



WANGUARD 5.4 User Guide

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

Copyright ©2014 Andrisoft SRL
All rights reserved.
Revision 2.50

Copyright & trademark notices

This edition applies to version 5.x 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. 2014. 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 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.

Sales: sales@andrisoft.com

Technical Support: support@andrisoft.com

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

© Copyright ANDRISOFT S.R.L. 2014. All rights reserved.

Table of Contents

1. IP Traffic Monitoring, Anomalies Detection & DDoS Mitigation with WANGUARD.....	5
WANGUARD Key Features & Benefits.....	5
WANGUARD Components.....	6
2. A first look at the WANGUARD Console.....	7
Side Region – used for navigation throughout the Console.....	7
Central Region – contains tabs with reports and dashboards.....	7
South Region – provides a quick look at the latest events, live statistics and graphs.....	7
Upper-right Menus – Help menu and User menu.....	7
3. Reports » Alarms & Tools.....	8
Anomalies.....	8
Active Anomalies.....	8
Anomalies Archive.....	10
Anomalies Overview.....	10
BGP Prefixes.....	10
BGP Operations.....	11
BGP Logs.....	11
Flow Collectors.....	12
Flows List.....	12
Flows Tops.....	12
Packet Analyzers.....	13
Active Captures.....	13
Captures Archive.....	15
4. Reports » Components.....	16
Overview.....	16
Console.....	16
Servers.....	17
Active Virtual Sensors.....	17
Active Sniffing Sensors.....	17
Active Flow Sensors.....	18
Active Filters.....	19
Sensors.....	20
Sensor Dashboard.....	20
Sensor Tops.....	20
Sensor Graphs.....	21
Flows List.....	22
Flows Top.....	23
AS Graphs.....	23
Country Graphs.....	23
Sensor Events.....	24
Anomalies Overview.....	24
Filters.....	24
Filter Dashboard.....	24
Filter Graphs.....	25
Filter Events.....	26
Filtering Rules Archive.....	26
5. Reports » Dashboards.....	27

6. Reports » IP Addresses & Groups.....	28
IP Dashboard.....	28
IP Graphs.....	28
IP Accounting.....	29
Flows List.....	29
Flows Tops.....	30
Profile Graphs.....	30
Anomalies Overview.....	30
7. Reports » Servers.....	31
Server or Console Dashboard.....	31
Server Graphs.....	31
Server or Console Events.....	32
Server Commands.....	32
8. Installation Guide.....	33
System Requirements.....	33
Sniffing Sensor – Minimum Hardware Requirements.....	33
Flow Sensor – Minimum Hardware Requirements.....	34
Filter – Minimum Hardware Requirements.....	34
Console Hardware Requirements.....	35
Software Installation & Download.....	35
Opening the Console for the First Time.....	35
Licensing Procedure.....	36
Quick Configuration Steps.....	36
9. Storage & Graphs Configuration.....	37
10. Anomalies Configuration.....	38
11. Response Configuration.....	39
List of Conditional Parameters and Dynamic Parameters.....	40
12. IP Zone Configuration.....	45
Anomaly Detection Settings & Thresholds Templates.....	46
13. Choosing a Method of Traffic Monitoring.....	48
Comparison between Packet-Based and Flow-Based Monitoring.....	48
14. Sniffing Sensor Configuration.....	50
Troubleshooting the Sniffing Sensor.....	51
15. Flow Sensor Configuration.....	53
Troubleshooting the Flow Sensor.....	55
16. Virtual Sensor Configuration.....	57
17. BGP Connection Configuration.....	58
18. Filter Configuration.....	60
19. Scheduled Reports.....	65
20. Events Reporting.....	66
21. Users Management.....	67
22. Appendix 1 – Network Basics You Should Be Aware Of.....	68
IPv4 Subnet CIDR Notation.....	70
23. Appendix 2 – Configuring NetFlow Data Export.....	71
Configuring NDE on an IOS Device.....	71
Configuring NDE on a CatOS Device.....	72
Configuring NDE on a Native IOS Device.....	72
Configuring NDE on a 4000 Series Switch.....	73

Configuring NDE on a Juniper Router (non-MX).....	73
24.Appendix 3 – Configuring Traffic Diversion.....	75
Understanding the BGP Diversion Method.....	75
BGP Configuration Guidelines.....	76
Filter System BGP Configuration.....	76
Filter System BGP Configuration Example.....	77
Cisco Router BGP Configuration.....	78
Cisco Router BGP Configuration Example.....	78
Understanding Traffic Forwarding Methods.....	79
Static Routing – Layer 2 Forwarding Method.....	79
GRE / IP over IP Tunneling – Layer 3 Forwarding Method.....	79
Configuring Static Routing – Layer 2 Forwarding Method.....	79
Configuring GRE / IP over IP Tunneling – Layer 3 Forwarding Method.....	80
25.Appendix 4 – Software Changelog.....	81
WANGUARD 5.4.....	81
WANGUARD 5.3.....	82
WANGUARD 5.2.....	83
WANGUARD 5.1.....	83
WANGUARD 5.0.....	84

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 interference of best-effort traffic with critical or real-time traffic. WANGUARD's network-wide surveillance of complex, multilayer, switched or 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

- **FULL NETWORK VISIBILITY** – Supports the latest IP traffic monitoring technologies: packet sniffing at 10 Gbps; NetFlow v5, v7 and v9; sFlow, IPFIX, NetStream, jFlow, cflowd.
- **DDOS DETECTION** – A fast traffic anomaly detection engine detects volumetric attacks by profiling the online behavior of users and by comparing over 120 live traffic parameters against user-defined thresholds.
- **DDOS MITIGATION** – Protects networks and services by detecting and cleaning malicious traffic on packet-scrubbing servers deployed in-line or out-of-line, or by using BGP black hole routing.
- **POWERFUL REACTION TOOLS** – Automate responses to threats using predefined or custom actions: send notification emails, announce prefixes in BGP, generate SNMP traps, modify ACLs, execute your own scripts with access to over 70 operational parameters through an easy-to-use API, etc.
- **DETAILED FORENSICS** – Samples of packets and flows for each attack are captured for forensic investigation. Detailed attack reports can be emailed to you, to affected customers or to attacker's ISP.
- **ADVANCED WEB CONSOLE** – Consolidated management and reporting through a single, interactive and configurable HTML5 web portal with customizable dashboards, user roles, remote authentication, etc.
- **PACKET SNIFFER** – Includes a distributed packet sniffer that can save packet dumps from different parts of the network. View packet details in a Wireshark-like web interface.
- **FLOW COLLECTOR** – Provides a fully featured NetFlow, sFlow, and IPFIX Analyzer and Collector that saves flows in a compressed format for long term storage. Flows can easily be searched, filtered and sorted.
- **COMPLEX ANALYTICS** – Generates the most complex reports with aggregated data for hosts, departments, interfaces, applications, ports, protocols, countries, autonomous systems, and more.
- **REAL-TIME REPORTING** – Bandwidth graphs are animated and have a short-term accuracy of just 5 seconds.
- **HISTORICAL REPORTING** – You can view reports from the last 5 seconds to the last 10 years by selecting any custom time period. Bandwidth histograms contain 95th-percentile values for burstable billing.

- **SCHEDULED REPORTING** – Any report can be generated automatically and emailed to interested parties at preconfigured intervals of time.
- **FAST & SCALABLE** – The software was designed to run on commodity hardware. Its components can be distributed on clustered servers.
- **THE LOWEST T.C.O.** – The most affordable on-premise DDoS protection solution on the market!

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

A first look at the WANGUARD Console

If you are an administrator seeking instruction on how to configure WANGUARD, see the “Installation Guide” chapter on page 33.

Please read the following chapters for 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 following cover the configuration of the solution.

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

Side Region – used for navigation throughout the Console

The Side Region is located on the east or west (default) edge of the window, according to the user's preference. If it is not visible, then has been 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 or by pressing Ctrl+R. Both sections contain multiple panels that can also collapse or expand, with such state being maintained 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 – contains tabs with reports and dashboards

The Console offers various ways to look at historic and live collected data. Each Report and Dashboard you open from the Side Region opens a tab (page) in the Central Region. You may switch between (sub-)tabs with (Alt+) Ctrl+→ and (Alt+) Ctrl+←. You can close all tabs except for the Landing Tab as defined by user preference. Initially, the Landing Tab is the Configuration Wizard.

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

The South Region is located at the bottom of the browser window. By default, it is collapsed; to expand it, click the small bottom edge or press Ctrl+E. This provides a quick way to view live data: events (system logs), animated graphs, anomalies, and statistics from all components.

Upper-right Menus – Help menu and User menu

The Help menu contains the User Manual, a few select tools and the About window. Dependent on 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 and the Console theme, and provides a Log Out option.

Reports » Alarms & Tools

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

Anomalies

The Anomalies tab provides live and historical data related to DoS and DDoS attacks or to other traffic anomalies.

The number of active traffic anomalies is displayed within the Reports » Alarms & Tools panel, and it is refreshed every 10 seconds. The color of the number reflects the highest severity of the active anomalies.

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

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	<p>The IP address or IP class of the traffic anomaly and the reverse DNS.</p> <p>In 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.</p> <p>Click it to open a new tab with data specific for the prefix.</p>
IP Group	<p>The IP group of the Prefix.</p> <p>Click it to open a new tab with data specific to the IP group.</p>
Anomaly	A short description of the anomaly.
Value	The peak value of the anomalous traffic. The latest value is displayed between parentheses.
Sensor	<p>The name of the Sensor that detected the anomaly.</p> <p>Click it to open a new tab with data specific to the Sensor.</p>
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	<p>Actions available for administrators, operators and unprivileged users with proper permissions:</p> <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 with the prefix • <i>Classify/Set Comment</i> – add or modify comments, or classify the impact of anomalies • <i>Manual Actions</i> – execute the response actions that have the condition: Manual Actions=yes. • <i>Expire the Anomaly</i> – force the Sensor to clear the anomaly. The Sensor must be running for the action to take effect.
Dropped	The percentage of anomalous traffic filtered by one or more WANGUARD 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 anomalous traffic and the overall traffic of the link (Sensor or interface).</p> <p>The exact rule severity and link severity are 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 percentage value between the anomaly traffic and the overall traffic.
Threshold	The threshold 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. This field is hidden if there are no comments.

When one or more Filters are activated to detect attackers and filtering rules, a new table appears in the same row as the traffic anomaly. The rows of the Filter table have a red background when the attack pattern is active, and a yellow background when the attack patterns are inactive.

Filter	The name of the Filter that detected the attack pattern. Click it to open a new tab with data specific to the Filter.
Filtering Rule	The filtering rule applied to drop the attack pattern. The Filter can dynamically apply the following filtering rules: <i>Source IP, Source Port, Destination Port, Packet Length, TimeToLive, IP Protocol</i> .

	Filtering rules are applied only when the filtering policy allows dropping of traffic. If the filtering rule conflicts with the Filter's Whitelist, then a white flag appears in the same row.
Firewall	Indicates if the filtering rule was applied by the software-based firewall or by the hardware-based firewall.
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 has 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 contains 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. The columns are explained in the previous paragraph.

The <Clear Active Anomalies> button clears all active anomalies directly from the Console, without any Sensor or Filter interaction.

Anomalies Overview

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

BGP Prefixes

The BGP Prefixes tab shows the prefixes announced by WANGUARD in the Border Gateway Protocol (BGP).

The number of active BGP announcements is displayed within the Reports » Alarms & Tools panel, and is refreshed every 10 seconds. The number is red when there is at least one active BGP announcement sent by a BGP connection configured for black-holing/null-routing, or blue when all active BGP announcements were sent by BGP connections configured for traffic diversion/off-ramping.

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

BGP Operations

BGP Operations provides live insights on BGP announcements sent by Sensors, Filters or Console.

Administrators and operators can send or withdraw BGP announcements manually. To send a new BGP announcement, click the top left button, enter the prefix and select a previously configured BGP connection.

The table below is visible when there is at least one active BGP announcement. The columns are:

BGP Connection	The BGP connection name as defined in the BGP connection's configuration (see page 58). When the grouping is set to "By BGP Connection" clicking the BGP connection's name will allow you to delete all announcements for that BGP connection with a single click.
Prefix	The prefix included the announcement. Individual IPv4 hosts have a /32 CIDR, IPv6 hosts have a /128 CIDR. When the grouping is set to "By IP/Mask" clicking the prefix will allow you to delete all announcements for that prefix with a single click.
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, the field contains the link to the detailed anomaly report. If there are multiple anomalies for the same prefix and BGP connection, they will be shown separately, but a single announcement is sent.
Comments	This field may contain user comments about the BGP announcement. If the field contains the word "ERROR" look for BGP Connection errors in Events (see page 66). If the field contains the word "Orphan" then the anomaly that triggered the announcement is no longer active, but the announcement still is. In this case you should withdraw the announcement manually.
Action	It contains a link for the manual removal of the BGP announcement. The Action field is visible for administrators and operators.

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. All columns are explained in the previous paragraph, except for the hidden User column that shows the account name of the user that sent the announcement.

You can modify the status of announcements manually by double-clicking the rows, or by clicking the <Clear Active Announcements> button. Both options only affect only the UI, not the BGPd configuration.

Flow Collectors

The **Reports » Alarms & Tools** panel contains a link to the **Flow Collectors** tab if at least one Flow Sensor was configured.

The number of active flow collectors is displayed within the Reports » Alarms & Tools panel, and it is refreshed every 10 seconds.

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

The Flow Collectors tab contains 3 sub-tabs located on the left bottom side of the window:

Flows List

You can list and filter the saved flow data by entering the fields below:

- **Sensors** – Select the Flow Sensor interfaces that captured the traffic you are 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 Filtering Expression** – Here you can enter a filtering expression for flows. Click the light bulb icon on the right to open a window that will show you the correct syntax. Frequently-used flows filters can be saved there and reused at a later time.
- **Output** – You can select several output formats, or you can type your own format. Click the light bulb icon on the left to open a window that will show you the correct syntax.

For better readability, the IPv6 addresses are shortened, such as that the middle nibbles are cut and replaced by ellipses (“...”). This is usually sufficient for recognition of a desired IPv6 address. If you need the full IPv6 address, check the option “IPv6 long”.

- **Export** – If the output is not very large, it can be emailed or printed.

If you need to list huge amounts of flow data, doing so solely from within the web browser may not be the best idea. In this case, select the “Dump” option to view the CLI command used to list the flows. 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 appropriate checkboxes, you can select how you want to have your flows aggregated. You may also aggregate entire subnets by 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 generate tops from the saved flow data by entering the fields below:

- **Sensors** – Select the Flow Sensor interfaces that capture the traffic you are 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 Filtering Expression** – Here you can enter a filtering expression for flows. Click the light bulb icon on the right to open a window that will show you the correct syntax. Frequently-used flows filters can be saved there and reused at a later time.
- **Output** – You can select several output formats, or you can type your own format. Click the light bulb icon on the left to open a window that will show you the correct syntax.

For better readability, the IPv6 addresses are shortened, such as that the middle nibbles are cut and replaced by ellipses (“...”). This is usually sufficient for recognition of a desired IPv6 address. If you need the full IPv6 address, check the option “IPv6 long”.

- **Export** – If the output is not very large, it can be emailed or printed.

If you need to list huge amounts of flow data, doing so solely from within the web browser may not be the best idea. In this case, select the “Dump” option to view the CLI command used to list the flows. 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 by selecting srcIPv4/<subnet bits>.
- **Limit** – Limit the output to only those statistic lines whose packets or bytes match the specified limit.
- **Top** – Limit the statistics to the first top N.

Packet Analyzers

The **Reports » Alarms & Tools** panel contains a link to the **Packet Analyzers** tab if at least one Sniffing Sensor was configured.

The number of active captures is displayed within the Reports » Alarms & Tools panel, and it refreshed every 10 seconds.

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 left side of the window:

Active Captures

Administrators, operators and unprivileged users with packet capturing privileges can generate packet dumps by clicking the <Add Traffic Capture> button. The options are:

- **Description** – A short description to help you identify the traffic capture.
- **Sniffing Sensors** – Select the Sniffing Sensors that will capture the traffic you are interested in. Multiple

selections can be made. Administrators can filter what Sensors are available to users.

- **BPF Expression** – Click the light bulb icon on the right to open a window containing the correct Berkley Packet Filter (BPF) syntax. Frequently used BPF expressions can be saved there for use at later time.

The use of a BPF expression is mandatory, but you can use the “ip” string when you want to capture all IP traffic.
- **Max Running Time** – The maximum running time.
- **Stop Capture On** – When 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 for 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, this 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, 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 are interested in.
- **Comments** – This field may contain comments about the traffic capture.

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

- **Description [BPF]** – The traffic capture's description and BPF expression.
- **Sampling** – The type of sampling being used.
- **From** – The date when the Sniffing Sensor started capturing packets.
- **Until** – The time or the conditions that will cause the Sniffing Sensor to stop capturing the traffic.
- **Status** – Indicates the status of the capture. It is green if the capturing thread is active.
- **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 to your computer. Click the third icon to stop the capture.

Captures Archive

The 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. The columns are explained in the previous paragraph.

Reports » Components

The **Reports » Components** panel contains links to the **Overview** tab and **Device Groups** tabs, and to detailed **Sensor** and **Filter** tabs.

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

The Device Groups tabs provide a real-time view of the status of the Sensor(s) and Filter(s) assigned to each device group. Administrators can restrict what device groups are available to unprivileged users.

Sensor tabs provide data specific to the selected Sensor, and Filter tabs provide data specific to the selected Filter.

Overview

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

Console

The Console System table has the following format:

Status	If the Console is functioning properly, a green check mark is displayed. If a red “X” is displayed instead, (re)start the WANsupervisor daemon on the Console server.
Online Users	The number of active Console sessions.
Free Graphs Disk	The disk space available on the partition that is 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.
Start Time	The date when the Console's database server started.

Servers

The Servers table has the following format:

Status	If the server is functioning properly, a green check mark is displayed. If a red “X” is displayed instead, (re)start the WANsupervisor daemon from the server and check the time synchronization with the Console's clock.
Server Name	Displays the name of the server and a colored box with the graph color as defined in configuration. Click to open a new tab with data specific to the server. Administrators and operators can right-click to open the server configuration.
Free RAM	The available RAM. The swap memory is not counted.
Load	The load reported by the operating system for the last 5 minutes.
CPU% Userspace	The percentage of CPUs used by the user space processes. Can be >100% on multiple cores/CPU's.
CPU% System	The percentage of CPUs used by the system. Can be >100% on multiple cores/CPU's.
CPU% Idle	The percentage of idle CPUs. Can be >100% on multiple cores/CPU's.
Free Flows Disk	The disk space available on the partition that is configured to store flows.
Free Dumps Disk	The disk space available on the partition that is configured to store packet dumps.
Contexts/IRQs/SoftIRQs	The number of context switches, hardware interrupts and software interrupts per second.
Uptime	The uptime of the server.

Active Virtual Sensors

The Virtual Sensors table has the same fields as the Sniffing Sensors table explained below. The table is not displayed if no Virtual Sensors are 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 check mark is displayed. If the Console cannot manage or reach the Sniffing Sensor, then a red “X” is displayed. In this case,
---------------	---

	make sure that the Sniffing Sensor is configured correctly, and the WANsupervisor daemon is running, and look for errors in the events log (see page 66).
Sensor Name	Displays the name of the Sniffing Sensor and a colored box with the graph color as defined in 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.
Pkts/s (In / Out)	The inbound and outbound packets/second throughput after validation.
Inbound Bits/s	The inbound bits/second throughput after validation and the inbound usage percent.
Outbound Bits/s	The outbound bits/second throughput after validation and the outbound usage percent.
Received Pkts/s	The rate of sniffed packets before validation.
IPs (Int / Ext)	The number of IP addresses that sent or received traffic. The Int(ernal)/Ext(ernal) IPs are the IPs inside/outside the IP Zone. The top generator value from the Sensor configuration enables or disables the monitoring of external IPs.
Dropped	The rate of packets dropped in the capturing process. A high number indicates a sniffing performance problem.
CPU%	The percentage of CPUs used by the Sniffing Sensor process.
RAM	The amount of memory used by the Sniffing Sensor process.
Start Time	The date when the Sniffing Sensor started.
Server	The server that runs the Sniffing Sensor. Click to open a new tab with data specific to the server. Administrators and operators can right-click to open the server configuration.

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 check mark is displayed. If the Console cannot manage or reach the Flow Sensor, then a red "X" is displayed. In this case, make sure that Flow Sensor is configured correctly and the WANsupervisor daemon is running, and look for errors in the events log (see page 66).
Sensor Name	Displays the name of the Flow Sensor. Click to open a new tab with data specific to the Sensor. Administrators and operators can right-click to open the sensor configuration.
Interface	The interface name and a colored box with the configured graph color. If the interface names are missing more than 2 minutes after the Sensor started, please check that the flow exporter's clock is synchronized with the server.
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.
IPs (Int / Ext)	The number of IP addresses that sent or received traffic. The Int(ernal)/Ext(ernal) IPs are the IPs inside/outside the IP Zone. The top generator value from the Sensor configuration enables or disables the monitoring of external IPs.
Flows/s	The rate of flows per second received by the Flow Sensor.
Flows Delay	<p>Because traffic data must be aggregated first, flow devices export flows with a configured delay. Some devices export flows much later than the configured delay and this field contains the maximum flows delay detected by the Flow Sensor.</p> <p>The Flow Sensor cannot run with delays over 5 minutes. To minimize the RAM usage and optimize the performance of the Flow Sensor process, the flows must be exported as soon as possible.</p>
Dropped	The number of unaccounted flows. A high number indicates a performance problem with the Sensor or a network connectivity issue with the flow exporter.
CPU%	The percentage of CPUs used by the Flow Sensor process.
RAM	The amount of memory used by the Flow Sensor process.
Start Time	The date when the Flow Sensor started.
Server	The server that runs the Flow Sensor. Click it to open a new tab with data specific to the server. Administrators and operators can right-click to open the server configuration.

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 check mark is displayed. If the Console cannot manage or reach the Filter, then a red "X" is displayed. In this case, make sure that the Filter is configured correctly and that the WANsupervisor daemon is running, and look for errors in the events log (see page 66).
Filter Name	Displays the name of the Filter. Click to open a new tab with data specific to the Filter. Administrators and operators can right-click to open the Filter configuration.
Anomaly#	When the Filter mitigates an anomaly, it contains the link to the anomaly report. Otherwise, the field displays the message "Filter offline".
Prefix	The IP address/mask from your network involved in the traffic anomaly. When the prefix is clicked, a new tab opens with data specific to the IP block or IP address.
IP Group	The IP group of the IP address.
Decoder	The decoder that is analyzing the traffic.

Pkts/s	The packets/second throughput towards the attacked IP address.
Bits/s	The bits/second throughput towards the attacked IP address.
IPs	The number of unique IP addresses making traffic with the attacked IP address.
Dropped	Represents the rate of packets dropped in the capturing process. A high number indicates a performance problem.
Peak CPU%	The maximum percentage of CPUs 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.
Server	The server that runs the Filter. Click to open a new tab with data specific to the server. Administrators and operators can right-click to open the server configuration.

Sensors

When you click a Sensor's name anywhere in the Console, the Sensor's tab opens. That tab includes a few sub-tabs located on the bottom side of the window. All sub-tabs have the following common toolbar fields:

- **Sensors** – Select the Sensors you are interested in, or select “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 into a single tab.

The Sensor Dashboard's configuration does not apply to a particular Sensor. The changes you make there will be visible on all Sensor dashboards.

The operation of dashboards is documented in the Reports » Dashboards chapter (page 27).

Sensor Tops

Sensor Tops allow you to generate various traffic tops for the selected Sensor(s). The top generator value from the Sensor configuration enables or disables data collection for Sensor tops.

- **Top Type**
 - *Talkers* – the IPs of your network that sent or received the most traffic for the selected decoder.
 - *IP Groups* – the IP groups that sent or received the most traffic for the selected decoder.
 - *External IPs* – the external IPs that sent or received the most traffic for the selected decoder.

- *Autonomous Systems* – the autonomous systems that sent or received the most traffic.
- *Countries* – the countries that sent or received the most traffic.
- *TCP Ports* – the most-used TCP ports.
- *UDP Ports* – the most-used UDP ports.
- *IP Protocols* – the most-used IP protocols.
- *IP Versions* – the most-used IP versions: IPv4 or IPv6.
- **Decoder** – Selects the decoder that analyzes the traffic you are interested in.
- **Direction** – The direction of the traffic (*Inbound* or *Outbound*).
- **Group Sensors** – If unchecked, each Sensor generates a different top. If checked, a single top includes the tops of all the selected Sensors.
- **DNS** – Check this if you need reverse DNS resolution for IP addresses. In some cases, this slows the top generation.

The number of top items and decoders can be modified in the Storage & Graphs Configuration (see page 37).

Generating tops for many Sensors or for large time frames may take several minutes. It may also require an increased *max_execution_time* parameter in *php.ini*.

Sensor Graphs

Sensor Graphs allows you to view a variety of Sensor-related histograms for the selected Sensor(s):

- **Data Units** – Select one or more parameters:
 - *Most Used* – Shows the most used data units, each in a different graph.
 - *Packets* – The packets/second rate.
 - *Bits* – The bits/second throughput.
 - *Applications* – Sensors can collect application-specific distribution data for: HTTP, HTTPS, SMTP, POP3, IMAP, SNMP, FTP, SSH, TELNET, SQL, NETBIOS, MS-DS, MS-RDP, DNS, ICMP and OTHERS.
 - *Bytes* – The bytes/second throughput.
 - *Internal/External IPs* – The number of IP addresses that sent or received traffic. The Internal/External IPs are the IPs inside/outside the IP Zone. The top generator value from the Sensor configuration enables or disables monitoring of External IPs. A spike in the Internal IPs graph usually means that an IP class scan was performed against your IP block. A spike in the External IPs graph usually means that you received a spoofed attack.
 - *Received Frames* – For Sniffing Sensors, represents the rate of packets received before validation. For Flow Sensors, represents the rate of flows received before validation.
 - *Dropped Frames* – For Sniffing Sensors, represents the number of packets dropped in the capturing process. A high number indicates a sniffing performance problem. For Flow Sensors, represents the number of unaccounted flows. A high number indicates a performance problem with the Flow

Sensor or a network connectivity issue with the flow exporter.

- *Unknown Frames* – For Sniffing Sensors, represents the rate of invalidated packets. For Flow Sensors, represents the rate of invalidated flows.
- *Unknown Sources* – The number of source IP addresses that did not pass validation.
- *Unknown Destinations* – The number of destination IP addresses that did not pass validation.
- *Avg Packet Size* – The average packet size in bits/packet.
- *CPU%* – The percentage of CPUs 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 updated IP accounting records.
- *HW Graphs* – The number of updated 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.
- *IP Structure RAM* – The number of RAM bytes used by each IP structure.
- **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 or no title for the “None” option, or you can enter your own text to be rendered as a title.
- **Graph Legend** – Select how detailed the graph's legend should be.
- **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 a graph. If checked, all selected Sensors generate a single graph that contains the combined data.
 - *Show Totals* – If multiple Sensors are used, render the data units stacked inside the graph.

Flows List

This is available only for Flow Sensors.

You can list and filter the flow data saved by the Flow Sensor. The options are listed on page 12, in the Flow Collectors chapter.

Flows Top

This is available only for Flow Sensors.

You can generate tops and process and filter the flow data saved by the Flow Sensor. The options are documented on page 12, in the Flow Collectors chapter.

AS Graphs

Flow Sensors and Sniffing Sensors can generate traffic and bandwidth histograms for autonomous systems. To enable this, set the top generator parameter to “Full” for Sniffing Sensors and to “Extended” for Flow Sensors.

The parameters are:

- **Sensors** – Select the Sensors you are 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. Frequently-used AS numbers can be saved there, for use at a later time.

If you don't know what AS number(s) a particular ISP has, go to Help » IP & AS Information » AS Numbers List. There you can apply different filters by clicking a table header's down arrow.
- **Export** – You can print, save as PDF or email the generated graphs.
- **Refresh** – The graphs are refreshed 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 or no title for the “None” option, or you can enter your own text to be rendered as a title.
- **Graph Options**
 - *Stack Sensors* – If unchecked, a different AS graph is generated for every Sensor. Otherwise, a single AS graph that contains summed traffic data is generated for all selected Sensors.
 - *Stack ASNs* – If you entered multiple AS Numbers, then you can sum all of them in a single AS graph. This is useful with ISPs and AS owners that have more than 1 allocated AS number.

Country Graphs

Flow Sensors and Sniffing Sensors can generate bandwidth histograms for Countries. To enable them, set the Top Generator parameter from the Sensor Configuration to “Extended” or to “Full”.

The parameters are:

- **Sensors** – Select the Sensors you are 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...”
- **Countries** – Select the country or countries from the drop-down list, or click the light bulb icon on the right to open a window with saved selections.
- **Export** – You can print, save as PDF or email the generated graphs.
- **Refresh** – The graphs are refreshed 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 or no title for the “None” option, or you can enter your own text to be rendered as a title.
- **Graph Options**
 - *Stack Sensors* – If unchecked, a different graph is generated for every Sensor. Otherwise, a single graph that contains the summed traffic data is generated for all selected Sensors.
 - *Stack Countries* – If you selected multiple countries, then you can sum all of them in a single graph, or you can see a separate graph for each country.

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 66).

Anomalies Overview

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

Filters

When you click a Filter's name anywhere in the Console, the Filter's tab opens. This tab includes a few sub-tabs located on the bottom side. All sub-tabs have the following common toolbar fields:

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

Filter Dashboard

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

The Filter Dashboard's configuration does not apply to a particular Filter. The changes you make there will be visible on any Filter dashboard.

The operation of dashboards is documented in the Reports » Dashboards chapter on page 27.

Filter Graphs

Filter Graphs allows you to view a variety of Filter-related histograms for the selected Filter(s):

- **Data Units** – Select one or more parameters:
 - *Most Used* – Shows most-used data units, each in a different graph.
 - *Anomalies* – The number of anomalies mitigated by the Filter's instances.
 - *Filtering Rules* – The number of filtering rules detected by the Filter's instances.
 - *SW Firewall Rules* – The number of filtering rules that were applied using the software firewall.
 - *HW Firewall Rules* – The number of filtering rules that were applied using the hardware firewall.
 - *Source IPs* – The number of IPs that sent traffic towards the attacked IP address.
 - *CPU%* – The percentage of CPUs used by the Filter's instances.
 - *Used RAM* – How much RAM was used by the Filter's instances.
 - *Filtered Packets* – How many packets were filtered by the software firewall.
 - *Filtered Bits* – How many bits were filtered by the software firewall.
 - *Dropped Packets* – The number of packets dropped in the capturing process.
 - *Received Packets* – The rate of packets received by the Filter's instances.
 - *Packets/s* – The number of packets/s analyzed by the Filter's instances.
 - *Bits/s* – The rate of bits/s analyzed by the Filter's instances.
 - *Filtering Rules* – The number of filtering rules for each filtering rule type.
 - *Total Excepted Rules* – The filtering rules detected but white-listed by the Filter's configuration.
- **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 or no title for the “None” option, or you can enter your own text to be rendered as a title.
- **Graph Legend** – Select how detailed the graph's legend should be.
- **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 Filters* – If unchecked, each selected Filter generates a different graph. If checked, all selected

Filters generate a single combined graph.

- *Show Totals* – If multiple Filters are selected, the total of data units are rendered in a stacked graph.

Filter Events

The list of events generated by the selected Filter(s) for the selected time frame. Events are explained in the Events Reporting chapter on page 66.

Filtering Rules Archive

The list of filtering rules detected by the Filter(s) for the selected time frame. Most fields are explained in the Reports » Alarms & Tools » Anomalies chapter on page 8.

Reports » Dashboards

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

A few sample dashboards are included by default in the Console. You can create your own by going to **Reports » Dashboards » + » Add Dashboard**.

In the dashboard configuration, you can choose to override the time frame of widgets with the time frame of the dashboard.

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 its title bar. To delete a widget, click the third icon from its title bar.

Along with the specific fields, every widget has a configurable title and height. Leave the widget's height parameter set 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.

Widget options are self-explanatory or are described in other chapters.

Only the administrator and operator are able to create, delete or edit dashboards. The user cannot make modifications to dashboards.

Reports » IP Addresses & Groups

This chapter describes how to generate complex traffic reports for IP addresses, IP blocks/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 groups defined in IP Zones. You can search or filter them by entering a sub-string contained in the name of the IP group you are interested in. You can use IP groups to generate reports for customers with multiple IP blocks allocated to them. To do that, use the same IP group name for all the customer's IP blocks.

If the reports are empty, set the IP accounting parameter or the IP graphs parameter to “Yes” in the corresponding IP Zone.

Clicking IP addresses or IP groups anywhere in the Console opens the same type of tab that contains sub-tabs on the bottom left side of the window. All sub-tabs use the following common toolbar fields:

- **Sensor** – Select the Sensors you are interested in or select “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 there will be visible for each IP, subnet or IP Group Dashboard.

The operation of Dashboards is documented in the Reports » Dashboards chapter on page 27.

IP Graphs

IP graphs allows you to view traffic histograms for the IP block, host or group:

- **Decoders & Data Unit** – Select the decoders that analyze the traffic you are 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 or no title for the “None” option, or you can enter your own text to be rendered as a title.
- **Graph Legend** – Select how detailed the graph's legend should be.
- **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 a different graph. If checked, all selected Sensors generate a single combined graph.
- *Stack Decoders* – If unchecked, each selected decoder generates a different graph.
- *Stack IPs* – Uncheck 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, 256 traffic graphs will be 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 be displayed as TOTAL OTHER and TCP as TCP OTHER. However, when you select TCP, HTTP and TCP+SYN as decoders, then the TCP+SYN decoder can be included in both TCP and HTTP, thus generating a decoder conflict. Check this to stop detection of conflicts between decoders, but keep in mind that graphs may be less accurate.
- *Stack Recursively* – When checked, subnet graphs can be created from the IPs graphs that are contained in the subnet.

The number of decoders, Data Units and aggregation types can be modified in the Storage & Graphs Configuration, see page 37.

IP Accounting

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

- **Decoders & Data Unit** – Select the decoders that analyze the traffic you are interested in. Data units available: *Packets*, *Bits* and *Bytes*.
- **Report Type** – Select the interval you want to use to aggregate the accounting data: *Daily*, *Weekly*, *Monthly*, or *Yearly*. The minimum time interval for accounting reports is 1 day, so if you select a shorter time interval, you will still see the data collected for the whole day.
- **Sum IPs** – Uncheck 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 for a large IP block like a /24 CIDR, 256 traffic accounting reports will be 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 37).

Flows List

You can list and filter the saved flow data for the IP class, host or group. The options are documented on page 12 in the Flow Collectors chapter.

This sub-tab is visible only when there is at least one Flow Sensor configured in the Console.

Flows Tops

You can generate tops from the saved flow data for the IP class, host or IP Group. The options are documented on page 12 in the Flow Collectors chapter.

This sub-tab is visible only when there is at least one Flow Sensor configured in the Console.

Profile Graphs

Profile graphs are used by WANGUARD to detect anomalies by traffic profiling. To view them, profile anomalies must be enabled in the Anomalies Configuration (see page 38) and activated in the IP Zone for the IP subnet.

Anomalies Overview

Here you can view trends and summarizations of attacks sent or received by the IP address, class or group for the selected time frame and decoders.

Reports » Servers

When you click a server's name anywhere in the Console, the server tab opens. This tab includes a few sub-tabs located on the bottom left side of the window. All sub-tabs have the following common toolbar fields:

- **Servers** – Select the servers you are interested in, or select “All” to select all servers. Multiple selections can be made. Administrators can filter what servers are available to users.
- **Time Frame** – Select predefined time frames or enter your own by selecting “Custom...”

Server or Console Dashboard

The Server or Console Dashboard allows you to group the most relevant data into a single tab.

The Server Dashboard's configuration does not apply to a particular server. The changes you make there will be visible on each server dashboard.

The operation of dashboards is documented in the Reports » Dashboards chapter on page 27.

Server Graphs

Server Graphs allows you to generate many server-related histograms for the selected server(s):

- **Data Units** – Select one or more parameters:
 - *Most Used* – Shows most-used data units, each in a different graph.
 - *System Load* – The load reported by the operating system.
 - *Free RAM* – The available RAM, not counting the swap space.
 - *Database/Graphs/SSD/Flow Collector/Package Dumps Disk - Free space* – How much disk space is available for each path.
 - *Uptime* – The uptime of the operating system.
 - *CPU% system/userspace/niced/idle* – The percentages of CPUs used by the system, used by the userspace processes, used by the processes with increased (nice) priorities, and idle.
 - *Number of processes* – The total number of processes that are running on the system.
 - *Hardware/Software CPU Interrupts* – The number of CPU interrupts made by hardware events, drivers, etc.
 - *Context Switches* – Indicates how much time the system spends on multi-tasking.
 - *Running Components* – The number of Sensor or Filter processes.
 - *Clock Delta* – The difference of time between the server and the Console, in seconds. Must be 0.
 - *Database/Graphs/SSD/Package Dumps/Flow Collector Disk - Total* – How much disk space is allocated

for the partitions that hold the path.

- *Database/Graphs/SSD/Packet Dumps/Flow Collector Disk - Free Inodes* – The number of free inodes held by the partitions that hold the paths.
- *Database/Graphs/SSD/Packet Dumps/Flow Collector Disk - Ops/s* – The number of reads and writes for the partitions that hold the paths.
- *Database/Graphs/SSD/Packet Dumps/Flow Collector Disk - Bytes/s* – The number of bytes/s for the partitions that hold the paths.
- *Server Interfaces - Packets/Bits/Errors/Dropped* – Various statistics collected for the interfaces defined in the Server Configuration window.
- **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 or title for the “None” option, or you can enter your own text to be rendered as a title.
- **Graph Legend** – Select how detailed the graph's legend should be.
- **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 Servers* – If unchecked, each selected server generates a different graph. If checked, all selected servers generate single combined graph.
 - *Show Totals* – If multiple servers are used, render the total of data units in a stacked graph.

Server or Console Events

The list of events generated by the selected server(s) or by the Console for the selected time frame. Events are explained in the Events Reporting chapter (see page 66).

Server Commands

Administrators can execute commands on the selected server(s) and see their output in the Console. The commands are executed by the WANsupervisor with the “andrisoft” account privileges, if the WANsupervisor service was not started with the “-n” option.

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 basic Linux or FreeBSD operation skills then no training is required for the software installation. Feel free to contact our support team with any issues.

Installing WANGUARD will 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 or 13.x** (free, Novell Enterprise Linux-based distribution), **Debian Linux 6.0 or 7.0** (free, community-supported distribution), **Ubuntu 12.x**. Other distributions may work but have not yet been tested.

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, etc.). You can also install the components on multiple servers distributed across your network.

We highly recommend running the software on physical servers and not on virtual machines, for the following reasons:

- Virtual machines don't have a stable clock source. This is a critical requirement for the Sensors.
- Virtual machines often suffer from disk I/O bottlenecks. This is a critical issue for the Console.
- The software needs resource allocation predictability.

Sniffing Sensor – Minimum Hardware Requirements

Sniffing Capacity	1 Gigabit Ethernet	10 Gigabit Ethernet
Architecture	x86 (32 or 64 bit)	x86 (64 bit)
CPU	2.0 GHz dual-core Xeon	2.8 GHz quad-core Xeon
RAM	1 GB	2 GB
Network Cards	Gigabit Ethernet with NAPI support Fast Ethernet for management	10 GbE card. Intel 82599 chipset recommended Fast Ethernet for management
Operating System	RHEL 5 / CentOS 5 or 6, Debian 6 or 7,	RHEL 5 / CentOS 5 or 6, Debian 6 or 7,

	Ubuntu Server 12, OpenSUSE 12 or 13	Ubuntu Server 12, OpenSUSE 12 or 13
Disk Space	10 GB (including operating system)	10 GB (including operating system)

Flow Sensor – Minimum Hardware Requirements

Flow-Processing Capacity	20 monitored interfaces, 15k active endpoints*
Architecture	x86 (32 or 64 bit)
CPU	2.0 GHz Xeon
RAM	4 GB
Network Cards	Fast Ethernet
Operating System	RHEL / CentOS 5 or 6, Debian 6 or 7, Ubuntu Server 12, OpenSUSE 12 or 13
Disk Space	15 GB (including operating system)

* The Flow Sensor is scalable. The number of monitored interfaces and endpoints grows linearly with the performance of the server.

Filter – Minimum Hardware Requirements

Filtering Capacity	1 Gbps*	10 Gbps*
Architecture	x86 (32 or 64 bit)	x86 (64 bit)
CPU	2.4 GHz Xeon	2.4 GHz quad-core Xeon
RAM	2 GB	4 GB
Network Cards	Gigabit Ethernet with NAPI support Fast Ethernet for management	10 GbE Cards with 82599 chipset Fast Ethernet for management
Operating System	RHEL 5 / CentOS 5 or 6, Debian 6 or 7, Ubuntu Server 12, OpenSUSE 12 or 13	RHEL 5 / CentOS 5 or 6, Debian 6 or 7, Ubuntu Server 12, OpenSUSE 12 or 13
Disk Space	10 GB (including operating system)	10 GB (including operating system)

* The Filter doesn't use the connection tracking system specific to stateful firewalls. This ensures a much better filtering and routing performance. However, we are unable to estimate the filtering performance you will be able to achieve, as the packet filter's capacity depends on many parameters: CPU type/speed/cache, Linux kernel version, NIC chipset, NIC driver, attack type, server load, routed traffic size, multi-core balance of hardware interrupts, number of existing rules, etc.

Having a dedicated filtering server for each monitored link is not always necessary. 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 Components*
Architecture	x86 (32 or 64 bit)
CPU	2.4 GHz dual-core Xeon
Memory	1 GB
Network Cards	Gigabit Ethernet
Operating System	RHEL / CentOS 5 or 6, Debian 6 or 7, Ubuntu Server 12, OpenSUSE 12 or 13
Software Packages	Apache 2.x+, php 5.2+, mysql 5.x, rrdtool 1.3+, ping, whois, traceroute, telnet, wireshark, tcpdump
Disk Space	10 GB (including operating system) + additional storage to store IP graphs data

* The Console is scalable. The number of monitored components grows linearly with the performance of the server.

To access the web interface provided by the Console, one of the following web browsers is required (others should also work but have not been tested): Google Chrome, Firefox 3.5 or later, Safari 3.0 or later, Internet Explorer 8 or later. Older versions of Internet Explorer contain a very slow javascript engine and won't work well. For the best Console experience, we highly recommend 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 or Adobe Flash are not required. To access the contextual help you must install the Adobe PDF Reader.

Software Installation & Download

The latest software installation instructions for RedHat-based, SuSE-based and Debian-based Linux distributions are listed on the Andrisoft website. The download link is included in the email with the trial key.

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 the Console for the First Time

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

To log into the 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. You can also access the Console securely by HTTPS if the Apache web server was configured with SSL/TLS support.

Licensing Procedure

If you have not yet licensed WANGUARD you will be asked to do so. Upload the *trial.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 into the Console using the default username/password combination: **admin/changeme**.

To understand how to navigate within the Console, please read the chapter beginning on page 7.

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 at the top-right part of the browser window and select <Change Password>.

Quick Configuration Steps

- Estimate storage requirements, review decoders and graphs parameters – page 37
- Setup anomalies detection parameters and decoders – page 38
- Configure the reaction to traffic anomalies – page 39
- Add your IP address ranges and important IPs to an IP Zone – page 45
- Add anomaly detection for prefixes, create thresholds templates – page 46
- Add and configure a Sensor, then start it – page 48
- Watch for errors in events log. Receive error notifications by email – page 66
- Generate reports and send them periodically by email – page 65
- Create your own dashboards and add useful widgets – page 27
- Configure Console accounts for your staff or customers – page 67
- Configure new BGP connections, if needed – page 58
- Configure Filters, if needed – page 60

Storage & Graphs Configuration

An important initial step in configuring WANGUARD is to make sure that the server(s) the software runs on have enough resources to process and store traffic information. Most resource-related options are found in Configuration » Global Settings » Storage & Graphs.

The default paths for **collected flows** and **packet dumps** exist only on the Console's file system. When the Sensors are installed on different systems, you should export these paths towards the Console's file system using an NFS share. If you do not, the Console won't be able to display the data saved on remote servers.

In a later chapter, you'll be able to configure the Sensors to generate traffic graphs for every IP that belongs to the monitored network. If you intend to use this feature, look carefully at the IP graphs options. Changing these options later will reset all existing IP graphs data.

IP graph files are stored on the Console's file system. There are 2 mutually exclusive methods for updating IP graph files, so select the appropriate one for you:

- **Write IP graph files directly on disk** – This method creates one file for every IP address directly in the defined Graphs Disk Path. If the RRDCache daemon is being used to speed up graphs I/O, add its socket path (unix:/var/rrdtool/rrdcached/rrdcached.sock on Redhat). The Andrisoft Knowledge Base contains an article on configuring RRDCacheD.
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 the length of time the data is stored.
This method is not suited for updating hundred of thousands of IP graphs with a 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 onto a larger but much slower disk.

Decoders determine the underlying protocols of each packet or flow. Enabling many decoders might cause a performance penalty for Sensors, 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 options 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.

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

Anomalies Configuration

Configure the anomaly detection engine by setting up the detection options and decoders in Configuration » Global Settings » Anomalies.

The Sensors are able to detect traffic anomalies by 2 methods:

- **Threshold Anomalies** – 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** – 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 anomaly detection for a subnet, the Sensor builds a behavioral traffic graph for about 25 hours. Note that profile anomaly detection is suited for hosts and subnets that have a predictable traffic pattern. Larger subnets are usually more predictable. False positives can be limited by adjusting the deviation percent and minimum packets and bits rates.

Traffic anomaly detection will be enabled individually for each subnet when configuring IP Zones (see page 45).

Response Configuration

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

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

- **Anomaly Actions** – Executed by Sensors for each traffic anomaly while the anomaly is active or when it expires.
- **Filtering Rule Actions** – Executed by Filters for each filtering rule while the rule is active or when it expires. Filtering rules expose malicious packets that share common OSI layer 3-7 fields (attacker IPs, TCP/UDP ports, length, protocols, TTL, content, 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.

Each action panel contains specific fields. The following fields are common:

- **Active** – Selects whether the action is enabled or disabled.
- **Priority** – Selects the order of execution relative to the other actions that are defined in the same panel. Lower numerical values correspond to increased priority.
- **Description** – The name or short description of the action.
- **Execution** – The actions in the “While... is detected” panel can be executed once for each anomaly or filtering rule, or can be executed periodically. The interval of execution is 5 seconds for Sniffing Sensors and Virtual Sensors, and 5-60 seconds for Flow Sensors (configurable graphs accuracy).
- **Log Execution** – The name of the action is visible in anomaly reports only when this is checked.
- **Conditions** – 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 (e.g. { and }) that can be included in the body of most actions. Every conditional parameter corresponds to a dynamic parameter.

By using the custom script action together with dynamic parameters, you can extend the reaction to anomalies and filtering rules. The script is executed on the Sensor or Filter system if it is accessible by the “andrisoft” user account.

The **List Prefixes** button allows you to see what IP classes are configured to use the response.

List of Conditional Parameters and 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. This is {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 Devices Group field.
7	Sensor IP	String	{sensor_ip}	The IP of the server running the Sensor.
8	Sensor Type [sniff,flow,virtual]	String	{sensor_type}	Type is "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 packets captured 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.

18	Manual Actions [yes,no]	String	{manual_actions}	When this condition is equal to “yes,” administrators and operators can manually trigger the action from Reports » Alarms & Tools » Anomalies.
20	Classification [Unclassified, etc..]	String	{classification}	Administrators and operators can manually classify anomalies in Reports » Alarms & Tools » Anomalies. The possible values are: Unclassified, False Positive, Possible Attack, Trivial Attack, Verified Attack, Crippling Attack.
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 the Console for the anomaly by administrators.
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}	Domain is “ip” when CIDR = 32 for IPv4 or 128 for IPv6; “subnet” in all other cases.
6	Anomaly Class [thresholds,profile]	String	{class}	Class is “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}	Comparison reads “above” for thresholds exceeding expectations or “under” for thresholds below expectations.
10	Unit [pkts/s,bits/s]	String	{unit}	Unit is “pkts/s” for packets per second anomalies or “bits/s” for bits per second anomalies.
11	Threshold Value	Number*	{rule_value}	The threshold value configured for the threshold.
12	Computed Threshold	Number*	{computed_threshold}	Threshold of the anomaly, 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 “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 ratio between the peak anomalous traffic rate and the threshold value.
17	Latest Rule Severity	Number	{latest_severity}	The ratio between the latest anomalous traffic rate and the threshold value.
18	Peak Link Severity	Number	{link_severity}	The ratio between the peak anomalous traffic rate and the interface's traffic rate.
19	Latest Link Severity	Number	{latest_link_severity}	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 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 filtering rule.
2	Filter Type (<i>ip, source, dest, proto, len, ttl</i>)	String	{filter_type}	The filtering rule 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 filtering rule's value.
		String	{filter_ip_dns}	If the filtering rule is for an IP, the dynamic parameter provides the reverse DNS of the IP.
4	Filter ISP	String	{filter_ip_isp}	If the filtering rule is for an IP, the dynamic parameter provides corresponding organization/ISP/autonomous system.
5	Filter Country	String	{filter_ip_country}	If the filtering rule is for an IP, the dynamic parameter provides the country the IP comes from.
6	Filter Pkts/s	Number*	{filter_pps}	The filtering rule's latest packets/second throughput.
7	Filter Bits/s	Number*	{filter_bps}	The filtering rule's latest bits/second traffic throughput.
8	Filter Peak Pkts/s	Number*	{filter_max_pps}	The maximum packets rate matched by the filtering rule's traffic.

9	Filter Peak Bits/s	Number*	{filter_max_bps}	The maximum bits rate matched by the filtering rule's traffic.
10	Filter Severity	Number	{filter_severity}	The severity field represents the ratio between filtering rule's traffic and threshold value.
11	Filter Packets	Number*	{filter_packets}	The number of packets matched by the filtering rule.
12	Filter Bits	Number*	{filter_bits}	The number of bits matched by the filtering rule.
13	Filter Time Interval (seconds)	Number	{filter_difftime}	The duration of the filtering rule.
14	Filter Whitelist	Number	{filter_whitelisted}	If the filtering rule is whitelisted, the value is 1. Otherwise, it is 0.
15	Filter Traffic Sample Size (bytes)	Number *	{filter_log_size}	The size of the traffic captured by the filtering rule's Capture Traffic action.
		String	{attacker_isp}	If the filtering rule is for an IP, the dynamic parameter provides the email address 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 in which you should include your IP address ranges. To add an IP Zone go 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.

To change the name of an IP Zone, open the IP Zone configuration window, provide a new description and press <Change Name>.

To copy an IP Zone, click the <Duplicate> button. A new IP Zone will be created that will have the same information and the same description as the original, but with the word “(copy)” attached. In some cases, when you have multiple Sensor systems, you may have to create multiple IP Zones that contain the same prefixes but have different settings for them. It is easier to duplicate an existing IP Zone than to add the same IP classes for each new IP Zone.

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

The IP Zone configuration window is divided in two vertical sections. The buttons that manage prefixes (IP address ranges or individual IPs) are located in the upper part of the left-hand section. When a new prefix is added, the tree below automatically updates itself. The right-hand section contains panels with user-provided settings for the selected prefix.

To enter IP addresses or IP blocks, you must use the CIDR notation. To enter individual hosts in IP Zones you must use the /32 CIDR for IPv4, or /128 for IPv6. For more information about the CIDR notation, see Appendix 1 from page 68.

Every IP Zone contains at least the 0.0.0.0/0 network. Because it has the /0 CIDR mask, it contains all IP addresses available for IPv4 and IPv6. For an easier configuration, every new prefix that you define will inherit by default the properties of the closest (having the biggest CIDR) IP class that includes it.

The **IP Settings** panel in the right-hand section provides the following options:

- **IP Group** – This combo box should contain a short description of 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,” the Console will collect graph data for every IP contained in the selected prefix.
- **IP Accounting** – If set to “Yes,” the Console will save daily accounting data for each IP contained in the selected prefix.

The **Storage Requirements** column indicates the necessary disk space for each Sniffing Sensor or Flow Sensor interface. Enabling IP graphs and IP accounting for very large prefixes (e.g. 0.0.0.0/0) is probably going to generate useless data that can potentially overload the system.

The **Comments** panel allows you to write a comment for the selected prefix. It is not visible elsewhere.

Anomaly Detection Settings & Thresholds Templates

Define traffic threshold rules by adding them to the **Thresholds Anomalies** panel from the IP Zone Configuration window. To ease the addition of identical threshold for multiple prefixes, add them to a thresholds template instead, by going to Configuration » Network & Policy » + » Add Thresholds Template.

A threshold rule is composed from:

- **Domain** – Sensors can detect anomalies to/from an IP contained in the subnet, or to/from the whole subnet.
- **Direction** – The direction of the 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 percentage of the total traffic received by the Sensor. Absolute values can be multiples of 1000 with K (kilo) appended, multiples of 1 million with M (mega) appended, or multiples of 1 billion with G (giga) appended.
- **Decoder** – Select one of the decoders enabled in the Anomalies Configuration window (see page 38).
- **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 prefixes. You can cancel inherited thresholds by selecting “Unlimited” in the Value field.
- **Inheritance** – Shows the parent prefix, if it exists.

Adding a threshold rule on 0.0.0.0/0 that reads, “Any IP receives over 5% TCP+SYN pkts/s” will catch port scans and all significant SYN attacks.

A threshold rule on 0.0.0.0/0 that reads, “Subnet receives under 5M TOTAL bits/s” will execute the response when the monitored link goes down.

You can configure illegal IP address ranges that should never be seen in normal traffic, like unallocated IP addresses or part of your internal IP address range that is unoccupied. Then add small thresholds to these to catch malicious activities such as scans and worms.

Adding similar threshold rules is not allowed, even if the rules have different values or responses. You can enable different actions for different threshold values by using conditional parameters inside the response. If you want to activate the Filter for UDP attacks stronger than 100 Mbps but also to null-route by BGP when those attacks reach 1 Gbps, add only the “Any IP receives over 100M UDP bits/s” rule. Then, inside the response add 2 actions: one that activates the Filter, and another that executes a BGP announcement with the condition “Peak Value” “over” “1G”.

The **Profile Anomalies** panel contains the Profiling Data parameter, which can have the following values:

- *Inherit* – inherit the value from the parent prefix
- *No* – do not generate profiling data for the selected prefix
- *For Subnet* – generate profiling data for all traffic received by the prefix
- *For IPs* – use carefully as it will generate profiling data for every IP contained in the prefix. This is not recommended for use on large subnets.
- *For All* – activate both of the above options.

Choosing a Method of Traffic Monitoring

This chapter explains the IP traffic monitoring methods supported by Sensors. There are 2 types of traffic-monitoring Sensors that differ only in the way they capture traffic:

- The **Sniffing Sensor** analyzes packets. It is used for in-line appliances, port mirroring or passive network tap.

In switched networks, only the packets for a specific device reach the device's network card. If the server running the Sniffing Sensor is not deployed in-line (in the main data-path), then a network TAP, or a switch or router that offers a “monitoring port” or “mirroring port,” must be used. In this case, the network device sends a copy of data packets traveling through selected ports or VLANs to the monitoring port. The Sniffing Sensor inspects every packet it receives and conducts packet-based traffic analysis.

- The **Flow Sensor** analyzes flows. It is used for monitoring NetFlow® (jFlow, NetStream, cflowd), sFlow® or IPFIX.

Many routers and switches can collect IP traffic statistics on monitored interfaces, and later export those statistics as flow records to the Flow Sensor. Because the flow protocol already performs pre-aggregation of traffic data, the flows of data sent to the Flow Sensor are much smaller than the monitored traffic. This makes the Flow Sensor a good option for monitoring remote or high-traffic networks. The main downside for flow-based traffic analysis is that computing pre-aggregation of traffic data adds at least a 30-second delay to the traffic statistics, so statistics are not in real time.

- The **Virtual Sensor** aggregates Sniffing Sensor and Flow Sensor data into a single anomaly detection domain.

The Virtual Sensor is not a “real” Sensor because it doesn't do any traffic capturing by itself. It is only used for anomaly detection in the summed-up traffic data collected by the Sniffing Sensors or Flow Sensors.

You can run Sniffing Sensors and Flow Sensors at the same time to achieve high availability and redundancy, and to be able to generate packets dumps and flow dumps.

Comparison between Packet-Based and Flow-Based Monitoring

We recommend using the Sniffing Sensor when the speed of detecting attacks is very important and when capturing raw packets for forensics is necessary. Since the Sniffing Sensor deals with every packet entering your network, it needs to run on a server with a powerful CPU.

The Flow Sensor receives pre-aggregated traffic information from routers, so its CPU usage is low. This enables the Flow Sensor to analyze the traffic from multiple 10GE or even 40GE interfaces, even if it runs on low-end hardware. Flows are saved in a compressed binary format and can be queried at a later date. The disadvantages of using the Flow Sensor are that it needs more RAM than the Sniffing Sensor, results in increased CPU usage on the

network device, and exhibits reduced speed in detecting attacks caused by the flow exporting technology.

The feature list is identical for both Sensor types. This is why only the generic term “Sensor” is used in the Console reports and throughout the documentation.

The table below lists the main differences between Sensors:

Sensor Type	Sniffing Sensor	Flow Sensor
Capturing Technology	<ul style="list-style-type: none"> - Port Mirroring (SPAN, Roving Analysis Port) - Network TAP - Sniffing packets passing an in-line appliance 	<ul style="list-style-type: none"> - NetFlow version 5, 7, 9 (jFlow, NetStream, cflowd) - sFlow version 4, 5 - IPFIX
Maximum Traffic Capacity per Sensor	10 GigE >150,000 endpoints*	multiples of 10 Gbps >100,000 endpoints*
Anomaly Detection Time	≤ 5 seconds	> flow export time (> 30 seconds) + 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
Packet Analyzer	Yes	No
Flow Collector	No	Yes

* 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 packets for a specific device reach the device's network card. If the server running the **Sniffing Sensor** is not deployed in-line (in the main data-path), then a network TAP, or a switch or router that offers a “monitoring port” or “mirroring port,” must be used. In this case, the network device sends a copy of data packets traveling through selected ports or VLANs to the monitoring port. The Sniffing Sensor inspects every packet it receives and conducts packet-based traffic analysis.

For instructions on how to configure switches or routers for port mirroring, consult the network device's documentation.

To configure an existing Sniffing Sensor, go to Configuration » Components and click the Sensor's name. To add a Sniffing Sensor, click the <+> button from the title bar of the Configuration » Components panel.

The Sniffing Sensor Configuration window contains the following fields:

- **Sensor Name** – A short name to help you identify the Sniffing Sensor.
- **Devices Group** – Optional description used within the Console to group multiple components by location, role, etc.
- **Graph Color** – The color used in graphs for this Sensor. The default color is a random one, but you can change it by entering another HTML color code or by clicking the drop-down menu.
- **Sensor License** – The license used by the Sensor. WANGUARD provides all features; WANSIGHT does not provide traffic anomaly detection and reaction.
- **Sensor Server** – The server running the Sensor. To define a new server, go to Configuration » Servers » + » Add Server.
- **Sniffing Interface** – The network interface listened by the Sniffing Sensor. The Linux network interface naming convention is eth0 for the first Ethernet interface, eth1.900 for the second Ethernet interface with 802.1Q VLAN 900, etc.

If the Sniffing Sensor server is deployed in-line, then this field must contain the network interface that receives the traffic entering your network.

- **Link Speed IN / OUT** – The speed of the monitored link. If set, it is used to generate reports based on usage percentage.
- **IP Zone** – The Sensor needs an IP Zone from which to learn about your network and to extract per-subnet settings.

IP Zones are documented in the IP Zones Setup chapter on page 45.

- **IP Validation** – This option can be used to distinguish the direction of the packets or to ignore certain IPs:
 - *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 by decreasing the CPU usage of the Sniffing Sensor.
- **Top Generator** – Allows generation of traffic tops:
 - *Basic* – Enables tops for Internal IPs, IP protocols, versions and TCP/UDP ports. It is the recommended value because it adds a very small performance penalty to the Sniffing Sensor.
 - *Extended* – Enables all tops from *Basic* as well as tops for External IPs (IPs not included in the IP Zone). It adds performance penalty of over 20%, especially during spoofed attacks.
 - *Full* – Enables all tops from *Extended* as well as tops and graphs for autonomous systems. It adds a performance penalty of over 20%, especially during spoofed attacks.
- **Comments** – Comments about the Sensor configuration can be saved here. They are not visible elsewhere.

To start the Sniffing Sensor, click the gray square button next to the Sensor's name from Configuration » Components. Check that the Sniffing Sensor starts properly by watching the events log (details on page 66).

If the Sniffing Sensor starts without errors, but you can't see any data collected by it in Reports » Components » Overview, follow the troubleshooting guide below.

Troubleshooting the Sniffing Sensor

- ✓ Look for warnings or errors produced by the Sniffing Sensor in the events log (details on page 66).

- ✓ Check that you have correctly configured the Sniffing Sensor. Each configuration field is explained in detail in the previous paragraph.
- ✓ Check that the sniffing interface is up using the “ifconfig <ethX>” command.
- ✓ Check that you have correctly configured the switch/TAP to send packets to the server on the configured interface.
- ✓ You can verify whether the server is receiving the packets through the configured interface with a tool like tcpdump. The syntax is “tcpdump -i <interface_usually_eth0> -n -c 100”.
- ✓ When IP Validation is not disabled, make sure that the IP Zone contains 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 to the **Flow Sensor**. Because the flow protocol already performs pre-aggregation of traffic data, the flows of data sent to the Flow Sensor are much smaller than the monitored traffic. This makes the Flow Sensor a good option for monitoring remote or high-traffic networks.

For detailed instructions on how to enable NetFlow, sFlow or IPFIX on your network device, please consult the appropriate user guide from your vendor. Appendix 2 from page 71 contains an example on how to configure NetFlow on Cisco IOS, CatOS and Juniper.

To configure an existing Flow Sensor, go to Configuration » Components and click the Sensor's name. To add a Flow Sensor click the <+> button from the title bar of the Configuration » Components panel.

The Flow Sensor Configuration window contains the following fields:

- **Sensor Name** – A short name to help you identify the Flow Sensor.
- **Devices Group** – Optional description used within the Console to group multiple components by location, role, etc.
- **Sensor Server** – The server running the Sensor. To define a new server, go to Configuration » Servers » + » Add Server.
- **Listener IP:Port** – The IP address of the network interface that receives flows and the destination port.
- **Repeater IP:Port** – Send all incoming flows to another host/collector by enabling the embedded packet repeater (optional).
- **Flow Collector** – All flows can be stored in an efficient binary format and queried in Reports » Alarms & Tools » Flow Collectors.
- **Sensor License** – The license used by the Sensor. WANGUARD provides all features; WANSIGHT does not provide traffic anomaly detection and reaction.
- **Flow Protocol** – The type of flows exported towards the Sensor: NetFlow, IPFIX or sFlow. For IPFIX implementations that retain the start time of the flows (e.g. Juniper MX), select the second value.
- **Flow Exporter IP** – The IP address of the router, switch, probe etc. Usually the loopback0 address of the router. For sFlow exporters, enter the IP that sends flow packets, not the Agent IP.
- **Sampling (1/N)** – Must contain the sampling rate configured on the flow exporter, or 1 when no sampling rate is configured. For NeFlow v9 and sFlow the value provided here is ignored because the sampling rate is automatically adjusted by the protocol.
- **Flow Timeout (s)** – The flow-active/inactive-timeout value from the IPFIX exporter that keeps the start time of the flows.
- **Time Settings** – The time offset between the time zone of the Flow Sensor's server and the flow exporter. Run NTP on both devices to keep their clocks synchronized. This is a critical requirement for the Flow Sensor.

- **SNMP Community** – The read-only SNMP community of the flow exporter allows the Console to gather interface information. If this field is left empty, you must enter the SNMP index, speed, etc. manually for each interface.
- **Monitored Interfaces** – The list of interfaces that should be monitored. Add only upstream interfaces if possible, to avoid producing duplicate flow entries. Settings per interface:
 - *Interface Name* – A short description used to identify the monitored interface. Descriptions longer than 10 characters clutter some reports.
 - *Graph Color* – The color used in graphs for this interface. The default color is a random one, but you can change it by entering another HTML color code or by clicking the drop-down menu.
 - *SNMP Index* – The interfaces are identifiable in flows only when their SNMP indexes is known.
 - *Traffic Direction* – The direction of the traffic entering the interface:
 - “Inbound” – For upstream interfaces, e.g. peering interfaces.
 - “Outbound” – For downstream interfaces, e.g. customer interfaces.
 - “Mixed” – Establishes the direction by IP/AS Validation.
 - “Null” – Null interface traffic is discarded by the router and also by the Flow Sensor.
 - *Link Speed In & Link Speed Out* – The capacity of the interface. If set, it is used to generate reports based on usage percentage.
 - *Top Generator* – Allows generating traffic tops:
 - “Basic” – Enables tops for Internal IPs, IP protocols, versions and TCP/UDP ports. It adds a very small performance penalty to the Flow Sensor.
 - “Extended” – Enables all tops from “Basic” as well as tops and graphs for autonomous systems and countries, but increases the CPU usage of the Flow Sensor by a few percentage points. If the router doesn't export AS information (e.g. non-BGP router), the Sensor uses an internal GeoIP database to get ASNs. Live stats for autonomous systems and countries are not very accurate.
 - “Full” – Enables all tops from “Extended” as well as tops for external IPs (IPs not included in the IP Zone), but increases the RAM usage several times over, especially during spoofed attacks. Live stats for autonomous systems and countries are very accurate.
- **IP Zone** – The Sensor needs an IP Zone from which to learn about your network and to extract per-subnet settings. For more information about IP Zones, please consult the IP Zones Setup chapter on page 45.
- **IP Validation** – This option can be used to distinguish the direction of the traffic or to ignore certain flows:
 - *Off* – The Flow Sensor analyzes all flows, but the traffic direction must be established per interface.
 - *On* – The Flow Sensor analyzes the flows that have the source and/or the destination IP in the selected IP Zone.
 - *Strict* – The Flow Sensor analyzes only the flows that have either the source or the destination IP in the IP Zone.
 - *Exclusive* – The Flow Sensor analyzes only the flows that have the destination IP in the 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 number is set to 0, then the IP address belongs to your autonomous system.

If enabled, only flows having the AS number set to “0” (your AS) are processed. This is rarely used to establish traffic direction.

AS validation has three options:

- *Off* – Disables 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.
- **Graphs Accuracy** – Low values increase the accuracy of Sensor graphs, at the expense of RAM usage. Setting this to under 20 seconds is not recommended.
- **Comments** – Comments about the Sensor configuration can be saved here. These are not visible elsewhere.

To start the Flow Sensor, click the gray square button next to the Sensor's name from Configuration » Components. Check that The Flow Sensor starts properly by watching the events log (details on page 66).

If the Flow Sensor starts without errors, but you can't see any data collected by it in Reports » Components » Overview after more than 5 minutes, follow the troubleshooting guide below.

Troubleshooting the Flow Sensor

- ✓ Look for warnings or errors produced by the Flow Sensor in the events log (details on page 66).
- ✓ Check that you have correctly configured the Flow Sensor. Each configuration field is explained in detail in the previous paragraph.
- ✓ You can verify that the server is receiving the flow packets on the configured Listener IP and Port with a tool like *tcpdump*. The syntax is “*tcpdump -i <interface_usually_eth0> -n -c 100 <flow_exporter_ip> and udp and <destination_port>*”.
- ✓ You can check if the local firewall, which allows the Flow Sensor to receive the flow packets, is enabled with the *iptables* command. The syntax is “*iptables -L -n -v*”.
- ✓ The clocks of both devices are synchronized with NTP. If the devices don't reside in the same time zone, adjust the time zone offset in the Flow Sensor configuration.
- ✓ The flow exporter's active/inactive flow timeout settings are set to less than 300 seconds. Flows sent with a delay of more than 300 seconds are automatically discarded and a warning is sent to the events log.

- ✓ Check that you have correctly configured the flow exporter to send flows to the server for each of the monitored interfaces. To see the interfaces that send flows, go to Reports » Alarms & Tools » Flow Collectors » Flows Top, select the Flow Sensor, set Output to Debug, set Top Type to Any interface and generate the top for the last 10 minutes. The In/Out_If column shows the SNMP index of every interface that exports flows.
- ✓ If you see statistics for only one traffic direction (inbound or outbound), go to Reports » Alarms & Tools » Flow Collectors » Flows List, and generate a listing for the last 10 minutes. If all your IPs are in a single column, check the flow exporter's configuration and feature list. Not all devices can export flows in both directions (e.g. Brocade XMR) or with the same SNMP interface index.
- ✓ If you defined interfaces with the Traffic Direction parameter set to “Mixed,” then make sure that the IP Zone you have selected for the Flow Sensor contains all your IP blocks.
- ✓ If the Flow Sensor does not show the correct statistics after upgrading the router's firmware, the SNMP index of interfaces has probably changed. In this case, enter the new SNMP index for each interface.

Virtual Sensor Configuration

The **Virtual Sensor** aggregates the Sniffing Sensor and Flow Sensor interfaces into a single anomaly detection domain. It disables the anomaly detection features of Sensors, and provides anomaly detection for the summed-up traffic data.

To configure an existing Virtual Sensor, go to Configuration » Components, and click the Sensor's name. To add a Virtual Sensor, click the <+> button from the title bar of the Configuration » Components panel.

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 the Virtual Sensor. The default color is a random one, but you can change it by entering another HTML color code or 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 Sensor interfaces.
- **IP Zone** – The Virtual Sensor must use an IP Zone to extract per-subnet settings. For more information about IP Zones consult the IP Zones Setup chapter on page 45.
- **Comments** – Comments about the Sensor configuration can be saved here. These are not visible elsewhere.

To start the Virtual Sensor, click the gray square button next to the Sensor's name from Configuration » Components. Check that the Virtual Sensor starts properly by watching the events log (details on page 66).

BGP Connection Configuration

Operators and administrators can view, send and withdraw BGP announcements manually from the Console. BGP announcement records are stored in Reports » Alarms & 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 your network by announcing upstream providers using a special BGP community, so that your side no longer routes the attacked addresses and the addresses are null-routed. This network protection technique is called black-holing or RTBH (Remote Triggered Black Hole).
- To divert (BGP off-ramping) the traffic to DDoSed destinations through a Filter system that will filter the attacker's traffic and then re-inject the cleaned traffic back into the network.

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

Before adding a BGP Connection, install and configure the BGPd daemon from the quagga package. Some BGPd configuration steps can be found on Appendix 3 – Configuring Traffic Diversion on page 75.

The BGP Connection Configuration window contains the following fields:

- **BGP Connection Name** – A short name or a description for the BGP Connection.
- **BGP Connection Role** – Set the correct role to see the number of BGP announcements in Reports » Alarms & Tools » BGP Prefixes in red for “Black-holing” and in blue for “Diversion”.
- **BGPd Server** – The server running the BGPd daemon. If the server is not the Console, make sure the BGPd daemon is accessible by telnet from the Console. To add a new server, go to Configuration » Servers » + » Add Server.
- **AS Number** – The AS number must match the one from your 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 parameter that should be appended to each announcement. This is not mandatory.
- **AS View** – If multiple AS views are defined in the BGPd configuration, 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 provided by zebra. This is a rarely-used feature.
- **Zebra Login & Enable Passwords** – The passwords for the zebra daemon.
- **Reject External IPs** – When this is selected, no BGP announcement can be sent for an IP that is not present in any of the configured IP Zones' subnets, excluding 0.0.0.0/0.

- **Reject IPv4 under /** – You can restrict sending prefixes that have the an 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 prefixes that have the an 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 Connection configuration can be saved here. These are not visible elsewhere.

Enable the BGP connection by clicking the gray square button next to the Sensor's name from Configuration » Components.

You can manually send a test BGP announcement with an unused IP address in Reports » BGP Prefixes » BGP Operations » Add BGP Announcements. Check that it functions properly by watching the events log (details on page 66).

If you see telnet connection errors in the events log, verify that the BGPd daemon is accessible by telnet on port tcp/2605 from the Console's server. You can clear BGP prefix errors from Reports » BGP Prefixes » BGP Archive.

Filter Configuration

The **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 worms and virus outbreaks. It includes a sophisticated traffic analysis engine that is able to detect and side-filter malicious traffic in a granular manner, without impacting the user experience or resulting in downtime.

If you do not plan to use the Filter you can safely skip this chapter.

The Filter can run on:

- **Out-of-line servers** – The Filter sends a BGP announcement to the border router or route reflector that sets the Filter's server as the next hop for the suspect traffic. The Filter then blocks the malicious traffic, and the cleaned traffic is routed back into the network.

The technique used to send only the traffic received by attacked destinations to the filtering server for cleaning is called traffic diversion, BGP off-ramping, sink hole routing, side filtering etc. Traffic diversion is explained in detail in Appendix 3 – Configuring Traffic Diversion (page 75).

To run the Filter in this mode, set the Traffic Diversion parameter. Other parameters must be set as in the routing mode explained below.

- **In-line servers acting as routers** – The Filter runs on a server that resides in the main data path, configured as an OSI Layer 3 router. The necessary steps for enabling IP forwarding are found in your Linux distribution manual.

To run the Filter in this mode, set the interface connected to the peering/border router as Inbound Interface. To inject the packets back into the network, set a core router as the default gateway, reachable through the Outbound Interface either directly or through a GRE/IP in IP tunnel.

- **In-line servers acting as network bridges** – The Filter runs on a server that resides in the main data path, configured as an OSI Layer 2 network bridge. The necessary steps for configuring a network bridge are found in your Linux distribution manual.

To run the Filter in this mode, set the Inbound Interface to the bridged interface, usually br0.

- **Servers connected to network taps or mirroring ports** – The Filter runs on a server that receives a copy of packets from a network tap or a mirroring port. Direct filtering is not possible, but the Filter is able to generate filtering rules that improve the visibility of attacks and can be applied on other in-line appliances or firewalls.

To run the Filter in this mode, set the Inbound Interface to be the same as the Sniffing Interface configured in the Sniffing Sensor.

The main function of the Filter is to find **attack patterns** and to translate them into **filtering rules**.

Each attack pattern is formed by malicious packets that share some common OSI Layer 3/Layer 4/Layer 5 fields. When an attack is launched from a non-spoofed IP address, the attack pattern is the IP of the attacker. When

the attack is spoofed and comes from random IPs, the attack pattern can be the source TCP or UDP port, the destination TCP or UDP port, IP protocol number, a common packet length, TTL, etc. When the Filter detects multiple attack patterns, it generates only the filtering rule(s) that have the least negative impact on normal customer traffic.

To configure an existing Filter, go to Configuration » Components and click the Filter's name. To add a Filter, click the <+> button from the title bar of the Configuration » Components panel.

The Filter Configuration window contains the following fields:

- **Filter Name** – A short name that will help you to identify the Filter system.
- **Devices Group** – Optional description used within the Console to group multiple components by location, role, etc.
- **Graph Color** – The color used in graphs for the Filter. The default color is a random one, but you can change it by entering another HTML color code or by clicking the drop-down menu.
- **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 may have to configure them first, using the *vconfig* command.

- **Sniffing Interface** – This switch configures the interface listened 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 very high because the Filter continuously inspects the 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 when the sniffing 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 may have to 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 filter rules are applied when the Filter generates a filtering rule.

The Filter does inbound software-based packet filtering and packet rate limiting using the Netfilter framework included in the Linux kernel. The software-based packet filter is very flexible. The Filter doesn't need the connection tracking mechanism specific to stateful firewalls, making the software-based packet filter very fast as well.

Available Software Filtering Policies:

- *No software filtering* – The Filter detects and reports filtering rules. The Linux firewall API is not used. You can implement other filtering commands using custom scripts executed by responses.
- *Drop filtering rules and forward valid traffic* – The Filter detects, reports and applies filtering rules. If the filtering rule is not whitelisted, then the traffic matched by it is dropped. The rest of the traffic is forwarded through the outbound interface.
- *Drop filtering rules and forward rate-limited valid traffic* – The Filter detects, reports and applies filtering rules and forwards rate-limited valid traffic. If the filtering rule is not whitelisted, then the traffic matched by it is dropped. The Filter system will not forward traffic that exceeds the anomaly's decoder packets/second threshold value.
- *Rate-limit filtering rules and forward valid traffic* – The Filter detects and reports filtering rules and rate-limits traffic to threshold values. The Filter forwards only the traffic matching the filtering rule that does not exceed the anomaly's decoder packets/second threshold value.
- *Apply the default FORWARDING policy* – The Filter detects and reports the filtering rules, and the default forwarding policy is applied. The Netfilter framework is still being used, but the rules have the “RETURN” target. This is used only for debugging Netfilter rules.
- *Drop filtering rules and accept valid local traffic* – The Filter detects, reports and applies filtering rules. If the filtering rule is not whitelisted, then the traffic matched by it is dropped. The rest of the traffic can be accepted by local services.
- *Drop filtering rules and accept rate-limited local valid traffic* – The Filter detects, reports and applies filtering rules and accepts rate-limited traffic to local services. If the filtering rule is not whitelisted, the traffic matched by it is dropped. Local services will not receive traffic that exceeds the anomaly's decoder packets/second threshold value.
- *Rate-limit filtering rules and accept local valid traffic* – The Filter detects and reports filtering rules, and rate-limits traffic to the threshold values. The Filter accepts only the traffic matching the filtering rule that does not exceed the anomaly's decoder packets/second threshold value.
- *Apply the default INPUT policy* – The Filter detects and reports filtering rules, and the default Netfilter INPUT policy is applied. The Netfilter framework is still being used, but all rules have the “RETURN” target. This is used only for debugging Netfilter rules.
- **HW Filtering Policy** – The Hardware Filtering Policy lets you select which hardware filters are applied when the Filter detects a filtering rule.

Available Hardware Filtering Policies:

- *No hardware filtering* – The Filter detects and reports filtering rules, but no hardware-based filters are applied.
- *Use Intel x520 or x540 10 Gbps NIC to block IPv4 sources* – The Filter blocks the sources of the attacks, but only if the sources are non-spoofed IPv4 addresses. For protecting against attacks from random IP sources, also define a SW filtering policy.
- *Use Intel x520 or x540 10 Gbps NIC to block IPv4 destinations* – The Filter blocks the IPv4 destinations of the attacks for any attack patterns it finds. Similar to BGP-based black-holing.
- *Use Silicom Director 10 Gbps NIC with PF_RING HW filters* – 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, but the Silicom Director states that it can filter them.

- **Use PF_RING** – Enable if you have PF_RING installed on the server. PF_RING provides high-speed packet analysis by decreasing the CPU usage of the Filter.
- **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 58. 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 no attack pattern detected will be deleted until the attack stops and the Filter becomes inactive. Usually, an attack pattern is associated with a filter (see Filtering Policy below).
- **Sampling (1/x)** – Must be equal to the number of filtering servers activated for the same anomaly. The value must be 1 when the Filter is not used in a clustered architecture where each filtering server receives traffic from a round-robin packet scheduler.
- **Whitelists** – A Filter Whitelist is a collection of user-created rules that prevent the filtering of critical traffic types. If the filtering policy permits, the Filter might filter attack patterns that you really don't want to be filtered.

The Filter drops destination ports and destination IP addresses only in worst-case scenarios when no other attack pattern is detected. In some cases, it is better 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, the Filter might filter port 53 UDP traffic towards your DNS server, making it partially unreachable. In this case, it is 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
- **Decoder** – You can choose what decoder the rule will match.
- **Rule Type** – Which filtering rules 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*.
- **Rule 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 with a white flag icon on its side.

- **Comments** – Comments about the Filter configuration can be saved here. These are not visible elsewhere.

Enable the Filter by clicking the gray square button next to the Filter's name from Configuration » Components. The Filter will run when the “Activate a Filter...” action is executed by a response to a traffic anomaly.

Check that the Filter runs properly by watching the events log (details on page 66).

Scheduled Reports

One of the greatest strengths of the Console is the ease with which it can generate complex Reports. Most reports created by clicking items from the Side Region can be printed, exported as PDFs or sent by email.

If you want to receive periodic reports by email without having to log into the Console, 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 by separating addresses with a 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 preconfigured time, enter a description and your email address, and then click the <Save & Execute Now> button. You should receive an email with the report in few seconds.

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

Events Reporting

“Events” are short text messages generated by WANGUARD components and logged by Console. You can see them in Reports » Components » [Component Name] » [Component Type] Event sub-tab. To search, sort or filter events, click the small down arrow that appears when hovering over the Event column header.

To see a live list with the **Latest Events**, click the small bottom edge of the window to raise the south region, or press Ctrl+E. The Latest Events tab displays on one side the 30 latest events, and on the other side the list of components that can generate events and the number of events generated in the last 24 hours by each component. The number's color indicates the maximum severity of the events: red means that there are ERRORS, blue is for INFO events, etc.

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

- **MELTDOWN** – Meltdown events are generated when a very serious error has occurred, such as a hardware error.
- **CRITICAL** – Critical events are generated when a significant software error is detected, like a memory exhaustion situation.
- **ERROR** – Error events are usually caused by misconfigurations or communication errors between components.
- **WARNING** – Warning events are generated when authentication errors occur, on I/O bottlenecks or when there are time synchronization issues.
- **INFO** – Informational events are generated when configurations are changed or when users log into the Console.
- **DEBUG** – Debug events are generated only to help with troubleshooting coding errors.

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

Users Management

To manage Console accounts or to configure authentication mechanisms, go to the Side Region and select Configuration » Global Settings » Users Management.

Each Console account must belong to one of the 3 available access levels or “roles”:

- **Administrator** – Has all privileges. Can manage accounts and reset passwords. Cannot view plain-text passwords because all passwords are encrypted. Is the only role able to access Configuration » Global Settings » License Manager.
- **Operator** – Can change any configuration but is not allowed to modify other accounts.
- **User** – Has read-only access to the Console, and all configurations are hidden. Provides **permission-based access** to reports, dashboards, sensors, IP groups, regions, etc.

To modify an account, double-click on it, or first select it and then press <Modify User>.

The **Full Name**, **Company**, **Position**, **Email**, **Phone** and **Comments** fields are optional and not used anywhere.

The **Landing Tab** list shows the tab that will be opened immediately after logging in. The list is dynamic and expands as you add sensors, dashboards, IP groups etc. Change the Landing Tab to a relevant dashboard or report.

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

The **Side Region Position** field lets you switch the Side Region's position to east or west.

The **Console Theme** field lets you change the Console's theme after re-logging in. Blue and gray are the most popular themes.

The **Authentication & Login** button provides LDAP and RADIUS-based authentication settings, and lets you set a MOTD message visible on the Login page.

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 networking, read about the technical basics in this appendix. 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 IP 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 the first two bits set to “1” and the 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 the first three bits set to “1” and the 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 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 on Cisco 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 the server running the Flow Sensor 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. The Flow Sensor drops flows older than 5 minutes!

```
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 the server running the Flow Sensor 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. The Flow Sensor drops flows older than 5 minutes!

```
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. The Flow Sensor drops flows older than 5 minutes!

```
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 (non-MX)

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 the WANGUARD Filter. The information provided here regarding router configurations is for informational purposes only. Please refer to the appropriate router user guides for detailed information.

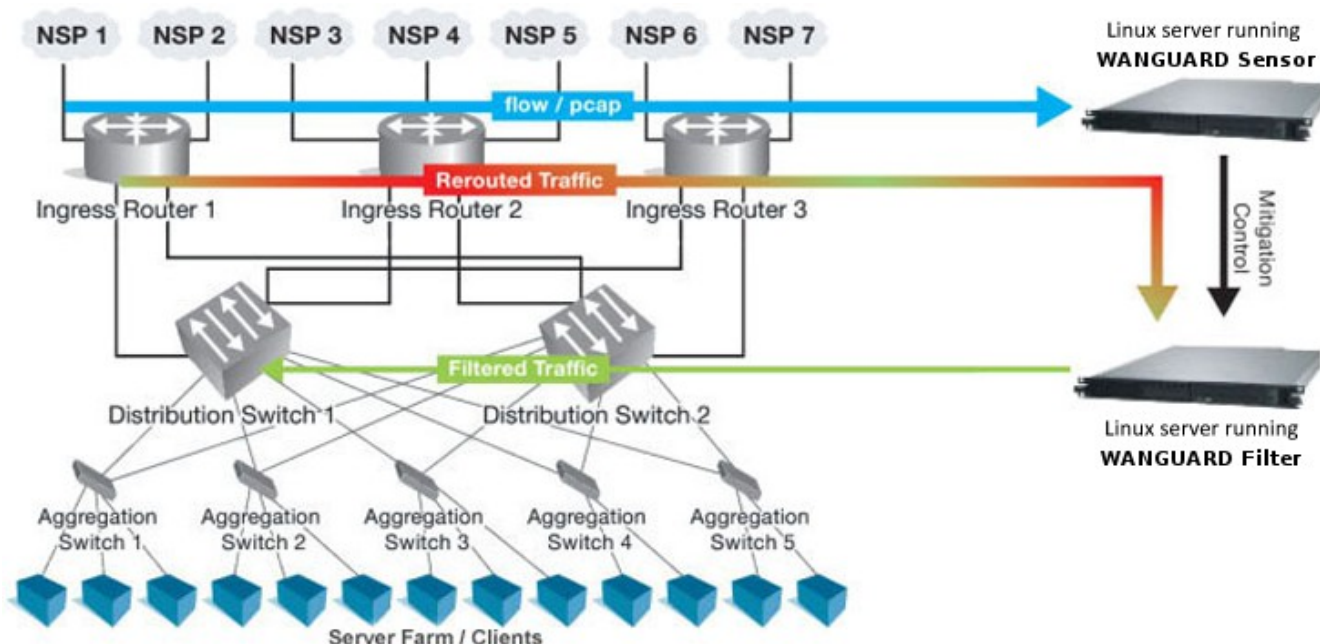
Understanding the BGP Diversion Method

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, the Filter sends a routing update where the Filter's server is listed as the best path for the attacked destinations.

The network prefix that the 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 to a Linux server running a WANGUARD Filter.

After BGP diversion is established, the router's routing tables point to the Filter server's IP address as the best route for 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 the Filter diverts traffic. The following examples are provided using common External Border Gateway Protocol v4 (eBGP). You should consider your own network configuration and determine whether eBGP or iBGP should be implemented in your network.

Follow these guidelines when the Filter server 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 appears in the router's routing tables. Using a recognizable value allows you to easily identify the Filter server 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 a 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 its configuration files in */etc/quagga* while Zebra keeps its 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 usually enough to get them started. You should replace “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 of 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 uses 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 configuring traffic diversion. The syntax in the commands is taken from the BGP configuration on a Cisco router.

The following configuration steps shows the commands used 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 section). The 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 to be 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 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 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 at the end point of the tunnel interface only, not at the zone's address.

To use this method, you must 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 60) with the Outbound Interface set to the virtual network interface created by the tunnel.

Appendix 4 – Software Changelog

WANGUARD 5.4

Released on 09 March 2014

- New decoders for the latest DDoS attacks: NTP, RDP and SNMP.
- Both Sensors can generate tops and graphs for Countries and tops for external IPs. To enable set the Top Generator parameter to Extended for the Sniffing Sensor and to Full for the Flow Sensor.
- Both Sensors can generate tops and graphs for Autonomous Systems based on GeoIP data.
- Add IPFIX compatibility with the Juniper MX and with other flow exporters that maintain the Start Time of flows.
- New API for manipulating IP Zones and Sensors from the CLI, located in /opt/andrisoft/api.
- All widgets in a dashboard can share a single Time Frame. This can be enabled from the Dashboard Configuration.
- Updated Autonomous Systems (Organizations) database with the latest entries from Internet Routing Registries.
- New FLOWSYN decoder that matches all packets from flows with the SYN flag set.
- New keyboard shortcuts: Ctr+E toggles South Region, Ctrl+R toggles Reports/Configuration, Ctrl+← and Ctrl+→ switches Central Region tabs, Alt+Ctrl+← and Alt+Ctrl+→ switches bottom tabs.
- New Flow Collector output mode called "Debug" that lists flows and tops from all flow exporting interfaces, not only from the interfaces monitored by the Flow Sensor.
- Communication between Console and remote components can be encrypted using SSL.
- New "Per All" option for the Profiling Data parameter in IP Zones that enables profiling for the included IPs and subnets.
- Unprivileged users can execute Anomaly Actions if the proper permission is enabled in the user profile.
- New roles for BGP Connections to see the number of BGP announcements from Reports » Alarms & Tools » BGP Prefixes in red for "Black-holing" and in blue for "Diversion".
- New Dashboard Widgets: Flows List and Flows Top
- New Dynamic Parameters: {filter_ip_isp} and {filter_ip_country} to get the ISP and the Country of attacking IPs.
- New button in Reports » Alarms » Anomalies » Anomalies Archive to force the clearing of active anomalies.
- New button in Reports » BGP Prefixes » BGP Archive to force the clearing of active BGP announcements.
- Reports » Tools » Flow Collector » Autonomous Systems moved to Reports » Components » Sensor » AS Graphs.

- Configuration » Components panel title button enables starting/stopping of all Components with a single click.
- Fix Console slowdown caused by the summarization of Events in the South Region » Latest Events when the Events Log is very big.
- Better User Manual & Admin Guide
- Fix for the Filter TTL bug, false positives on outbound attacks on decoders other than TOTAL, fix HTML code for Scheduled Reports
- Few other smaller bug fixes.

WANGUARD 5.3

Released on 23 December 2013

- Reports->Server tabs with dashboards, graphs, command outputs.
- Filter tabs with dashboards, graphs, filtering rules.
- Complex SQL Filter for all tables + Selectable Columns.
- Anomalies can be manually classified.
- Actions (eg. manual mitigation) can be manually activated for active anomalies.
- A new Anomalies Overview widget.
- The Quagga's bgpd.conf file is editable from the BGP Connection configuration window.
- Data Retention settings for packet dumps.
- Events logs are now also present in many Reports tabs.
- IP Zone shows storage requirements for IP graphs, IP accounting and Profiling Data.
- WANsupervisor generates warning events when the server has a disk full, no free RAM or a high System Load.
- Console users can be disabled.
- Added a new SIP decoder for VoIP traffic.
- Scheduled emailed reports for IP Groups, Sensors, Filters and Servers.
- Graphs showing totals can be stacked.
- Packet analyzer shows hex and ASCII dumps of packets.
- Few bugfixes.

WANGUARD 5.2

Released on 16 August 2013

- Full IPv6 support for the Sensors and for the Filter.
- The Filter now supports all decoders.
- The Filter supports In-NIC hardware filters present in Intel x520 10Gbps network cards.
- Extended the Packet Analyzer with auto-stop functions.
- Console GeoIP support to see the countries of the attackers.
- The Sniffing Sensor can use native PF_RING functions.
- Events logs are now also present in many Reports tabs.
- Added IPv4 and IPv6 mask restrictions to BGP announcements.
- Graphing IP sweeps can be disabled.
- Numerical Dynamic Parameters can now be obtained in a shorter form by appending `_kilo`, `_mega`, `_giga`, `_prefix` (auto).
- Added new decoders: Flows (Flow Sensor only), SSH, Youtube, NetFlix, Hulu.
- Redesigned the Configuration Side Area.
- Right clicking on Reports->Anomalies & Tools directly open Archives.
- Many bugfixes.
-

WANGUARD 5.1

Released on 16 March 2013

- Multiple CPU support for packet sniffing at 10 Gbps wire-speeds and beyond.
- The web interface automatically adjusts to time-zone differences.
- Console now supports multiple RADIUS and LDAP servers, and different remote user profiles.
- Improvements for Anomaly Reports. The histogram time interval can be changed. Mitigation logs can be included.
- New Dynamic Parameters {ip_dns}, {attacker_isp}, {anomaly_log}, {filter_log}, {filter_ip_dns}.
- Console provides loading feedback for time-consuming tasks and actions.
- The Flow Sensor can now analyze flows from routers in different time-zones.
- Percentage thresholds can be configured with minimal packets/second and bits/second values.
- A new decoder for WWW traffic.
- Various bug-fixes.

WANGUARD 5.0

Released on 20 October 2012

- A completely new Flow Collector interface. It allows easy navigation into flow data and provides powerful statistics and summaries
- A brand new Packet Analyzer. You can now capture packets using just a few clicks and then view the dumps in detail in a Wireshark-like web interface
- A new Configuration Wizard
- A new License Manager available for Administrators. It's flexible, allows rebranding and customisations
- Users can change their passwords, Themes and Side Region position
- Combined Reports for IPs, IP Groups and Sensors
- Two new separate Dashboards for IP Reports and Sensor Reports
- Most Reports can be sent by Email
- Sensor graphs can be summed
- Added lots of new Sensor graphs: average packet size, CPU%, RAM, no. of IP graphs, no. of IP Accounting records etc.
- Easier IP Zone configuration that can be listed
- Dashboard permissions
- Dashboard widget hierarchy
- A new HTML widget for Dashboards
- Custom fields for Live Sensor Stats Widget
- ASN graph widget
- Unified panels in an intuitive manner
- Bookmarks lets you save frequently used, manually entered data
- A new "Quick Search" button with full Reports functionality
- Live & archived tops by "IP Group" and "IP version"
- Perpetual sessions
- Sensors support configurable traffic decoders everywhere: stats, tops, graphs, accounting
- Flow engine rewrite. Now supports NetFlow v9, IPFIX and native sFlow
- Flows can be collected in an efficient binary format
- Flow Sensor consumes less RAM
- Flow Sensor supports "Mixed" traffic interfaces
- Flow Sensor supports multiple time-zones
- Flow Sensor restarts itself gracefully if it doesn't receive flows for a long time

- 32bit AS numbers support for Sensors
- Sniffing Sensor adds a new Traffic Capturing framework. Supports full captures, file rotations and advanced sampling
- Protocols distribution generator takes into account the traffic's direction
- The number of top items is configurable for Sensors
- Sensors compatible with the new Server-based configuration
- Filter compatible with the new thresholds system
- Filter compatible with the new Server-based configuration
- A.T.L.A.S. - Andrisoft's Threat Level Analysis System that enables managed security services and remote NOC supervision
- Console shows live attacks in much more detail and includes additional actions: add a comment, withdraw BGP announcement, generate a Report
- Detailed Anomaly Reports. Can be sent automatically by email
- Decoder-based thresholds
- Prediction-based thresholds
- Percentage-based thresholds
- Minimum thresholds
- Non-inheritable thresholds
- Per subnet thresholds
- Each threshold rule can have it's own Response
- Anomalies Overview Report
- BGP announcements grouping by Router or IP/Mask
- Configurable BGP announcements timeouts
- Customizable Anomalies expiration time
- Added a link severity parameter
- Severity bar is colored to indicate the link's severity