



WANGUARD 5.2

User Manual & Administrator's Guide

Console + Sniffing Sensor + Flow Sensor + Virtual Sensor + Filter

Copyright & trademark notices

This edition applies to version 5.0 of the licensed program WANGUARD and to all subsequent releases and modifications until otherwise indicated in new editions.

Notices

References in this publication to ANDRISOFT S.R.L. products, programs or services do not imply that ANDRISOFT S.R.L. intends to make these available in all countries in which ANDRISOFT S.R.L. operates. Evaluation and verification of operation in conjunction with other products, except those expressly designated by ANDRISOFT S.R.L., are the user's responsibility. ANDRISOFT S.R.L. may have patents or pending patent applications covering subject matter in this document. Supplying this document does not give you any license to these patents. You can send license inquiries, in writing, to the ANDRISOFT S.R.L. marketing department, sales@andrisoft.com.

Copyright Acknowledgment

© ANDRISOFT S.R.L. 2013. All rights reserved.

All rights reserved. This document is copyrighted and all rights are reserved by ANDRISOFT S.R.L. No part of this document may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying and recording, or by any information storage and retrieval system without the permission in writing from ANDRISOFT S.R.L.

The information contained in this document is subject to change without notice. If you find any problems in the documentation, please report them to us in writing. ANDRISOFT S.R.L. will not be responsible for any loss, costs or damages incurred due to the use of this documentation.

WANGUARD is a SOFTWARE PRODUCT of ANDRISOFT S.R.L. WANGUARD and WANSIGHT are trademarks of ANDRISOFT S.R.L. Other company, product or service names may be trademarks or service marks of others.

ANDRISOFT S.R.L.

Str. Lunei L30 Ap. 11, 300109 Timisoara, Timis, Romania

Phone: +40721250246

Sales: sales@andrisoft.com

Technical Support: support@andrisoft.com

Website: <http://www.andrisoft.com>

© Copyright ANDRISOFT S.R.L. 2013. All rights reserved.

Table of Contents

1. IP Traffic Monitoring, Anomalies Detection & DDoS Mitigation with WANGUARD	4
WANGUARD Key Features & Benefits	4
WANGUARD Components	5
2. A first look at WANGUARD Console	6
Side Region – used for navigation throughout the Console	6
Central Region – home of tabbed Reports and Dashboards	6
South Region – provides a quick look on the latest events, live statistics and graphs	6
Upper-right Menus – Help menu and User menu	6
3. Reports » Anomalies & Tools	7
Anomalies	7
Active Anomalies	7
Anomalies Archive	9
Anomalies Overview	9
BGP Prefixes	9
BGP Operations	9
BGP Logs	10
Flow Collector	10
List Flows	10
Flows Tops	11
Autonomous Systems	12
Packet Analyzer	13
Active Captures	13
Captures Archive	15
4. Reports » Dashboards	16
5. Reports » Interfaces	17
Overview	17
Console	17
Active Virtual Sensors	18
Active Sniffing Sensors	18
Active Flow Sensors	18
Active Filters	19
Sensors	20
Sensor Dashboard	20
Sensor Graphs	21
Sensor Tops	22
List Flows	23
Flows Tops	23
Sensor Events	23
Anomalies Overview	24
6. Reports » IP Addresses & Groups	25
IP Dashboard	25
IP Graphs	25
IP Accounting	26
List Flows	27
Flows Tops	27
Profile Graphs	27
Anomalies Overview	27

7. Installation Guide.....	28
System Requirements.....	28
Sniffing Sensor Hardware Requirements.....	28
Flow Sensor Hardware Requirements.....	29
Filter Hardware Requirements.....	29
Console Hardware Requirements.....	29
Software Installation & Download.....	30
Opening Console for the first time.....	30
Licensing Procedure.....	30
Quick Configuration Steps.....	31
8. Storage & Graphs Configuration.....	32
9. Anomalies Configuration.....	33
10. Response Configuration.....	34
Conditional & Dynamic Parameters.....	35
11. IP Zone Configuration.....	40
Anomaly detection settings & Thresholds Templates.....	41
12. Choosing a method of traffic monitoring.....	42
Comparison between Packet Sniffing and Flow Monitoring.....	43
13. Sniffing Sensor Configuration.....	44
14. Flow Sensor Configuration.....	47
15. Virtual Sensor Configuration.....	50
16. BGP Connection Configuration.....	51
17. Filter Configuration.....	53
18. Scheduled Reports.....	57
19. Events Reporting.....	58
20. Users Management.....	59
21. Appendix 1 – Network Basics You Should Be Aware Of.....	60
IPv4 Subnet CIDR Notation.....	62
22. Appendix 2 – Configuring NetFlow Data Export.....	63
Configuring NDE on an IOS Device.....	63
Configuring NDE on a CatOS Device.....	64
Configuring NDE on a Native IOS Device.....	64
Configuring NDE on a 4000 Series Switch.....	65
Configuring NDE on a Juniper Router.....	65
23. Appendix 3 – Configuring Traffic Diversion.....	67
Understanding the BGP Diversion Method.....	67
BGP Configuration Guidelines.....	68
Filter System BGP Configuration.....	68
Filter System BGP Configuration Example.....	69
Cisco Router BGP Configuration.....	70
Cisco Router BGP Configuration Example.....	70
Understanding Traffic Forwarding Methods.....	71
Static Routing – Layer 2 Forwarding Method.....	71
GRE / IP over IP Tunneling – Layer 3 Forwarding Method.....	71
Configuring Static Routing – Layer 2 Forwarding Method.....	71
Configuring GRE / IP over IP Tunneling – Layer 3 Forwarding Method.....	72

IP Traffic Monitoring, Anomalies Detection & DDoS Mitigation with WANGUARD

Unforeseen traffic patterns affect user satisfaction, pressure over-subscription plans, and clog costly transit links. Providing high performance and reliable network services is central to the success of today's organizations. As the business cost of network malfunctions continues to increase, rapid identification and mitigation of threats to network performance and reliability becomes critical in order to meet expected SLAs and network availability requirements. Such threats can include propagating worms, botnet attacks, Denial Of Service attacks (SYN flood, UDP flood etc.), misuse of services, and data traffic interfering with real-time traffic. WANGUARD's network-wide surveillance of complex, multilayer, switched and routed environments together with its unique combination of features is specifically designed to meet the challenge of pin-pointing and resolving any such threats.

WANGUARD Key Features & Benefits

- **DDOS DETECTION & MITIGATION** – It contains an innovative anomaly detection engine that you can use to define traffic policies, detect attacks and filter them.
- **POWERFUL ALERTING** – You can automate responses to threats using pre-defined, extensible actions: send emails, announce prefixes in BGP, null-routing, send SNMP traps etc.
- **DETAILED ATTACK INFORMATION** – View attack details with attackers and packet samples. Attack reports can be emailed automatically to you or to your customers.
- **TRAFFIC MONITORING** – Supports the latest traffic monitoring technologies: 10 Gbps packet sniffing, NetFlow v5, v7 and v9, sFlow, IPFIX, NetStream, jFlow and more.
- **FULLY-FEATURED CONSOLE** – Consolidated management through a single, interactive and configurable web portal with custom Dashboards and user Roles.
- **COMPLEX ANALYTICS** – Provides the most complex Reports with aggregated data for hosts, departments, interfaces, applications, ports, protocols etc.
- **REAL-TIME REPORTING** – The fastest solution on the market with an accuracy of just 5 seconds. The high accuracy makes traffic graphs appear animated.
- **HISTORICAL REPORTING** – You view the last half hour to last 10 years Reports, and also select any custom time period. Supports 95th percentile.
- **SCHEDULED REPORTING** – You can generate Scheduled Reports and email them to you or to your customers at preconfigured intervals of time.
- **NETFLOW ANALYZER** – Provides a fully featured NetFlow Analyzer and Collector. Also works with sFlow, jFlow, cFlow, NetStream and IPFIX.
- **PACKET SNIFFER** – A distributed Packet Sniffer can save packet dumps from different parts of your network. Access the dumps from a Wireshark-like web interface.
- **ADVANCED CONFIGURATION** – You can fine-tune most parameters, from the accuracy of IP graphs and authentication methods to the data retention intervals.

- **OUTSTANDING SUPPORT** – Standard Support inquiries sent by email are answered by experienced engineers in 24 hours or less. We can use Skype or TeamViewer.
- **CONTEXTUAL HELP** – Includes a Contextual Help system, an Installation Wizard and a User Manual and Administrator's Guide in PDF format for easy printing.

The recorded data is stored in an internal SQL database that can be easily queried and referenced. The recorded monitoring statistics can be viewed through a rich, easy-to-use Ajax-based web interface.

WANGUARD Components

Andrisoft WANGUARD is an enterprise-grade Linux-based solution that delivers the functionality NOC, IT & Security teams need to effectively monitor and protect their network through a single, integrated package. The components have been built from the ground up to be high performing, reliable and secure.

WANGUARD relies on the **Sniffing Sensor** or on **Flow Sensor** to provide in-depth traffic analysis, traffic accounting, bandwidth monitoring and traffic anomaly detection. The collected information enables you to generate complex traffic reports, graphs and tops, instantly pin down the cause of network incidents, understand patterns in application performance and make the right capacity planning decisions. When DoS, DDoS or DrDOS attacks occur, the **Filter** detects attack patterns and scrubs off anomalous traffic in a granular manner without impacting the user experience or resulting in downtime.

The **Console** offers single-point management and reporting by consolidating data received from all WANGUARD components deployed within the network.

A first look at WANGUARD Console

If you're an administrator and you look on how to configure WANGUARD, skip to the Installation Chapter on page 28.

Please read the following chapters in order to get a clear overview of the basic premises required for the proper operation of the software. The next 5 chapters cover all reporting features, while the latter cover the configuration of the solution.

To understand the operation of the Console please be aware of the structure of the web application:

Side Region – used for navigation throughout the Console

It is located on the east or west edge of the window, according to the user's preference. If it's not visible then it's either collapsed or hidden by an Administrator. Clicking the edge of regions expands or collapses them.

The Side Region contains 2 sections – Reports and Configuration – that can be collapsed or expanded by clicking the title bars. Both sections contain multiple panels than can also collapse or expand, their state being kept between sessions. Panels are refreshed automatically every 5 to 10 seconds.

The Reports section title bar contains a “Quick Search” functionality button. Shortcut: Ctrl-S

Central Region – home of tabbed Reports and Dashboards

The Console offers various ways to look at historic and live collected data. Each Report or Dashboard you request through the Side Region opens a tab (page) in the Central Region. You may switch between tabs or close them all except for the Landing Tab defined in the user's preference. Initially the Landing Tab is this Configuration Wizard.

South Region – provides a quick look on the latest events, live statistics and graphs

It's located on the bottom of the browser window. It's collapsed by default so to expand click the small bottom edge. It provides a quick way to view live data: animated graphs, events, anomalies, statistics from all components.

Upper-right Menus – Help menu and User menu

The Help menu contains the User Manual, some useful tools and the About window. Depending on the context, the User Manual will open at the chapter describing the last opened window or tab. The Contextual Help works only with Adobe PDF Reader.

The User menu lets you quickly change the password, the Console theme and provides a Log Out option.

Reports » Anomalies & Tools

The **Reports » Anomalies & Tools** panel contains links to the **Anomalies** tab, to the **BGP Prefixes** tab, to the **Flow Collector** tab and to the **Packet Analyzer** tab.

Anomalies

The Anomalies tab contains live and historical data that relates to DoS and DDoS attacks or other traffic anomalies. The number of active traffic anomalies is displayed within the Anomalies panel and it's refreshed every 10 seconds. It's not displayed if it's zero.

The Anomalies tab contains 3 sub-tabs located on the bottom side:

Active Anomalies

Active Anomalies contains a table visible only while Sensors detect active traffic anomalies. The table's rows represent active anomalies, sorted by start time in descending order. The table's columns are:

No	The unique index of the anomaly. Click it to open a detailed Anomaly Report.
Prefix	The IP address or IP class of the traffic anomaly and the reverse DNS. In the front of the Prefix, the graphic arrow indicates the direction of the traffic: inbound when the arrow is pointing towards the Prefix, or outbound when the arrow is pointing away from the Prefix. Click it to open a new tab with data specific for the Prefix.
IP Group	The IP Group of the Prefix. Click it to open a new tab with data specific for the IP Group.
Anomaly	A short description of the anomaly.
Value	The peak value of the anomalous traffic. Between parenthesis, the latest value.
Sensor	The name of the Sensor that detected the anomaly. Click it to open a new tab with data specific to the Sensor.
From	The time and date when the anomaly started.
Latest Alarm	How much time passed since the last detection of the anomaly.
Pkts/s – Bits/s	The latest packets/second and bits/second throughput of the TOTAL traffic.
Actions	Actions available for Administrators and Operators:

	<ul style="list-style-type: none"> • <i>View Traffic Graph</i> – available if IP Graphs is enabled for the prefix • <i>View Traffic Log</i> – available if the Response contains a Traffic Capturing action • <i>Delete BGP Route</i> – available if a BGP announcement was sent for the prefix • <i>Set Comment</i> – add or modify comments • <i>Stop Traffic Anomaly</i> – force the Sensor to clear the anomaly
Dropped	The percent of the anomalous traffic filtered by one or more Filter systems.
Severity	<p>The severity field represents graphically the ratio between the anomalous traffic and the threshold value. Every bar represents 100% of the threshold value.</p> <p>The color of the severity indicates the link's severity: 0-25% blue, 25%-50% yellow, 50%-75% orange, 75%-100% red. The link's severity is the ratio between the anomaly traffic and the overall traffic of the link (Sensor or Interface).</p> <p>The exact rule's severity and link's severity is displayed as a tool-tip.</p>
PARAMETERS VISIBLE ONLY WHEN DISPLAY IS SET TO "FULL"	
Total Pkts	The number of packets from the total traffic during the anomaly.
Total Bits	The number of bits from the total traffic during the anomaly.
Overall Traffic	The percent between the anomaly traffic and the overall traffic.
Threshold	The threshold's value.
IP Zone	The IP Zone of the Sensor. Click it to open the Prefix settings from the IP Zone.
Template	The Thresholds Template that contained the anomaly's rule, if any.
Expiration	The number of seconds between the latest alarm and the time the anomaly becomes inactive.
Response	The name of the Response executed for the anomaly.
Comments	User comments. The field is hidden if there are no comments.

If one or more Filters are activated to detect attackers and attack patterns then a new table will show up in the same traffic anomaly row. The rows in the table will have red background for active attack patterns and yellow background for inactive attack patterns.

Filter	The name of the Filter that detected the attack pattern.
Filtering Rule	<p>The filtering rule applied to drop the attack pattern's traffic. The Filter dynamically applies the following filtering rules: <i>Source IP, Source Port, Destination Port, Packet Length, TimeToLive, IP Protocol</i>.</p> <p>Filtering rules are applied only when the filtering policy allows dropping of traffic. If the filter conflicts with the Filter's Whitelist, then a red exclamation point shows up.</p>

Firewall	Indicates if there was a software filter installed or a firewall filter, or both.
From	The date and time when the attack pattern was first detected
Until	The date and time when the attack pattern was last detected.
Latest Alarm	How much time passed since the last detection of the attack pattern.
Max Pkts/s	The maximum packets/second throughput for the traffic matching the attack pattern.
Max Bits/s	The maximum bits/second throughput for the traffic matching the attack pattern.
Pkts	The number of packets counted in the traffic matching the attack pattern.
Bits	The number of bits counted in the traffic matching the attack pattern.
Log	When clicking the icon, a new tab opens with a packet-level capture of the attack pattern. Available only if the Response contains a Traffic Capturing action.

Anomalies Archive

The Anomalies Archive shows all traffic anomalies sorted by time in descending order. By clicking the down arrow on any column header, you can apply filters, change sorting direction and hide or show columns.

The <+> sign from the first column expands the row with additional information about the anomaly, mitigation information etc. Other columns are explained on previous paragraphs.

Anomalies Overview

Here you can view trends and summarizations of attacks for the selected time-frame, Sensors and decoders.

BGP Prefixes

The BGP Prefixes tab contains live and historical data that relates to BGP announcements. The number of active BGP announcements is displayed within the Anomalies panel and it's refreshed every 10 seconds. It's not displayed if it's zero.

The BGP Prefixes tab contains 2 sub-tabs located on the bottom side:

BGP Operations

BGP Operations provides live insight on BGP announcements made by Sensors, Filters or through Console.

Administrator or Operator roles can add and remove BGP announcements manually. To add a new BGP announcement you must enter the Prefix and select a previously configured BGP Connection.

The table containing BGP announcements is visible only while announcements are active. The table's columns are:

BGP Connection	The BGP Connection name as defined in the BGP Connection's configuration – see page 51.
Prefix	The Prefix included the announcement. When announcing individual hosts use the /32 CIDR for IPv4 and /128 CIDR for IPv6.
From	The date when the BGP announcement was sent.
Until	The date when the BGP announcement will be withdrawn.
Anomaly	If the BGP announcement was triggered by a Response to an anomaly, it contains the link to the detailed Anomaly Report.
Comments	This field may contain comments about the BGP announcement. If it contains the word ERROR then look for details in Events – see page 58.
Action	It contains a link for the manual removal of the BGP announcement. The Action field is visible only for Administrator or Operator roles.

BGP Logs

BGP Logs shows all BGP announcements sent by WANGUARD, sorted by time in descending order. By clicking the down arrow on any column header, you can apply filters, change sorting direction and hide or show columns. The columns are explained in the previous paragraph.

You can modify the status of announcements by double-clicking the rows.

Flow Collector

The **Reports » Anomalies & Tools** panel contains links to the **Flow Collector** tab if at least one Flow Sensor was configured.

Here you can list, aggregate, filter and sort individual flows, generate traffic tops and statistics, and view traffic graphs for Autonomous Systems.

The Flow Collector tab contains 3 sub-tabs located on the bottom side:

List Flows

You can list and filter the flow data according to your needs by entering the fields below:

- **Sensors**

Select the Flow Sensor Interfaces that captured the traffic you're interested in. Multiple selections can be made. Administrators can filter which Sensors are available to users.

- **Time-frame**

Select predefined time-frames or enter your own by selecting “Custom...”.

- **Flows Filter**

Here you can enter a filter for flows. Click the lightbulb icon on the right to open a window containing the correct syntax. Frequently used flows filters can be saved there and reused at any time later.

- **Output**

You can select several output formats, or you can type your own format that conforms to the format specification of nfdump.

For better readability IPv6 addresses are shortened, such as that the middle nibbles are cut and replaced by dots '...'. Most often this is good enough to recognize a wanted IPv6 address your are looking for. If you need the full IPv6 address, check the option 'IPv6 long'.

- **Export**

If the output is small you can send it by Email or you can Print it.

But when you need to generate huge amounts of flow data, doing that solely through the browser may not be the best idea. In this case, select the “Dump” option to view the CLI command used to generate the data. You can execute the command locally, forward the output to a file etc.

- **Aggregation**

By default the flows are not aggregated. By clicking on the checkboxes, you can select how you want to have your flows aggregated. You may also aggregate entire subnets when selecting srcIPv4/<subnet bits>.

- **Limit Flows**

List only the first N flows of the selected time slot.

- **Sorting**

When listing flows from different Flow Sensors you may sort them according to the start time of the flows. Otherwise the flows are listed in sequence of the selected Flow Sensors.

Flows Tops

You can process and filter the flow data to generate tops by entering the fields below:

- **Sensors**

Select the Flow Sensor Interfaces that captured the traffic you're interested in. Multiple selections can be made. Administrators can filter which Sensors are available to users.

- **Time-frame**

Select predefined time-frames or enter your own by selecting “Custom...”.

- **Flows Filter**

Here you can enter a flows filter. Click the lightbulb icon on the right to open a window containing the correct syntax. Frequently used flows filters can be saved there and reused at any time later.

- **Output**

You can select several output formats, or you can type your own format that conforms to the format specification of nfdump.

For better readability IPv6 addresses are shortened, such as that the middle nibbles are cut and replaced by dots '...'. Most often this is good enough to recognize a wanted IPv6 address your are looking for. If you need the full IPv6 address, check the option 'IPv6 long'.

- **Export**

If the output is small you can send it by Email or you can Print it.

But when you need to generate huge amounts of flow data, doing that solely through the browser may not be the best idea. In this case, select the "Dump" option to view the CLI command used to generate the data. You can execute the command locally, forward the output to a file etc.

- **Top Type**

Select the statistics you want from the menu and the order option.

- **Aggregation**

By default the flows are not aggregated. By clicking on the checkboxes, you can select how you want to have your flows aggregated. You may also aggregate entire subnets when selecting srcIPv4/<subnet bits>.

- **Limit**

Limit the output only to those statistic lines whose packets or bytes match the specified limit.

- **Top**

Limit the statistics to the first top N.

Autonomous Systems

If you are using the Flow Sensor, you can generate traffic and bandwidth histograms for Autonomous Systems. You can use this option if you have BGP-enabled flow exporters that are configured to include AS information in exported flows.

The parameters are:

- **Sensors**

Select the Flow Sensors you're interested in. Multiple selections can be made. Administrators can filter which Sensors are available to users.

- **Time Frame**

Select predefined time-frames or enter your own by selecting "Custom...".

- **AS Numbers**

Click the lightbulb icon on the right to open a window containing the correct syntax. Often used AS numbers can be saved there and used at any time later.

If you don't know what AS number(s) is a particular ISP having then you can click the upper-right side of

the window: Help » AS Information » AS Numbers List. There you can apply different filters by clicking the table header's down icon.

- **Export**

You can print, save as PDF or email the generated AS graphs.

- **Refresh**

The report is refreshed only when you press the <<Generate>> button. If you select a refresh interval then the graphs will be constantly refreshed.

- **Graphs Size**

You can select predefined sizes or you can enter your own size in the “<X pixels> x <Y pixels>” format.

- **Graphs Title**

Graphs can have an automatically-generated title for the “Default” option, no title for the “None” option, or you can enter your own text that will be rendered as a title.

- **Stack Sensors**

If unchecked, a different AS graph is generated for every Flow Sensor. Otherwise a single AS graph that contains summed traffic data is generated for all Flow Sensors.

- **Stack ASNs**

If you entered multiple AS Numbers then you can sum all of them in a single AS graph. Useful with ISPs and AS owners that have more than 1 allocated AS number.

Packet Analyzer

The **Reports » Anomalies & Tools** panel contains links to the **Packet Analyzer** tab if at least one Sniffing Sensor was configured.

The Packet Analyzer allows you to easily capture packets using distributed Sniffing Sensors. You can view packet dumps directly from the Console using an integrated Wireshark-like interface.

The tab contains 2 sub-tabs located on the bottom:

Active Captures

Administrators, Operators and Users with Packet Capturing privileges can generate packet dumps by clicking the <<Add Capture>> button. The options are:

- **Description**

A short description to help you identify the capture.

- **Sniffing Sensors**

Select the Sniffing Sensors that could capture the traffic you're interested in. Multiple selections can be made. Administrators can filter what Sensors are available to users.

- **BPF Expression**

Click the lightbulb icon on the right to open a window containing the correct BPF – Berkley Packet Filter syntax. Often used BPF expressions can be saved there and used at any time later.

The use of a BPF expression is mandatory but you can use the “ip” string to capture all IP traffic.

- **Max Running Time**

The maximum running time.

- **Stop Capture On**

When the Max Running Time is set to “Unlimited” you can set an exact date when the capture will stop.

- **Max File Size (MB)**

Before writing a raw packet to a file, check whether the file is currently larger than the <number> and, if so, close the current file and open a new one.

- **Max Packets**

The capture stops after receiving <number> packets.

- **Max File Number**

Setting this will limit the number of files created to the specified <number>, and begin overwriting files from the beginning, thus creating a 'rotating' buffer. In addition, it will name the files with enough leading 0s to support the maximum number of files, allowing them to sort correctly.

- **Time Rotation (s)**

If specified, it rotates the file every <number> seconds.

- **Sampling Type & Value**

Select “None” when no packet sampling is required. Select “1 / Value” to save just one packet every <value> packets. Select “Value / 5s” to save maximum <value> packets every 5 seconds.

- **Filename Prefix**

The name of the capture file. If any file-rotation options are used then a number will be appended to the filename.

- **Snapshot (bytes/pkt)**

Snarf <number> bytes of data from each packet rather than the default of 65535 bytes. Note that taking larger snapshots both increases the amount of time it takes to process packets and, effectively, decreases the amount of packet buffering. This may cause packets to be lost. You should limit <number> to the smallest number that will capture the protocol information you're interested in.

- **Comments**

This field may contain comments about the traffic capture.

Active Captures are listed as a table with in following format:

- **Description [BPF]**

The capture's description and the BPF expression.

- **Sampling**
The type of sampling that is being used.
- **From**
The date when the Sniffing Sensor started capturing packets.
- **Until**
The time or the conditions that will cause the stopping of the capture.
- **Status**
It indicates the status of the capture. It's green if the capturing thread still runs.
- **Interface**
The Sniffing Sensor or the Filter that captures packets.
- **Files / Size**
The number of dump files generated, and the size of the latest dump file.
- **Packets**
The number of packets captured.
- **Actions**
Click the first icon to view the latest dump file in a Wireshark-like web interface. Click the second icon to download the latest dump file. Click the third icon to stop the capture.

Captures Archive

Captures Archive lists all captures sorted by time in descending order. By clicking the down arrow on any column header, you can apply filters, change sorting direction and hide or show columns.

The <+> sign from the first column expands the row with additional information about the capture, and provides access to every capture file. Other columns are explained on previous paragraphs.

Reports » Dashboards

Wouldn't it be nice to see all your relevant data in a single tab? The **Dashboard** allows you to group data according to your needs.

Few sample Dashboards are included in the Console, but you can create more by going to **Reports » Dashboards**. Open an existing Dashboard and click <<Add Dashboard>>.

Then add some **widgets** to your Dashboard. To sort them, click the title bar and move them around. To collapse a widget, click the first icon on the widget title bar. To edit a widget, click the second icon from it's title bar. To delete a widget, click the third icon from it's title bar.

Along with specific fields, every widget has a configurable title and height. Leave the widget's height parameter to "Auto" for the widget to take all the vertical space it needs. To restrict the height of a widget enter a number of pixels instead.

Widgets options are self-explanatory or are described in other Reports chapters.

Only "Administrator" and "Operator" roles are able to create, delete or edit Dashboards. The "User" role doesn't allow modifications on Dashboards.

Reports » Interfaces

The **Reports » Interfaces** panel contains links to the **Overview** tab, **Interface Groups** tabs and to detailed **Sensor** tabs.

The Overview tab provides a real-time view on the status of all WANGUARD components.

Interface Groups tabs provide a real-time view on the status of belonging Sensor(s) and Filter(s). Administrators can restrict what Interface Groups are available to users.

Sensor tabs provide data specific to the selected Sensor.

Overview

The Overview tab contains a self-refreshing table with real-time system parameters collected from all active WANGUARD components.

Console

The Console System table has the following format:

Status	If the Console is functioning properly, a green "checked" arrow is displayed. If there's a red cross instead, (re)start the WANsupervisor daemon from the Console server.
Online Users	The number of active Console sessions.
Free Graphs Disk	The disk space available on the partition configured to store IP graphs.
Free DB Disk	The disk space available on the partition that is configured to store the database.
DB Size	The amount of disk space used by the database.
DB Active Clients	The number of clients that are currently using the SQL server.
DB Active Connections	The number of active connections on the SQL server.
Avg DB Queries/s	The average number of database queries per second reported by the SQL server.
Load	The load of the operating system for the last 5 minutes.
RAM	The amount of RAM used by PHP processes.
Started	The date when the Console's database server started.

Active Virtual Sensors

The Virtual Sensors table has the same fields as the Sniffing Sensors table explained below. The table is not displayed if there are no Virtual Sensors running.

Active Sniffing Sensors

The Active Sniffing Sensors table is not displayed if there are no Sniffing Sensors running. The table has the following format:

Status	If the active Sniffing Sensor is functioning properly then a green “checked” icon is displayed. If Console cannot manage or reach the Sniffing Sensor then a red “X” icon is displayed. In this case make sure that Sniffing Sensor is configured correctly, make sure that the WANsupervisor daemon is running and look for errors in the Events – see page 58.
Sensor Name	Displays the name of the Sniffing Sensor and a colored box with the Graph Color as defined in its configuration. Click it to open a new tab with data specific to the Sensor. Administrators and Operators can right-click it to open the Sensor's configuration.
IPs	The number of IP addresses that sent or received traffic. Only your network's IP addresses are counted.
Pkts/s (In / Out)	The inbound and outbound packets/second throughput after validation.
Inbound Bits/s	The inbound bits/second throughput after validation, and inbound usage percent.
Outbound Bits/s	The outbound bits/second throughput after validation, and outbound usage percent.
Received Pkts/s	The rate of sniffed packets before validation.
Dropped	The rate of packets dropped in the capturing process. When the number is high, it indicates a sniffing performance problem.
Load	The load of the operating system for the last 5 minutes.
CPU%	The CPU percent used by the Sniffing Sensor process.
RAM	The amount of memory used by the Sniffing Sensor process.
Started	The date when the Sniffing Sensor started.

Active Flow Sensors

The Active Flow Sensors table is not displayed if there are no Flow Sensors running. The table has the following format:

Status	If the active Flow Sensor is functioning properly then a green “checked” icon is displayed. If Console cannot manage or reach the Flow Sensor then a red “X” icon is displayed. In this case make sure that Sniffing Sensor is configured correctly, make sure that the WANsupervisor daemon is running and look for errors in the Events – see page 58.
Sensor Name	Displays the name of the Flow Sensor. Click it to open a new tab with data specific to the Sensor. Administrators and Operators can right-click it to open the Sensor's configuration.
Interface	The interface name and a colored box with the configured Graph Color. If the interface names are missing after more than 2 minutes after the Sensor started, please check that the flow exporter's clock is synchronized with the server.
IPs	The number of IP addresses that sent or received traffic through the interface. Only your network's IP addresses are counted.
Pkts/s (In / Out)	The inbound and outbound packets/second throughput after validation.
Inbound Bits/s	The inbound bits/second throughput after validation, and inbound usage percent.
Outbound Bits/s	The outbound bits/second throughput after validation, and outbound usage percent.
Flows/s	The rate of flows per second received by the Flow Sensor.
Flows Delay	Because traffic data must be aggregated first, flow devices export flows with a configured delay. Some devices export flows much later than the configured delays, and this field contains the maximum flows delay detected by the Flow Sensor. Flow Sensor cannot run with delays over 5 minutes. To minimize the RAM usage and the performance of the Flow Sensor process, the flows must be exported as soon as possible.
Dropped	The number of unaccounted flows. If the number is high it indicates a performance problem on the Sensor or a network connectivity issue with the flow exporter.
Load	The load of the operating system for the last 5 minutes.
CPU%	The CPU percent used by the Flow Sensor process.
RAM	The amount of memory used by the Flow Sensor process.
Started	The date when the Flow Sensor started.

Active Filters

The Active Filters table is not displayed if there are no Filters running. The table has the following format:

Status	If the active Filter is functioning properly then a green “checked” icon is displayed. If Console cannot manage or reach the Filter then a red “X” icon is displayed. In this case make sure that Filter is configured correctly, make sure that the WANsupervisor daemon is running and look for errors in the Events – see page 58.
Filter Name	Displays the name of the Filter

Anomaly#	If the Filter mitigates an anomaly it contains the link to the Anomaly Report. Otherwise will display the message "Filter offline".
IP Address	The IP address from your network involved in the traffic anomaly. If the IP address is clicked then a new tab opens with data specific to the IP address.
IP Group	The IP Group of the IP address.
Decoder	The traffic type that exceeded the threshold: <i>TCP+SYN, TCP, UDP, ICMP, OTHER</i> .
IPs	The number of unique IP addresses detected making traffic with the attacked IP address.
Pkts/s	The total packets/second throughput towards the attacked IP address.
Bits/s	The total bits/second throughput towards the attacked IP address.
Dropped	It represents the rate of packets dropped in the capturing process. When the number is high it indicates a performance problem.
Load	The load of the operating system for the last 5 minutes.
Peak CPU%	The maximum CPU percent used by the Filter process.
RAM	The amount of memory used by the Filter process.
Start Time	The date when the Filter started to mitigate the anomaly.

Sensors

When you click a Sensor's name anywhere in the Console, the Sensor's tab is opened. The Sensor tab includes few sub-tabs located on the bottom side. All sub-tabs use the following common toolbar fields:

- **Sensors**
Select the Sensors you're interested in or "All" to select all Sensors. Multiple selections can be made. Administrators can filter what Sensors are available to users.
- **Time Frame**
Select predefined time-frames or enter your own by selecting "Custom...".

Sensor Dashboard

The Sensor Dashboard allows you to group the most relevant data a Sensor can give you to a single tab.

The Sensor Dashboard's configuration does not apply to a particular Sensor. The changes you make here will be visible for each Sensor.

The operation of Dashboards is documented in the Reports » Dashboards chapter on page 16.

Sensor Graphs

Sensor Graphs allows you to generate various Sensor-related histograms for the selected Sensor(s):

- **Data Units**

Select one or more parameters:

- *Default* – Shows most used parameters, each one in a different graph.
- *Packets* – The packets/second rate.
- *Bits* – The bits/second throughput.
- *Distribution* – Sensors can collect protocols distribution data for: HTTP, HTTPS, SMTP, POP3, IMAP, SNMP, FTP, SSH, TELNET, SQL, NETBIOS, MS-DS, MS-RDP, DNS, ICMP, OTHERS.
- *Bytes* – The bytes/second throughput.
- *IPs* – The number of IP addresses that sent or received traffic. Only your network's IP addresses are counted. A spike in this graph usually means that an IP class scan was performed.
- *Received frames* – For Sniffing Sensors it represents the rate of received packets before validation. For Flow Sensors it represents the rate of received flows before validation.
- *Dropped frames* – For Sniffing Sensors it represents the number of packets dropped in the capturing process. When the number is high, it indicates a sniffing performance problem. For Flow Sensors it represents the number of unaccounted flows. If the number is high it indicates a performance problem on the Flow Sensor or a network connectivity issue with the flow exporter.
- *Unknown frames* – For Sniffing Sensors it represents the rate of invalidated packets. For Flow Sensors it represents the rate of invalidated flows.
- *Unknown Sources*
The number of source IP addresses that didn't pass validation.
- *Unknown Destinations*
The number of destination IP addresses that didn't pass validation.
- *Avg Packet Size*
The average packet size: bits/packet.
- *CPU%*
The CPU percent used by the Sensor process.
- *RAM*
The amount of memory used by the Sensor process.
- *Load*
The load of the operating system for the last 5 minutes.
- *IP Graphs*
The number of updated IP graphs files.

- *IP Accounting*
The number of IP accounting records updated.
- *HW Graphs*
The number of files updated for traffic profiling files.
- *IP Graphs Time*
The number of seconds needed to update the IP graphs files.
- *HW Graphs Time*
The number of seconds needed to update the traffic profiling files.
- *Processing Time*
The number of seconds needed to perform traffic analysis functions.
- *IP Structures*
The number of internal IP structures.
- **Graphs Size**
You can select predefined sizes or you can enter your own size in the “<X pixels> x <Y pixels>” format.
- **Graphs Title**
Graphs can have an automatically-generated title for the “Default” option, no title for the “None” option, or you can enter your own text that will be rendered as a title.
- **Graph Legend**
Select the details of the graph's legend.
- **Consolidation**
If you are interested in spikes, select the *MAXIMUM* aggregation type. If you are interested in average values, select the *AVERAGE* aggregation type. If you are interested in low values, select the *MINIMUM* aggregation type.
- **Graph Options**
 - *Stack Sensors* – If unchecked, each selected Sensor generates different graphs. If checked, all selected Sensors generate single graphs that contain combined data
 - *Show Totals* – If multiple sensors are used, render the total of Data Units

Sensor Tops

Sensor Tops allows you to generate various Sensor-related tops for the selected Sensor(s):

- **Top Type**
 - *Talkers* – the IPs of your network that generate most traffic for the selected Decoder
 - *IP Groups* – the IP Groups that generate most traffic for the selected Decode
 - *TCP Ports* – the most used TCP ports

- *UDP Ports* – the most used UDP ports
- *IP Protocols* – most used IP protocols
- *IP Versions* – IPv4 and IPv6
- *AS Numbers* – the Autonomous Systems that generate most traffic. Available only for Flow Sensors
- **Decoder**
Select the decoder that analyzes the traffic you're interested in.
- **Direction**
The direction of the traffic: *Inbound* or *Outbound*.
- **Group Sensors**
If unchecked, each Sensor generates a different top. If checked, all selected Sensors are combined in a single top instead.
- **DNS**
Check this if you need reverse DNS resolution for IP addresses. This might slow down the top generation.

The number of top items and decoders can be modified in the Storage & Graphs Configuration, see page 32.

Generating tops for many Sensors and large time-frames may take minutes. It may require the increase of *max_execution_time* parameter from *php.ini*.

List Flows

This is available only for Flow Sensors.

You can list and filter the flow data for the Flow Sensor. The options are documented on page 10 in the Flow Collector chapter.

Flows Tops

This is available only for Flow Sensors.

You can process and filter the flow data to generate tops for the Flow Sensor. The options are documented on page 10 in the Flow Collector chapter.

Sensor Events

The list of events generated by the selected Sensor(s) for the selected time-frame. Events are explained in the Events Reporting chapter, see page 58.

Anomalies Overview

Here you can view trends and summarizations of attacks detected by Sensor(s), for the selected time-frame and decoders.

Reports » IP Addresses & Groups

This chapter describes how to generate complex traffic reports for IP addresses, IP subnets and IP Groups.

The **Reports » IP Addresses** panel allows the quick generation of IP traffic reports by entering the IP / CIDR in the upper side of the Panel, or by selecting an IP class or host from the expandable tree below.

The **Reports » IP Groups** panel lists all IP Group names that exists in IP Zones. You can search or filter them by entering a sub-string contained in the IP Group's name you're interested in. Use IP Groups to generate reports for clients that have multiple allocated IP classes. You just have to define those IP classes with the same IP Group name.

If the reports are empty, check if the IP addresses and subnets have in IP Zones the "IP Accounting" parameter and "IP Graphs" parameter set to *Yes*.

Clicking IP Addresses or IP Groups opens the same type of tab that contains few sub-tabs on the bottom side. All sub-tabs use the following common toolbar fields:

- **Sensors**
Select the Sensors you're interested in or "All" to select all Sensors. Multiple selections can be made. Administrators can filter which Sensors are available to users.
- **Time Frame**
Select predefined time-frames or enter your own by selecting "Custom...".

IP Dashboard

The IP Dashboard allows you to group the most relevant data for IPs, subnets and IP Groups to a single tab.

The IP Dashboard's configuration does not apply to a particular IP, subnet or IP Group. The changes you make here will be visible for each IP, subnet or IP Group.

The operation of Dashboards is documented in the Reports » Dashboards chapter on page 16.

IP Graphs

IP Graphs allows you to generate various histograms for the IP classes, host or IP Group:

- **Decoders & Data Unit**
Select the decoders that analyze the traffic you're interested in. Data Units available: *Packets*, *Bits* and *Bytes*.
- **Graphs Size**
You can select predefined sizes or you can enter your own size in the "<X pixels> x <Y pixels>" format.
- **Graph Title**

Graphs can have an automatically-generated title for the “Default” option, no title for the “None” option, or you can enter your own text that will be rendered as a title.

- **Graph Legend**

Select the details of the graph's legend.

- **Consolidation**

If you are interested in spikes, select the *MAXIMUM* aggregation type. If you are interested in average values, select the *AVERAGE* aggregation type. If you are interested in low values, select the *MINIMUM* aggregation type.

- **Graphs Stacking**

- *Stack Sensors* – If unchecked, each selected Sensor generates different graphs. If checked, all selected Sensors generate single graphs that contain combined data
- *Stack Decoders* – When is checked, the graphs will contain data from all selected decoders
- *Stack IPs* – Un-check this option if you want a different traffic graph displayed for every IP address contained in the IP class or IP Group. Use carefully because when this option is used with a /24 CIDR then 256 traffic graphs are displayed, one for each IP address in the “C” class
- *Stack Conflicts* – If decoders can be included one within the other (e.g. TOTAL contains TCP that contains HTTP and HTTPS), the graph will display stacked decoders to show the most specific ones. This generates both accurate and intuitive traffic graphs. In the example, TOTAL will show as TOTAL OTHER and TCP as TCP OTHER. But when you select TCP, HTTP and TCP+SYN as decoders, then the TCP+SYN decoder can be included both in TCP and HTTP, thus generating a decoder conflict. Check this to stop detecting conflicts between decoders, but keep in mind that graphs might not be as accurate
- *Stack Recursively* – When is checked, subnet graphs can be created contained IPs graphs.

The number of decoders, Data Units and aggregation types can be modified in the Storage & Graphs Configuration, see page 32.

IP Accounting

IP Accounting allows you to generate various traffic accounting reports for the IP class, host or IP Group:

- **Decoders & Data Unit**

Select the decoders that analyze the traffic you're interested in. Data Units available: *Packets*, *Bits* and *Bytes*.

- **Report Type**

Select the interval you want to aggregate the accounting data: *Daily*, *Weekly*, *Monthly*, *Yearly*.

- **Sum IPs**

Un-check this option if you want a different traffic accounting report displayed for every IP address contained in the IP class or IP Group. Use carefully because when this option is used with a /24 CIDR then 256 traffic accounting reports are displayed, one for each IP address in the “C” class.

- **Sum Sensors**

If unchecked, each Sensor generates a different traffic accounting report. If checked, all selected Sensors generate a single traffic accounting report that contains the summed traffic accounting data.

The number of decoders can be modified in the Storage & Graphs Configuration, see page 32.

List Flows

You can list and filter the flow data for the IP class, host or IP Group. The options are documented on page 10 in the Flow Collector chapter.

This is available only if there is at least one configured Flow Sensor.

Flows Tops

You can process and filter the flow data to generate tops for the IP class, host or IP Group. The options are documented on page 10 in the Flow Collector chapter.

This is available only if there is at least one configured Flow Sensor.

Profile Graphs

Profile graphs are used by WANGUARD to detect anomalies through traffic profiling. To view them, profile Anomalies must be enabled in the Anomalies Configuration – see page 33 and activated for the IP subnet.

Anomalies Overview

Here you can view trends and summarizations of attacks from/to the IP addresses, classes or IP Groups, for the selected time-frame and decoders.

Installation Guide

WANGUARD can be installed on common server hardware, provided that the system requirements listed later in this chapter are met. If you have some basic Linux or FreeBSD operation skills then no training is required for the software installation. Feel free to contact our support team for any issues.

Installing WANGUARD does not generate any negative side effects on your network's performance. Installation and configuration may take less than an hour; after that your network will be monitored and protected immediately. No baseline data gathering is required.

System Requirements

WANGUARD 5.x has been tested with the following distributions: **Red Hat Enterprise Linux 5.x or 6.x** (commercial Linux distribution), **CentOS 5.x or 6.x** (free, Red Hat Enterprise Linux based distribution), **OpenSUSE 12.x** (free, Novell Enterprise Linux based distribution), **Debian Linux 6.0** (free, community supported distribution), **Ubuntu 12.x**. Other distributions may work but haven't been tested yet.

The WANGUARD architecture is completely **scalable**. By installing the software on better hardware, the number of monitored and protected endpoints and networks increases. All WANGUARD components can be installed on a single server if enough resources are provided (RAM, CPU, Disk Space, Network Cards). You can also install the components on multiple servers distributed across your network.

Sniffing Sensor Hardware Requirements

Sniffing Capacity	1 Gigabit Ethernet	10 Gigabit Ethernet
Architecture	x86 (32 or 64 bit)	x86 (64 bit)
CPU	1 x Xeon 2.0 GHz	2 x Xeon 2.8 GHz
RAM	500 MB	1 GB
Network Cards	1 x Gigabit Ethernet with NAPI support 1 x Fast Ethernet for management	1 x 10 GbE card. Chipset 82599 is recommended 1 x Fast Ethernet for management
Operating System	RHEL 5 / CentOS 5, RHEL / CentOS 6, Debian 6, Ubuntu Server 12, OpenSUSE 12	RHEL 5 / CentOS 5, RHEL / CentOS 6, Debian 6, Ubuntu Server 12, OpenSUSE 12
Disk Space	10 GB (including Operating System)	10 GB (including Operating System)

Flow Sensor Hardware Requirements

Flow-processing Capacity	20 monitored interfaces, 15k active endpoints
Architecture	x86 (32 or 64 bit)
CPU	1 x Xeon 2.0 GHz
RAM	4 GB
Network Cards	1 x Fast Ethernet
Operating System	RHEL / CentOS 5, RHEL / CentOS 6, Debian 6, Ubuntu Server 12, OpenSUSE 12
Disk Space	15 GB (including Operating System)

Filter Hardware Requirements

Filtering Capacity	1 Gbps link	10 Gbps link
Architecture	x86 (32 or 64 bit)	x86 (64 bit)
CPU	1 x Xeon 2.0 GHz or 1 x Opteron 1.8 GHz	4 x Xeon 2.4 GHz
RAM	2 GB	4 GB
Network Cards	1 x Gigabit Ethernet with NAPI support 1 x Fast Ethernet for management	1 x 10 GbE Cards with 82599 chipset 1 x Fast Ethernet for management
Operating System	RHEL 5 / CentOS 5, RHEL / CentOS 6, Debian 6, Ubuntu Server 12, OpenSUSE 12	RHEL 5 / CentOS 5, RHEL / CentOS 6, Debian 6, Ubuntu Server 12, OpenSUSE 12
Disk Space	10 GB (including Operating System)	10 GB (including Operating System)

Having a dedicated filtering server for each monitored link is not always required. You can deploy a single filtering server that will protect multiple links, as long as you can re-route the traffic towards it and re-inject the cleaned traffic to a downstream router. For very large networks, a dedicated filtering server for each upstream link is highly recommended.

Console Hardware Requirements

Capacity	< 5 Managed Sensors
Architecture	x86 (32 or 64 bit)
CPU	1 x Xeon 2.4 GHz or 1 x Opteron 1.8 GHz

Memory	1 GB
Network Cards	1 x Gigabit Ethernet
Operating System	RHEL / CentOS 5, RHEL / CentOS 6, Debian 6, Ubuntu Server 12, OpenSUSE 12
Software Packages	apache 2.x+ php 5.2+ mysql 5.x rrdtool 1.3+ perl-rrdtool perl-DBD-MySQL ping, whois, traceroute, telnet, wireshark, tcpdump
Disk Space	10 GB (including Operating System) + additional storage when storing IP graphs data

To access the web interface provided by Console, one of the following web browsers is required (other should also work but have not been tested): Google Chrome, Firefox 3.5 or later, Safari 3.0 or later. Internet Explorer has a slow javascript engine and a non-standard behavior so it's not recommended. For the best Console experience we highly recommend the Chrome or Firefox and a 1280 x 1024 pixels or higher resolution display.

The web browser must have javascript and cookies support activated. Java support and Flash are not required. To access the Contextual Help you must install Adobe PDF Reader.

Software Installation & Download

Software installation instructions are listed and updated on the Andrisoft website for RedHat-based, SuSE-based and Debian-based Linux distributions.

You can try a fully functional version of WANGUARD for 30 days. You can switch to a full-time, registered version by applying a purchased license key.

Opening Console for the first time

Console is the web interface through which you will control and monitor all other components. If you followed correctly the installation instructions, from now on you will only need to log in to Console to manage WANGUARD.

To log in to Console open <http://<hostname>/wanguard>. If the page cannot be displayed, make sure the Apache web server is running and the firewall does not block incoming traffic on port 80.

Licensing Procedure

If you haven't licensed WANGUARD yet, you will be asked to do so. You must upload the *andrisoft.key* file we sent you by email by clicking the key icon. The license key contains encrypted information about the licensed capabilities of the software. You can switch between WANGUARD and WANSIGHT solely by changing the license key.

Log in to Console using the default username / password combination of **admin / changeme**.

To understand how to navigate within the Console, please read the chapter from page 6.

If the Console is installed on a public server, you should immediately change the default password for the "admin" account. To do that click the **Admin** menu on the top-right part of the window and select <<Change Password>>.

Quick Configuration Steps

- Estimate storage requirements, review decoders and graphs parameters – page 32
- Setup the anomalies detection parameters and decoders – page 33
- Configure the reaction to traffic anomalies – page 34
- Add your IP address ranges and important IPs to an IP Zone – page 40
- Add anomaly detection for prefixes, create Thresholds Templates – page 41
- Add and configure a Sensor then set it as Active – page 42
- Watch for errors in Events. Receive error notifications by email – page 43
- Generate Reports and send them periodically by email – page 57
- Create your own Dashboards and add relevant widgets – page 16
- Create new accounts for your staff or customers – page 59
- Configure new BGP Connections, if needed – page 51
- Configure Filters, if needed – page 53

Storage & Graphs Configuration

An important step in configuring WANGUARD is to make sure that the involved servers have enough resources to process and withhold traffic information. Most resource-related parameters are found in Configuration » Global Settings » Storage & Graphs.

The default paths for **collected flows** and **packet dumps** exist only on the Console's filesystem. When the Sensors are installed on different systems, you should export these paths towards the Console through NFS. If you don't, the Console won't be able to display the collected data.

In a later chapter you'll be able to configure the Sensors to generate traffic graphs for lots of IPs, depending on the size of the monitored network. If you intend to use this feature then look carefully at the IP graphs parameters. Changing these parameters later requires the recreation of all IP graphs.

IP graph files are stored on the Console's filesystem. There are 2 different methods for updating IP graph files, so select the appropriate tab:

- **Write IP graph files directly on disk**

This method creates one file for every IP address directly on the defined Graphs Disk Path. RRDCache daemon optimizes I/O access but you must add it's path (usually it's `unix:/var/rrdtool/rrdcached/rrdcached.sock`) and you must configure it first.

The first accuracy parameter or "Archive" (default is 5 minutes) specifies the granularity of the graphs for recent data. It can be set as high as 5 seconds and as low as 10 minutes. The averages and intervals values specify the accuracy / granularity and for how long do you want the data to be stored.

This method is not suited for updating tens of thousands of IP graphs with very high granularity.

- **Write IP graph files in RAM or SSD first**

This method is suited for high granularity IP graphs. It creates a file for every IP address on RAM or SSD and updates it there. The files are moved periodically on a larger but much slower disk.

Decoders determine the underlying protocols of each packet or flow. Enabling too many decoders might cause a performance penalty, but you will be able to better differentiate the traffic.

Consolidation functions build consolidated values for Archives. If you are interested in traffic spikes, check MAXIMUM. If you are interested in average values, check AVERAGE. For low traffic values, check MINIMUM.

All IP graph parameters have a direct impact on the storage space required in the Console's file system. The *Disk space required for each IP graph file* value will be updated when you click the <<Update>> button. If you change the graphs parameters, make sure you delete old .rrd files from the defined **Graphs Disk Path**.

The deletion of old data can be automated in Configuration » Global Settings » Data Retention.

Anomalies Configuration

An important initial step in configuring WANGUARD is setting up anomalies detection parameters and decoders.

Anomalies detection parameters are located in Configuration » Global Settings » Anomalies. The Sensors are able to detect many types of traffic anomalies.

There are 2 categories / classes of traffic anomalies:

- **Threshold Anomalies**

Are detected when a previously user-defined packets/s or bits/s rate threshold (absolute value or percentage) has been reached. The traffic can be differentiated through decoders. Enable only the decoders for the traffic for which you will apply thresholds. Decoders determine the underlying protocols of each packet or flow.

- **Profile Anomalies**

Are detected through a behavioral recognition approach. The Sensors detect any activity that deviates from the "normal" traffic received by the protected subnets.

After enabling Profile Anomalies detection for a subnet, the Sensor builds a behavioral traffic graph for about 25 hours. Note that Profile Anomalies detection is suited for hosts and subnets that have a predictable traffic pattern. Larger subnets are usually the more predictable.

False positives can be limited by adjusting the deviation percent and minimum packets and bits rates.

Traffic anomalies detection will be enabled individually for each subnet when configuring IP Zones on page 40.

Response Configuration

Responses provide a unique and powerful way to automate reactions to traffic anomalies and attack patterns. To add a Response, go to Configuration » Network & Policy » Add Response. If you don't plan to use this feature, you may safely skip this chapter.

When invoked by a Sensor or Filter, a Response runs the contained **Actions**. These are modules that provide means to execute various commands, send notifications, write logs and more. There are 2 types of Actions:

- **Anomaly Actions**

Are executed by Sensor, for each traffic anomaly while the anomaly is active and once when it expires.

- **Pattern Actions**

Are executed by WANGUARD Filter, for each attack pattern while the pattern is active and once when it expires. Attack patterns are malicious packets that share common layer 3-7 fields (attacker IPs, TCP / UDP ports, length, protocols, TTL, contents etc.).

To modify, delete or rename an Action you must select the Action's description in the left section. The <<List Prefixes>> button allows you to see what IP classes are configured to use the Response.

All Actions have their specific fields together with the following common fields:

- **Active** selects if the Action is enabled or disabled.
- **Priority** selects the order of execution relative to the other Actions that are defined within the same panel. Lower numerical values correspond to increased priority.
- **Description** is the name or description of the Action.
- **Execution**. The Actions in the "While... is detected" panel can be executed every 5 seconds, depending on the Sensor's Accuracy parameter. The name of the Action is visible in the Console if **Log Execution** is checked.
- **Conditions** are rules that must be passed before the Action is executed. Each Condition is formed from a **Conditional Parameter**, a comparison function and a user-defined value. Conditional Parameters are dynamic, internal parameters that are updated every 5 seconds by Sensors and Filters.

Dynamic Parameters are parameters defined within curly brackets - { and } that can be included in the body of most Actions. Every Conditional Parameter has a correspondence with a Dynamic Parameter.

Using the Custom Script Action together with Dynamic Parameters you can extend and customize the reaction to anomalies and patterns.

Conditional & Dynamic Parameters

#	CONDITIONAL PARAMETER	TYPE	DYNAMIC PARAMETER	DESCRIPTION
GENERAL PARAMETERS				
1	IP Address	String	{ip}	The IP or Subnet involved in the anomaly.
		String	{ip_dns}	The reverse DNS of the IP involved in the anomaly. It's {ip} if the lookup is not successful.
2	CIDR	Number	{cidr}	The CIDR (prefix mask) of the IP or Subnet involved in the anomaly.
3	Prefix	String	{prefix}	The IP/CIDR involved in the anomaly.
4	IP Group	String	{ip_group}	The IP Group of the IP or Subnet involved in the anomaly.
5	Sensor Name	String	{sensor}	The Sensor's name.
6	Sensor Group	String	{sensor_group}	The Sensor's Interface Group.
7	Sensor IP	String	{sensor_ip}	The IP of the server running the Sensor.
8	Sensor Type [sniff,flow,virtual]	String	{sensor_type}	It's "sniff" for the Sniffing Sensor, "flow" for the Flow Sensor, or "virtual" for the Virtual Sensor.
9	Sensor ID	Number	{sensor_id}	The unique ID of the Sensor.
10	Flow Exporter IP	String	{router_ip}	The Flow exporter's IP. Empty when using the Sniffing Sensor.
11	IP Zone Name	String	{ipzone}	The IP Zone used by the Sensor.
12	Response Name	String	{response}	The Response used for the anomaly.
13	Template Name	String	{template}	The Template that defined the anomaly's triggering rule, if any.
14	Expiration Delay (seconds)	String	{expiration}	The number of seconds between the last time the anomaly is detected and the time the anomaly is expired.
15	Captured Packets	Number	{captured_pkts}	The number of captured packets during the Response, if any.
16	BGP Log Size (bytes)	Number	{bgplog_bytes}	The size of the BGP announcements logs.
17	Unique Dynamic Parameters	String	{exclusive}	The Unique Dynamic Parameters contain Dynamic Parameters that must be unique for the validation of an Action.
ANOMALY PARAMETERS				
1	Anomaly Description	String	{anomaly}	A description of the anomaly.

2	Anomaly ID	Number	{anomaly_id}	The unique identification number of the anomaly.
3	Anomaly Comment	String	{comment}	The comment added in Console by Administrators for the Anomaly.
4	Direction [incoming,outgoing]	String	{direction}	The direction of the rule that triggered the anomaly. Can be "incoming" or "outgoing".
5	Domain [ip,subnet]	String	{domain}	It's "ip" when CIDR = 32 for IPv4 or 128 for IPv6. It's "subnet" otherwise.
6	Anomaly Class [thresholds,profile]	String	{class}	It's "thresholds" for threshold-based anomalies or "profile" for profiling-based anomalies.
7	Threshold Type [absolute,percentage]	String	{threshold_type}	Threshold-based anomalies can be defined as "absolute" values or as a "percentage" of the total traffic received by the Sensor.
8	Anomaly Decoder (Protocol) [TOTAL,...]	String	{decoder}	The traffic decoder (protocol) for the detected anomaly.
9	Comparison [above,under]	String	{operation}	It's "above" for exceeding thresholds or "under" for the less-than-expected thresholds.
10	Unit [pkts/s,bits/s]	String	{unit}	It's "pkts/s" for packets per second anomalies or "bits/s" for bits per second anomalies.
11	Threshold Value	Number*	{rule_value}	It's the threshold value configured for the threshold.
12	Computed Threshold	Number*	{computed_threshold}	It's the threshold of the anomaly. It's dynamically adjusted for profiling-based and percentage-based anomalies.
13	Peak Value	Number*	{value}	The highest value of the traffic decoder for "above" thresholds or the lowest value for the "under" thresholds.
14	Latest Value	Number*	{latest_value}	The latest value given by the traffic decoder that detected the anomaly.
15	Sum Value	Number*	{sum_value}	The sum of the values given by the traffic decoder as long as the anomaly is active.
16	Peak Rule Severity	Number	{severity}	The field contains the ratio between the peak anomalous traffic rate and the threshold value.
17	Latest Rule Severity	Number	{latest_severity}	The field contains the ratio between the latest anomalous traffic rate and the threshold value.
18	Peak Link Severity	Number	{link_severity}	The field contains the ratio between the peak anomalous traffic rate and the interface's traffic rate.

19	Latest Link Severity	Number	{latest_link_severity}	The field contains the ratio between the latest anomalous traffic rate and the interface's traffic rate.
		String	{anomaly_log_10}	The first 10 packets or flows of the anomalous traffic.
		String	{anomaly_log_50}	The first 50 packets or flows of the anomalous traffic.
		String	{anomaly_log_100}	The first 100 packets or flows of the anomalous traffic.
		String	{anomaly_log_500}	The first 500 packets or flows of the anomalous traffic.
		String	{anomaly_log_1000}	The first 1000 packets or flows of the anomalous traffic.
OVERALL TRAFFIC PARAMETERS				
1	Peak TOTAL Pkts/s	Number*	{total_pps}	The peak packets/s throughput of the IP or Subnet, for all traffic.
2	Peak TOTAL Bits/s	Number*	{total_bps}	The peak bits/s throughput of the IP or Subnet, for all traffic.
3	Latest TOTAL Pkts/s	Number*	{latest_total_pps}	The latest packets/s throughput of the IP or Subnet, for all traffic.
4	Latest TOTAL Bits/s	Number*	{latest_total_bps}	The latest bits/s throughput of the IP or Subnet, for all traffic.
5	TOTAL Packets	Number*	{sum_total_pkts}	The sum of packets of the IP or Subnet, for all traffic, during the anomaly.
6	TOTAL Bits	Number*	{sum_total_bits}	The sum of bits of the IP or Subnet, for all traffic, during the anomaly.
TIME-RELATED PARAMETERS				
1	From (unixtime)	Number	{from_unixtime}	The time in unixtime format when the traffic anomaly started.
2	Until (unixtime)	Number	{until_unixtime}	The time in unixtime format when the traffic anomaly expired.
3	From (ISO 8601)	String	{from}	The time in iso8601 format when the traffic anomaly started.
4	Until (ISO 8601)	String	{until}	The last time in iso8601 format when the traffic anomaly expired.
5	Duration (seconds)	Number	{duration}	The number of seconds the anomaly was active.
6	Internal Ticks	Number	{tick}	Internal tick parameter. The Sniffing Sensor increments the value every 5 seconds while the anomaly is being detected.
FILTERS PARAMETERS				
1	Number of Filters	Number	{filters}	The number of WANGUARD Filters activated for the anomaly.

2	Filters Pkts/s	Number*	{filters_pps}	The latest packets/second throughput recorded by active Filter(s) in the anomalous traffic.
3	Filters Bits/s	Number*	{filters_bps}	The latest bits/second throughput recorded by active Filter(s) in the anomalous traffic.
4	Filters Max Pkts/s	Number*	{filters_max_pps}	The maximum packets/second throughput recorded by active Filter(s) in the anomalous traffic.
5	Filters Max Bits/s	Number*	{filters_max_bps}	The maximum bits/second throughput recorded by active Filter(s) in the anomalous traffic.
6	Filtered Packets	Number*	{filters_filtered_packets}	The number of packets filtered by active Filter(s).
7	Filtered Bits	Number*	{filters_filtered_bits}	The number of bits filtered by active Filter(s).
8	Filters CPU Usage	Number	{filters_max_cpu_usage}	The maximum CPU% used by Filter(s).
FILTER PARAMETERS				
1	Filter #	Number	{filter_id}	The unique ID of the attack pattern.
2	Filter Type (<i>ip, source, dest, proto, len, ttl..</i>)	String	{filter_type}	The attack pattern type: - ip (attacker's IP address) - source (source port of the attacker) - dest (destination port of the victim) - proto (the IP Protocol field) - len (the size of the packets) - ttl (the TimeToLive field) - others.
3	Filter Value	String	{filter_value}	The attack pattern's value.
		String	{filter_ip_dns}	If the attack pattern is an IP, the Dynamic Parameter provides the reverse DNS of the IP.
4	Filter Pkts/s	Number*	{filter_pps}	The attack pattern's latest packets/second throughput.
5	Filter Bits/s	Number*	{filter_bps}	The attack pattern's latest bits/second throughput.
6	Filter Peak Pkts/s	Number*	{filter_max_pps}	The maximum packets rate matched by the attack pattern.
7	Filter Peak Bits/s	Number*	{filter_max_bps}	The maximum bits rate matched by the attack pattern.
8	Filter Severity	Number	{filter_severity}	The severity field represents the ratio between attack pattern traffic and threshold values.
9	Filter Packets	Number*	{filter_packets}	The number of packets matched by the attack pattern.
10	Filter Bits	Number*	{filter_bits}	The number of bits matched by the attack pattern.
11	Filter Time Interval (seconds)	Number	{filter_difftime}	The duration of the attack pattern.
12	Filter Whitelist	Number	{filter_whitelisted}	If the attack pattern is whitelisted, the value is 1. Otherwise it's 0.
13	Filter Traffic Sample Size (bytes)	Number *	{filter_log_size}	Attack pattern traffic sample size.

	String	{attacker_isp}	If the attack pattern is an IP, the Dynamic Parameter provides the email of the attacker's ISP.
	String	{filter_log_10}	The first 10 packets of the attack pattern's traffic.
	String	{filter_log_50}	The first 50 packets of the attack pattern's traffic.
	String	{filter_log_100}	The first 100 packets of the attack pattern's traffic.
	String	{filter_log_500}	The first 500 packets of the attack pattern's traffic.
	String	{filter_log_1000}	The first 1000 packets of the attack pattern's traffic.

* Numerical values can be returned in multiples of 1,000 by appending `_kilo` to the Dynamic Parameter. The same goes for 1,000,000 by appending `_mega` and 1000,000,000 by appending `_giga`. To get the biggest multiplier (k,M,G) for the value, append `_prefix`.

IP Zone Configuration

IP Zones are hierarchical, tree-like structures that must include your IP address ranges and important IPs. Add an IP Zone by going to Configuration » Network & Policy » Add IP Zone. Sensors use IP Zones to learn about your network and to extract per-subnet settings. An IP Zone may be used by multiple Sensors, but a Sensor can only use one IP Zone.

To change the name of an IP Zone you must first open the IP Zone Configuration window, provide a new description and then press <<Change Name>>.

To copy the selected IP Zone you must click the <<Duplicate>> button. A new IP Zone will be created and it will have the same information and the same description with the word “(copy)” attached. In some cases when you have multiple Sensor systems, you may have to create multiple IP Zones that share the same prefixes. Instead of recreating the same IP classes for each new IP Zone you can duplicate an existing IP Zone and modify only few parameters.

To delete an IP Zone you must first open the IP Zone Configuration window, press <<Delete>> button and then confirm the deletion.

The IP Zone Configuration window is divided in two vertical sections. In the upper side of the left section there are buttons to manage Prefixes (IP address ranges or individual IPs). When adding a new Prefix, the tree below is automatically updated. The right section contains panels with user-provided settings for the selected Prefix.

WANGUARD understands IPs and IP classes entered in the CIDR notation. To enter individual hosts in IP Zones you must use the /32 CIDR for IPv4 and /128 for IPv6. For more about CIDR notation you can consult the Appendix 1 – Network Basics You Should Be Aware Of on page 60.

Every IP Zone contains at least the 0.0.0.0/0 network. Because it has the /0 CIDR it contains all IP addresses available for both IPv4 and IPv6. To ease the configuration, every new Prefix that you define, inherits by default the properties of the closest (having the biggest CIDR) IP class that includes it.

The **IP Settings** panel on the right section contains the following configurable parameters:

- **IP Group** combo box should contain a short description for the selected Prefix. Setting the same IP Group for more than one subnet will allow you to easily generate combined Reports.
- **IP Graphs**. If set to “Yes”, then the Console will collect graphs data for every IP contained in the selected IP class.
- **IP Accounting**. If set to “Yes”, then the Console will save daily accounting data for every IP contained in the selected IP class.

Enabling IP Graphs and IP Accounting for very large Prefixes (e.g. 0.0.0.0/0) is probably going to generate useless data and overload the system.

The **Comments** panel allows you to write a comment for the selected Prefix. It's not visible elsewhere.

Anomaly detection settings & Thresholds Templates

Thresholds Anomalies panel in the IP Zone Configuration window can contain user-defined traffic thresholds. To ease the addition of thresholds with the same values for multiple prefixes, use Thresholds Templates (Configuration » Network & Policy » Add Thresholds Template).

Thresholds have the following parameters:

- **Domain.** Sensors can detect anomalies to / from an IP contained in the subnet, or to / from the whole subnet.
- **Direction.** The direction of the analyzed traffic: can be "receives" for inbound traffic or "sends" for outbound traffic.
- **Comparison.** Select "over" for volumetric anomalies (e.g. DrDoS, DDoS) or "under" to detect the lack of traffic towards a monitored subnet or server.
- **Value.** Write the threshold value as an absolute number or as a percent of the total traffic received by the Sensor. Absolute values can be multiples of 1000 with K (kilo) appended, multiple of 1 million with M (mega) appended, or multiple of 1 billion with G (giga) appended.
- **Decoder.** Select one of the Decoders enabled in the Anomalies Configuration window – see page 33.
- **Unit.** DDoS attacks reach an unusually high number of packets per second, so select "pkts/s" to detect them. For bandwidth-related anomalies select "bits/s".
- **Response.** Select a previously defined Response or select "None" if you're not interested in reacting to the anomaly.
- **Parent.** Select "Yes" if the threshold should be inherited by more specific subnets. You can cancel inherited thresholds by selecting "Unlimited" in the Value field.
- **Inheritance.** Shows who's the parent Prefix, if any.

Adding a threshold on 0.0.0.0/0 that reads "Any IP receives over 5% TCP+SYN pkts/s" will catch ports scans and all SYN attacks. Threshold on 0.0.0.0/0 "Subnet receives under 5M TOTAL bits/s" will trigger the Response when the monitored link is down. You can configure "illegal" IP address ranges that should never be seen in normal traffic, for example, unallocated IP addresses or part of your internal IP address range that is unoccupied. You then add small thresholds to them to catch malicious activities such as scans and worms.

Profile Anomalies panel contains the Profiling Data parameter that can have the following values:

- *Inherit* – inherit the value from the parent Prefix
- *No* – don't generate profiling data for the selected Prefix
- *For Subnet* – generate profiling data for the whole traffic received by the Prefix
- *For IPs* – use carefully as it will generate profiling data for every IP contained in the Prefix. Don't use this on large Prefixes, and never for 0.0.0.0/0

Choosing a method of traffic monitoring

This section explains the available methods you can use for traffic monitoring. Reading this chapter is strongly recommended, as it will help you understand how to deploy Sensor in your network.

The Sensor was designed to monitor the traffic from the smallest branch office with tens of endpoints to the largest enterprises with hundreds of thousands of endpoints.

Depending on your network topology and configuration, your needs and your hardware, you must choose between the 2 types of Sensors:

- **Sniffing Sensor for Port Mirroring (SPAN, Roving Analysis Port) or Network TAP or In-line deployment**

In switched networks only the traffic for a specific device is sent to the device's network card. If the Sensor system is not deployed in-line (in the main data-path) then a network TAP, or a switch or router that offers a "monitoring port" must be used. In this case, the network device sends a copy of data packets traveling through a port or VLAN to the monitoring port. A Sniffing Sensor inspects every packet it receives to do the traffic analysis.

Packet sniffing provides extremely fast and accurate traffic analysis and accounting results. The downside is that it needs fast CPUs and good NICs.

- **Flow Sensor for NetFlow® (v5,v7,v9 – jFlow, NetStream, cflowd) or sFlow® (v4,v5) or IPFIX**

Many routers and switches can collect IP traffic statistics on monitored interfaces, and later export those statistics as flow records, towards the Flow Sensor to do the actual traffic analysis.

Because the Flow protocol already perform pre-aggregation of traffic data, the flows of data sent to the monitoring server are much smaller than the monitored traffic. This makes the Flow Sensor a good option for monitoring remote or high-traffic networks. The downsides are that computing pre-aggregation of traffic data requires large amounts of RAM, it has up to 5 minutes delays, and the accuracy of traffic parameters is lower than when inspecting packets (especially when sampling is used).

Virtual Sensor aggregates Sniffing Sensors and Flow Sensors' Interfaces into a single anomaly detection domain. It disables the anomaly detection features of containing Sensors, and provides anomaly detection for the summed up traffic data.

In high availability scenarios it's recommended to use both methods of traffic capturing. Add a new Sensor by going to Configuration » Components » Add Sensor.

Comparison between Packet Sniffing and Flow Monitoring

The table below provides a quick comparison between the two available traffic capturing technologies. The hardware requirements for each method are different. We keep an updated hardware requirements list on our website.

	WANGUARD Sensor	
	Sniffing Sensor	Flow Sensor
Capturing Technology	- Port Mirroring (SPAN, Roving Analysis Port) - Network TAP - In-line Deployment	- NetFlow version 5, version 7, version 9 – including jFlow, NetStream, cflowd* - sFlow version 4, version 5* - IPFIX*
Maximum Traffic Capacity per Sensor	10 GigE >150,000 endpoints**	multiples of 10 Gbps >10,000 endpoints**
Anomaly Detection Time	<= 5 seconds	< flow export time + 5 seconds
IP Graphs Accuracy	>= 5 seconds	>= 20 seconds
Traffic Validation Options	IP classes, MAC addresses, VLANs, BPF	IP classes, Interfaces, AS Numbers, Ingress / Egress

* **Cisco Systems** (ASR 1000, ISR G1, ISR G2 - 800, 1800, 1900, 2800, 2900, 3800, 3900 -, 1700, 3660, Catalyst 4500, Catalyst 6500, Cisco 7500, 7600, 10000, 12000, ASR 9000, CRS-1, C3KX-SM-10G, XR12000), **Adtran** (NetVanta 3200, 3305, 4305, 5305, 1524, 1624, 3430, 3448, 3130, 340, and 344), **AlaxalA Networks** (AX7800R , AX7800S , AX7700R , AX5400S), **Alcatel** (OmniSwitch 6850 , OmniSwitch 9000), **Allied Telesis** (SwitchBlade 7800R series , SwitchBlade 7800S series , SwitchBlade 5400S series), **Brocade** (BigIron series, FastIron series, IronPoint series, NetIron series, SecureIron series, ServerIron series), **Barracuda** (Barracuda NG Firewall), **Comtec Systems** (!-Rex 16Gi & 24Gi & 24Gi-Combo), **Dell - Force 10 Networks** (PowerConnect 6200 series, PowerConnect 8200 series, E series), **D-Link** (DGS-3600 series), **Enterasys**, **Extreme Networks** (Alpine 3800 series, BlackDiamond 6800 series, BlackDiamond 8800 series, BlackDiamond 10808, BlackDiamond 12804C , BlackDiamond 12804R , Summit X450 Series , Summit i series), **Fortigate** (FortiSwitch series, FortiGate series), **Huawei**, **H3C**, **Hewlett-Packard** (ProCurve 2610 series, ProCurve 2800 series , ProCurve 2900 series, ProCurve 2910al series, ProCurve 3400cl series , ProCurve 3500yl series , ProCurve 4200vl series , ProCurve 5300xl series , ProCurve 5400zl series , ProCurve 6200yl series , ProCurve 6400cl series , ProCurve 6600 series, ProCurve 8212zl, ProCurve Wireless Edge Services xl Module, ProCurve Wireless Edge Services zl Module, ProCurve Access Point 530ProCurve 9300m series , ProCurve Routing Switch 9408sl), **Hitachi** (GR4000 , GS4000 , GS3000), **Juniper Networks**, **Maipu** (S3300 Series, S3400 Series, S3900 Series), **MikroTik**, **NetGear** (GSM7352S-200, GSM7328S-200), **Nortel** (5500 & 8600 Series), **NEC** (IP8800/R400 series , IP8800/S400 series , IP8800/S300 series), **Palo Alto**, **Riverbed**, **Sonicwall** (SonicWall NSA E5500), **Vyatta** (Vyatta 514, Vyatta 2500 series, Vyatta Virtual Router, Firewall, VPN)

** An endpoint is an IP address that belongs to your network. The software is not limited by the number of connections between IPs.

Sniffing Sensor Configuration

In switched networks only the traffic for a specific device is sent to the device's network card. If the Sensor system is not deployed in the main data-path then a network TAP, or a switch or router that offers a "monitoring port" must be used. In this case, the network device sends a copy of data packets traveling through a port or VLAN to the monitoring port. The Sniffing Sensor inspects every packet it receives to do the traffic analysis.

For configuring Cisco switches please consult Catalyst Switched Port Analyzer (SPAN) Configuration Example on <http://www.cisco.com/warp/public/473/41.html>. To configure TAPs or other devices that support port mirroring, please consult the producer's documentation.

The Sniffing Sensor Configuration window contains the following fields:

- **Sensor Name**

A short name to help you identify the Sniffing Sensor.

- **Interface Group**

Optional description used within the Console to group multiple interfaces by location, roles etc.

- **Graph Color**

The color used in graphs for this Sensor. The default color is a random one, but you can change it. To change the color you can enter the color as a HTML Color Code or you can manually select the color by clicking the drop-down menu.

- **Sensor Type**

If the license key permits it, you can change the type of the Sensor.

- **Sensor Server**

The Server running the Sensor. To add a new one go to Configuration » Servers » Add Server.

- **Sniffing Interface**

The network interface that receives the traffic. The Linux network interface naming convention: eth0 for the first ethernet interface, eth1.900 for the second ethernet interface with 802.1Q VLAN 900 and so on.

If the Sniffing Sensor server is deployed in-line then it must contain the network interface that receives the traffic towards your network.

- **Link Speed IN / OUT**

The speed of the monitored link. If set, it is used to generate Reports based on usage percent.

- **IP Zone**

The Sensor must use an IP Zone to learn about your network and to extract per-subnet settings.

For more information about IP Zones please consult IP Zones Setup chapter on page 40.

- **IP Validation**

This option can be used to distinguish the direction of the packets or to ignore unwanted IP traffic:

- *Off* – The Sensor analyzes all traffic, but you must enable MAC Validation to distinguish the direction of traffic
- *On* – The Sensor analyzes the traffic that has the source and/or the destination IP in the selected IP Zone
- *Strict* – The Sensor analyzes the traffic that has either the source or the destination IP in the selected IP Zone
- *Exclusive* – The Sensor analyzes the traffic that has the destination IP in the selected IP zone but not the source IP

- **MAC Validation / Address**

This option can be used to distinguish the direction of the packets or to ignore unwanted OSI Layer 2 traffic:

- *None* – The Sensor analyzes all traffic, but you must enable IP Validation to distinguish the direction of traffic
- *Upstream MAC* – MAC validation is active and the MAC Address belongs to the upstream router
- *Downstream MAC* – MAC validation is active and the MAC Address belongs to the downstream router

The MAC Address must be written using the Linux convention – six groups of two hexadecimal values separated by colons (:).

- **BPF Expression**

You can filter the type of traffic the Sensor receives. Use BPF expressions or tcpdump-style syntax.

- **Use PF_RING**

Enable if you have PF_RING installed on the server. PF_RING provides high-speed packet analysis.

- **Top N**

Allows extended traffic tops generation. Enabling this causes a minor performance penalty.

- **Comments**

Comments about the Sensor configuration can be saved here. Not visible elsewhere.

To start the Sniffing Sensor click gray square button from the Side Region.

After setting a Sensor as Active, you should see if it starts properly by watching the Events – see page 58.

If the Sniffing Sensor starts without errors, but you can't see any data after more than 10 seconds, please check the following:

- ✓ You have correctly configured the switch/TAP to send packets to the server on the configured interface.
- ✓ The server is receiving the packets packets through the configured interface. You can verify this with a tool

like *tcpdump*. The syntax is “`tcpdump -i <interface_usually_eth1> -n -c 100`”.

- ✓ If the IP Validation is not disabled, then the IP Zone must contain all your subnets.

Flow Sensor Configuration

Many routers and switches can collect IP traffic statistics on monitored interfaces, and later export those statistics as flow records, towards the Flow Sensor to do the actual traffic analysis.

For detailed instructions on how to enable NetFlow, sFlow or IPFIX on your network device, please consult the vendor's website. An example on how to configure NetFlow on IOS, CatOS & Juniper can be found in Appendix 2 at page 63.

The Flow Sensor Configuration window contains the following fields:

- **Sensor Name**
A short name to help you identify the Flow Sensor.
- **Interface Group**
Optional description used within the Console to group multiple interfaces by location, roles etc.
- **Flow Protocol**
The type of flows exported towards the Sensor: NetFlow or IPFIX or sFlow.
- **IP Address**
The IP address of the router, switch, probe etc. Usually the Loopback0 address of the router. Each server running the Flow Sensor must have its system time synchronized with the flow exporter.
- **Sampling (1/N)**
Must contain the sampling rate as configured on the flow exporter. It's 1/1 if no sampling is used. NeFlow v9 contains the sampling value so in that case the value of this field will be ignored.
- **Time Zone**
The time-zone difference between the Console's server and the flow exporter.
- **Sensor Server**
The Server running the Sensor. To add a new one, go to Configuration » Servers » Add Server.
- **Listener IP/Port**
The IP address of the network interface that receives flows and the destination port.
- **Sensor Type**
If the license key permits it, you can change the type of the Sensor.
- **SNMP Community**
The read-only SNMP community of the network device allows Console to connect to the flow exporter and request SNMP indexes and other useful information for adding new interfaces.
- **Monitored Interfaces**
The list of interfaces that should be monitored. Add as few as possible. Settings per interface:

- *Interface Name* – A short description used to identify the monitored interface
- *Graph Color* – The color used in graphs for this interface. The default color is random, but you can change it. To change the color you can enter the color as a HTML Color Code or you can manually select the color
- *SNMP Index* – The interfaces are identifiable in flows only after their SNMP indexes
- *Traffic Direction* – The direction of the traffic entering the interface:
 - Select "Inbound" for upstream interfaces
 - Select "Outbound" for downstream interfaces
 - Select "Mixed" to establish the direction by IP / AS Validation
 - Traffic entering the "Null" interface is discarded by the router and by the Flow Sensor
- *Link Speed IN & Link Speed Out* – The speed of the interface. If set, it is used to generate Reports based on usage percent
- *Top N* – Allows extended traffic tops generation. Enabling this, RAM usage slightly increases
- **IP Zone**

The Sensor must use an IP Zone to learn about your network and to extract per-subnet settings. For more information about IP Zones please consult IP Zones Setup chapter on page 40.
- **IP Validation**

This option can be used to distinguish the direction of the traffic or to skip unwanted flows:

 - *Off* – The Sensor analyzes all traffic but the traffic direction must be established per interface
 - *On* – The Sensor analyzes the traffic that has the source and/or the destination IP in the selected IP Zone
 - *Strict* – The Sensor analyzes the traffic that has either the source or the destination IP in the selected IP Zone
 - *Exclusive* – The Sensor analyzes the traffic that has the destination IP in the selected IP zone but not the source IP
- **AS Validation**

Flows from BGP-enabled routers might contain the source and destination AS (Autonomous System) number. In most configurations, if the AS is set to 0 then the IP address belongs to your Autonomous System.

If enabled, only flows having AS number set to "0" (your AS) are processed. Rarely used to establish traffic direction.

AS Validation has three options:

 - *Off* – Will disable AS Validation
 - *On* – Only flows that have the source ASN and / or the destination ASN set to 0 are analyzed
 - *Strict* – Only flows that have either the source ASN or the destination ASN set to 0 are analyzed
- **Flow Collector**

All received flows can be stored in an efficient binary format and queried in Reports » Collections.

- **Graphs Accuracy**

Low values increase the Sensor's accuracy but the Flow Sensor will use more RAM.

- **Repeater IP:Port**

Send all incoming flows to another host and port by enabling the packet repeater.

- **Comments**

Comments about the Sensor configuration can be saved here. Not visible elsewhere.

To start the Sensor click gray square button from the Side Region.

After setting a Sensor as Active you should see if it starts properly by watching the Events – see page 58.

If the Flow Sensor starts without errors, but you can't see any data after more than 5 minutes, please check the following:

- ✓ You have correctly configured the flow exporter to send flows to the server for each of the configured interfaces.
- ✓ The server is receiving the flow packets on the configured port. You can verify this with a tool like *tcpdump*. The syntax is “`tcpdump -i <interface_usually_eth0> -n -c 100 udp and <destination_port>`”.
- ✓ The local firewall is allowing the Flow Sensor to receive the flow packets. You can check if the firewall is enabled with the *iptables* command. The syntax is “`iptables -L -n -v`”.
- ✓ Both the server and the flow exporter reside in the same time-zone. The clocks must be synchronized with NTP.
- ✓ The flow exporter's active/inactive flow timeout settings are less than 300 seconds. Flows sent with a delay of more than 300 seconds are automatically discarded.
- ✓ If you have “Mixed” interfaces then the IP Zone you have selected for the Flow Sensor must contain all your subnets.

Virtual Sensor Configuration

Virtual Sensor aggregates Sniffing Sensors and Flow Sensors' Interfaces into a single anomaly detection domain. It disables the anomaly detection features of containing Sensors, and provides anomaly detection for the summed up traffic data.

The Virtual Sensor Configuration window contains the following fields:

- **vSensor Name**
A short name to help you identify the Virtual Sensor.
- **Graph Color**
The color used in graphs for this Virtual Sensor. The default color is a random one, but you can change it. To change the color you can enter the color as a HTML Color Code or you can manually select the color by clicking the drop-down menu.
- **vSensor Server**
The Server running the Virtual Sensor. To add a new one go to Configuration » Servers » Add Server.
- **Link Speed IN / OUT**
The summed up speeds of the aggregated Sensors.
- **IP Zone**
The Virtual Sensor must use an IP Zone to extract per-subnet settings. For more information about IP Zones please consult IP Zones Setup chapter on page 40.
- **Comments**
Comments about the Sensor configuration can be saved here. Not visible elsewhere.

BGP Connection Configuration

Operators and Administrators can view, send and withdraw BGP announcements manually from the Console. The BGP announcements records are stored in Reports » Anomalies & Tools » BGP Prefixes » BGP Archive.

The Sensors and Filters can be configured to send and withdraw BGP announcements automatically, in the following cases:

- To protect networks by announcing upstream providers using a special BGP community, that your side does not route the attacked addresses anymore, or that they should null-route the announced addresses. This network protection technique is called black-holing.
- To divert DoS, DDoS and DrDoS traffic through a Filter system that will clean the traffic.

If you do not need any of those features you can safely skip this chapter.

BGP Connections can be used only after you have previously installed and configured the bgpd daemon included in the quagga (<http://www.quagga.net>) package. Some bgpd configuration steps can be found on Appendix 3 – Configuring Traffic Diversion at page 67.

The BGP Connection Configuration window contains the following fields:

- **BGP Connection Name**
A short name or a description for the BGP Connection.
- **BGPd Server**
The Server running the bgpd daemon. To add a new Server go to Configuration » Servers » Add Server.
- **AS Number**
Your AS number must match the one from the bgpd configuration.
- **Login Password**
The password needed to connect to the bgpd daemon.
- **Enable Password**
Configuration mode password of the bgpd daemon.
- **Route Map**
The route-map that should be appended to each announcement. This is not mandatory.
- **AS View**
If multiple AS views are defined in the bgpd configuration, then you must enter which view you want to use for this configuration. This is not mandatory.
- **Zebra Local Blackhole**
Check if you need the local black-hole feature quagga zebra. This is a rarely used feature.
- **Zebra Login & Enable Passwords**

The passwords for the zebra daemon.

- **Reject External IPs**

If this is selected no BGP announcement will be sent for an IP that's not present in any of the configured IP Zones' subnets, excluding 0.0.0.0/0.

- **Reject IPv4 under /**

You can restrict sending BGP prefixes that have the IPv4 CIDR mask less than the configured value. For example, a value of 32 rejects all prefixes that are not hosts.

- **Reject IPv6 under /**

You can restrict sending BGP prefixes that have the IPv6 CIDR mask less than the configured value. For example, a value of 128 rejects all prefixes that are not hosts.

- **Comments**

Comments about the BGP Router configuration can be saved here. Not visible elsewhere.

To enable the BGP Connection click gray square button from the Side Region.

Filter Configuration

WANGUARD Filter was designed to protect networks from internal and external threats (availability attacks on DNS, VoIP, Mail and similar services, unauthorized traffic resulting in network congestion), botnet attacks, zero-day worm and virus outbreaks. It includes sophisticated traffic analysis algorithms that are able to detect and side-filter malicious traffic in a granular manner without impacting the user experience or resulting in downtime.

It can be deployed in-line or out-of-line by diverting the malicious traffic towards the filtering server. Traffic diversion uses the `bgpd` daemon included in the `quagga` software package. An example on how to configure `bgpd` and network devices for traffic diversion is found on Appendix 3 – Configuring Traffic Diversion at page 67.

If you don't plan to use the Filter you can skip this chapter.

To manage Filters go to Configuration » Components. The Filter Configuration window contains the following fields:

- **Filter Name**

A short name that will help you to identify the Filter system.

- **Interface Group**

Optional description used within the Console to group multiple interfaces by location, roles etc.

- **Graph Color**

The color used in graphs for the Filter. The default color is a random one, but you can change it.

- **Filter Type**

Can only be “Activated by Response”, for now.

- **Filter Server**

The Server running the Filter. To add a new Server go to Configuration » Servers » Add Server.

- **Inbound Interface**

The network interface that receives the malicious traffic. If the Filter system is deployed in-line, then this is the interface that receives the traffic entering your network.

The network interface's name must use the interface naming conventions of the Linux operating system: `eth0` for the first interface, `eth1` for the second, `eth0.900` for the first interface with VLAN 900 and so on. If VLANs are used then you should configure them first using the `vconfig` command.

- **Sniffing Interface**

This switch configures the interface monitored by the Filter.

- *Inbound* – The Filter analyzes the traffic coming towards the Inbound Interface. The generated statistics are very accurate but the CPU usage is higher because the Filter continuously inspects malicious packets even if they are not being forwarded
- *Outbound* – The Filter analyzes only the traffic passing the Outbound Interface. Choosing this

option makes the Filter consume less CPU, because the malicious packets that are dropped do not reach the Outbound Interface. The disadvantage of this option is that the Filter will not record traffic statistics for the dropped traffic

- **Outbound Interface**

This field is optional if the Monitor Interface is set to “Inbound”. The cleaned traffic is sent to a downstream router through this network interface. The gateway must be reachable through this interface.

If GRE or IP over IP tunneling is being used, then you must firstly configure a virtual network interface with the *ip* command, part of the *iproute2* package.

- **SW Filtering Policy**

The Software Filtering Policy lets you select which software filters will be applied when the Filter detects an attack pattern. An attack pattern is formed by malicious packets that share some common OSI Layer 3, OSI Layer 4 or OSI Layer 5 fields. When an attack comes from a non-spoofed IP address, the attack pattern is the source IP address of the attacker. In case of a spoofed attack, the attack pattern could be the source TCP or UDP port, the destination TCP or UDP port, IP protocol number, packets size, TTL etc.

The Filter does inbound software-based traffic filtering and packet rate limiting using the Netfilter framework.

Available Filtering Policies are:

- *No software filtering* – The Filter detects and reports attack patterns. The Linux firewall API is not used, but you can implement other filtering commands using Responses Actions
- *Filter the attack patterns and forward valid traffic* – The Filter detects, reports and filters attack patterns. If an attack pattern is not whitelisted, then the whole traffic matched by the attack pattern is dropped. The rest of the traffic is forwarded through the Outbound Interface
- *Filter the attack patterns and forward limited valid traffic* – The Filter detects, reports and filters the attack patterns and forwards rate-limited valid traffic. If an attack pattern is not whitelisted, then the whole traffic matched by the attack pattern is dropped. The Filter system will not forward traffic that exceeds the anomaly's traffic type packets/second threshold value
- *Limit the attack patterns and forward valid traffic* – The Filter detects, reports and rate-limits the attack patterns to the threshold values. The Filter only forwards attack patterns traffic that does not exceed the anomaly's traffic type packets/second threshold value for the attacked IP address
- *Apply default FORWARDING policy* – The Filter detects and reports the attack patterns, and the default forwarding policy is applied. Netfilter is still being used, but all the rules have the “RETURN” target. This is used only mainly when debugging Netfilter rules
- *Filter the attack patterns and accept valid local traffic* – The Filter detects, reports and filters the attack patterns. If an attack pattern is not whitelisted, then the whole traffic matched by the attack pattern is dropped. The rest of the traffic is accepted towards the Inbound Interface
- *Filter the attack patterns and accept limited local valid traffic* – The Filter detects, reports and filters the attack patterns and accepts rate-limited local valid traffic. If an attack pattern is not whitelisted, then the whole traffic matched by the attack pattern is dropped. The Filter system will not accept traffic that exceeds the anomaly's traffic type packets/second threshold value for the attacked IP address

- *Limit the attack patterns and accept local valid traffic* – The Filter detects, reports and rate-limits the attack patterns to the threshold values. The Filter only accepts attack patterns traffic that does not exceed the anomaly's traffic type packets/second threshold value for the attacked IP address
- *Apply default INPUT policy* – The Filter detects and reports the attack patterns, and the default Netfilter INPUT policy is applied. Netfilter is still being used, but all the rules have the “RETURN” target. This is used only mainly when debugging Netfilter rules

- **HW Filtering Policy**

The Hardware Filtering Policy lets you select which hardware filters will be applied when the Filter detects an attack pattern.

- *No hardware filtering* – The Filter detects and reports attack patterns but no hardware-based filters are applied.
- *Intel X520 10 Gbps NIC - block only IPv4 sources* – The Filter blocks the sources of the attacks but only if the sources are non-spoofed IPv4 addresses. For protecting against attacks with random IP sources, also define a SW Filtering Policy.
- *Intel X520 10 Gbps NIC - block only IPv4 destinations* – The Filter blocks the IPv4 destinations of the attacks, for any attack patterns it finds. Also known as black-holing.

Intel X520 10 Gbps NIC - PF_RING – The Filter uses the PF_RING framework to apply the following hardware-base filtering rules: Source IP, Destination IP, TCP/UDP Source Port, TCP/UDP Destination Port, IP Protocol. Other attack patterns cannot be filtered by the X520 NIC.

- **Use PF_RING**

Enable if you have PF_RING installed on the server. PF_RING provides high-speed packet analysis.

- **Traffic Diversion**

The Traffic Diversion field provides a selection of currently defined BGP Connections that may be used for traffic diversion. When a BGP Connection is selected, the Filter sends a BGP announcement through it, so that the Filter system becomes the next-hop for the attacked IP address. When the attack ends, the Filter automatically withdraws the BGP announcement and the traffic towards the IP address will be routed normally.

For more information about defining BGP Connections, please consult the BGP Connection Setup chapter on page 51. If the Filter system is deployed in-line, or you don't plan to use traffic diversion, you can leave the BGP Connection field set to “None”.

- **Filters Timeout**

This field contains the number of seconds of inactivity required for the deletion of an attack pattern. If set to 0 then every attack pattern detected will not be deleted until the attack stops and the Filter becomes inactive. Usually, an attack pattern is associated with a filter (see Filtering Policy below).

- **Whitelists**

A Filter Whitelist is a collection of user-created rules that prevents the filtering of critical traffic types. If the filtering policy permits, the Filter could filter attack patterns that you really don't want to be filtered.

The Filter filters destination ports and destination IP addresses only in worst-case scenarios when no other attack pattern is detected. In some cases, it's best to let the malicious traffic enter the network than to filter some critical destination IPs and destination ports. For example, if your DNS server is being attacked by spoofed addresses on port 53 UDP, then the Filter might filter port 53 UDP traffic towards

your DNS server making your DNS partially unreachable. In this case, it's best to configure a Whitelist that will prevent this behavior.

To add a new rule to the Whitelist you must enter the following fields:

- **Subnet**
The attacked IP address should be included in this subnet. You can set this to 0.0.0.0/0 for generic whitelists.
- **Description**
Add a description, explanation or comment for the exception
- **Protocol**
You can choose what type of traffic the rule will match: *ANY, TCP, UDP, ICMP*.
- **Parameter**
Which traffic parameter should be compared: *IP Address, Source Port, Destination Port, Packet Length, IP Packet TimeToLive, IP Protocol Type*.
- **Operator**
Operators for strings and numbers: *equal, non-equal*. Operators for numbers: *less than, greater than*.
- **Value**
The user-defined value that should be compared.

When an attack pattern cannot be filtered because it conflicts with the Filter's whitelists, then the attack pattern is reported near a red exclamation point.

- **Comments**

Comments about the Filter configuration can be saved here. Not visible elsewhere.

To start a Filter click gray square button from the Side Region.

Scheduled Reports

One of the greatest strengths of the Console is the ease to generate complex Reports. Most Reports created through the Side Region can be printed, exported as PDF or sent by email. But if you want to create periodic Reports, go to Configuration » Schedulers » Add Report.

Through **Scheduled Reports** you can configure the Console to automatically generate Reports and send them by email to you or to your customers at preconfigured intervals of time.

You can include more than one email address in the **Email To** field separated by comma.

The emails are sent periodically according to the settings in the **Scheduler** tab.

To see how the email would look like without waiting for the pre-configured time, enter a description, your email address and then click the <<Save & Execute Now>> button. You should receive an email with the Report immediately.

The emails are formatted as HTML messages and include MIME attachments, so make sure to use compatible email clients.

Events Reporting

An "event" is a text message generated by a WANGUARD component and logged by the Console. To see the latest Events, raise the **South Side** by clicking the small bottom edge on the window.

The **Latest Events** tab lists the latest 30 events. On the side it contains a list of components that can generate events. The number on the right margin indicates the number of events in the last 24h, but no more than 30. The number's color indicates the maximum severity in the events: red means that there are ERRORS, blue is for INFO events etc.

Click a component to open a dedicated tab, then click the small down arrow on any column header to search, sort and filter Events.

Event's **severity** indicates the importance of the event:

- **MELTDOWN** – Meltdown events are generated when a very serious error is detected, such as a hardware error.
- **CRITICAL** – Critical events are generated when a significant software error is detected such as a memory exhaustion.
- **ERROR** – Error events are caused by misconfigurations or communication errors between WANGUARD components.
- **WARNING** – Warning events are generated when authentication errors occur, when there are errors updating files or when there are synchronization issues.
- **INFO** – Informational events are generated when configurations are changed or when users log in to the Console.
- **DEBUG** – Debug events are generated only for troubleshooting coding errors.

As an Administrator, you should keep Events with high severities under surveillance. Configure the Console to send Events periodically by email, syslog or SNMP in Configuration » Schedulers » Events Reporting.

Users Management

In the Side Region select Configuration » Global Settings » Users Management.

Currently there are three available access levels or "roles" for users:

- **Administrator** – Has all privileges. Can manage other accounts and can reset passwords. Passwords are immediately encrypted.
- **Operator** – Has all privileges except modifying other accounts or viewing the License Manager.
- **User** – Denies modifications and hides all configurations. Provides **permission-based access** to Reports, Dashboards, Interfaces, IP Groups, Regions etc.

To modify an account you must double-click it or select it and then press <<Modify User>>.

The **Full Name**, **Company**, **Position**, **Email**, **Phone** and **Comments** fields are optional.

The **Landing Tab** list contains the tabs that can be opened immediately after logging in. The list is dynamic and expands as you add Interfaces, Dashboards, IP Groups etc. You should change the Landing Tab to a relevant Dashboard or Report.

The **Minimum Severity** field selects the minimum severity level of the events that will be displayed in the Console.

The **Side Region Position** field lets you switch the Side Region's position between East and West.

The **Console Theme** field lets you switch the Console theme after re-logging in. Blue and Gray are the most popular themes.

Console Users » **Authentication** contains LDAP and RADIUS-based authentication settings.

You can enable cookie-based authentication by clicking the **Persistent Sessions** checkbox.

Appendix 1 – Network Basics You Should Be Aware Of

If you are new to network administration and network monitoring, read about the technical basics in this section. It will help you understand how WANGUARD works. If you are already used to IP addresses and IP classes you can safely skip this appendix.

IP Addresses

In order for systems to locate each other in a distributed environment, nodes are given explicit addresses that uniquely identify the particular network the system is on and uniquely identify the system to that particular network. When these two identifiers are combined, the result is a globally-unique address. This address, known as “IP address”, as “IP number”, or merely as “IP” is a code made up of numbers separated by three dots that identifies a particular computer on the Internet. These addresses are actually 32-bit binary numbers, consisting of the two sub addresses (identifiers) mentioned above which, respectively, identify the network and the host to the network, with an imaginary boundary separating the two.

An IP address is, as such, generally shown as 4 octets of numbers from 0-255 represented in decimal form instead of binary form.

For example, the address 168.212.226.204 represents the 32-bit binary number 10101000.11010100.11100010.11001100.

The binary number is important because that will determine which class of network the IP address belongs to. The Class of the address determines which part belongs to the network address and which part belongs to the node address (see IP address Classes further on).

The location of the boundary between the network and host portions of an IP address is determined through the use of a subnet mask. This is another 32-bit binary number which acts like a filter when it is applied to the 32-bit IP address. By comparing a subnet mask with an IP address, systems can determine which portion of the IP address relates to the network and which portion relates to the host. Anywhere the subnet mask has a bit set to “1”, the underlying bit in the IP address is part of the network address. Anywhere the subnet mask is set to “0”, the related bit in the IP address is part of the host address. The size of a network is a function of the number of bits used to identify the host portion of the address. If a subnet mask shows that 8 bits are used for the host portion of the address block, a maximum of 256 host addresses are available for that specific network. If a subnet mask shows that 16 bits are used for the host portion of the address block, a maximum of 65,536 possible host addresses are available for use on that network.

An Internet Service Provider (ISP) will generally assign either a static IP address (always the same) or a dynamic address (changes every time one logs on). ISPs and organizations usually apply to the InterNIC for a range of IP addresses so that all clients have similar addresses. There are about 4.3 billion IP addresses. The class-based, legacy addressing scheme places heavy restrictions on the distribution of these addresses. TCP/IP networks are inherently router-based, and it takes much less overhead to keep track of a few networks than millions of them.

IP Classes

Class A addresses always have the first bit of their IP addresses set to “0”. Since Class A networks have an 8-bit network mask, the use of a leading zero leaves only 7 bits for the network portion of the address, allowing for a maximum of 128 possible network numbers, ranging from 0.0.0.0 – 127.0.0.0. Number 127.x.x.x is reserved for loopback, used for internal testing on the local machine.

Class B addresses always have the first bit set to “1” and their second bit set to “0”. Since Class B addresses have a 16-bit network mask, the use of a leading “10” bit-pattern leaves 14 bits for the network portion of the address, allowing for a maximum of 16,384 networks, ranging from 128.0.0.0 – 181.255.0.0.

Class C addresses have their first two bits set to “1” and their third bit set to “0”. Since Class C addresses have a 24-bit network mask, this leaves 21 bits for the network portion of the address, allowing for a maximum of 2,097,152 network addresses, ranging from 192.0.0.0 – 223.255.255.0.

Class D addresses are used for multicasting applications. Class D addresses have their first three bits set to “1” and their fourth bit set to “0”. Class D addresses are 32-bit network addresses, meaning that all the values within the range of 224.0.0.0 – 239.255.255.255 are used to uniquely identify multicast groups. There are no host addresses within the Class D address space, since all the hosts within a group share the group’s IP address for receiver purposes.

Class E addresses are defined as experimental and are reserved for future testing purposes. They have never been documented or utilized in a standard way.

WANGUARD uses extensively, throughout its components, IP Addresses and IP Classes with the CIDR notation.

IPv4 Subnet CIDR Notation

CIDR	CLASS	HOSTS NO.	MASK
/32	1/256 C	1	255.255.255.255
/31	1/128 C	2	255.255.255.254
/30	1/64 C	4	255.255.255.252
/29	1/32 C	8	255.255.255.248
/28	1/16 C	16	255.255.255.240
/27	1/8 C	32	255.255.255.224
/26	1/4 C	64	255.255.255.192
/25	1/2 C	128	255.255.255.128
/24	1 C	256	255.255.255.000
/23	2 C	512	255.255.254.000
/22	4 C	1024	255.255.252.000
/21	8 C	2048	255.255.248.000
/20	16 C	4096	255.255.240.000
/19	32 C	8192	255.255.224.000
/18	64 C	16384	255.255.192.000
/17	128 C	32768	255.255.128.000
/16	256 C, 1 B	65536	255.255.000.000
/15	512 C, 2 B	131072	255.254.000.000
/14	1024 C, 4 B	262144	255.252.000.000
/13	2048 C, 8 B	524288	255.248.000.000
/12	4096 C, 16 B	1048576	255.240.000.000
/11	8192 C, 32 B	2097152	255.224.000.000
/10	16384 C, 64 B	4194304	255.192.000.000
/9	32768 C, 128B	8388608	255.128.000.000
/8	65536 C, 256B, 1 A	16777216	255.000.000.000
/7	131072 C, 512B, 2 A	33554432	254.000.000.000
/6	262144 C, 1024 B, 4 A	67108864	252.000.000.000
/5	524288 C, 2048 B, 8 A	134217728	248.000.000.000
/4	1048576 C, 4096 B, 16 A	268435456	240.000.000.000
/3	2097152 C, 8192 B, 32 A	536870912	224.000.000.000
/2	4194304 C, 16384 B, 64 A	1073741824	192.000.000.000
/1	8388608 C, 32768 B, 128 A	2147483648	128.000.000.000
/0	16777216 C, 65536 B, 256 A	4294967296	000.000.000.000

Appendix 2 – Configuring NetFlow Data Export

This appendix is a brief guide to setting up the NetFlow data export (NDE) on Cisco and Juniper routers or intelligent Cisco Layer 2 / Layer 3 / Layer 4 switches. If you have problems with the configuration contact your network administrator or Cisco consultant. For devices that run hybrid mode on a Supervisor Engine (Catalyst 65xx series) it is recommended to configure IOS NDE on the MSFC card and CatOS NDE on the Supervisor Engine. For more information about setting up NetFlow please visit <http://www.cisco.com/go/netflow>.

Configuring NDE on an IOS Device

In the configuration mode on the router or MSFC, issue the following to start NetFlow Export.

First enable Cisco Express Forwarding:

```
router(config)# ip cef
router(config)# ip cef distributed
```

And turn on flow accounting for each input interface with the interface command:

```
interface
ip route-cache flow
```

For example:

```
interface FastEthernet0
ip route-cache flow
interface Serial2/1
ip route-cache flow
```

It is necessary to enable NetFlow on all interfaces through which traffic (you are interested in) will flow. Now, verify that the router (or switch) is generating flow stats – try command 'show ip cache flow'. Note that for routers with distributed switching (GSR's, 75XX's) the RP cli will only show flows that made it up to the RP. To see flows on the individual line cards use the 'attach' or 'if-con' command and issue the 'sh ip ca fl' on each LC.

Enable the exports of these flows with the global commands:

```
router(config)# ip flow-export version 5
router(config)# ip flow-export destination <ip address> 2000
router(config)# ip flow-export source FastEthernet0
```

Use the IP address of your Flow Sensor's server and the configured listening port. UDP port 2000 is used as an example. The 'ip flow-export source' command is used to set up the source IP address of the exports sent by the equipment.

If your router uses the BGP protocol, you can configure AS to be included in exports with command:

```
router(config)# ip flow-export version 5 [peer-as | origin-as]
```

The following commands break up flows into shorter segments: 1 minute for active traffic and 30 seconds for inactive traffic. Please use only this values as it decreases the RAM usage and increases performance of the Flow Sensor.

```
router(config)# ip flow-cache timeout active 1
router(config)# ip flow-cache timeout inactive 30
```

In enable mode you can see current NetFlow configuration and state.

```
router# show ip flow export
router# show ip cache flow
router# show ip cache verbose flow
```

Configuring NDE on a CatOS Device

In privileged mode on the Supervisor Engine enable NDE:

```
switch> (enable) set mls nde <ip_address> 2000
```

Use the IP address of your Flow Sensor's server and the configured listening port. UDP port 2000 is used only as an example.

```
switch> (enable) set mls nde version 5
```

The following command is required to set up flow mask to full flows.

```
switch> (enable) set mls flow full
```

The following commands break up flows into shorter segments: ~1 minute for active flows and ~ 30 seconds for inactive flows. Please use only this values as it decreases the RAM usage and increases performance of Flow Sensor.

```
switch> (enable) set mls agingtime long 8
switch> (enable) set mls agingtime 4
```

If you want to account all traffic within the specified VLANs rather than inter VLAN traffic use CatOS 7.2 or higher and issue the following command:

```
switch> (enable) set mls bridged-flow-statistics enable
```

And enable NDE:

```
switch> (enable) set mls nde enable
```

To see current NetFlow configuration and state issue the following commands:

```
switch> (enable) show mls nde
switch> (enable) show mls debug
```

Configuring NDE on a Native IOS Device

To configure NDE use the same commands as for the IOS device. In the enable mode on the Supervisor Engine, issue the following, to set up the NetFlow export version 5.

```
switch(config)# mls nde sender version 5
```

The following commands break up flows into shorter segments: ~1 minute for active flows and ~ 30 seconds for inactive flows. Please use only this values as it decreases the RAM usage and increases performance of Flow Sensor.

```
switch(config)# mls aging long 8
switch(config)# mls aging normal 4
```

On the Supervisor Engine 1 issue the following to put full flows into the NetFlow exports:

```
switch(config)# mls flow ip full
```

If you have a Supervisor Engine 2 or 720 running IOS version 12.1.13(E) or higher, issue the following commands instead:

```
switch(config)# mls flow ip interface-full
switch(config)# mls nde interface
```

Configuring NDE on a 4000 Series Switch

Configure the switch the same as an IOS device, but instead of command 'ip route cache flow' use command 'ip route-cache flow infer-fields'. This series requires a Supervisor IV with a NetFlow Services daughter card to support NDE.

Configuring NDE on a Juniper Router

Juniper supports flow exports by the routing engine sampling packet headers and aggregating them into flows. Packet sampling is done by defining a firewall filter to accept and sample all traffic, applying that rule to the interface and then configuring the sampling forwarding option.

```
interfaces {
  ge-0/1/0 {
    unit 0 {
      family inet {
        filter {
          input all;
          output all;
        }
        address 192.168.1.1/24;
      }
    }
  }
}
firewall {
  filter all {
    term all {
      then {
        sample;
      }
    }
  }
}
```

```
                accept;
            }
        }
    }
}

forwarding-options {
    sampling {
        input {
            family inet {
                rate 100;
            }
        }
        output {
            cflowd 192.168.1.100 {
                port 2000;
                version 5;
            }
        }
    }
}
}
```

Appendix 3 – Configuring Traffic Diversion

This appendix describes how to configure traffic diversion for Filter. Information provided here regarding router configurations is for informational purposes only. Please refer to the appropriate router user guides for detailed information.


Understanding the BGP Diversion Method

Following standard Border Gateway Protocol (BGP) routing definitions, routers select the routing path with the longest matching prefix (also known as the “most specific”). After establishing a BGP session with the router, Filter sends a routing update where the Filter system is listed as the best path for the attacked destinations.

The network prefix that Filter announces is longer than the one already listed in the router’s routing table, overriding the router's routing table definition.

To configure traffic diversion in Layer 2 or Layer 3 network topologies, perform the following:

1. Configure traffic diversion using BGP
2. Configure the appropriate traffic forwarding method



The figure above provides an example of traffic diversion from Ingress Router 1,2,3 towards a Linux server running the Filter software.

After BGP diversion is established, the router's routing tables points to the Filter server as the best route to the attacked addresses and the router forwards all traffic destined to those addresses to the Filter server.

BGP Configuration Guidelines

This section provides general guidelines for BGP configuration on the Filter server and on a divert-from router.

The guidelines provided in this section apply to the BGP configuration on any router from which Filter system diverts the traffic. The following examples are provided using common External Border Gateway Protocol v4 (eBGP). You should consider the network configuration and determine whether eBGP or iBGP should be implemented in your network.

Follow these guidelines when the Filter system and adjacent routers operate using common eBGP:

1. Configure bgpd with an easy recognizable Autonomous System number.

The bgpd sends routing information only when it diverts traffic. This route appear in the router's routing tables. Using a recognizable value allows you to easy identify the Filter system in the router's routing tables.

2. To ensure that the bgpd routing information is not redistributed to other internal and external BGP neighboring devices, perform the following:

- Configure the bgpd not to send routing information and to drop incoming BGP routing information
- Set the bgpd BGP community attribute values to *no-export* and *no-advertise*.

A match in the community attributes enables bgpd to filter BGP announcements on the router and enforce this policy.

3. Enter the *soft-reconfiguration inbound* command during the setup procedures. This command is useful for troubleshooting and allows you to restore a routing table without reconnecting to neighboring device.

Filter System BGP Configuration

You must configure the BGP using the Zebra software (<http://www.zebra.org>) or the Quagga software (<http://www.quagga.net>). Quagga is a fork of Zebra and the differences are minimal. Quagga keeps it's configuration files in */etc/quagga* while Zebra keeps it's configuration files in */etc/zebra*.

After installing Quagga or Zebra, you will have to create some basic configuration files, so both zebra and bgp daemons could start. Setting the passwords for the two daemons is enough to get them started. You should change "zebrapass" and "bgppass" with your own passwords.

```
[root@localhost ~]# echo 'password zebrapass' > /etc/quagga/zebra.conf
[root@localhost ~]# echo 'password bgppass' > /etc/quagga/bgpd.conf
[root@localhost ~]# /etc/init.d/zebra start
[root@localhost ~]# /etc/init.d/bgpd start
```

It is a good idea to tighten the security in the zebra daemon. You must connect to the zebra daemon with telnet on localhost port 2601 (default zebra port) with the previously defined password ("zebrapass") and issue the

following commands:

```
[root@localhost ~]# telnet 127.0.0.1 2601
localhost> enable
localhost# config terminal
localhost(config)# service password-encryption
localhost(config)# write
localhost(config)# exit
localhost# exit
```

To configure the bgpd daemon you must telnet to port 2605 and enter the previously defined password (“bgppass”). You must then switch to the privileged mode by entering the *enable* command.

```
[root@localhost ~]# telnet 127.0.0.1 2605
localhost> enable
localhost#
```

Switch to terminal configuration mode by entering the *config terminal* command. The prompt will change indicating that the system has entered the configuration mode:

```
localhost# config terminal
localhost(config)#
```

You should then enable encrypted passwords and set a new password for the configuration mode:

```
localhost(config)# service password-encryption
localhost(config)# enable password enablepass
```

Configure routing on bgpd using the commands shown in the following example. Please note that you can use the prefix-list, route-map, or distribute-list method for filtering outgoing routing information about the router. The following example describes the distribute-list method. You can use the prefix-list or route-map filtering method types as long as the routing information is not sent to bgpd.

```
localhost(config)# router bgp <WANGUARD-Filter-AS-number>
localhost(config-router)# bgp router-id <WANGUARD-Filter-IP-address>
localhost(config-router)# neighbor <Router-IP-address> remote-as <Router-AS-number>
localhost(config-router)# neighbor <Router-IP-address> description <description>
localhost(config-router)# neighbor <Router-IP-address> soft-reconfiguration inbound
localhost(config-router)# neighbor <Router-IP-address> distribute-list nothing-in in
localhost(config-router)# neighbor <Router-IP-address> route-map WANGUARD-Filter-out out
localhost(config-router)# exit
localhost(config)# access-list nothing-in deny any
localhost(config)# route-map WANGUARD-Filter-out permit 10
localhost(config-route-map)# set community x:x no-export no-advertise
localhost(config-route-map)# exit
localhost(config)# write
localhost(config)# exit
```

Filter System BGP Configuration Example

To display the router configuration, enter the *show running-config* command from the “enable” command level. In the following example, the router's AS number is 1000, and the bgpd AS number is 64000.

The following partial sample output is displayed:

```
localhost# show running-config
... ..
```

```

router bgp 64000
  bgp router-id 192.168.1.100
  neighbor 192.168.1.1 remote-as 1000
  neighbor 192.168.1.1 description divert-from router
  neighbor 192.168.1.1 soft-reconfiguration inbound
  neighbor 192.168.1.1 distribute-list nothing-in in
  neighbor 192.168.1.1 route-map WANGUARD-Filter-out out
!
access-list nothing-in deny any
!
route-map WANGUARD-Filter-out permit 10
  set community 1000:64000 no-export no-advertise
!
line vty
... ..

```

Cisco Router BGP Configuration

This section describes the router's BGP configuration used when you configure traffic diversion. The syntax in the commands is taken from the BGP configuration on a Cisco router.

The following configuration steps shows the commands to use to configure BGP on a Cisco router:

```

r7500(config)# router bgp <Router-AS-number>
r7500(config-router)# bgp log-neighbor-changes
r7500(config-router)# neighbor <WANGUARD-Filter-IP-address> remote-as <WANGUARD-Filter-AS-
number>
r7500(config-router)# neighbor <WANGUARD-Filter-IP-address> description <description>
r7500(config-router)# neighbor <WANGUARD-Filter-IP-address> soft-reconfiguration-inbound
r7500(config-router)# neighbor <WANGUARD-Filter-IP-address> distribute-list
routesToWANGUARDFilter out
r7500(config-router)# neighbor <WANGUARD-Filter-IP-address> route-map WANGUARD-Filter-in in
r7500(config-router)# no synchronization
r7500(config-router)# exit
r7500(config)# ip bgp-community new-format
r7500(config)# ip community-list expanded <WANGUARD-Filter-community-name> permit no-export
no-advertise
r7500(config)# route-map WANGUARD-Filter-in permit 10
r7500(config-route-map)# match community <WANGUARD-Filter-community-name> exact match
r7500(config-route-map)# exit
r7500(config)# ip access-list standard routesToWANGUARDFilter
r7500(config-std-nacl)# deny any

```

The *no synchronization* command prevents the distribution of the bgpd routing updates into Interior Gateway Protocol.

Cisco Router BGP Configuration Example

To display the router configuration, enter the *show running-config* command from the router global command level. In the following example, the router's AS number is 1000 and the bgpd AS number is 64000.

The following partial output is displayed:

```

r7500# show running-config
... ..
router bgp 1000
  bgp log-neighbor-changes
  neighbor 192.168.1.100 remote-as 64000

```

```

neighbor 192.168.1.100 description Filter appliance
neighbor 192.168.1.100 soft-reconfiguration inbound
neighbor 192.168.1.100 distribute-list routesToWANGUARDFilter out
neighbor 192.168.1.100 route-map WANGUARD-Filter-in
no synchronization
!
ip bgp community new-format
ip community-list expanded WANGUARD-Filter permit 1000:64000 no-export no-advertise
!
route-map WANGUARD-Filter-in permit 10
match community WANGUARD-Filter exact match
ip access-list standard routesToWANGUARDFilter
deny any
...

```

Understanding Traffic Forwarding Methods

This section provides details on traffic forwarding methods. Traffic forwarding methods are used to forward the cleaned traffic from the Filter system to a downstream router.

The following terminology is used in this section:

- Divert-from router – Router from which the bgpd diverts the attacked destinations traffic.
- Inject-to router – Router where bgpd forwards the cleaned traffic towards attacked destinations.
- Next-hop router – Router that is the next-hop to the destinations according to the routing table on the divert-from router before traffic diversion is activated.

Static Routing – Layer 2 Forwarding Method

In a Layer 2 topology, the Filter system, divert-from router, and next-hop router are on the same network or VLAN. In a Layer 2 topology, a divert-from router and an inject-to router are two different devices. The next-hop router and the inject-to router are the same device.

GRE / IP over IP Tunneling – Layer 3 Forwarding Method

In a Layer 3 topology, the divert-from and inject-to routers are the same router (referred to as the router in this chapter). Filter sends a BGP announcement that modifies the router's routing table to divert the zone traffic to the Filter system. Filter cleans the traffic and returns the cleaned traffic to the same router. The divert-from router then sends the traffic to the router that appears as the best path to the zone. This process may result in a malicious routing loop. In this case you may have to use a tunnel that is configured between the Filter system and the next-hop router to forward clean traffic. The inject-to router does not perform routing decisions according to the zone address and forwards the packets to the next-hop router.

Configuring Static Routing – Layer 2 Forwarding Method

The Layer-2 Forwarding (L2F) method is used in a Layer 2 topology when all three devices—the Filter system, the divert-from router, and the next-hop router—are located in one shared IP network. In a Layer 2 topology, a divert-from router and an inject-to router are two separate devices. The next-hop router and the inject-to router are the same device.

The Filter system issues an ARP query to resolve the MAC address of the inject-to/next-hop router and then forwards the traffic. For this reason, no configuration on the routers is required when using the L2F method. The only thing you have to configure when using this method is the default gateway on the Filter system so that it points to the inject-to/next-hop router.

Configuring GRE / IP over IP Tunneling – Layer 3 Forwarding Method

In the tunnel diversion method, you configure a tunnel between the Filter server and each of the next-hop routers. The Filter's server sends the traffic over the tunnel that ends in the next-hop router of the destined zone. Because the returned traffic goes over a tunnel, the inject-to router performs a routing decision on the end point of the tunnel interface only, not on the zone's address.

To use this method you have to run the standard Linux tool *ip* to create and route GRE / IP over IP tunnels that will be used to inject the cleaned traffic back into the network. You must then configure the Filter (page 53) with the Outbound Interface set to the virtual network interface created by the tunnel.