDoS
Create another certificate request
Create a self-signed certificate

- Send the Certificate Request to a Certificate Authority
  - In order for this certificate to be used, submit this file to a Certificate Authority (online authority). They will generate a certificate response file, which contains your public key and is digitally signed by the Certificate Authority.
- Upload the Signed Certificate
  - When you receive your response file from the online authority, you will need to install this on the FortiDDoS. The response file contains your public key that has been signed by the authority.

Figure 61: Uploading the certificate response file received from the CA

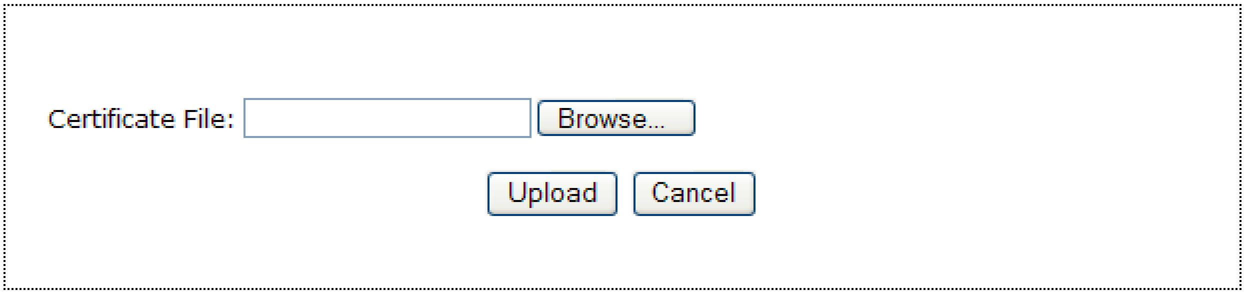

- Click *Upload Certificate.*
- Use the Browse button to locate the file containing the signed certificate response received from the CA.
- Click *Upload***.** Your browser will then be redirected to the login screen.

Managing your browser's trusted roots

To avoid getting a trusted root error during the connection to the management interface of the FortiDDoS, you must add the certificate's issuer as a trusted root. To do that follow the steps shown below.

- When you get the trusted root warning, click *View Certificate.* You will see a dialog similar to one shown in Figure 62.

Figure 62: Viewing the appliance certificate

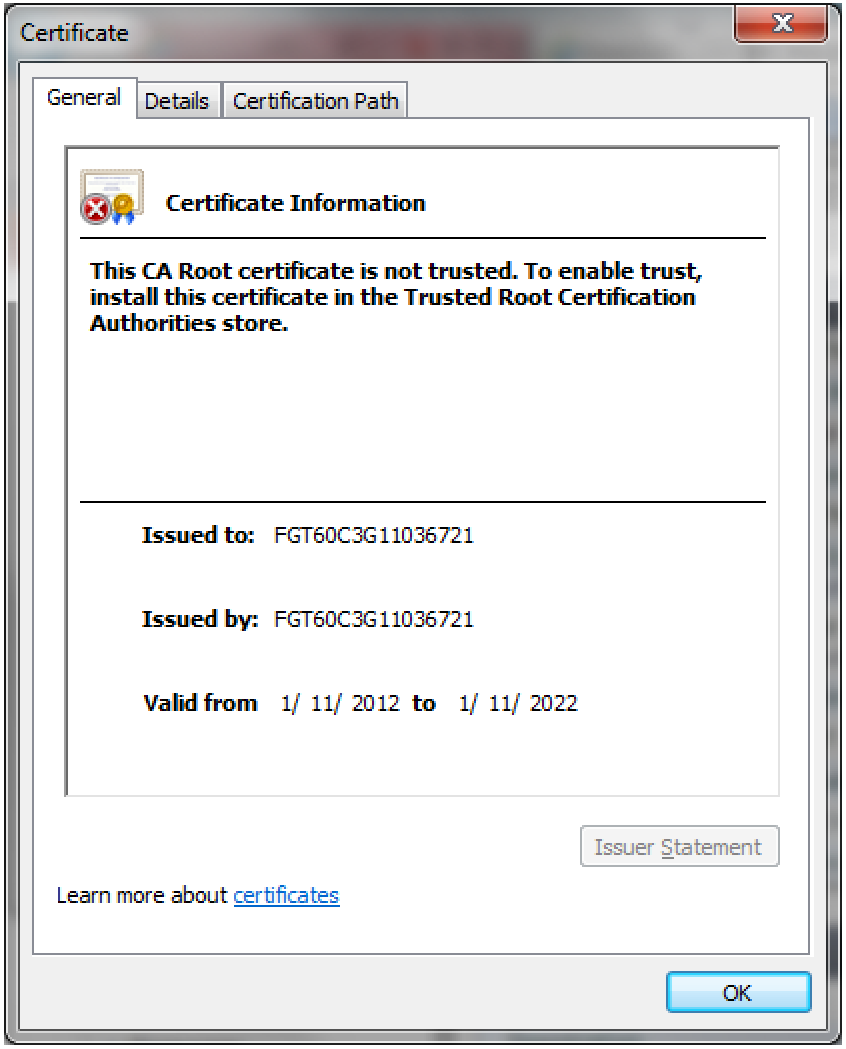

Certificate Information

This CA Root certificate is not trusted. To enable trust, install this certificate in the Trusted Root Certification Authorities store.

Issued to: FGT60C3G11036721

Issued by: FGT60C3G11036721

Valid from 1/ 11/ 2012 **to** 1/ 11/ 2022

Learn more about certificates

- Click on *Install Certificate.*
- You will see a *Certificate Import Wizard* dialog shown in Figure 63.

Figure 63: Certificate import wizard

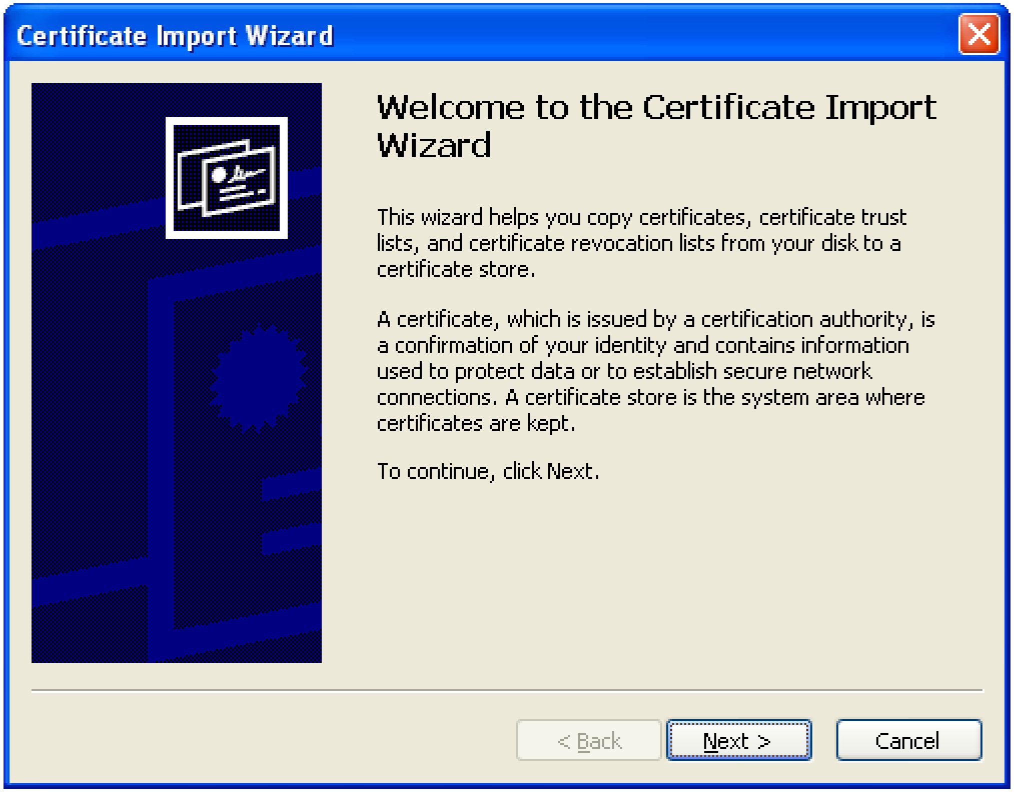

Welcome to the Certificate Import Wizard

This wizard helps you copy certificates, certificate trust lists, and certificate revocation lists from your disk to a certificate store.

A certificate, which is issued by a certification authority, is a confirmation of your identity and contains information used to protect data or to establish secure network connections. A certificate store is the system area where certificates are kept.

To continue, click Next.

- Click *Next.*
- You will then see a *Certificate Import Wizard* with *Certificate Store* Screen. Click *Place all certificates in the following store.* See Figure 64.

Figure 64: Certificate store

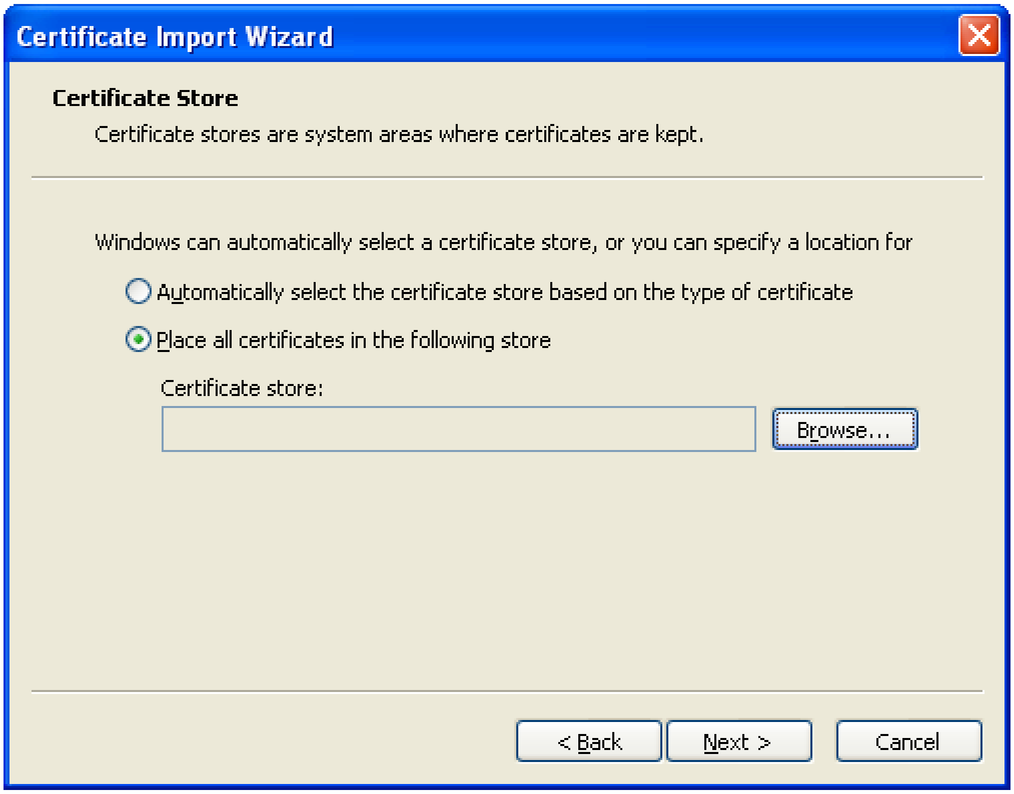

Certificate Import Wizard

Certificate Store

Certificate stores are system areas where certificates are kept.

Windows can automatically select a certificate store, or you can specify a location for

○ Automatically select the certificate store based on the type of certificate

● Place all certificates in the following store

Certificate store:

- In *Certificate Store*, click *Browse.*
- You will then see a dialog (see Figure 65) that allows you to select the store.

Figure 65: Select certificate store

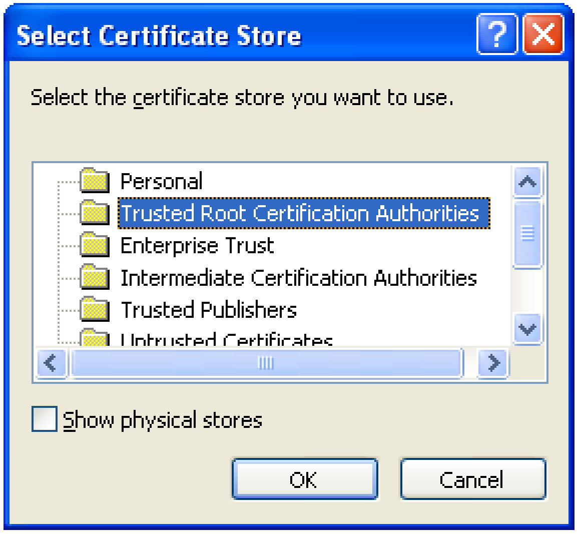

Select Certificate Store

Select the certificate store you want to use.

- Personal
- Trusted Root Certification Authorities
- Enterprise Trust
- Intermediate Certification Authorities
- Trusted Publishers
- Untrusted Certificates

□ Show physical stores

- Click on *Trusted Root Certification Authorities* and click *OK*.
- Click *Next*. This will finish the wizard through Figure 66. Click *Finish.*

Figure 66: Completing the certificate import wizard

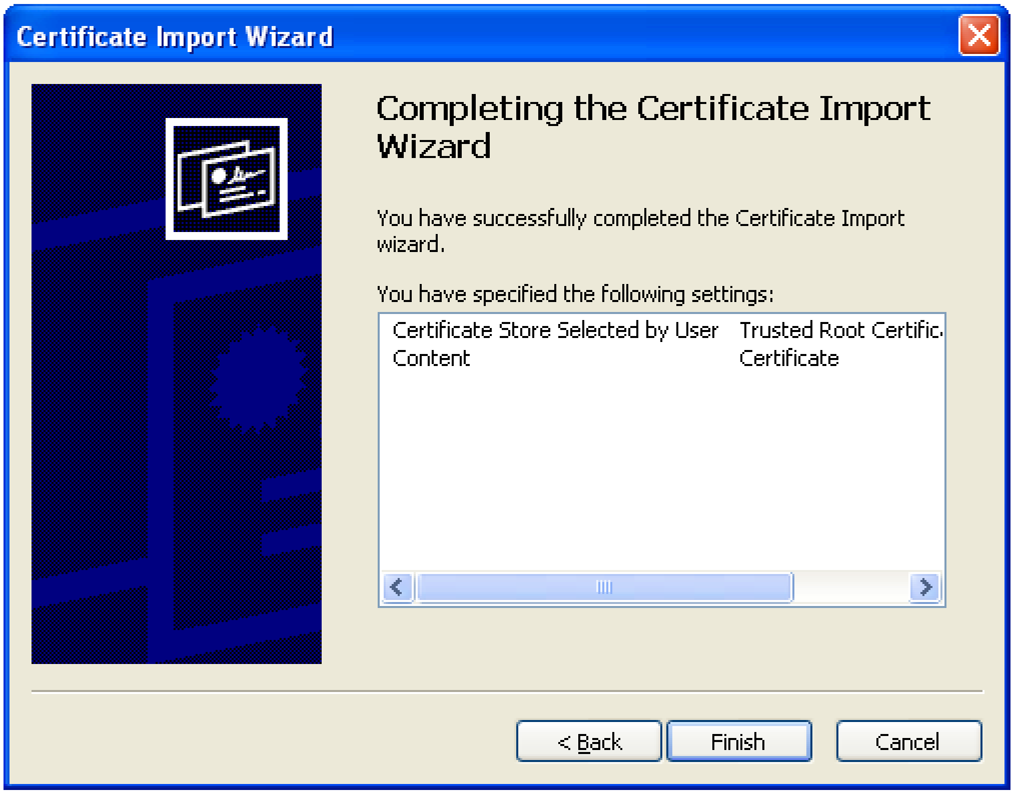

Certificate Import Wizard

Completing the Certificate Import Wizard

You have successfully completed the Certificate Import wizard.

You have specified the following settings:

Certificate Store Selected by User	Trusted Root Certific
Content	Certificate

- You will then see the final security warning shown in Figure 67. Click *Yes.*

Figure 67: Security warning

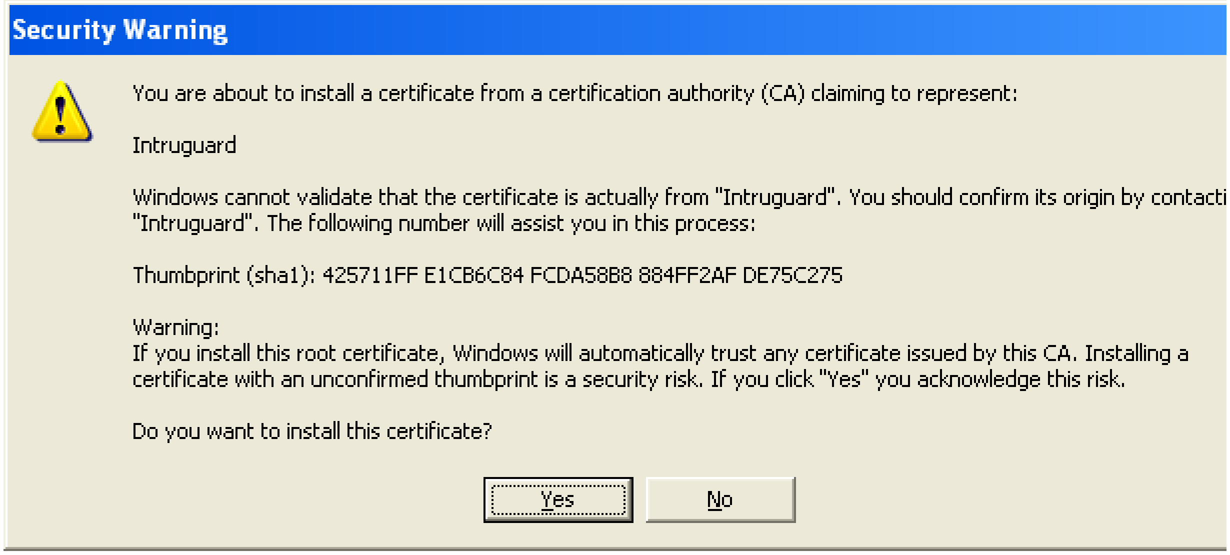

Security Warning

You are about to install a certificate from a certification authority (CA) claiming to represent:

Intruguard

Windows cannot validate that the certificate is actually from "Intruguard". You should confirm its origin by contacti "Intruguard". The following number will assist you in this process:

Thumbprint (sha1): 425711FF E1CB6C84 FCDA58B8 884FF2AF DE75C275

Warning:
If you install this root certificate, Windows will automatically trust any certificate issued by this CA. Installing a certificate with an unconfirmed thumbprint is a security risk. If you click "Yes" you acknowledge this risk.

Do you want to install this certificate?

Your browser is now ready to trust the FortiDDoS or current certificate authority. From your next browser session onwards, you will not see the trusted root warning.

Managing a RAID array

FortiDDoS has a Redundant Array of Intelligent Disks (RAID). This is a pair of hard disks which are mirrored automatically. This results in increased data integrity, fault tolerance.

Figure 68 shows a picture of the FortiDDoS-100A front panel where the RAID disks are located. The two disks are marked DISK1 and DISK2 respectively.

If any of the disks malfunctions it can be removed while the system is running. Due to the redundancy, the system can continue to run with the alternative disk. You can turn the knobs and pull the malfunctioning disk out.

Figure 68: RAID array on the front panel of a FortiDDoS-100A

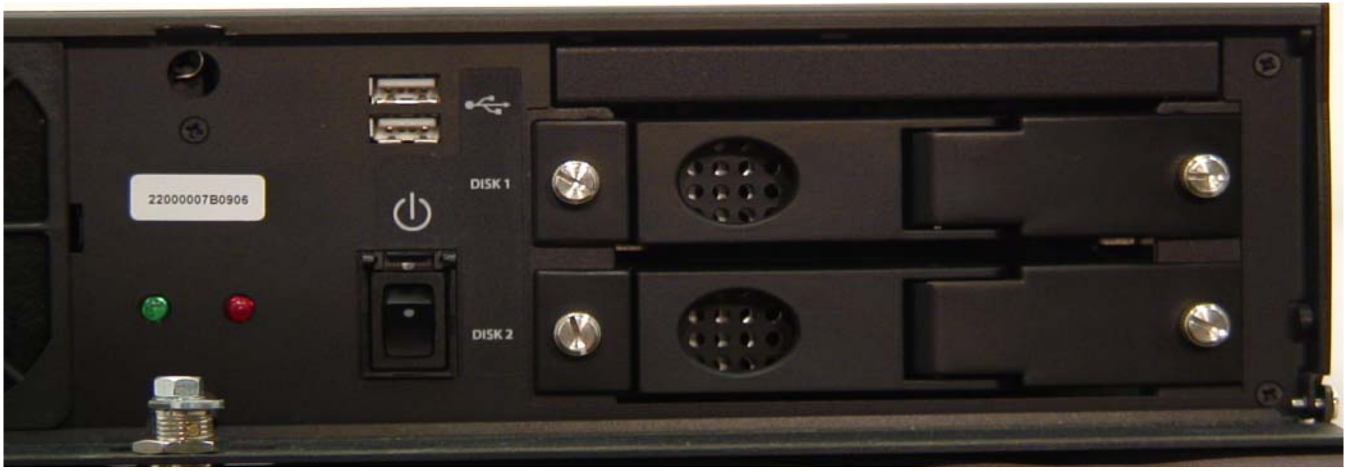

Monitoring event monitor for disk failures

As part of regular monitoring of the network, you must ensure that you keep an eye on the monitor screen and have e-mail notifications turned on. You must keep the *Event Categories* > *Device Events* checked in the Monitor screen. This is shown in the Figure 69.

Figure 69: Event categories for device events

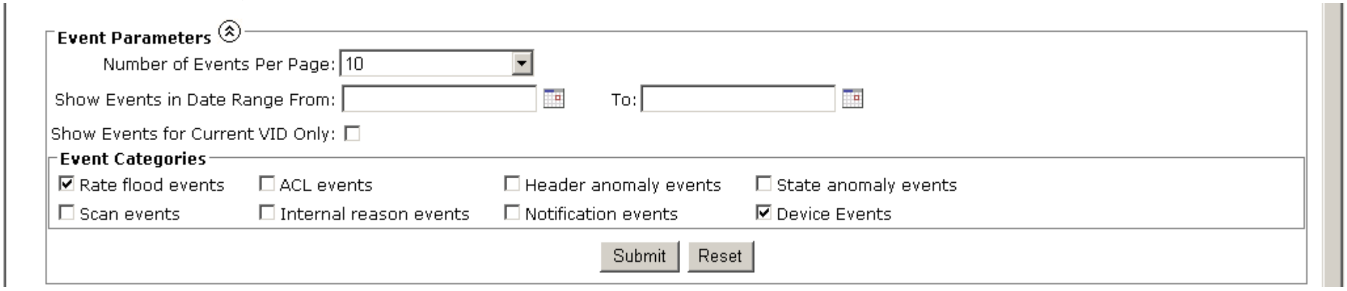

Event Parameters

Number of Events Per Page: 10

Show Events in Date Range From: To:

Show Events for Current VID Only: □

Event Categories

☑ Rate flood events	□ ACL events	□ Header anomaly events	□ State anomaly events
□ Scan events	□ Internal reason events	□ Notification events	☑ Device Events

Submit Reset

Once the Device Events are enabled, in case of disk failures, you will see an event similar to one shown in Figure 70.

Figure 70: RAID failure event in the monitor screen

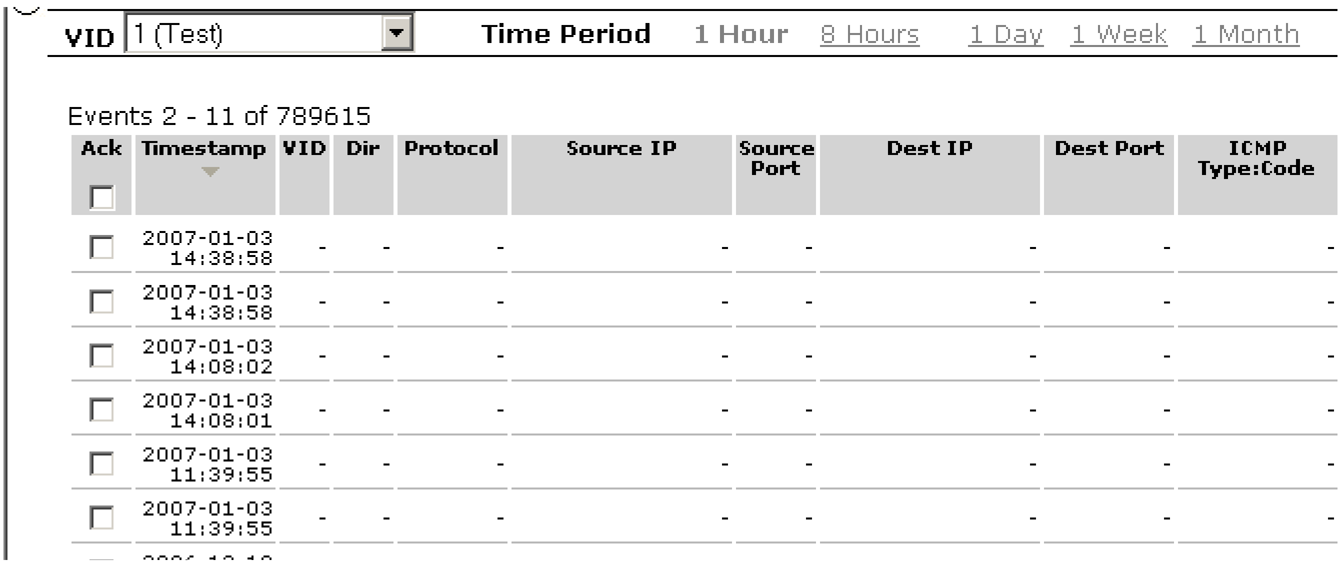

VID 1 (Test) Time Period 1 Hour 8 Hours 1 Day 1 Week 1 Month

Events 2 - 11 of 789615

Ack ☐	Timestamp	VID	Dir	Protocol	Source IP	Source Port	Dest IP	Dest Port	ICMP Type:Code
☐	2007-01-03 14:38:58	-	-	-	-	-	-	-	-
☐	2007-01-03 14:38:58	-	-	-	-	-	-	-	-
☐	2007-01-03 14:08:02	-	-	-	-	-	-	-	-
☐	2007-01-03 14:08:01	-	-	-	-	-	-	-	-
☐	2007-01-03 11:39:55	-	-	-	-	-	-	-	-
☐	2007-01-03 11:39:55	-	-	-	-	-	-	-	-

Once you know that one of your disks has failed, you must replace the disk or repair the disk after contacting Fortinet Technical Support.

After replacing the faulty disk, select *Manage* > *Global* > *Disk Drives.* This will show you the current state of the disks. Newly mounted disk will show as *degraded*. Please refer to Figure 71.

Figure 71: Disk details screen when one of the disks is degraded

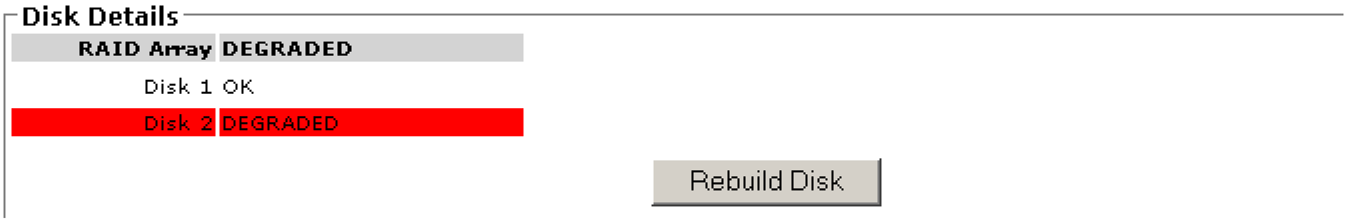

Disk Details

RAID Array	DEGRADED
Disk 1	OK
Disk 2	DEGRADED

You must then click *Rebuild Disk* to start the rebuild process.

You will then see a screen similar to Figure 72 where you will be asked for confirming the rebuild operation. Enter YES or yes to continue.

Figure 72: Confirming a rebuild disk operation

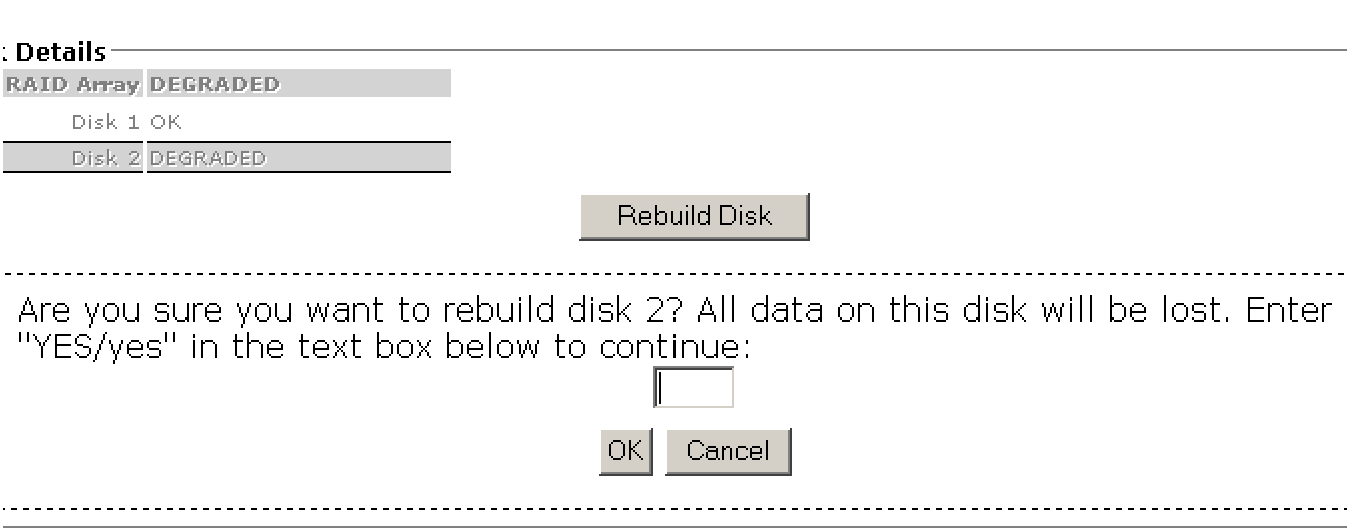

: Details

RAID Array	DEGRADED
Disk 1	OK
Disk 2	DEGRADED

Are you sure you want to rebuild disk 2? All data on this disk will be lost. Enter "YES/yes" in the text box below to continue:

Once the rebuild operation starts, you will see a screen similar to Figure 73 where you will see the rebuild progress in percentage. The system can operate normally as long as both the disks are not degraded simultaneously.

Figure 73: Rebuilding disk in progress

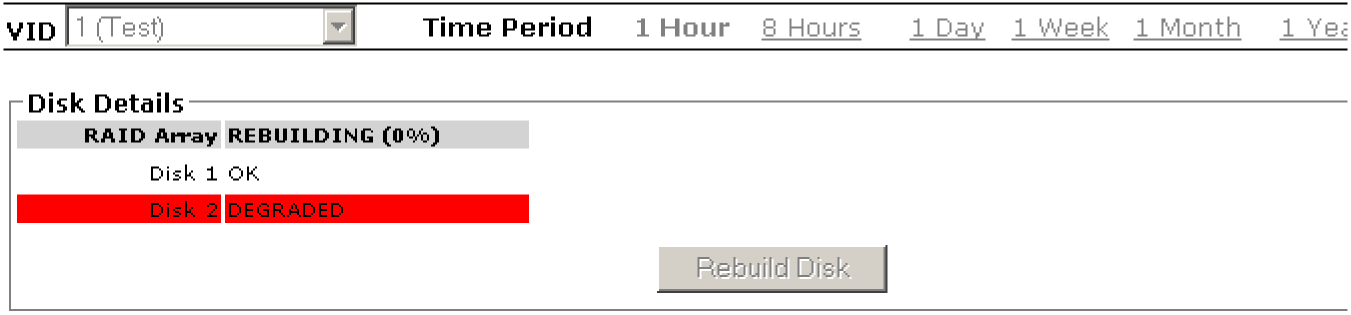

VID 1 (Test) Time Period 1 Hour 8 Hours 1 Day 1 Week 1 Month 1 Yea

Disk Details

RAID Array	REBUILDING (0%)
Disk 1	OK
Disk 2	DEGRADED

At the end of the rebuilding operation, the disks will become *OK* as shown in Figure 74.

Figure 74: Disk status after the rebuild operation is complete

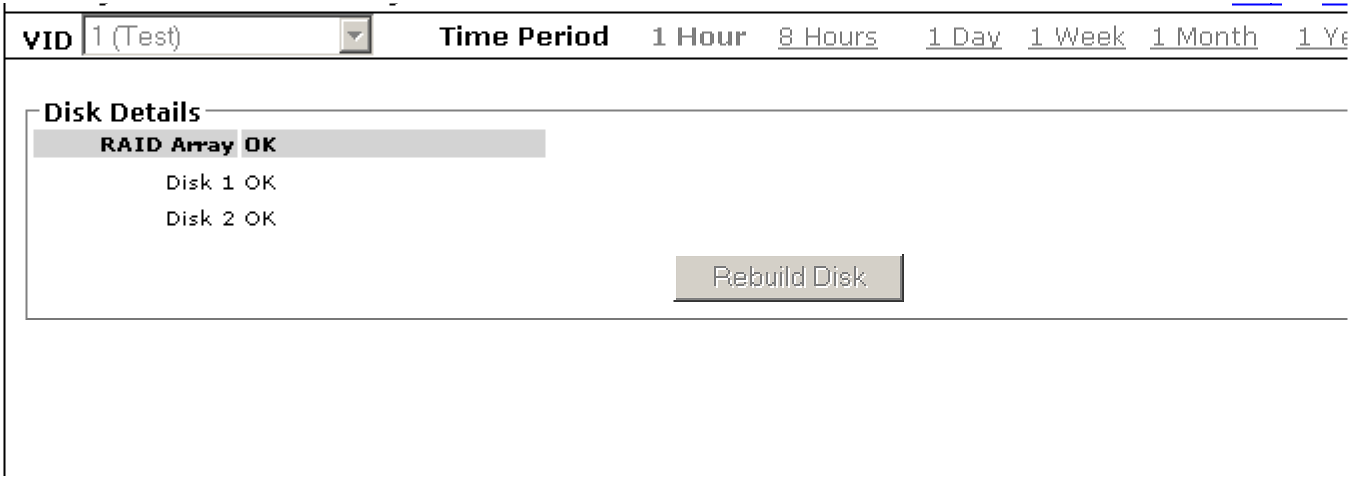

VID 1 (Test)

Disk Details

RAID Array OK

Disk 1 OK

Disk 2 OK

Displaying login, access and audit trails

FortiDDoS keeps information about following user actions in a database table that can be reviewed by an administrator with privileges:

- Login to the Web-based Manager- including failed attempts - known as Login Trail.
- Access to Web-based Manager menu items - known as Access Trail
- Changes done to the configuration - known as Audit Trail

Login trail Whenever a user logins to the Web-based Manager, the following information is recorded in a database table:

- Time of the attempt
- IP address from where logged in
- User name used for login purpose
- Status - whether success or failure.

You can view the Login Trail using *Show* > *Global* > *Reports* > *Login Trail*. Figure 75 shows tan example login trail report.

Access trail Whenever a user accesses a menu item in the Web-based Manager, the following information is recorded in a database table:

- Time of the access
- User name
- IP address from where logged in
- Page accessed
- Any relevant parameters that further give details of the access.

You can view the Login Trail using *Show* > *Global* > *Reports* > *Access Trail*.

Audit trail Whenever a user changes the configuration of the appliance through Web-based Manager, the following information is recorded in a database table:

- Time of the access
- User name
- IP address from where logged in
- Page accessed
- VID
- Action: showing relevant parameters that further give details of the change.

You can view the Audit Trail using *Show* > *Global* > *Reports* > *Audit Trail*.

Please note that all the details of the configuration are not recorded. In some cases, these changes are substantial and may not fit in the screen.

Figure 75: Login trail example report

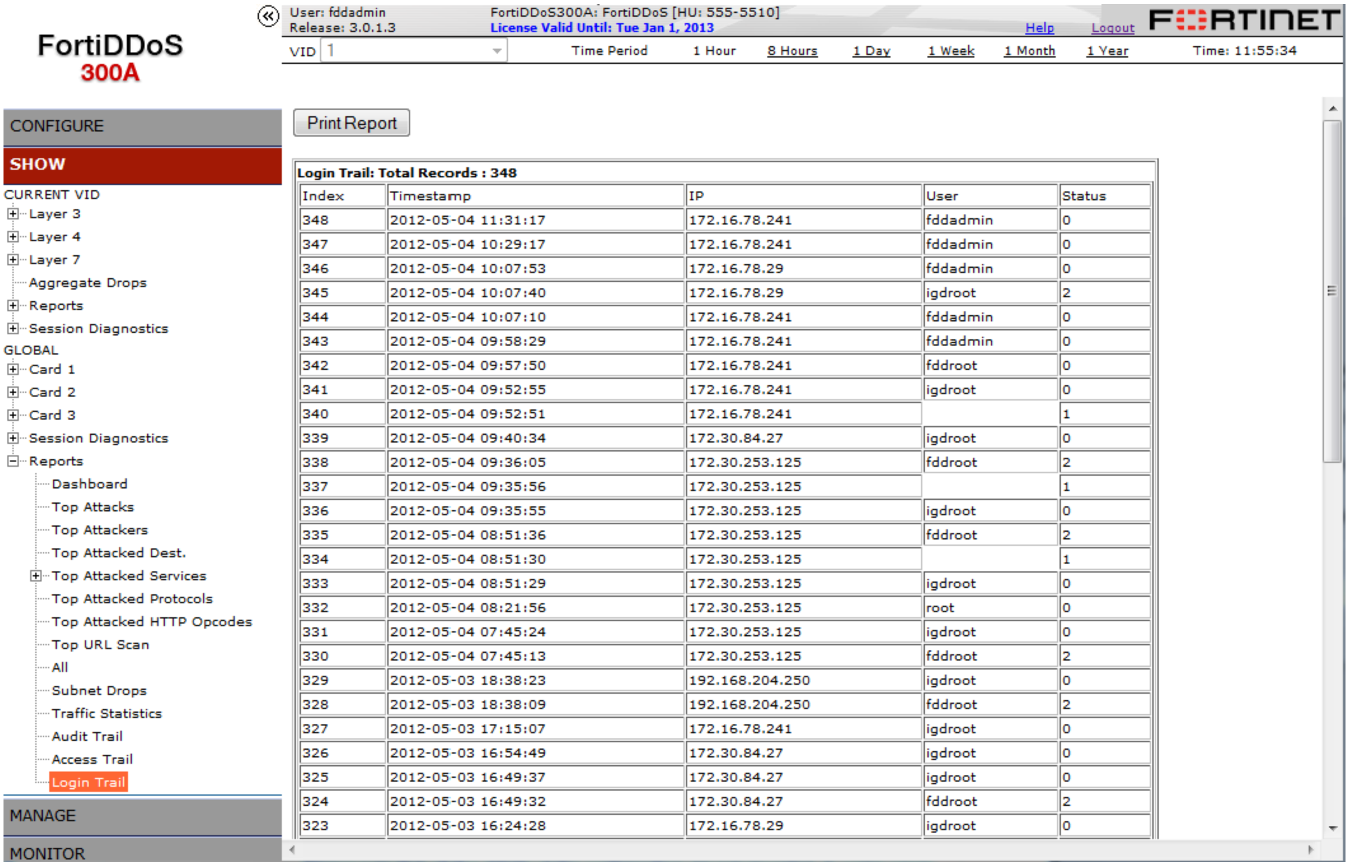

User: fddadmin
Release: 3.0.1.3

FortiDDoS300A: FortiDDoS [HU: 555-5510]
License Valid Until: Tue Jan 1, 2013

VID 1

Time: 11:55:34

Login Trail: Total Records : 348

Index	Timestamp	IP	User	Status
348	2012-05-04 11:31:17	172.16.78.241	fddadmin	0
347	2012-05-04 10:29:17	172.16.78.241	fddadmin	0
346	2012-05-04 10:07:53	172.16.78.29	fddadmin	0
345	2012-05-04 10:07:40	172.16.78.29	igdroot	2
344	2012-05-04 10:07:10	172.16.78.241	fddadmin	0
343	2012-05-04 09:58:29	172.16.78.241	fddadmin	0
342	2012-05-04 09:57:50	172.16.78.241	fddroot	0
341	2012-05-04 09:52:55	172.16.78.241	igdroot	0
340	2012-05-04 09:52:51	172.16.78.241		1
339	2012-05-04 09:40:34	172.30.84.27	igdroot	0
338	2012-05-04 09:36:05	172.30.253.125	fddroot	2
337	2012-05-04 09:35:56	172.30.253.125		1
336	2012-05-04 09:35:55	172.30.253.125	igdroot	0
335	2012-05-04 08:51:36	172.30.253.125	fddroot	2
334	2012-05-04 08:51:30	172.30.253.125		1
333	2012-05-04 08:51:29	172.30.253.125	igdroot	0
332	2012-05-04 08:21:56	172.30.253.125	root	0
331	2012-05-04 07:45:24	172.30.253.125	igdroot	0
330	2012-05-04 07:45:13	172.30.253.125	fddroot	2
329	2012-05-03 18:38:23	192.168.204.250	igdroot	0
328	2012-05-03 18:38:09	192.168.204.250	fddroot	2
327	2012-05-03 17:15:07	172.16.78.241	igdroot	0
326	2012-05-03 16:54:49	172.30.84.27	igdroot	0
325	2012-05-03 16:49:37	172.30.84.27	igdroot	0
324	2012-05-03 16:49:32	172.30.84.27	fddroot	2
323	2012-05-03 16:24:28	172.16.78.29	igdroot	0

FORTINET®