FORTINET®

FortiDDoS v3.2
Installation Guide

April 1, 2013

28-320-183686-20130401

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	docs.fortinet.com
Knowledge Base	kb.fortinet.com
Customer Service & Support	support.fortinet.com
Training Services	training.fortinet.com
FortiGuard	www.fortiguard.com
Document Feedback	techdocs@fortinet.com

Table of Contents

Introduction

Scope of this document

This document gives the details in installing and configuring FortiDDoS devices.

This document covers:

- Package contents
- Simple deployment overview
- Installing the physical system
- Setting up network properties
- Configuring interface settings
- Checking system status
- Configuring the operating mode
- Assigning Virtual Identifiers (VIDs) to protect systems
- Performing a sanity test
- Monitoring events
- Showing traffic
- Using bypass switches for fail-over
- Using traffic diversion in service provider environment
- Using load balancing to support higher bandwidth in service provider environment
- Using FortiGuard IP Reputation Service

Introduction

This document explains the tasks required to initially install a FortiDDoS device in a network. We assume that you have already read the *FortiDDoS Fundamentals Guide*, and are familiar with the fundamental concepts related to FortiDDoS devices. This document explains package contents, system overview, selecting a mode of operation, the physical installation, how to change the IP address for the management port, and how to assign Virtual Identifiers to protect specific systems. It also shows you how to perform a ping test, and an overview of how to monitor events and traffic.

Package contents

Before you begin, please be sure the following items are in the package. If the package is not complete, contact your supplier.

- FortiDDoS device
- Power cord
- Documentation CD
- Brackets to mount the chasis on a rack

Note: Please retain the carton, including the original packing materials in case there is a need to return the product.

Simple deployment overview

The simple configuration and location of interfaces on the FortiDDoS devices are described below.

Physical interfaces

Refer to Figure 1 on page 7 for physical interfaces on FortiDDoS devices.

Traffic Processing (TP) Boards

FDD-100A contains one Traffic Processing Board.

FDD-200A contains two Traffic Processing Boards.

FDD-300A contains three Traffic Processing Boards.

Data ports on each TP Board

There are two pairs of Ethernet ports located on the back panel of the FortiDDoS device. There are copper and SFP ports. At a given time, you can use either copper or fiber for a link.

For the FDD-100A, ports are marked LAN 1, WAN 1, LAN 2 and WAN 2.

The FDD-200A has additional ports that are marked LAN 3, WAN 3, LAN 4 and WAN 4.

The FDD-300A has additional ports that are marked LAN 5, WAN 5, LAN 6 and WAN 6.

USB keyboard port

Use of the keyboard port is optional and is to be used during diagnostics on the console.

Serial Interface through USB port

A serial console can be connected using a USB to serial adapter. The console can be used for Command Line Interface (CLI) access for advanced usage.

Monitor port

Use of monitor port is optional and is to be used during diagnostics on the console. This is a DVI port. If you have VGA monitor connector, you will need a DVI to VGA adapter for connecting your monitor.

Ethernet port for management interface

The management of FortiDDoS devices is normally done over IP over Ethernet. The Gigabit Ethernet Port can be connected to a private or public network. The device Web-based Manager can then be accessed over SSL using the connection. For diagnostics, the CLI can be accessed over secure shell (ssh) over the same network.

Figure 1: Back panel of a FortiDDoS 100A device with copper and fiber interfaces and the management Interfaces

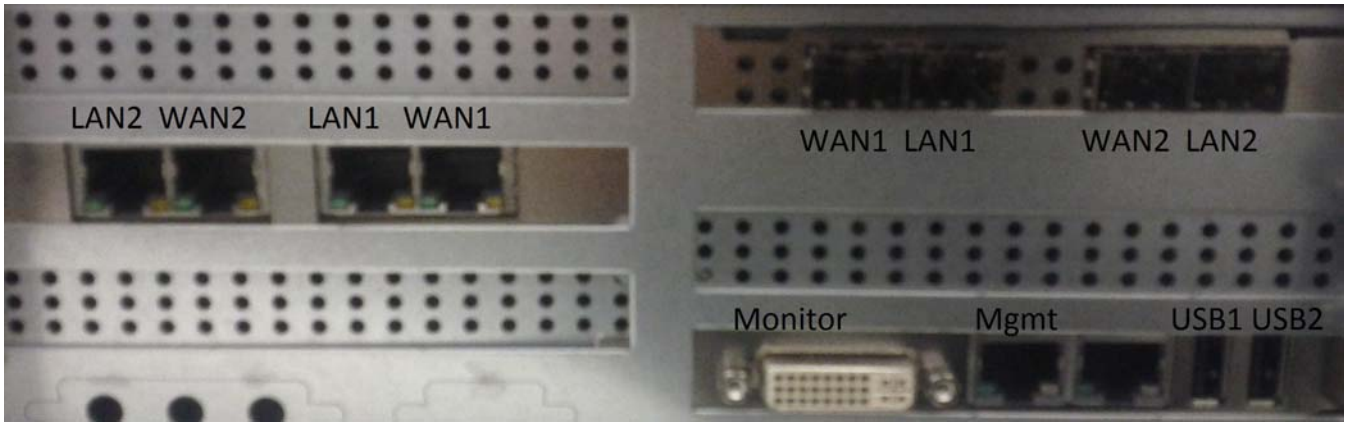

Simple deployment

The FortiDDoS device is designed to protect a system or a network of systems from rate-based attacks and anomaly attacks. If multiple systems or workgroups are protected by a FortiDDoS device, a switch will be required between the FortiDDoS appliance and the protected systems.

Figure 2: A simple network prior to installation of a FortiDDoS device

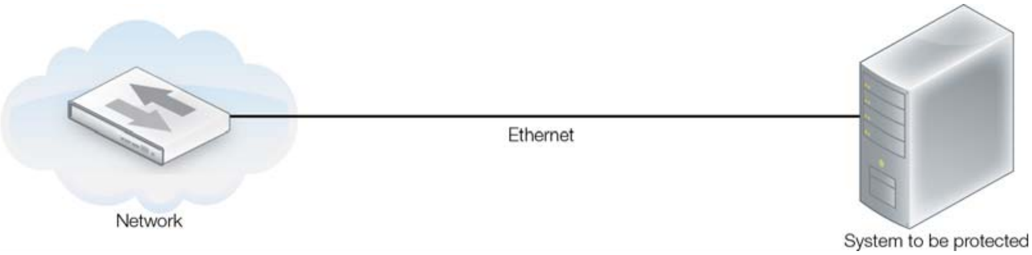

In a simple network shown in Figure 2, a system is connected to an Ethernet local area network.

In the simplest configuration, you can install a FortiDDoS unit as an inline device, as shown in Figure 3.

Figure 3: Network with a FortiDDoS device protecting a single system

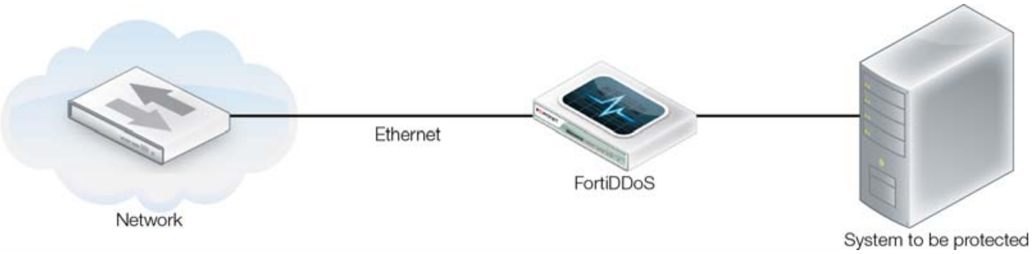

The appliance is stateful and bidirectional, so a concept of ‘direction’ must be introduced to differentiate between inbound and outbound traffic.

As shown in Figure 4 below, in a typical installation, the Ethernet segment connected to the protected systems is connected to LAN 1. The Ethernet segment connected to the rest of the network (typically the Internet) is connected to WAN 1.

Figure 4: Recommended directionality of FortiDDoS devices

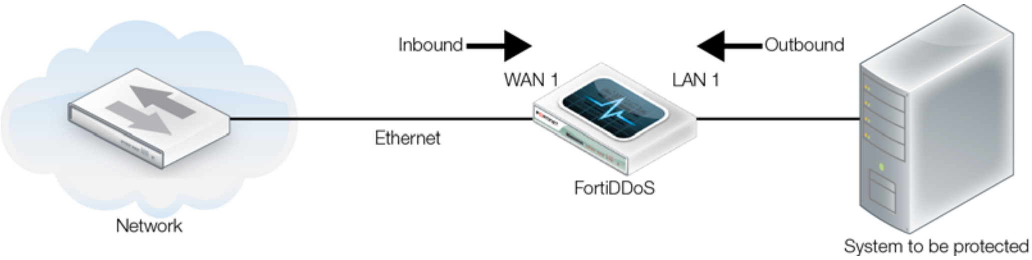

Basic web hosting deployment

More complex setups can protect multiple systems. In a basic web hosting deployment a FortiDDoS device can protect systems in multiple customer cages as shown in Figure 5. You can either use a single VID system or a multiple VID system. Please refer to the *FortiDDoS Fundamentals Guide* for concepts related to VID and the *FortiDDoS Web-Based Manager Guide* for the actual configuration of VIDs.

Figure 5: Basic web hosting deployment of FortiDDoS devices

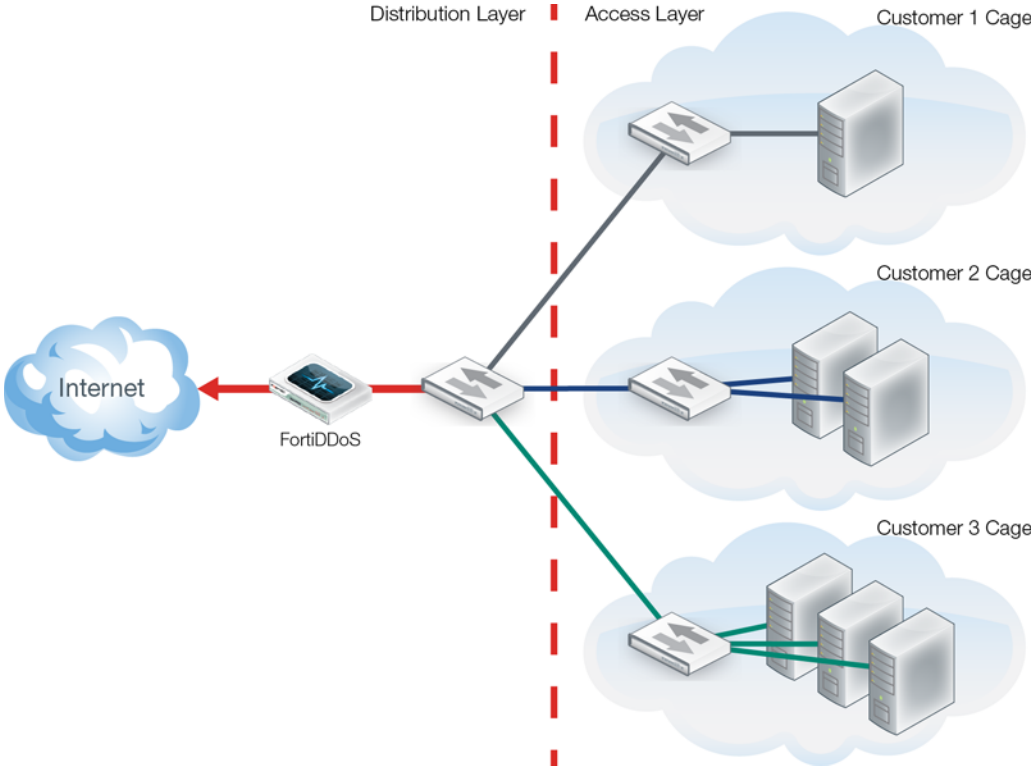

Managed hosting deployment with high availability

Figure 6 shows another setup protecting multiple systems in a data center environment.

In this case two FortiDDoS devices independently protect the routers and the subsequent networks from DoS and DDoS attacks.

Figure 6: Managed hosting deployment with high availability

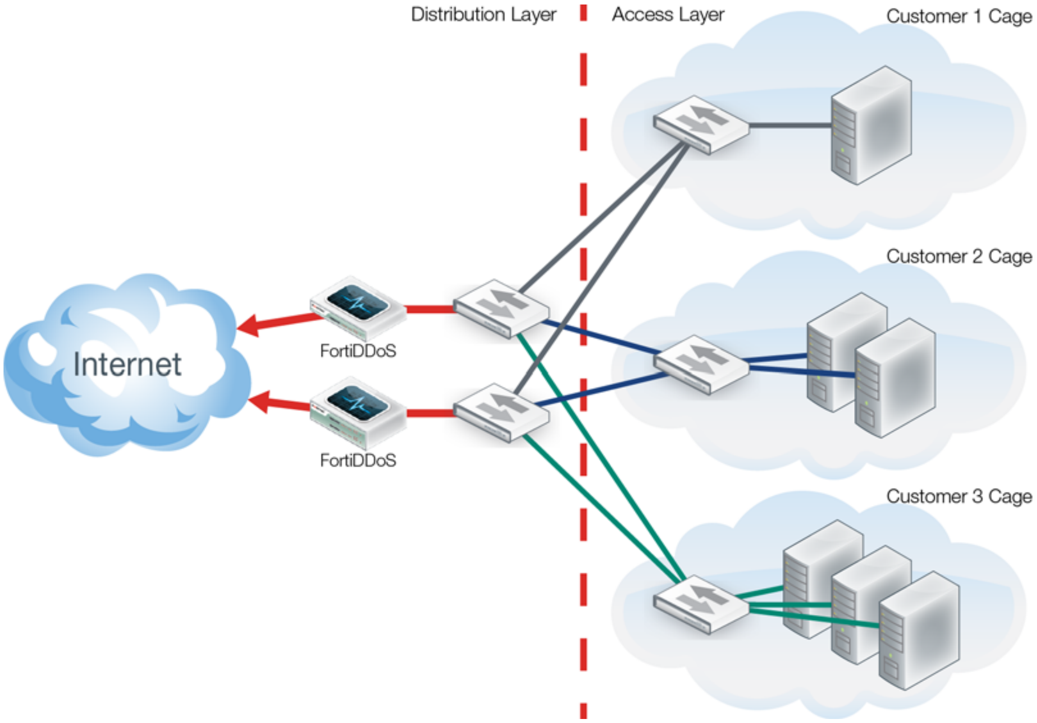

Installation & Initial Configuration

There are nine main steps to install and configure FortiDDoS devices. They are:

1. Installing the physical system
2. Setting up network properties
3. Configuring interface settings
4. Checking system status
5. Configuring the operating mode
6. Assigning Virtual Identifiers (VIDs) to protect systems.
7. Performing a sanity test
8. Monitoring events
9. Showing traffic

Installing the physical system

Follow these steps to install the system:

Connecting the power cord

1. Take the FortiDDoS device out of the box and make sure the power switch is off.
2. Connect one end of the power cord to an appropriate 110/220 outlet and the other end to the appliance itself.

Caution: The appliance must be switched off for at least 30 seconds before restarting.

Connecting the management ports

To manage the FortiDDoS device via a web browser:

1. Connect the 10/100 ethernet port to a workgroup switch/router or use a crossover Ethernet cable to a computer with an HTML web browser. The IP address of the management port is preset to 192.168.1.1.
2. You must first access the FortiDDoS device using this IP address, but you may change it by clicking *Manage* > *Global* > *Device Configuration* > *IP Address* after you connect.

Setting up network properties

To set the network properties:

1. From a workstation or PC, access the graphical user interface on the FortiDDoS device over the management Ethernet by using the default address `https://192.168.1.1` as the URL.

Note: You must use https when entering the address. The system will not respond to access requests using http (without the ‘s’).

2 Log in using the default user ID *fddroot* and the default password *rootpasswd*.

3 You can change the IP address to one that is appropriate for your domain using the *Manage* > *Global* > *Device Configuration* > *IP address* menu. The DNS and gateway settings are used to send E-mail summaries of events. E-mail cannot be sent until valid addresses are configured for these fields.

4 The host name is used to logically name the FortiDDoS system for easy reference.

Following table contains the default IP addresses and name assignments of your FortiDDoS device.

Table 1: Default IP Addresses & Hostname

IP Address for FortiDDoS device	192.168.1.1
Netmask	255.255.255.0
DNS 1 Address	UNDEFINED
DNS 2 Address	UNDEFINED
Gateway Address	UNDEFINED
Hostname	A unique character string assigned by the factory

Configuring interface settings

Every network is different and the interfaces to which the main ports and auxiliary ports of the FortiDDoS device are connected have to be described clearly to the device so that it can communicate with the networks without any errors.

You must know the network settings before installing FortiDDoS device. The existing switches/routers/firewalls have their ports set to certain speed, duplexity, and flow control mode.

With those settings in mind, you must set the values in the *Configure*> *Global*> *Card 1* > *Interface Settings Menu* shown in Figure 7.

1 Configure LAN 1 to Copper or Fiber and then to Auto or Forced depending on the port connected to LAN 1. If you set it to Forced, configure the speed, the duplex value and the flow control value. Some of the settings will be enabled or disabled depending on the Interface or Mode.

2 Repeat step 1 for WAN 1, LAN 2 and WAN 2.

3 Repeat the same steps for other cards in the appliance.

Figure 7: Configuring interface settings for ports

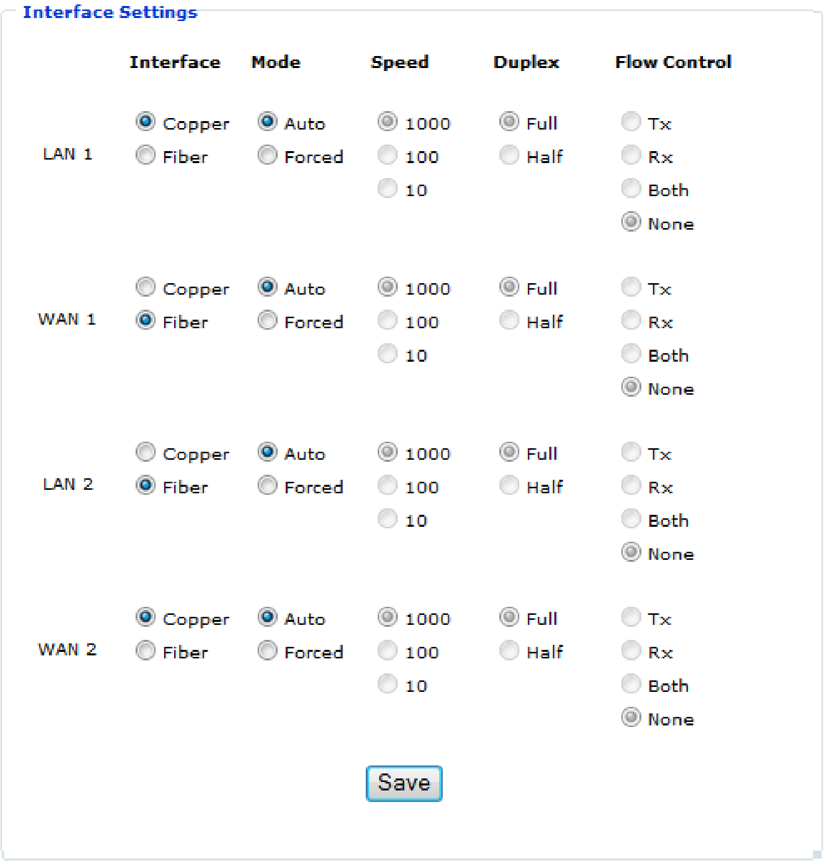

Interface Settings

	Interface	Mode	Speed	Duplex	Flow Control
LAN 1	● Copper ○ Fiber	● Auto ○ Forced	● 1000 ○ 100 ○ 10	● Full ○ Half	○ Tx ○ Rx ○ Both ● None
WAN 1	○ Copper ● Fiber	● Auto ○ Forced	● 1000 ○ 100 ○ 10	● Full ○ Half	○ Tx ○ Rx ○ Both ● None
LAN 2	○ Copper ● Fiber	● Auto ○ Forced	● 1000 ○ 100 ○ 10	● Full ○ Half	○ Tx ○ Rx ○ Both ● None
WAN 2	● Copper ○ Fiber	● Auto ○ Forced	● 1000 ○ 100 ○ 10	● Full ○ Half	○ Tx ○ Rx ○ Both ● None

Checking system status

The System Status page (accessible from the main menu), shown in Figure 8 and Figure 9 show the overall health of the system.

These pages shows the port status. You must make sure that the Configured, Actual and Link Partner Ability correspond correctly to your network and expectation.

These pages will also tell you if the *sendmail* service is operational. This service can be used in conjunction with the Event Monitor to notify you (or other email recipients) of system events. This can be configured under the *Configure* > *Current VID* > *Event Notification* menu.

For the FortiDDoS device to send a mail message, it must be able to contact a Domain Name Server (DNS) to resolve the domain name of the email addresses. The status page will indicate whether the system is able to reach a DNS server.

In case you are having trouble establishing connectivity, you must carefully study the values for configured, actual and link partner ability.

Note: In case of forced settings, the link partner abilities must be ignored.

Figure 8: Status page for FortiDDoS devices with copper connections - Part 1

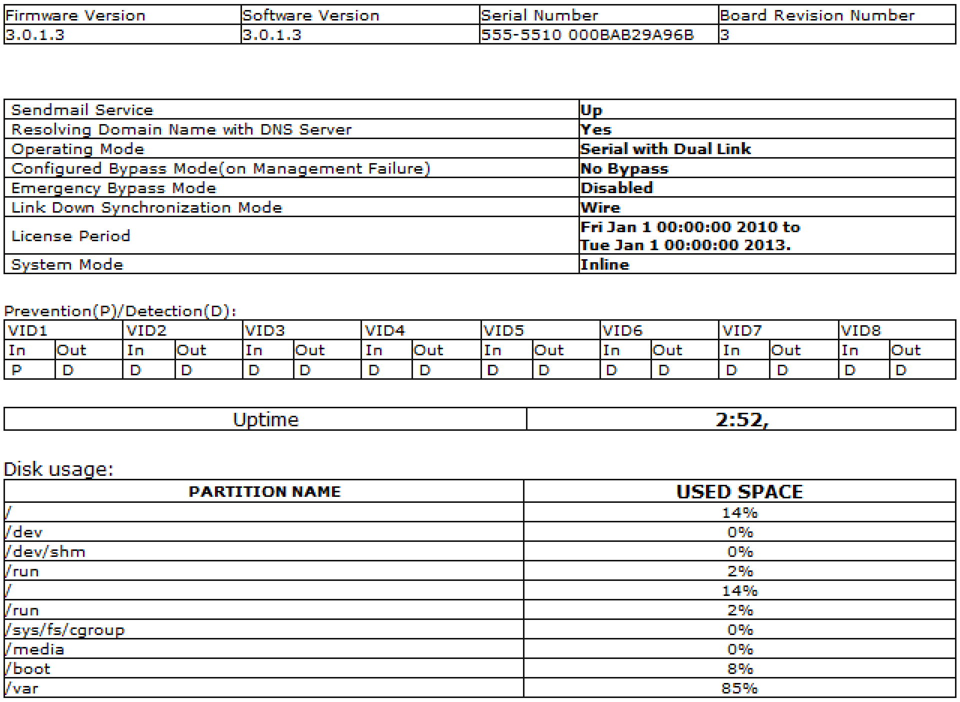

Firmware Version	Software Version	Serial Number	Board Revision Number
3.0.1.3	3.0.1.3	555-5510 000BAB29A96B	3

Sendmail Service	**Up**
Resolving Domain Name with DNS Server	**Yes**
Operating Mode	**Serial with Dual Link**
Configured Bypass Mode(on Management Failure)	**No Bypass**
Emergency Bypass Mode	**Disabled**
Link Down Synchronization Mode	**Wire**
License Period	**Fri Jan 1 00:00:00 2010 to** **Tue Jan 1 00:00:00 2013.**
System Mode	**Inline**

Prevention(P)/Detection(D):

VID1		VID2		VID3		VID4		VID5		VID6		VID7		VID8	
In	Out	In	Out	In	Out	In	Out	In	Out	In	Out	In	Out	In	Out
P	D	D	D	D	D	D	D	D	D	D	D	D	D	D	D

Uptime	**2:52,**

Disk usage:

PARTITION NAME	USED SPACE
/	14%
/dev	0%
/dev/shm	0%
/run	2%
/	14%
/run	2%
/sys/fs/cgroup	0%
/media	0%
/boot	8%
/var	85%

Figure 9: Status page for FortiDDoS devices with copper connections - Part 2

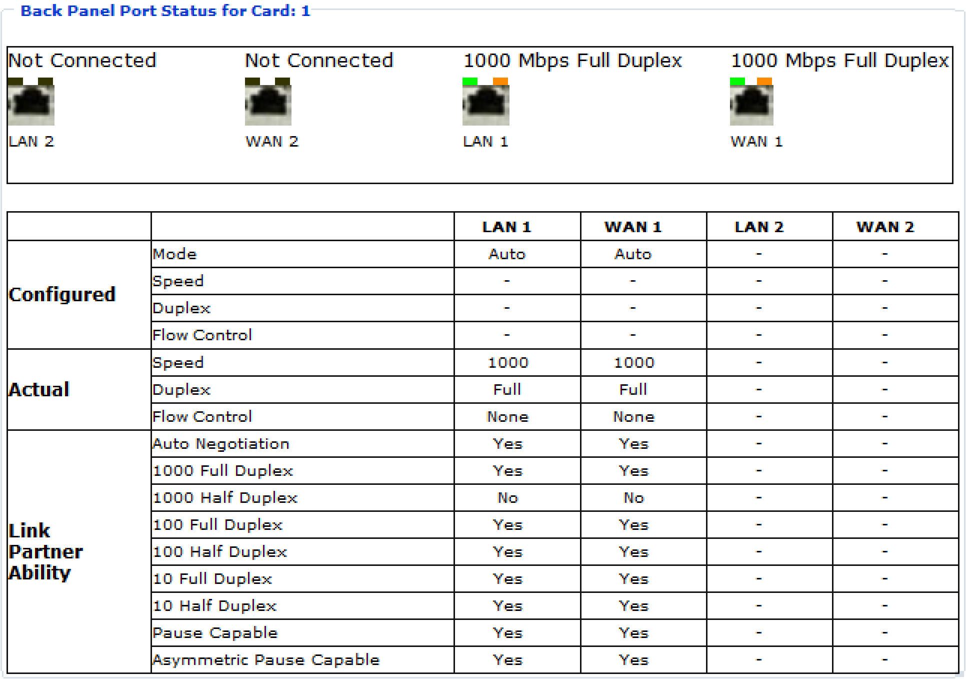

Back Panel Port Status for Card: 1

Not Connected	Not Connected	1000 Mbps Full Duplex	1000 Mbps Full Duplex
LAN 2	WAN 2	LAN 1	WAN 1

		LAN 1	WAN 1	LAN 2	WAN 2
Configured	Mode	Auto	Auto	-	-
	Speed	-	-	-	-
	Duplex	-	-	-	-
	Flow Control	-	-	-	-
Actual	Speed	1000	1000	-	-
	Duplex	Full	Full	-	-
	Flow Control	None	None	-	-
Link Partner Ability	Auto Negotiation	Yes	Yes	-	-
	1000 Full Duplex	Yes	Yes	-	-
	1000 Half Duplex	No	No	-	-
	100 Full Duplex	Yes	Yes	-	-
	100 Half Duplex	Yes	Yes	-	-
	10 Full Duplex	Yes	Yes	-	-
	10 Half Duplex	Yes	Yes	-	-
	Pause Capable	Yes	Yes	-	-
	Asymmetric Pause Capable	Yes	Yes	-	-

Configuring the operating mode

Serial mode

Serial mode is the default mode of operation. In the Default mode, the FortiDDoS device is positioned ‘inline’, meaning it is between the protected system(s) and the rest of the network. Figure 10 shows this.

Direction-based VID-based detection mode

Detection mode for a set of chosen VIDs in a specific direction is a mode in which the appliance does not perform any blocking of data, but it does log events and build traffic profiles. Data passes through the FortiDDoS device as it travels to and from the protected system(s) and the rest of the network. After a sufficient learning period of 2-14 days, the FortiDDoS device should be placed inline (in Prevention mode).

Figure 10: Logical network configuration for Detection Mode

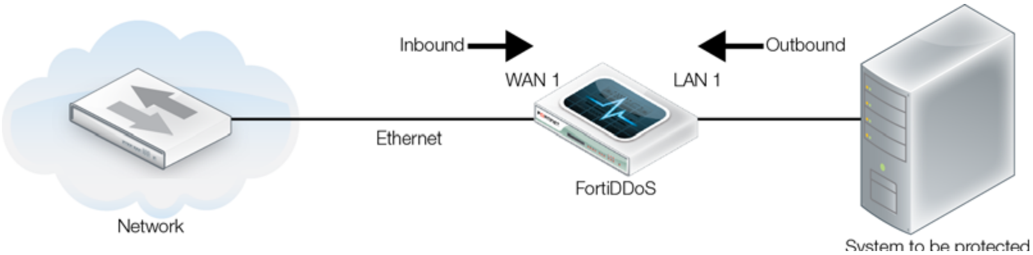

FortiDDoS devices can be simply placed in series (in-line) or can be placed in tandem with a bypass switch to avoid failures.

Fortinet recommends FortiBridge bypass switches for failover protection. For other bypass switches available in the market, please contact your Sales Engineer to check if it is qualified to work with FortiDDoS appliances. Refer to “Configuration Options” on page 22.

Direction-Based VID-Based Prevention Mode

Prevention Mode for a set of chosen VIDs in a specific direction is the full-function operating mode of the FortiDDoS device. Place the unit inline between the protected system(s) and the rest of the network. Any anomalous traffic or traffic that exceeds threshold values is blocked. You can configure the unit to send any blocked traffic out the corresponding auxiliary ports to a forensic capture device for further analysis.

Note: For connecting and configuring the bypass switches, the procedure remains same as the Detection Mode. Please refer to the sections above.

Configuring additional modes

To set the function of the auxiliary ports Click *Configure* > *Global* > *Operating Mode*.

1 **Asymmetric Pair - Internal**: LAN 2 and WAN 2 will be connected in such a way that asymmetric traffic between two networks can be combined. This mode is useful in case you want to connect 2 FortiDDoS devices in an asymmetric network. Traffic from 2 uplinks is combined in both the FortiDDoS device using copies on auxiliary ports.

Figure 11: Ports

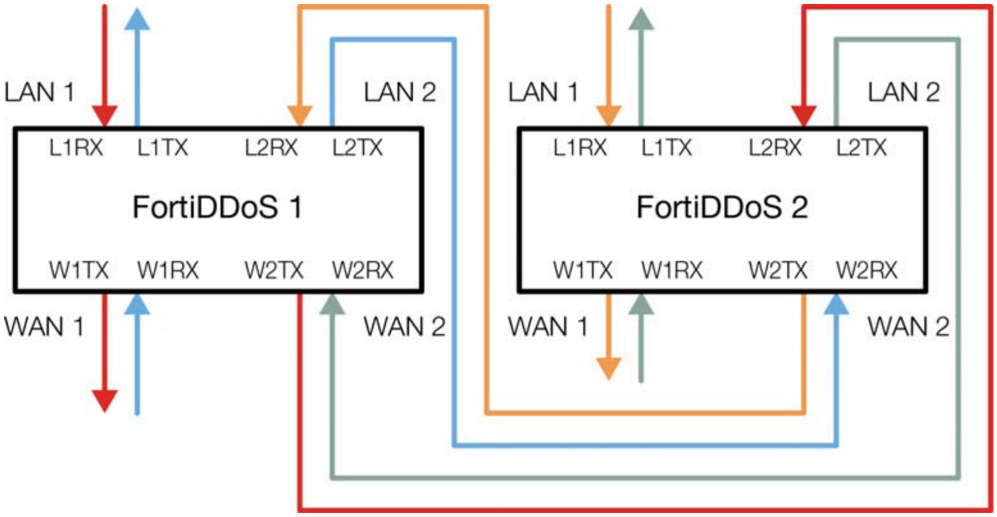

2 **Asymmetric Pair - External**: Connect LAN 2 and WAN 2 to external ports that copy the traffic for the other link. The task of copying is entrusted to an external source. Traffic from 2 uplinks is combined in both the FortiDDoS device using copies available on the auxiliary ports.

3 **Default Mode**: Connect LAN 2 to internal network and WAN 2 to the second Internet link. This mode is useful in case you want to connect 1 FortiDDoS device in an asymmetric network or a network having two Internet links. Traffic from 2 links is combined internally in the device. However, at the egress port, the traffic corresponds to the corresponding link. E.g. WAN 1 receives traffic from LAN 1 and vice versa. Similarly LAN 2 receives traffic from WAN 2 and versa.

Configuring prevention or detection mode for a set of VIDs in a specific direction

To set the Prevention/Detection Mode of a set of VIDs, click *Configure* > *Global* > *Operating Mode*. Please refer to Figure 11 above.

In *Prevention/Detection Mode section***,** click the VIDs you want in Prevention Mode and leave the VIDs unchecked if you want them in Detection Mode. You can choose the modes in Inbound or Outbound or both directions.

Click *Save*.

Configuring bypass mode

Bypass is relevant in case of appliance management path failure. It is assumed that the data path failures are handled separately - in some cases using an external bypass switch.

In case of Management Path failure, the user can choose one of the following:

1 Extrinsic Bypass

2 Intrinsic Bypass

3 No Bypass

Choose Extrinsic Bypass in case you want the external bypass switch to be enabled - implies no prevention.

Choose Intrinsic Bypass in case you want the intrinsic bypass to be triggered - in case you do not have external bypass switches and also implies no prevention.

Choose No Bypass in case you want the existing mode to continue without updated thresholds - implies no continuous learning and adaptive prevention/detection.

To set the Bypass Mode of the appliance, click *Configure* > *Global* > *Operating Mode*. Please refer to Figure 11 above.

In *Bypass Mode section***,** select one of the above bypass modes.

Click *Save*.

Configuring emergency bypass mode

At certain times, to eliminate the possibility of malfunction of the FortiDDoS device, you may want to bypass the device logic while keeping the device inline. To achieve such a functionality, you can keep the appliance in Emergency Bypass Mode. This ensures that the packets which arrive at ingress ports are simply transferred to the corresponding egress ports - just like a wire.

To set the Emergency Bypass Mode of the appliance, click *Configure* > *Global* > *Operating Mode*. Please refer to Figure 11 above.

In *Emergency Bypass Mode* section**,** click on the checkbox for *Emergency Bypass*.

Click *Save*.

Configuring link down synchronization or link state propagation

Link Down Synchronization lets you configure FortiDDoS device to force the partner link down on a segment when one of the links goes down. The device monitors the link state for a pair of ports which are protecting a segment. These correspond to LAN 1 (connected to LAN) or WAN 1 (connected to the Internet). Similarly for Dual WAN Link mode, these ports correspond to LAN 2 and WAN 2.

If the link goes down on either port, the partner port is disabled. Link Down Synchronization once enabled, propagates the link state across the FortiDDoS device. This is the default functionality. If you want to disable this functionality, you must select Hub mode.

This feature is not useful when using bypass switches and must be set to HUB mode instead of default WIRE mode.

To enable Link Down Synchronization, you don’t have to make any changes. It is set as the factory default.

To set the Link Down Synchronization to Hub Mode, click *Configure* > *Global* > *Link Down Synchronization*.

In *Link Down Synchronization* section**,** click on the radio button for *Hub*.

Type `yes` in the text box and press *OK*.

Note: Changes to Link Down Synchronization requires restarting the services - which leads to some downtime. Please plan for the downtime.

Assigning Virtual Identifiers (VIDs) to protect systems

Virtual Identifiers (VIDs) enable you to “virtualize” the device to behave as if it were multiple physical appliances with each appliance conforming to a single server/network.

Because each networked system has different traffic characteristics, the FortiDDoS device allows you to build a unique profile for each server/network you want to protect.

These servers/networks may be specific to individual departments, or may be used for different applications such as web hosting, E-mail or DNS queries.

Your FortiDDoS device is pre-configured to accommodate either 1, 4, or 8 VIDs. Each VID can be measured against a unique set of threshold parameters that are independent of other VIDs. This is shown in Figure 12 Not all VIDs need to be configured; you may leave blank those that you do not need.

Figure 12: Network with FortiDDoS protecting multiple VIDs

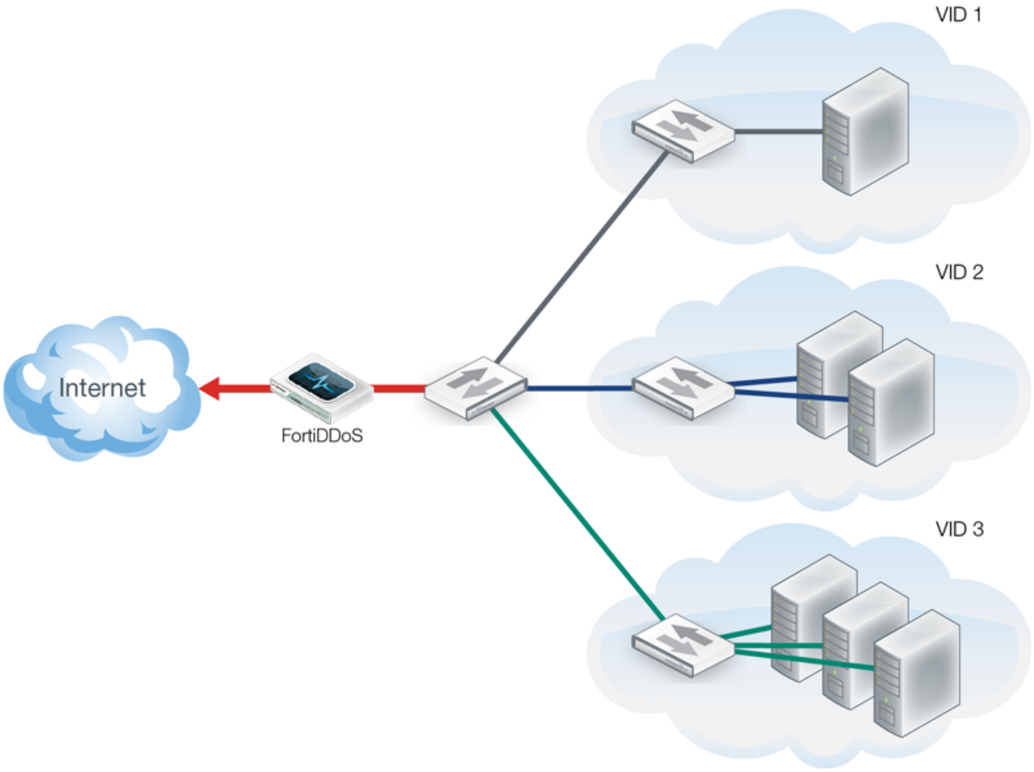

Note: It is recommended that you use a single network switch between the FortiDDoS device and protected systems. The goal is to avoid inserting any potential source of attack traffic that does not pass through the device.

Configuring VIDs

To configure a VID:

1 From the main menu, click *Configure* > *Global* > *VIDs*

2 Simply enter the following information:

- Subnet ID

  This ID is used to for subnet-based reporting. Administrator can generate attack event report for individual subnets.

  You can enter up to 512 subnets. Please refer to the datasheet of your appliance.

- IP Address

  This corresponds to the IP address of the subnet you want to add to a VID.

- Netmask

  This corresponds to the Netmask of the subnet you want to add to a VID.

- VID Number

  This corresponds to the actual VID number.
- Alternate VID Number

  This corresponds to the alternate VID number. An alternate VID is a VID where the subnet is transferred once the traffic to the subnet exceeds the Threshold (below)
- Threshold

  The threshold corresponds to the packet rate beyond which the subnet is moved to Alternate VID number. If the traffic goes below this threshold for a preconfigured time period, it goes back to the (original) VID Number. The timeout is defined in the same screen in a different fieldset. This feature can be enabled or disabled by switching the check-box *Allow VID Switching based on thresholds* to off or on.
- Comment

For a detailed description of VID configuration, please refer to the *Web-based Manager Administration Guide*.

Performing a sanity test

The following steps can serve as a simple demonstration of how FortiDDoS devices block traffic. To run the demo, the network configuration should be in **serial prevention** mode as shown in Figure 13. The protected server should respond to ICMP Echo (ping) packets, and a connected system upstream must be capable of generating a series of ICMP Echo Request packets.

Figure 13: Ping test configuration

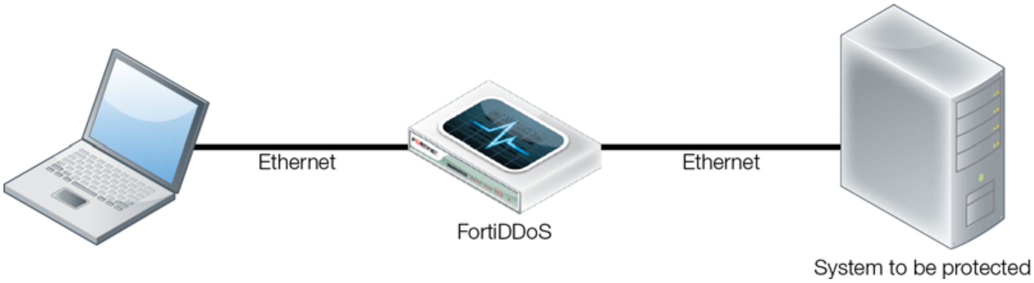

Steps for performing a ping test

1 Configure the FortiDDoS device threshold for ping to 5 packets per second.

To do this, click *Configure* > *Current VID* > *Blocking Threshold* > *Layer 3* > *Protocols* from the main menu. You set the ICMP threshold here because it is in the Layer 3 packet that you determine the type of protocol to use. In this case, you will set a threshold for ICMP packets, which corresponds to protocol number 1. Refer to Figure 14.

Set your inbound and outbound thresholds for ICMP to 5 packets per second.

Figure 14: Blocking Conditions for ICMP for Ping Test

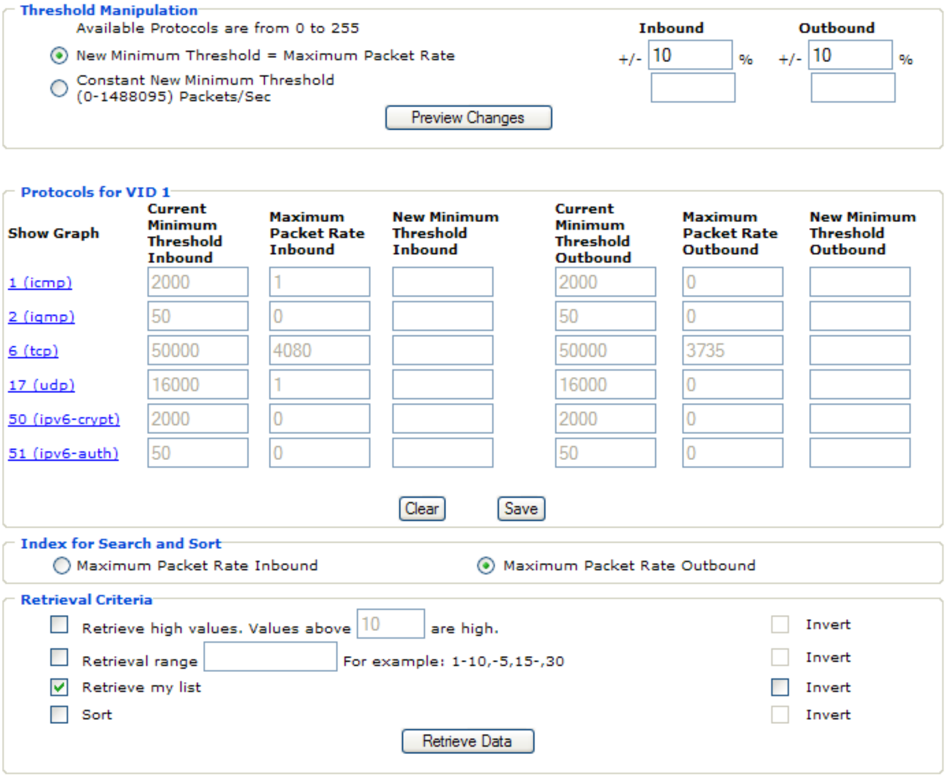

Threshold Manipulation

Available Protocols are from 0 to 255

● New Minimum Threshold = Maximum Packet Rate

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

Inbound +/- 10 %

Outbound +/- 10 %

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)	2000	1		2000	0	
2 (igmp)	50	0		50	0	
6 (tcp)	50000	4080		50000	3735	
17 (udp)	16000	1		16000	0	
50 (ipv6-crypt)	2000	0		2000	0	
51 (ipv6-auth)	50	0		50	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

You must click **Save** in the screen above the Layer 3 Classifier table to record your settings in the system. The system may pause for a few seconds before confirming that the new values have been updated.

2 Generating ICMP (ping) traffic

From the PC/Workstation, generate a small, controlled flood of 100 ICMP Echo (ping) packets directed to the protected system. In UNIX/LINUX, the command line input will look like this:

> `ping -c 100 -i 0.1 AA.BB.CC.DD` (where `AA.BB.CC.DD` represents the IP address of the protected system)
>
> The command above will generate an ICMP Echo Request (ping) packet to the specified address every 0.1 seconds until 100 packets are sent. This is the equivalent of 10 packets per second for 10 seconds.

Following is a screen capture from an actual ping flood test. Notice that the first few pings are allowed to pass and receive a response. As soon as the rate per second rises above the threshold, (somewhere in the first 11 packets) the FortiDDoS device blocks all ICMP packets for the 10 second threshold. After the blocking period, ICMP packets are again allowed until the threshold is reached.

In the sequence below, this is reflected by responses to the first 7 ping requests, followed by no response to the next 80 packets (blocked by the appliance). Then packets 87-93 are allowed before the threshold is again reached.

```
[root@client1 win]# ping -c 100 -i 0.1 172.16.0.50
PING 172.16.0.50 (172.16.0.50) 56(84) bytes of data.
64 bytes from 172.16.0.50: icmp_seq=1 ttl=64 time=0.503 ms
64 bytes from 172.16.0.50: icmp_seq=2 ttl=64 time=0.307 ms
```

```
64 bytes from 172.16.0.50: icmp_seq=3 ttl=64 time=0.220 ms
64 bytes from 172.16.0.50: icmp_seq=4 ttl=64 time=0.314 ms
64 bytes from 172.16.0.50: icmp_seq=5 ttl=64 time=0.260 ms
64 bytes from 172.16.0.50: icmp_seq=6 ttl=64 time=0.281 ms
64 bytes from 172.16.0.50: icmp_seq=7 ttl=64 time=0.206 ms

64 bytes from 172.16.0.50: icmp_seq=87 ttl=64 time=0.275 ms
64 bytes from 172.16.0.50: icmp_seq=88 ttl=64 time=0.336 ms
64 bytes from 172.16.0.50: icmp_seq=89 ttl=64 time=0.192 ms
64 bytes from 172.16.0.50: icmp_seq=90 ttl=64 time=0.192 ms
64 bytes from 172.16.0.50: icmp_seq=91 ttl=64 time=0.247 ms
64 bytes from 172.16.0.50: icmp_seq=92 ttl=64 time=0.172 ms
64 bytes from 172.16.0.50: icmp_seq=93 ttl=64 time=0.284 ms
--- 172.16.0.50 ping statistics ---
100 packets transmitted, 14 received, 86% packet loss, time
   11253ms
rtt min/avg/max/mdev = 0.172/0.270/0.503/0.082 ms
```

The line above mentions that 14 responses were received, indicating 86 packets were not received.

Note: The number of blocked requests may vary between 80 and 90 depending on when the flood is started relative to the FortiDDoS device one second boundary.

Monitoring events

The Monitor button on the screen shows the properties of all events that have occurred for a selected period of time.

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

The events can be viewed at various levels as a table.

When packets are dropped by the appliance, you can see the cause of the drops and other details as events in the event monitor.

Refer to the *DDoS Fundamentals Guide* for further details.

Showing traffic

The FortiDDoS user interface provides several granular traffic graphs. You can see the traffic through each VID independently. The detailed description of these graphs is available in the *FortiDDoS Web-based Manager Guide*. Corresponding to the ping test in "Performing a sanity test" on page 18, activity will appear in the following and several other graphs:

1. *Show > Global > Card 1 > LAN 1 and WAN 1*
2. *Show > Current VID > Layer 3 > My Graphs > Protocols*
3. *Show > Current VID > Layer 4 > My Graphs > ICMP Types and Codes*.

Showing event reports

the FortiDDoS device user interface provides several granular event reports to summarize the past attack events. You can see the reports for each VID independently. The detailed description of these reports is available in the *FortiDDoS Web-based Manager Guide*. Corresponding to the ping test, activity will appear in the following reports:

- Top Attacked Services and Top Attacked ICMP Type and Code
- Top Attacked Protocols
- Top Attacks.

Using bypass switches for fail-over

Bypass switches are useful for fail-over purpose. They can be used for the occassional maintenance required for FortiDDoS devices. Passive bypass switches are useful in case of power failure. If both the FortiDDoS device and the failover switch share the same power, external connectivity can still be maintained.

As shown in Figure 15, when the bypass switch is in disabled mode, the in-line traffic continues to flow through the FortiDDoS device. This is the default mode.

Figure 15: Bypass Switch in Disabled Mode

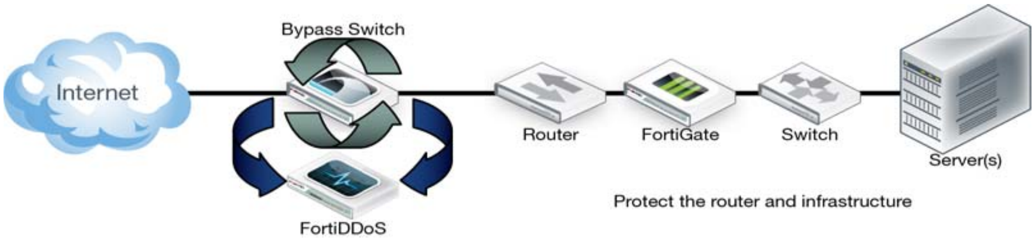

As shown in Figure 16, when the bypass switch is in bypass enabled mode, all in-line traffic is routed through the bypass switch. In this mode, the switch allows the FortiDDoS device to be removed and replaced without network downtime. Once power is restored to the bypass switch, network traffic is seamlessly diverted to the FortiDDoS device, allowing it to resume its critical functions.

Figure 16: Bypass switch in enabled mode

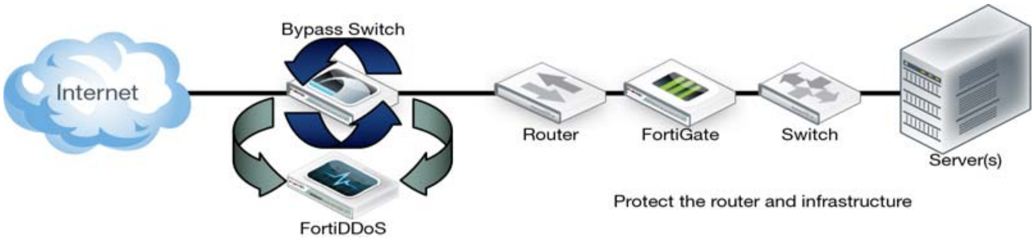

Using an optical bypass switch with heartbeat

Fortinet recommends using a FortiBridge devices as your optical bypass switches to provide a permanent and trouble-free access.

Figure 17: The optical bypass device used in serial detection mode

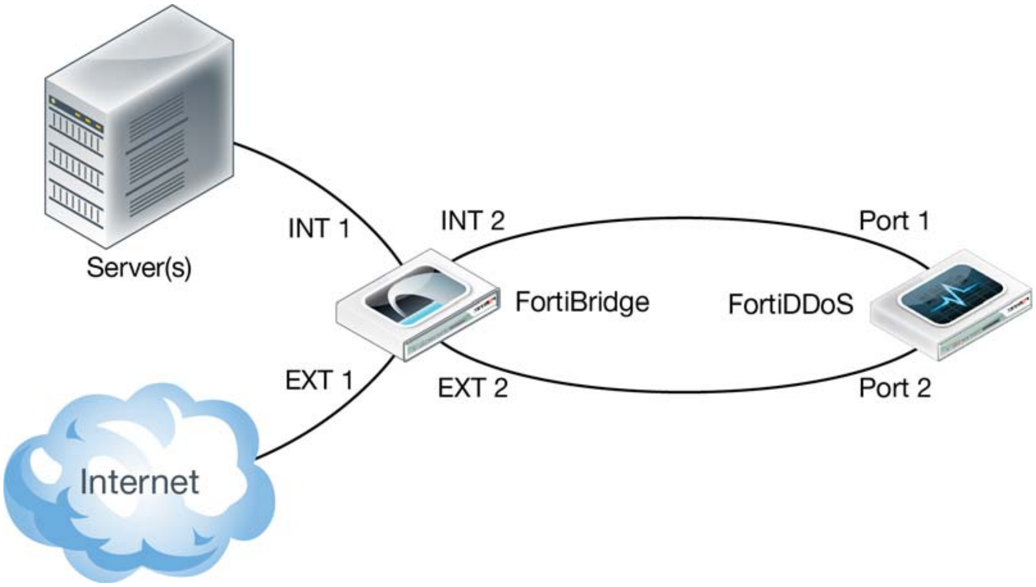

The optical bypass switch with heartbeat, monitors the link to the attached FortiDDoS device by sending a heartbeat packet to the device once every second. If the optical bypass switch does not receive the heartbeat back, it automatically switches network traffic to bypass the unresponsive FortiDDoS device - even if the device is still receiving power. The optical bypass continues to send the heartbeat and restores the traffic through the FortiDDoS device as soon as the link is restored.

Configuring the optical bypass switch

Refer to the *FortiBridge QuickStart Guide* and *FortiGate Hardware Guide* to set the following parameters:

- Input timeout period
- Input retry count

Connecting the optical bypass switch to the network

1 Connect the INT 1 port to the Server side.

2 Connect the EXT 1 port to the Internet side.

3 Connect the INT 2 port to the Server Port of the FortiDDoS device.

4 Connect the EXT 2 port to the Internet Port of the FortiDDoS device.

Using copper 10/100/1000 bypass switch with heartbeat

When the bypass switch loses power, in-line traffic continues to flow through the network link, but is no longer routed through the FortiDDoS device. This switch also allows the FortiDDoS device to be removed and replaced without network downtime. Once power is restored to the bypass switch, network traffic is seamlessly diverted to the FortiDDoS device, allowing it to resume its critical functions.

Configuring the bypass switch

A 10/100/1000 bypass switch allows you to set communication parameters. Refer to the *FortiGate Hardware Guide* to set the following parameters:

- Auto-negotiation
- Line speed
- Link Fault Detect (LFD)
- Input timeout period
- Input retry count

Connecting the 10/100/1000 bypass switch to the network

1. Connect the INT 1 port to the Server side.
2. Connect the EXT 1 port to the Internet side.
3. Connect the INT 2 port to the Server Port of the FortiDDoS device.
4. Connect the EXT 2 port to the Internet Port of the FortiDDoS device.

Configuring MAC Addresses for Bypass Switch Heatbeat 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.

To configure, please refer to section “Configuring MAC Addresses for Bypass Switch Heatbeat Packets” in the *FortiDDoS Web-based Manager v3.1 Reference Guide*.

Using traffic diversion in service provider environment

Traffic diversion

A FortiDDoS device can be deployed in Service Provider environment where the total bandwidth may be much more than what the appliance can support, but the attack traffic to a specific subnet or server is within its capacity. In such cases, normal traffic is sent through its regular path while the attack traffic is manually diverted by Service Provider staff during attack through the FortiDDoS device. The FortiDDoS device in turn cleans the traffic and injects it back to the network.

The FortiDDoS device is a layer-2 bridge and therefore does not have either a MAC address or an IP address in the data path (path of the packets.). To allow such diversions, you must therefore connect the device to interfaces on the routers or switches that have a routeable IP address.

Refer to Figure 18. Following terminology is used in this section:

- Divert-from router—Router from which the FortiDDoS device diverts the attacked customer traffic.
- Inject-to router—Router to which the FortiDDoS device forwards the legitimate traffic to the attacked customer.

A very simple deployment is explained in Figure 18. This involves Layer 2 forwarding through the FortiDDoS device.

One additional interface on the Divert-from Router - Router 1 is used to divert the traffic to the attacked destination. This traffic passes through the FortiDDoS device. The traffic is then forwarded to the Inject-to Router - Router 2. These two interfaces are in the same network (192.168.1.x) and therefore an ARP request from Router 1 for 192.168.1.2 passes through the FortiDDoS device and reaches Router 2 and Router 2 can respond back with an ARP reply and vice versa.

A static route is added on Router 1 for addresses for the attacked customer network. Having the longest matching prefix, the rule matches first and therefore all traffic to attacked customer network is diverted from Router 1 to Router 2 through the FortiDDoS device network rather than going straight from Router 1 to Router 2. The return path for traffic should preferably be via the FortiDDoS appliance. The solution will work even if the traffic is unidirectional through the FortiDDoS device. Bidirectional traffic helps the FortiDDoS device determine the statefulness within connections.

Figure 18: Traffic Diversion and a FortiDDoS device

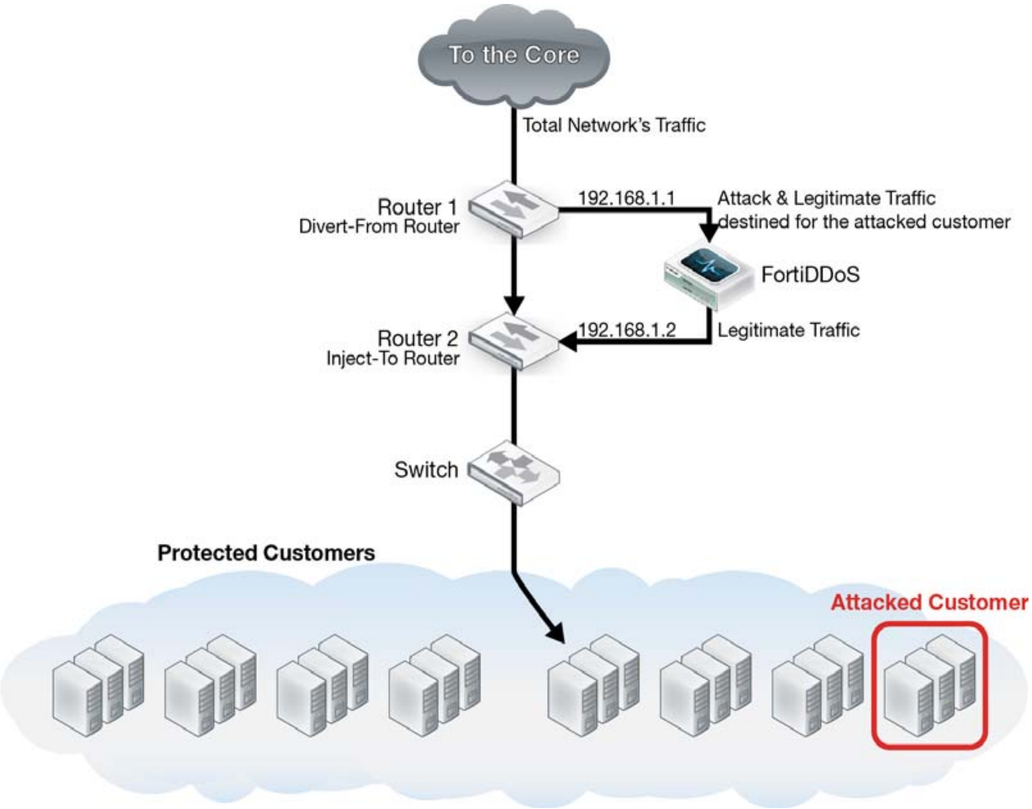

Traffic diversion using a single divert-from and inject-to router and a switch

Refer to Figure 19. A single router acts a Divert-from and Inject-to router. A very simple deployment is explained in Figure 19. This involves Layer 2 forwarding through the FortiDDoS device.

One interface on the Internet side of the Router is used to divert the traffic to the attacked destination. This traffic passes through the FortiDDoS device through a switch. The traffic is then forwarded to the Inject-to interface on the same Router through the same switch.

To ensure that the traffic is symmetric and both incoming and outgoing traffic to/from the attacked destination go through the FortiDDoS device, the LAN interface of the Router is used to divert the traffic from the attacked destination. This traffic passes through the FortiDDoS device through a switch. The traffic is then forwarded to the Inject-to interface on the same Router through the same switch.

A static route is added on the Router for addresses for the attacked customer network. Having the longest matching prefix, the rule matches first and therefore all traffic to attacked customer network is diverted from Router to the L3 Switch through the FortiDDoS device network rather than going straight from Router to Distribution Switch.

The return path for traffic should preferably be via a FortiDDoS appliance. The solution will work even if the traffic is unidirectional through the FortiDDoS appliance. Bidirectional traffic helps the FortiDDoS device determine the statefulness within connections.

To ensure that the return traffic passes through the FortiDDoS device, use Policy Based Routing (PBR) available in most routers. This allows routing based on source address of the packets and interface to be routed via an address.

Figure 19: Traffic diversion using a single divert-from and inject-to router and a FortiDDoS unit

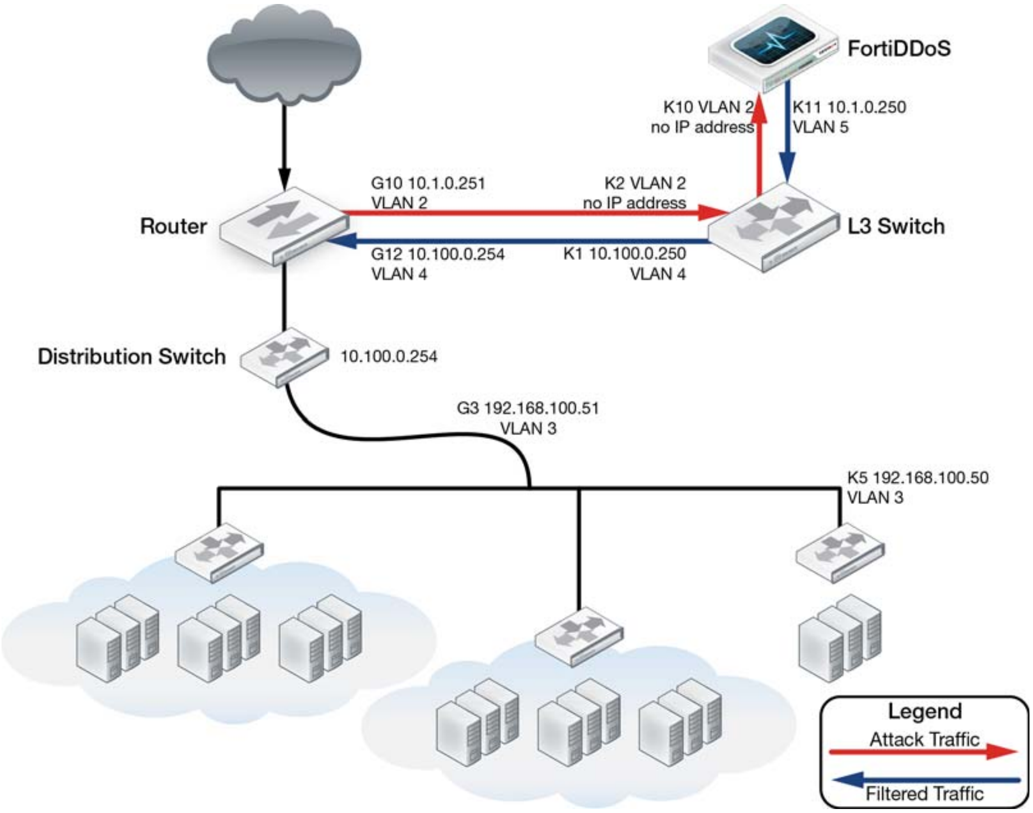

Router configuration for diversion

```
vlan internal allocation policy ascending
!
interface GigabitEthernet1/0/1
  no switchport
  no ip address
!
interface GigabitEthernet1/0/2
  switchport access vlan 3
  switchport trunk encapsulation dot1q
!
interface GigabitEthernet1/0/3
!
interface GigabitEthernet1/0/4
!
interface GigabitEthernet1/0/5
!
interface GigabitEthernet1/0/6
!
interface GigabitEthernet1/0/7
!
interface GigabitEthernet1/0/8
!
interface GigabitEthernet1/0/9
!
interface GigabitEthernet1/0/10
  switchport access vlan 2
!
interface GigabitEthernet1/0/11
ip address 10.100.0.250 255.255.255.0
no ip directed-broadcast
ip policy route-map FDD-X00A-PBR
!
interface GigabitEthernet1/0/12
!
interface Vlan2
  ip address 10.1.0.251 255.255.255.0
!
interface Vlan3
  ip address 192.168.100.51 255.255.255.0
!
!
ip classless
ip route 207.117.1.0 255.255.255.0 10.1.0.250
!
!
ip access-list extended zone-A
permit ip any  207.117.0.0 0.0.0.255
!
route-map FDD-X00A-PBR permit 100
match ip address zone-A
set ip next-hop 10.200.0.254
```

```
!
route-map FDD-X00A-PBR permit 101
description let thru all other packets without modifying next-hop
```

Switch configuration for traffic diversion

```
interface GigabitEthernet1/0/1
  no switchport
  no ip address
  channel-group 1 mode on
!
interface GigabitEthernet1/0/2
  switchport access vlan 3
  switchport trunk encapsulation dot1q
!
interface GigabitEthernet1/0/3
  switchport access vlan 3
!
interface GigabitEthernet1/0/4
  switchport access vlan 3
  switchport trunk encapsulation dot1q
!
interface GigabitEthernet1/0/5
!
interface GigabitEthernet1/0/6
!
interface GigabitEthernet1/0/7
!
interface GigabitEthernet1/0/8
!
interface GigabitEthernet1/0/9
!
interface GigabitEthernet1/0/10
  switchport access vlan 2
!
interface GigabitEthernet1/0/11
  switchport access vlan 5
!
interface GigabitEthernet1/0/12
  switchport access vlan 4
!
!

interface Vlan1
  no ip address
!
interface Vlan3
  ip address 192.168.100.50 255.255.255.0
!
interface Vlan4
  ip address 10.100.0.250 255.255.255.0
!
!
interface Vlan5
```

```
  ip address 10.1.0.250 255.255.255.0
!
ip classless
ip route 0.0.0.0 0.0.0.0 10.100.0.254
```

Threshold setting using predefined profiles

When traffic for different attacked customers are diverted through the FortiDDoS device during attacks, the device may not have the granular traffic thresholds set correctly corresponding to the traffic level normally experienced by the customer network.

To solve this issue, you can take two approaches:

1. **Learning Mode:** During normal times, train customer network traffic in different VIDs, archive profile for future use. Restore the threshold configuration during attack to a specific VID. Divert the traffic and configure the VIDs so that the FortiDDoS device uses the thresholds corresponding to that VID for the traffic.
2. **Predefined Profiles Mode:** Predefine different traffic levels - say 1 Mbps, 10 Mbps, 20 Mbps, 100 Mbps, etc. in various VIDs. Use additional predefined parameters such as SYN/second, SYNs/Src, Concurrent Connections/Source, Packets/second, etc. Use such predefined traffic level configurations for different VIDs and send the attack traffic to a VID that corresponds to the customer traffic level based on past historical knowledge of the data.

Using load balancing to support higher bandwidth in service provider environment

Load balancing

Many data center architectures require protecting network infrastructure, and server farms. With these requirements becoming more prevalent, traffic requirements on some networks may exceed the capabilities of the FortiDDoS appliance. Furthermore, the FortiDDoS devices in such network topologies could potentially become a network bottleneck. FortiDDoS appliances are restricted by interface speeds and support only 1 Gbps full duplex throughput. Thus to increase the overall throughput, you require some type of load balancing solution using multiple FortiDDoS appliances.

This leads to the requirement that the load-balancing device must exceed the throughput of numbers of multiple FortiDDoS devices.

Load Balancer intercepts all traffic between the LAN and the WAN, and dynamically distributes the load among the available FortiDDoS appliances, based on Load Balancer configuration. Load Balancing utilizes all the appliances concurrently, providing overall improved performance, scalability and availability.

The FortiDDoS device is a layer-2 bridge and therefore does not have either a MAC address or an IP address in the data path (path of the packets.). For transparent bridges, the Load Balancer receives a packet, makes a load balancing decision, and forwards the packet to a FortiDDoS device. The FortiDDoS device does not perform NAT on the packets; the source and destination IP addresses are not changed.

The load balancer must perform the following:

- Balance traffic across two or more FortiDDoS devices in your network, allowing them to work in parallel.

- Maintains state information about the traffic flowing through it and ensures that all traffic between specific IP address source and destination pairs flows through the same FortiDDoS unit.
- Performs health checks on all paths through the FortiDDoS devices. If any path is not operational, the load balancer diverts traffic away from that path, maintaining connectivity across the FortiDDoS devices.

You can use an external load balancer such as Linux Virtual Server (LVS), Cisco Content Switching Module (CSM), or Avaya Load Balancing Manager.

Load Balancing allows you to:

- Maximize FortiDDoS productivity.
- Scale FortiDDoS performance.
- Eliminate the FortiDDoS device as a single point of failure.

You must use **Sandwich topology** for Load Balancing using FortiDDoS device.

Sandwich topology

Refer to Figure 20. This topology requires a load-balancing device before and after the FortiDDoS device cluster. This type of design ensures the highest level of security due to physical separation of the FortiDDoS interfaces across multiple switches.

Each Load Balancer load balances between IP address interfaces of the peer device behind the FortiDDoS device. For this to work, each FortiDDoS device must reside in a different VLAN and subnet, and the physical ports connected to the FortiDDoS device must be on different VLANs as well. In addition, for each VLAN, both load balancers must be in the same subnet. Each load balancer interface and the FortiDDoS device connected to it reside in a separate VLAN. This ensures persistency since all the traffic through a particular FortiDDoS device is contained in the device’s VLAN.

In typical load balancers, there are two hash predictors:

- **Bi-Directional hash**, which requires both load-balancing devices to share a common hash value that ultimately produces the same route. Accomplish bi-directional hashing by hashing the source and destination IP address along with the destination port of the given flow. The load balancers ensure that all packets belonging to a session pass through the same FortiDDoS device in both directions. The devices select a FortiDDoS device based on a symmetric hash function of the source and destination IP addresses. This ensures that packets traveling between the same source and destination IP addresses traverse the same FortiDDoS device.
- **Uni-Directional hash** produces the route in the same fashion as a bi-directional hash and also creates a TCP connection table with the reverse flow path defined. This allows you to match return path traffic against this connection table rather than being hashed.

Figure 20: Load balancing using FortiDDoS devices using sandwich topology

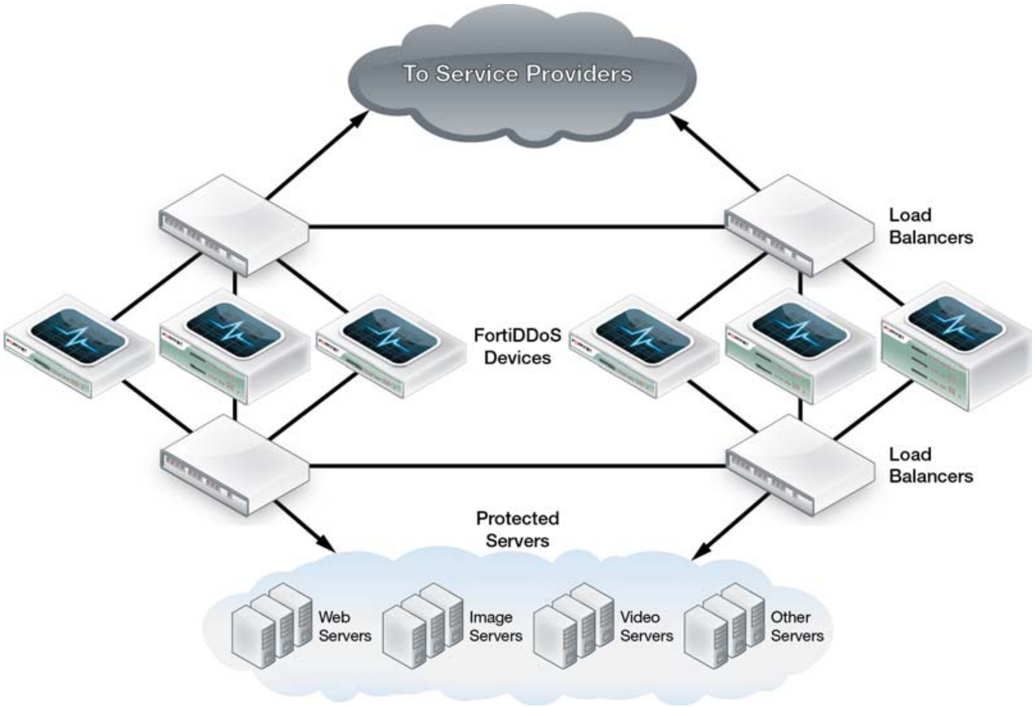

Refer to Figure 21, traffic flows through the FortiDDoS devices and the devices filter the traffic in both directions.

FortiDDoS devices do not have IP addresses on VLANs. Instead, you configure alias IP addresses on each switch interface to which the FortiDDoS device connects. The Load Balancing Switches use the alias IP addresses to direct traffic to the correct FortiDDoS device.

On the path to the intranet, Load Balancing Switch 1 (LBS1) balances traffic across VLANs 101, 102, and 103 through the firewalls to Load Balancing Switch 2. On the path to the Internet, Load Balancing Switch 2 (LBS2) balances traffic across VLANs 201, 202, and 203 through the FortiDDoS device to Load Balancing Switch 1. Each Load Balancing Switch uses the alias IP addresses configured on the other Load Balancing Switch as targets for the load-balancing process.

The LBS1 selects a FortiDDoS device based on source IP address using the hash address source predictor. The LBS2 selects a FortiDDoS device based on the destination IP address using the hash address destination predictor. This predictor allows the LBS2 to select the same FortiDDoS device for return flows and buddy connections such as in case of FTP.

Figure 21: Using VLANs and FortiDDoS devices in sandwich topology

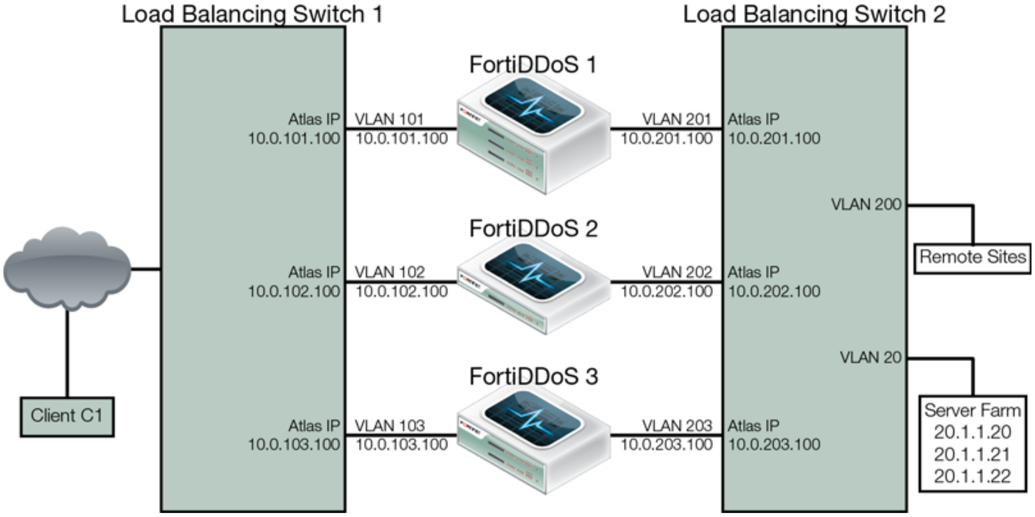

Switch Configuration for load balancing

```
(clientSide-84.82) #show run
!Current Configuration:
!
!System Description "FortiSwitch-248B-DPS 48x1G & 4x10G"
!System Software Version "5.2.0.2.4"

serviceport ip 192.168.22.98 255.255.255.0 0.0.0.0
vlan database
vlan name 10 "egress"
vlan name 11 "ingress"
exit

port-channel "egress" 1
interface 0/1
channel-group 1/1
exit
interface 0/3
channel-group 1/1
exit
interface 0/5
channel-group 1/1
exit
interface 0/7
channel-group 1/1
exit
interface 0/9
channel-group 1/1
exit
interface 0/11
channel-group 1/1
exit
interface 0/13
channel-group 1/1
exit
```

```
interface 0/15
channel-group 1/1
exit
port-channel "ingress" 2
interface 0/2
channel-group 1/2
exit
interface 0/4
channel-group 1/2
exit
interface 0/6
channel-group 1/2
exit
interface 0/8
channel-group 1/2
exit
interface 0/10
channel-group 1/2
exit
interface 0/12
channel-group 1/2
exit
interface 0/14
channel-group 1/2
exit
interface 0/16
channel-group 1/2
exit

mac-addr-table aging-time 60000

interface 0/1
no cdp run
switchport allowed vlan add 10
exit

interface 0/2
no cdp run
exit
interface 0/3
no cdp run
exit

interface 0/4
no cdp run
exit

interface 0/5
no cdp run
exit

interface 0/6
no cdp run
exit
```

```
interface 0/7
no cdp run
exit

interface 0/8
no cdp run
exit

interface 0/9
no cdp run
exit

interface 0/10
no cdp run
exit

interface 0/11
no cdp run
exit

interface 0/12
no cdp run
exit

interface 0/13
no cdp run
exit

interface 0/14
no cdp run
exit

interface 0/15
no cdp run
exit

interface 0/16
no cdp run
exit

interface 0/17
no cdp run
switchport allowed vlan add 10
switchport native vlan 10
exit

interface 0/18
no cdp run
switchport allowed vlan add 11
switchport native vlan 11
exit
```

```
interface 0/49
no cdp run
switchport allowed vlan add 10
switchport native vlan 10
exit

interface 0/50
no cdp run
switchport allowed vlan add 11
switchport native vlan 11
exit

interface 1/1
staticcapability
switchport allowed vlan add 10
switchport native vlan 10
lacp collector max-delay 0
exit

interface 1/2
staticcapability
switchport allowed vlan add 11
switchport native vlan 11
lacp collector max-delay 0
exit

interface 1/3
staticcapability
switchport allowed vlan add 10
switchport tagging 10
lacp collector max-delay 0
exit

interface 1/4
staticcapability
switchport allowed vlan add 11
switchport tagging 11
lacp collector max-delay 0
exit

router rip
exit
router ospf
exit
exit

(clientSide-84.82) #
(clientSide-84.82) #show load-balance
Hash Mode: src-dst-ip-ipport
```

Using FortiGuard IP Reputation Service

The FortiGuard IP Reputation Service aggregates data from around the world to provide up-to-date information about malicious sources. With breaking intelligence from distributed network gateways and world-class research by FortiGuard Labs, a subscription to FortiGuard IP Reputation Service helps block attacks.

Today's attackers infect and control hosts with botnets to launch phishing, spamming, and DDoS attacks. Threats can strike fast, disappear, and re-appear just as quickly in another form. These attacks are difficult to detect and investigate. They're often over by the time they're identified as DDoS attacks, and chasing non-existent attacks takes up valuable resources. If a malicious machine attacks a target in one location, the rest of the global network needs to find out fast in order to pre-empt the next wave. FortiGuard IP Reputation Service provides the updates.

FortiGuard IP Reputation Service:

- Protects against malicious sources associated with web attacks, phishing activity, web scanning, scraping etc.
- Blocks large scale DDoS attacks from known infected sources
- Protects against centrally managed and automated botnet attacks
- Blocks access from anonymous and open proxies (on FortiWeb platforms)
- Delivers daily IP reputation updates
- Automatically downloads
- Includes analysis tools to better understand origin of attack using Geo IP location

Configuring FortiGuard IP Reputation Service

After purchasing the service and registering your FortiDDoS serial number with FortiGuard, refer to the *FortiDDoS Web-based Manager Reference Guide* to configure access control lists using IP reputation and schedule IP reputation list updates.

FORTINET®