

# Wansight 8.0 User Guide

- Console
- Packet Sensor, Flow Sensor, SNMP Sensor, Sensor Cluster

## Copyright & Trademark Notices

This edition applies to version 8.0 of the licensed program Wansight and 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. sales department, [sales@andrisoft.com](mailto:sales@andrisoft.com).

## Copyright Acknowledgment

© 2020, ANDRISOFT S.R.L. All rights reserved.

All rights reserved. This document is copyrighted, and ANDRISOFT S.R.L reserves all rights. 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 and Wansight are SOFTWARE PRODUCTS 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.**

**Website:** <https://www.andrisoft.com>  
**Sales and pre-sales:** [sales@andrisoft.com](mailto:sales@andrisoft.com)  
**Technical support:** [support@andrisoft.com](mailto:support@andrisoft.com)

© 2020, ANDRISOFT. All rights reserved.

## Table of Contents

|   |    |
|---|----|
| 1. Traffic Monitoring and IP Accounting with Wansight.....                | 5  |
| Key Features & Benefits.....  | 5  |
| Software Components.....  | 5  |
| 2. Choosing a Method of Traffic Monitoring.....                           | 7  |
| Comparison between Packet Sniffing, Flow Monitoring and SNMP Polling..... | 8  |
| 3. System Requirements and Software Installation.....                     | 9  |
| System Requirements.....  | 9  |
| Software Installation.....  | 12 |
| Opening the Console.....  | 12 |
| Licensing Procedure.....  | 13 |
| Quick Configuration Steps.....  | 13 |
| 4. Basic Concepts of Wansight Console.....                                | 14 |
| 5. Configuration » General Settings » Graphs & Storage.....               | 16 |
| Sensor and Applications Graph Troubleshooting.....                        | 19 |
| IP/Subnet Graph Troubleshooting.....                                      | 19 |
| AS and Country Graph Troubleshooting.....                                 | 19 |
| 6. Configuration » General Settings » Custom Decoders.....                | 20 |
| 7. Configuration » Network & Policy » IP Zone.....                        | 22 |
| 8. Configuration » Servers.....   | 24 |
| Distributing the Software over Multiple Servers.....                      | 25 |
| Server Troubleshooting.....   | 25 |
| 9. Configuration » Components » Packet Sensor.....                        | 26 |
| Packet Sensor Troubleshooting.....  | 28 |
| Packet Sensor Optimization Steps for Intel 82599.....                     | 29 |
| 10. Configuration » Components » Flow Sensor.....                         | 30 |
| Flow Sensor Troubleshooting.....  | 33 |
| 11. Configuration » Components » SNMP Sensor.....                         | 35 |
| SNMP Sensor Troubleshooting.....  | 37 |
| 12. Configuration » Components » Sensor Cluster.....                      | 38 |
| 13. Configuration » Schedulers » Scheduled Reports.....                   | 39 |
| 14. Configuration » Schedulers » Event Reporting.....                     | 40 |
| 15. Configuration » General Settings » Outgoing Email.....                | 41 |
| 16. Configuration » General Settings » User Management.....               | 42 |
| 17. Configuration » General Settings » User Authentication.....           | 44 |
| 18. Reports » Tools » Flows.....  | 46 |
| Flow Records.....   | 46 |
| Flow Tops.....  | 46 |
| 19. Reports » Tools » Packets.....  | 48 |
| Active Packet Dumps.....  | 48 |
| Packet Dump Archive.....  | 49 |
| 20. Reports » Devices » Overview.....                                     | 50 |
| 21. Reports » Devices » Sensors.....                                      | 54 |

|   |    |
|---|----|
| Sensor Dashboard.....                                 | 54 |
| Sensor Graphs.....                                    | 54 |
| Sensor Tops.....                                      | 55 |
| Sensor Events.....                                    | 56 |
| AS Graphs.....  | 56 |
| Country Graphs.....                                   | 57 |
| Flow Records.....                                     | 58 |
| Flow Tops.....  | 58 |
| 22. Reports » Dashboards.....                         | 59 |
| 23. Reports » IP Addresses & Groups.....              | 60 |
| IP Dashboard.....                                     | 60 |
| IP Graphs.....  | 60 |
| IP Accounting.....                                    | 61 |
| Flow Records.....                                     | 62 |
| Flow Tops.....  | 62 |
| 24. Reports » Servers.....                            | 63 |
| Console / Server Dashboard.....                       | 63 |
| Console / Server Graphs.....                          | 63 |
| Console Events.....                                   | 64 |
| Server Events.....                                    | 64 |
| Server Commands.....                                  | 64 |
| 25. Appendix 1 – IPv4 Subnet CIDR Notation.....       | 65 |
| 26. Appendix 2 – Configuring NetFlow Data Export..... | 66 |
| Configuring NDE on older IOS Devices.....             | 66 |
| Configuring NDE on a CatOS Device.....                | 67 |
| Configuring NDE on a Native IOS Device.....           | 67 |
| Configuring NDE on a 4000 Series Switch.....          | 68 |
| Configuring NDE on IOS XE.....                        | 68 |
| Configuring NDE on IOS XR.....                        | 68 |
| Configuring NDE on IOS-XE Everest.....                | 69 |
| Configuring NDE on a Juniper Router (non-MX).....     | 70 |
| 27. Appendix 3 – DPDK Configuration.....              | 71 |

## Traffic Monitoring and IP Accounting with Wansight

Andrisoft Wansight is an award-winning enterprise-grade software which delivers to NOC and IT teams the functionality needed for monitoring the traffic of large WAN networks. Andrisoft Vanguard extends Wansight with advanced anomaly detection and DDoS mitigation capabilities. If these capabilities are also needed, the upgrade to Vanguard can be done by uploading a new license key.

### Key Features & Benefits

- ✓ **FULL NETWORK VISIBILITY** – Supports all major IP traffic monitoring technologies: packet sniffing, NetFlow version 5, 7 and 9; sFlow version 4 and 5; IPFIX and SNMP
- ✓ **FAST, SCALABLE & ROBUST** – Designed to run on commodity server hardware by leveraging high-speed packet capturing technologies such as DPDK, PF\_RING Vanilla, PF\_RING ZC and Netmap. Can run as a cluster with its software components distributed across multiple servers
- ✓ **ENTERPRISE-GRADE WEB CONSOLE** – Provides consolidated management and reporting through a highly-configurable multi-tenant web portal with customizable dashboards, user roles, and remote authentication
- ✓ **PACKET SNIFFER** – Saves packet dumps using a distributed packet sniffer that can be deployed on different network entry points. Displays packet details in a Wireshark-like web interface
- ✓ **FLOW COLLECTOR** – Contains a fully-featured NetFlow, sFlow, and IPFIX collector that saves flow data in a compressed format for long term storage. Flows can easily be searched, filtered, sorted, and exported
- ✓ **COMPLEX ANALYTICS** – Generates 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** – Generates PDF and HTML reports and sends them automatically by email to the interested parties at preconfigured intervals of time
- ✓ **COMPLETE REST API** – All configurations and collected data can be easily queried and referenced via a fully-featured RESTful API which exposes hundreds of internal parameters, anomaly data, graphs and tops

All configurations are stored in an SQL database that is easy to query, backup and restore.

### Software Components

**Wansight Sensor** provides in-depth traffic analysis, traffic accounting, and bandwidth monitoring. The collected information enables complex traffic reports with graphs and tops, which help you to instantly pin down the cause of network incidents, understand patterns in application performance and make the right capacity planning decisions.

**Wansight Console** is a multi-tenant web graphical user interface that functions as the administrative core of the software. It offers single-point management and reporting by consolidating the data from all Wansight Sensors deployed within the network.

For the sake of brevity, Wansight Sensor is hereby referred to as the Sensor, and Wansight Console as the Console.

## Choosing a Method of Traffic Monitoring

This chapter describes the traffic monitoring technologies supported by Wansight Sensor. There are four Wansight Sensor “flavors”, each having a different way of obtaining traffic information:

- **Packet Sensor** analyzes packets. It can be used on appliances that are either deployed in-line (servers, firewalls, routers, bridges, IDSes, load-balancers) or connected to a mirrored port or TAP.

*In switched networks, only the packets for a specific device reach the device's network card. If the server running a Packet 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 copies of data packets traveling through selected ports or VLANs to the monitoring port. Packet Sensor inspects every packet it receives and conducts packet-based traffic analysis*

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

*Many routers and switches can collect IP traffic statistics and periodically send them as flow records to a Flow Sensor. Because the flow protocol already performs pre-aggregation of traffic data, the flow data sent to Flow Sensor is much smaller than the monitored traffic, and this makes Flow Sensor a good option for monitoring remote or high-traffic networks. The main downside of flow-based traffic analysis is that pre-aggregating traffic data adds a delay of at least 30 seconds to collecting real-time traffic statistics*

- **SNMP Sensor** monitors the bandwidth usage of routers and switches on a port-by-port basis.

*When this technology is used, an SNMP Sensor queries the device (e.g. router, switch, server) for the traffic counters of each port with small data packets. These are triggering reply packets from the device. Compared to other bandwidth monitoring technologies, the SNMP option is very basic and offers no IP-specific information. SNMP creates the least CPU and network load*

- **Sensor Cluster** aggregates pre-existing Sensor traffic data into a single, unified IP graphing domain.

*Sensor Cluster sums up the traffic data collected by Packet Sensors, Flow Sensor and SNMP Sensor interfaces and performs the same tasks as the other Sensors (IP graphing, IP accounting, etc.)*

For redundancy, high availability and to be able to view packet dumps and flow data, you can deploy Flow Sensor(s) and Packet Sensor(s) simultaneously.

## Comparison between Packet Sniffing, Flow Monitoring and SNMP Polling

Packet Sensor is recommended when the speed of detecting attacks is critical, or when there is a need for capturing raw packets for forensics and troubleshooting. Because every packet entering the network is inspected, Packet Sensor needs to run on servers with powerful CPUs.

Flow Sensor analyzes pre-aggregated traffic information sent by routers and switches, so it can monitor traffic passing through multiple 10/40/100 GbE interfaces even when it runs on a low-end server. By comparison, Flow Sensor has some disadvantages:

- x It exhibits reduced speed in processing real-time traffic information. Flow exporters aggregate traffic data over time, making the traffic visible only after a delay (flow aging) that usually exceeds 30 seconds
- x It provides slightly less accurate traffic readings because in most cases the packets or flows are sampled
- x Enabling the flow exporter functionality may result in an increased CPU load on the network device when the flow collection is not performed in hardware
- x Flows can be dropped if a powerful spoofed DDoS attack fills the TCAM of the network device

SNMP Sensor is useful to monitor devices that cannot export flows or mirror packets, or to compare flow and SNMP-derived statistics in order to ensure the flow data's accuracy.

|   | Packet Sensor   | Flow Sensor   | SNMP Sensor   |
|---|---|---|---|
| <b>Traffic Monitoring Technology</b>        | <ul style="list-style-type: none"> <li>• Sniffing packets passing an in-line appliance</li> <li>• Port mirroring (SPAN, Roving Analysis Port)</li> <li>• Network TAP</li> </ul> | <ul style="list-style-type: none"> <li>• NetFlow version 5, 7 and 9 (jFlow, NetStream, cflowd)</li> <li>• sFlow version 4 and 5</li> <li>• IPFIX</li> </ul> | <ul style="list-style-type: none"> <li>• SNMP version 1</li> <li>• SNMP version 2c</li> <li>• SNMP version 3</li> </ul> |
| <b>Maximum Traffic Capacity per Sensor*</b> | 100 GigE  | multiples of 100 Gbps   | multiples of 100 Gbps   |
| <b>DDoS Detection Time**</b>                | ≤ 1 seconds   | ≥ flow aging time (≥ 30 seconds usually) + 5 seconds  | ≥ 5 seconds   |
| <b>IP Graph Granularity</b>                 | ≥ 5 seconds   | ≥ 20 seconds  | N/A (SNMP offers no details about IPs)  |
| <b>Traffic Validation Options</b>           | IP classes, MAC addresses, VLANs, BPF   | IP classes, Interfaces, AS Numbers, Ingress/Egress  | Interfaces  |
| <b>Packet Dumps</b>                         | Yes   | No  | No  |
| <b>Flow Collector</b>                       | No  | Yes   | No  |

\* The number of connections between IPs is not a limiting factor

\*\* Vanguard Sensor detects the IPs under attack. Vanguard Filter detects the sources of each attack



## System Requirements and Software Installation

Installing Wansight will not generate any negative side effects on the network's performance. Full installation and configuration may take less than an hour.

Wansight runs exclusively on Linux platforms. To install and configure the software you need basic Linux operation skills and at least medium computer networking skills. If you encounter software installation issues or if you have questions about the system requirements listed below, please contact [support@andrisoft.com](mailto:support@andrisoft.com).

### System Requirements

Wansight 8.0 can be installed on the following 64-bit Linux distributions: CentOS 7 or 8 (free, Red Hat-based), Red Hat Enterprise Linux 7 or 8 (commercial), Debian Linux 7 to 10 (free, community-supported), Ubuntu Server 14 to 20 (free, Debian-based).

Wansight was designed to be completely scalable, so it can be installed either on a single server with adequate hardware resources or on multiple servers distributed across the network. It is highly recommended to install the software on dedicated servers and not on Virtual Machines, mainly because:

- Having fast and uninterrupted access to the hard disk is a critical requirement of the Console
- The resources must be provisioned in a predictable and timely manner
- Some virtualized environments do not have a stable-enough clock source

| Importance of HW resources | CPU Speed (> GHz/core) | CPU Cores (> cores) | RAM Size (> GB) | HDD Size (> GB) | HDD/SSD Speed (> Mbytes/s) | Network Adapter (Vendor, Model) |
|----------------------------|------------------------|---------------------|-----------------|-----------------|----------------------------|---------------------------------|
| Console                    | High                   | High                | High            | Very High       | Very High                  | Very Low                        |
| Packet Sensor              | Very High              | High                | Medium          | Low             | Low                        | Very High                       |
| Flow Sensor                | Low                    | Low                 | High            | Medium          | High                       | Very Low                        |
| SNMP Sensor                | Very Low               | Low                 | Very Low        | Very Low        | Very Low                   | Very Low                        |
| Sensor Cluster             | Medium                 | Medium              | Medium          | Very Low        | Very Low                   | Very Low                        |

|               |                      |                 |                   |                |                     |
|---------------|----------------------|-----------------|-------------------|----------------|---------------------|
| <b>Legend</b> | Very High Importance | High Importance | Medium Importance | Low Importance | Very Low Importance |
|---------------|----------------------|-----------------|-------------------|----------------|---------------------|

## Minimum System Requirements for Console

| Capacity     | <10 components (Sensors, Filters, BGP Connectors)         |
|--------------|---|
| Architecture | 64-bit x86  |
| CPU          | 1x 2.4 GHz quad-core Xeon                                 |
| RAM          | 1 x 8 GB  |
| NICs         | 1 x Fast Ethernet for management                          |
| HDDs         | 2 x 7200 RPM HDD (SSD highly recommended), RAID 1, 350 GB |

The Console server stores the database and centralizes all operational logs, graphs and IP accounting data. Its performance is determined by its configuration, the performance of the I/O and the performance of the applications it relies on: MySQL/MariaDB, Apache HTTPD and PHP.

To access the web interface, use one of the following web browsers: Google Chrome 64+, Firefox 52+, Safari 10+, Microsoft Edge 12+, Opera 43+. JavaScript and cookies must be enabled. Java and Adobe Flash are not required. For the contextual help you may need to install Adobe PDF Reader.

For the best experience, we recommend using Google Chrome and a 1280x1024 or higher resolution display.

Firefox shows mismatched colors on some PCs. You can fix it by enabling color management:

1. Type in about:config on your Firefox address bar
2. Use the filter field to search for color\_management
3. Set gfx.color\_management.mode to 1
4. Set gfx.color\_management.enablev4 to true
5. Restart Firefox

## Minimum System Requirements for Packet Sensor

| Capacity     | 10 Gbit/s, 14 Mpkts/s (wire rate)   | 40 Gbit/s, ±30 Mpkts/s  |
|--------------|---|---|
| Architecture | Intel Xeon 64-bit, dedicated server   | Intel Xeon 64-bit, dedicated server   |
| CPU          | 1x 2.4 GHz Xeon E5-2640v4   | 1 x 2.4 GHz Xeon E5-2680v4  |
| RAM          | 4 x 2 GB DDR4 (quad channel)  | 4 x 8 GB DDR4 (quad channel)  |
| NICs         | 1 x 10 GbE adapter (Intel 82599+ or PF_RING/DPDK-supported chipset)<br>1 x Fast Ethernet for management | 1 x 40 GbE adapter (Intel XL710+ or other DPDK-supported chipset)<br>1 x Fast Ethernet for management |
| HDDs         | 2 x 5400 HDD, RAID 1, 10 GB (including OS)  | 2 x 5400 HDD, RAID 1, 10 GB (including OS)  |

Packet Sensor can run load-balanced over multiple CPU cores with the following hardware / Capture Engines:

- Intel 82599 chipset network adapters, such as Intel X520, Intel X540, HP X560 or Silicom PE310G4DBi9-T
- PF\_RING (with or without ZC) high-speed packet I/O framework

- Netmap high-speed packet I/O framework and its supported NICs
- Data Plane Development Kit (DPDK) and all its supported NICs

You can easily scale the Packet Sensor’s capacity above 100 Gbit/s by enabling packet sampling on the switch or TAP, or by defining a Sensor Cluster that aggregates multiple Packet Sensor instances running on different servers equipped with 10, 40 or 100 Gbit/s network adapters.

### Minimum System Requirements for Flow Sensor

| Capacity     | 15000 flows/s                    |
|--------------|----------------------------------|
| Architecture | 64-bit x86                       |
| CPU          | 1 x 2.0 GHz dual-core Xeon       |
| RAM          | 1 x 8 GB                         |
| NICs         | 1 x Fast Ethernet for management |
| HDDs         | 2 x 7200 RPM HDD, RAID 1, 60 GB  |

Flow Sensor can monitor an almost unlimited number of interfaces. On modern hardware, the processing of tens of thousands of flows/s is also not a problem. Each Flow Sensor can handle the flows of a single flow exporter. Any server with enough RAM can run tens of Flow Sensor instances. For this type of Sensor, the amount of RAM is much more important than the speed of the CPU.

Flow Sensor can store flow data on the local disk in a highly compressed binary format.

### Minimum System Requirements for SNMP Sensor

| Capacity     | 20+ devices                      |
|--------------|----------------------------------|
| Architecture | 64-bit x86                       |
| CPU          | 1 x 1.6 GHz dual-core Xeon       |
| RAM          | 1 x 1 GB                         |
| NICs         | 1 x Fast Ethernet for management |
| HDDs         | 2 x 5200 RPM HDD, RAID 1, 20 GB  |

Each SNMP Sensor can monitor a single device with an unlimited number of interfaces. Any server can run an almost unlimited number of SNMP Sensor instances.

### Minimum System Requirements for Sensor Cluster

The hardware requirements for Sensor Cluster are very low because the traffic information is pre-aggregated by the associated Flow Sensor, Packet Sensor or SNMP Sensor instances. It is best to run it on the Console server.

## Software Installation

The download link is listed in the email with the trial license key. The latest software installation instructions are listed on [www.andrisoft.com](http://www.andrisoft.com).

Each trial license key activates all features for 30 days. You can install the trial license key on any number of servers. To switch to a full, registered version, apply a license key purchased from the online store.

## Opening the Console

Wansight Console provides a web interface and centralized system through which you can control and monitor all the other components. If you have correctly followed the installation instructions, from now on you will only need to log in to Console to manage and monitor servers and software components. SSH access will be needed for updating the software.

Open the Console at [http://<console\\_hostname>/wansight](http://<console_hostname>/wansight). If the page cannot be displayed, make sure that the Apache web server is running and the firewall does not block incoming traffic on port 80 or 443. You can also access it securely via HTTPS if the Apache web server was configured to serve pages over SSL/TLS.

If you have not licensed Wansight, you will be asked to do so. Upload the *trial.key* file emailed to you by clicking the key icon. The license key contains encrypted information about the licensed capabilities of the software. You can replace the license key in Configuration » General Settings » License Manager.

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

If the Console is installed on a public server, you should immediately change the default password of the “admin” account. To do so, click the **Admin** menu at the top-right corner of the browser window and select [**Change Password**]. From the same menu you can change the Console layout and theme.

To understand how to navigate within the Console, read the dedicated chapter on page 14.

## Licensing Procedure

When the trial period is over you will have to purchase as many Sensor and Filter licenses, in form of annual subscriptions, as the number of Sensors and Filters configured and enabled in Configuration » Components.

- You will have to purchase as many Sensor licenses as the number of flow exporters (usually border or edge routers) monitored by Flow Sensors. There is no limit on the number of interfaces a Flow Sensor can monitor.
- You will have to purchase as many Sensor licenses as the number of interfaces (ports) listened by Packet Sensors. Multiple Packet Sensors listening to the same interface (e.g. when using a multi-queue NIC) use a single Sensor license. Packet Sensor can monitor an unlimited number of IPs/domains
- You can mix Vanguard Sensor licenses together with Wansight Sensor licenses
- Sensor Cluster does not require licensing
- Console does not require licensing

You can distribute the licensed Sensors on any number of servers without additional licensing costs. The license key must contain the hardware keys listed under Configuration » General Settings » License Manager » Requirements. The minimum licensing period is 12 months and the maximum licensing period is 48 months.

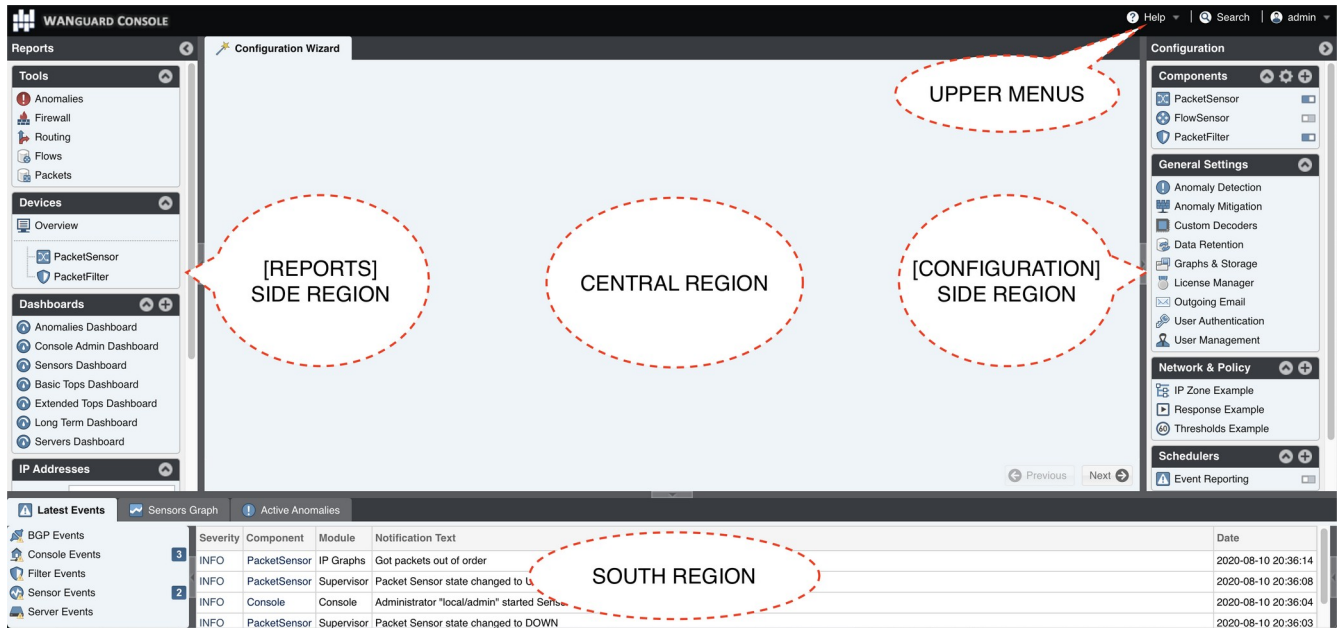
## Quick Configuration Steps

- ➔ Estimate storage requirements, review decoders and IP graph settings – page 16
- ➔ Add your IP address ranges and important hosts to an IP Zone – page 22
- ➔ Configure a Packet Sensor – page 26, Flow Sensor – page 30, or SNMP Sensor – page 35
- ➔ Watch the event log. Receive error notifications by email – page 40
- ➔ Create personalized Console accounts for your staff or customers – page 42
- ➔ Create dashboards and add widgets containing useful information – page 59

## Basic Concepts of Wansight Console

Please read this chapter in order to understand the basic premises required to properly operate the software. The next few chapters cover the software's configuration, while the last chapters cover the reporting features.

To understand how to operate the Console web interface you should be aware of its structure:



### Side Region

Is located at the east and/or west edge of the browser's window, according to the user's preference. If it is not visible, it has been either collapsed or hidden by an administrator. Clicking the edge of regions expands or collapses them.

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. These panels are constantly updated.

### Central Region

Each report, dashboard or tool you select in the Side Region opens a tab (page) in the Central Region. You can switch between (sub-) tabs with a mouse or with the keyboard shortcut (Alt+) Ctrl+ → and (Alt+)Ctrl+ ←. You can close all tabs except for the Landing Tab (initially set as the Configuration Wizard). To change the Landing Tab, edit your user profile in Configuration » General Settings » User Management. Tabs can be rearranged with drag-and-drop.

## **South Region**

South Region provides a quick way to view live data such as events (system logs) or animated traffic graphs. It is located at the bottom of the browser's window. By default, it is collapsed; to expand it, click the thin line near the lower edge or press Ctrl+E.

## **Upper Menus**

The Help menu contains links to the User Guide, to various helper tools and to Software Updates. Depending on context, the User Guide opens the chapter describing the last-opened window or tab.

The "Quick Search" button allows you to open tabs with the keyboard. Keyboard shortcut: Ctrl+S.


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

## Configuration » General Settings » Graphs & Storage

An important initial step in configuring Wansight is to make sure that the servers the software runs on have enough resources to collect IP graphs, flows and packet dumps.

In a later chapter, you will be able to configure Sensors to generate traffic graphs, tops and accounting data for every IP that belongs to the monitored network. If you intend to use this feature, you may want to change the IP storage settings, as changing these later might require a reset of all existing IP graph, top and accounting data.

All storage-related settings are located in Configuration » General Settings » Graphs & Storage.



The **Graph Storage Engine** parameter lets you choose which third-party software will store the graph data:

- **RRDTool** is installed and enabled by default. It is very stable but in some cases it can be slow (due to the disk seek time), it uses a lot of storage space (due to its preallocation of data), and some configuration changes (e.g. adding a decoder) resets the existing graph data
- **InfluxDB** is a database created specifically to store time series data. It uses more RAM than RRDTool but it needs less storage space, it is faster in many cases, it can be configured for High Availability, and existing data is not deleted by configuration changes

Wansight can use **Graph Storage Engine 1** as the only backend, or it can use two backends at the same time in the following combinations: RRDTool + InfluxDB, or InfluxDB (primary) + InfluxDB (secondary). Click on the options button from the right to configure the selected engine.

- **RRDTool (On Disk)** – This method optimizes the long-term storage of IP graph data by allowing up to three **Round Robin Archives**. The values within the Round Robin Archives determine the granularity of the graphs and the interval of time they are saved. These entries specify for how long, and how accurately the data should be stored. A smaller data average (5 seconds minimum) generates a very accurate graph,



but requires more disk space, while a bigger data average (5 minutes maximum) is less accurate but uses less disk space.

RRDTool stores each graph in a separate file, therefore on non-SSD drives the disk seek time may be too high to update thousands of IP graph files every few minutes. If this is the case, configure the **RRDCached Socket** to increase the I/O performance of the Console server ([KB article link](#)). If the speed of updating IP graph files is still not fast enough, consider the other options listed below. **Graph Data Units** lets you choose the data units stored inside IP graph files. **Graph Consolidation** lets you choose how to build consolidated values for Round Robin 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

- **RRDTool (In RAM)** – This method is not optimal for long-term storage because it allows a single Round Robin Archive per IP graph file. The files are created and updated in the **Graphs RAM Path**, and moved periodically onto a larger, albeit slower disk. Select this method when you can't use InfluxDB and when the previous method (configured with RRDCached) is not fast enough to sustain updating thousands of very high-granularity IP graphs. The other parameters are explained by the previous paragraph
- **InfluxDB** – This is the recommended graph storage engine. The configuration steps are mandatory:
  - In the **InfluxDB URL** field enter the full URL of the InfluxDB http(s) server (e.g. http://127.0.0.1:8086)
  - **InfluxDB Username** and **InfluxDB Password** are optional and should be set if `auth-enabled=true` in `influxdb.conf`
  - **InfluxDB Database** is optional. By default, the name of the InfluxDB database is "andrisoft". The field should be set only when using a single InfluxDB for multiple Consoles.
  - Click on the **<Initialize InfluxDB>** button to create a new InfluxDB database called "andrisoft"

The **Graphing Engine Backend** parameter lets you choose from which Graph Storage Engine to pull the data when rendering graphs.

**Decoders** represent internal functions that differentiate and classify the underlying protocols of each packet or flow. Each enabled decoder increases the size of IP graph, top and accounting data, and causes a small performance penalty. It is recommended to enable only the decoders you are interested in. You can define your own decoders in Configuration » General Settings » Custom Decoders. Built-in decoders:

| DECODER  | DESCRIPTION   |
|----------|---|
| IP       | Matches all IP packets, irrespective of higher protocols. Always enabled                    |
| TCP      | Matches TCP traffic   |
| TCP+SYN  | Matches TCP traffic with SYN flag set and ACK unset. Flow Sensor counts one packet per flow |
| UDP      | Matches UDP traffic   |
| ICMP     | Matches ICMP traffic  |
| OTHER    | Matches IP protocols that differ from TCP, UDP and ICMP                                     |
| INVALID  | Matches TCP or UDP port set to 0, or IP protocol set to 0                                   |
| FLOWS    | Matches flow records and replaces packets/s with flows/s. Works only with Flow Sensor       |
| FLOW+SYN | Matches flow records with SYN flag set. Flow Sensor counts all packets per flow             |
| FRAGMENT | Matches fragmented IP packets. Works only with Packet Sensor                                |

|                   |  |
|-------------------|--|
| <b>TCP-NULL</b>   | Matches TCP traffic without TCP flags, indicative of reconnaissance sweeps                     |
| <b>TCP+RST</b>    | Matches TCP traffic with RST flag set  |
| <b>TCP+ACK</b>    | Matches TCP traffic with SYN flag unset and ACK set  |
| <b>TCP+SYNACK</b> | Matches TCP traffic with SYN flag set and ACK flag set   |
| <b>NETBIOS</b>    | Matches TCP traffic on source or destination port 139  |
| <b>QUIC</b>       | Matches Google's QUIC protocol on UDP port 80 and 443  |
| <b>UDP-QUIC</b>   | Matches UDP traffic not part of the QUIC protocol  |
| <b>MEMCACHED</b>  | Matches UDP traffic on port 11211  |
| <b>HTTP</b>       | Matches TCP traffic on source or destination port 80   |
| <b>HTTPS</b>      | Matches TCP traffic on source or destination port 443  |
| <b>MAIL</b>       | Matches TCP traffic on source or destination ports 25, 110, 143, 465, 585, 587, 993, 995       |
| <b>DNS</b>        | Matches UDP traffic on source or destination port 53   |
| <b>SIP</b>        | Matches TCP or UDP traffic on source or destination port 5060                                  |
| <b>IPSEC</b>      | Matches IP traffic on IP protocol 50 or 51   |
| <b>WWW</b>        | Matches TCP traffic on source or destination ports 80, 443                                     |
| <b>SSH</b>        | Matches TCP traffic on source or destination port 22   |
| <b>NTP</b>        | Matches UDP traffic on source or destination port 123  |
| <b>SNMP</b>       | Matches UDP traffic on source or destination ports 161, 163                                    |
| <b>RDP</b>        | Matches TCP or UDP traffic on source or destination port 3389                                  |
| <b>YOUTUBE</b>    | Matches IP traffic going or coming from Youtube AS 43515, 36561, or from Youtube subnets       |
| <b>NETFLIX</b>    | Matches IP traffic going or coming from Netflix AS 55095, 40027, 2906, or from Netflix subnets |
| <b>HULU</b>       | Matches IP traffic going or coming from Hulu AS 23286, or from Hulu subnets                    |
| <b>FACEBOOK</b>   | Matches IP traffic going or coming from Facebook AS 54115, 32934, or from Facebook subnets     |

In most cases it is safe to ignore the parameters under the **Advanced Settings** section.

**Graph IP Sweeps** option allows storing IP graph data for IPv4 and/or IPv6 addresses that receive traffic without sending any traffic back. Do not set it to "Off" when monitoring unidirectional links or asymmetric traffic.

**Sensor Top N** (default: 20) specifies the maximum number of items stored for ordered sets of data, such as top Talkers, External IPs, ASNs, Countries, TCP/UDP ports, IP protocols, and so on.

Packet Sensor saves packet dumps on the local disk in the path defined in the **Packet Dump Path** field. Flow Sensor saves flow data on the local disk in the path defined in the **Flow Collector Path** field. When the Console is not installed on the same server that runs the Sensor, export these paths towards the Console's file system using an NFS share ([KB article link](#)). If you do not, the Console is not able to display data saved on remote servers.

RRDTool stores all graph files on the Console server, in the **Graphs Disk Path**. Graph files are optimized for storing time series data and do not grow over time. All IP graph options have a direct impact on the storage space required on the Console server. The size of each IP graph file is listed on the bottom of the window in the **Disk space required for each IP graph file** field. When Sensor Clusters are not used, the maximum number of IP graph files that

could be generated can be calculated with the formula: ((number of Packet Sensors) + (number of Flow Sensor interfaces)) x (number of IPs contained in subnets with IP Graphing set to "Yes" in the IP Zone).

It is highly recommended to automate the deletion of old data and to monitor the disk usage of IP graphs in Configuration » General Settings » Data Retention.

## Sensor and Applications Graph Troubleshooting

- ✓ Ensure that all Sensors run correctly by verifying the event log and by viewing live statistics from Reports » Devices » Overview. The troubleshooting guide for Packet Sensor is located on page 28, for Flow Sensor on page 33 and SNMP Sensor on page 37
- ✓ Discontinuous Sensor graphs can be caused by enabling IP Accounting for too many/large subnets when there is a slow connection between the Sensor and the MySQL/MariaDB running on the Console server

## IP/Subnet Graph Troubleshooting

- ✓ Ensure that all Sensors run correctly by verifying the event log and by viewing live statistics displayed in Reports » Devices » Overview. The troubleshooting guide for Packet Sensor is located on page 28, for Flow Sensor on page 33 and SNMP Sensor on page 37
- ✓ Generating IP graph data causes the biggest impact on the load of the Console server. Enable each feature (IP graphing, IP accounting) sequentially for each subnet, after making sure that the Console server can handle it. The storage requirements for each subnet are listed in the IP Zone, and the current disk usage in Configuration » General Settings » Data Retention
- ✓ The internal process used for saving IP graph data is `/opt/andrisoft/bin/genrrds_ip`. If it is overloading the Console server or the event log contains warnings such as "Updating IP graph data takes longer than 5 minutes", use InfluxDB or RRDCacheD, RAM/SSD updating method, faster disk drivers, enable IP graphing for fewer subnets, or deploy a Sensor Cluster configured to aggregate IP graph data
- ✓ The internal process used for generating IP or subnet graphs is `/opt/andrisoft/bin/gengraph_ip`. Console users launch the process for each requested IP or subnet graph. If the Console server gets too loaded by `gengraph_ip`, execute "killall gengraph\_ip" and configure InfluxDB or RRDCacheD. When launched, the process stops only when the graph is generated. This process can be slow when users request subnet graphs for subnets not specifically defined in the IP Zone. It is not possible to throttle the number of graphs requested by users

## AS and Country Graph Troubleshooting

- ✓ Ensure that all Sensors run correctly by verifying the event log and by viewing live statistics from Reports » Devices » Overview. The troubleshooting guide for Packet Sensor is located on page 28, for Flow Sensor on page 33 and SNMP Sensor on page 37
- ✓ To enable AS and Country graphs, set the Stats Engine parameter to either "Extended" for Flow Sensor, or "Full" for Packet Sensor
- ✓ SNMP Sensor is not able to generate AS graphs or Country graphs

## Configuration » General Settings » Custom Decoders

**Decoders** represent internal functions that differentiate and classify the underlying protocols of each packet and flow. The predefined decoders are listed in the “Graphs & Storage” chapter on page 16. If you do not need to define custom decoders, you can safely skip this section.

The screenshot shows the 'Edit Custom Decoder' window with the following configuration:

- Decoder Name:** MEMCACHE
- Decoder Color:** #0FFB84
- Decoder Description:** User-supplied decoder for memcache traffic
- Packet Matching Expressions:**
  - BPF Syntax:** udp and port 11211
  - ACL IPv4 Syntax:** 0.0.0.0/0 0.0.0.0/0 0-65535 11211-11211 17-17 0/0x0 0/0x0 0-65535 0.0.0.0/0 0.0.0.0/0 11211-11211 0-65535 17-17 0/0x0 0/0x0 0-65535
  - ACL IPv6 Syntax:** ::0 ::0 11211-11211 0-65535 17-17 0/0x0 0/0x0 0-65535 ::0 ::0 0-65535 11211-11211 17-17 0/0x0 0/0x0 0-65535
- Flow Matching Expressions:**
  - Flow Syntax:** proto 17 and port 11211
  - Flowspec Syntax:** protocol [6]
- Advanced Settings:**
  - Included Decoders:** IP, UDP
  - Conflicting Decoders:** UDP, INVALID
  - Filter Engine:** UDP
  - Netfilter Expression:** -m multiport --ports 11211

To manage user-defined decoders go to Configuration » General Settings » Custom Decoders.

Each custom decoder is defined by:

- **Decoder Name** – A short name to help you identify the decoder. This field is mandatory
- **Decoder Color** – The color used in graphs for the decoder. The default color is a random one, which can be changed by entering a different HTML color code or by clicking the drop-down menu
- **Decoder Description** – An optional short description of the decoder
- **BFP Syntax** – This syntax is used by Packet Sensor and Packet Filter when the Capture Engine option is not set to DPDK. Click the star icon on the right to see details about the correct syntax. Examples:
  - To match TCP packets with the SYN flag set, enter *tcp[tcplflags] & tcp-syn!=0*
  - To match UDP packets with the destination port under 1024, enter *proto 17 and dst portrange 1-1023*
  - To match memcached packets, enter *proto 17 and port 11211*
- **ACL IPv4/IPv6 Syntax** – This syntax is used by Packet Sensor and Packet Filter when the Capture Engine option is set to DPDK. Click the star icon on the right to see a few examples and get all the details about the correct syntax
- **Flow Syntax** – This syntax is used by Flow Sensor and Flow Filter. Click the star icon on the right to see the correct syntax. Examples:
  - To match TCP flows having only the SYN flag set, enter *flags S and not flags AFRPU*
  - To match flows with the MPLS label0 set to 2, enter *mpls label0=2*

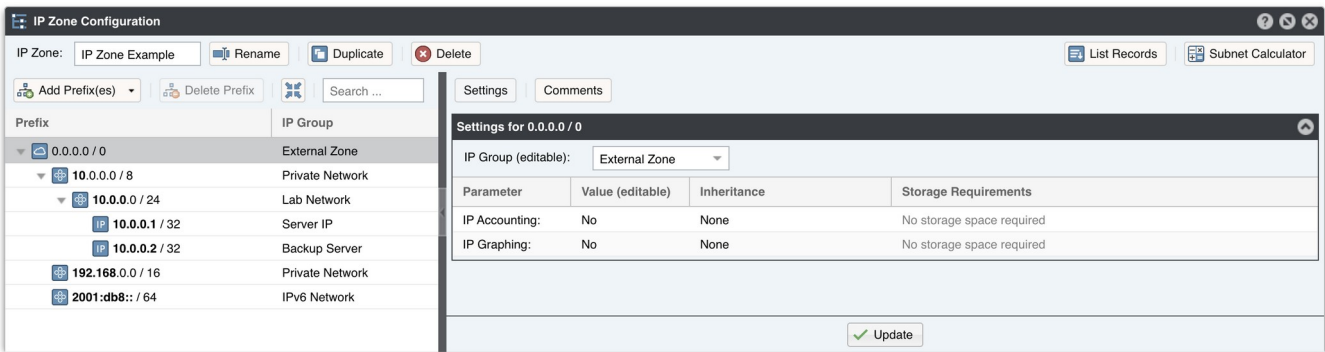
- To match memcached packets, enter *proto 17 and port 11211*
- **Flowspec Syntax** – Enter a Flowspec expression if you intend to use BGP Flowspec for traffic redirection or DDoS mitigation. Click the star icon on the right to open a window that shows you the correct syntax. Example:
  - To match memcached packets, enter *port 11211; protocol 17;*
- **Included Decoders** – Required if you plan to use the decoder for IP graphing. Select the decoders that include the matched traffic, or choose IP if not sure
- **Conflicting Decoders** – Required if you plan to use the decoder for IP graphing. Select the decoders that might match same traffic, but not always. The option is used only for displaying stacked decoders inside IP graphs

## Configuration » Network & Policy » IP Zone

**IP Zones** are hierarchical, tree-like data structures used by Sensor to extract per-subnet settings and to learn the protected network's boundaries.

In most configurations, you will have to add your IP blocks to the IP Zones listed in Configuration » Network & Policy. There are several ways to add prefixes (IPs/IP blocks/subnets/ranges): from the web interface, via the REST API by accessing [http://<console\\_ip>/wanguard-api-ui](http://<console_ip>/wanguard-api-ui), or by executing the command `php /opt/andrisoft/api/cli_api.php` on the Console server.

To add a new IP Zone, go to Configuration » Network & Policy » [+] and select [IP Zone]. You only need more than one IP Zone when you want to use different per-subnet settings for different Sensors. If this is the case, it may be easier to open an existing IP Zone that already includes your IP address ranges, and duplicate it by pressing the [**Duplicate**] button.



The IP Zone Configuration window is divided into two vertical sections. The buttons that manage prefixes are located in the upper part of the left section. When a new prefix is added the tree below automatically updates itself. The section on the right contains panels with user-provided settings for the selected prefix.

To enter IP addresses or IP blocks, use the CIDR notation. To enter individual hosts in IP Zones, use the /32 CIDR mask for IPv4 and /128 for IPv6. For more information about the CIDR notation consult Appendix 1 on page 65.

Every IP Zone contains the network 0.0.0.0/0. Because it's CIDR mask is /0, this "supernet" includes all IP addresses available for IPv4 and IPv6. For an easier configuration, every new prefix that you define inherits by default the properties of the most-specific (having the biggest CIDR mask) IP class that includes it.

The **IP Settings** panel from the section on the right provides the following parameters:

- **IP Group** – Set a short description of the selected prefix, or the name of the customer that uses it. When you set the same IP group on multiple prefixes you will be able to generate aggregated traffic reports. This combo box is editable
- **IP Graphing** – Set to "Yes" to permit the collection of graph data for every IP contained in the selected prefix. The **Graph IP Sweeps** option from Configuration » General Settings » Graphs & Storage can be used to prevent generating graph data for IPs that only receive traffic without sending traffic in return. IP Graphing is always enabled for the subnets explicitly defined in the IP Zone. Do not enable this option on many/large subnets without a performance impact assessment

- **IP Accounting** – Set to “Yes” for the Sensor to generate daily accounting data for each IP contained in the selected prefix. IP Accounting is always enabled for the subnets explicitly defined in the IP Zone. Do not enable on many/large subnets without a performance impact assessment

The **Storage Requirements** column indicates the disk space needed by each Packet Sensor and Flow Sensor interface to store the generated data. Enabling IP graphing and IP accounting for very large prefixes (e.g. 0.0.0.0/0) might generate data that could overload the Console server and fill the disk space. The storage requirements for IP graph data is possible to estimate only for RRD files. For InfluxDB this is not possible.

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

## Configuration » Servers

Any server running a Wansight component must be added under Configuration » Servers. The Console server is automatically added during installation.

To add a new server, click the [+] button from the title bar of the Configuration » Servers panel. To change the configuration of an existing server, go to Configuration » Servers and click its name.

The screenshot shows the 'Server Configuration' window with the following details:

- Server Name: Console
- Server Color: #4E7FA6
- Server ID: 1
- Device Group: (empty dropdown)
- IP Address: 192.168.242.1
- Linux Distribution: Generic Debian
- Hardware Key: Fw8k9EISoisw+xx4fKME8zCVRyaQ=

Network Interfaces (optional) section:

| Interface | Speed IN | Speed OUT | Color        |
|-----------|----------|-----------|--------------|
| eth0      | 1 Gbps   | 1 Gbps    | Light Purple |
| lo        | 10 Gbps  | 10 Gbps   | Dark Grey    |
| p1p2      | 10 Gbps  | 10 Gbps   | Dark Red     |
| p255p1    | 10 Gbps  | 10 Gbps   | Dark Purple  |
| ...       | ...      | ...       | ...          |

Buttons: Add Interface, Edit Interface, Delete Interface. A 'Save' button is at the bottom.

- **Server Name** – A short name to help you identify the server
- **Server Color** – The color used in graphs for this server. The default color is a random one, which can be changed by clicking the drop-down menu
- **Server ID** – Unique identifier of the server. This id is used when exporting NFS shares
- **Device Group** – Optional description used to group servers by location, role, etc.
- **IP Address** – An IP address defined on the server. Can be public or private, IPv4 or IPv6
- **Linux Distribution** – Select the Linux distribution installed on the server
- **Hardware Key** – Read-only string used for licensing purposes. The hardware key field is updated by the WANsupervisor service when the hardware, IP or hostname changes. If the hardware key is unregistered, send it to [sales@andrisoft.com](mailto:sales@andrisoft.com)
- **Network Interfaces (optional)** – The WANsupervisor service can monitor packets/s, bits/s, errors and dropped frames for each server interface. The data is available in Reports » Servers » [Server] » Server Graphs » Data Units = Server Interfaces. These stats are provided by the Linux kernel
- **Comments** – These observations are not visible elsewhere



## Distributing the Software over Multiple Servers

You can distribute Sensors over multiple servers for load and geographical distribution, or for high-availability and redundancy:

1. Add the new server in Console, under Configuration » Servers, set the IP and enter a Server Name
2. Install the software on the new server by following the installation instructions from the link contained in the response mail to the evaluation request
3. When executing `/opt/andrisoft/bin/install_supervisor` enter the IP of the Console server and the Console database password
4. Start the WANsupervisor service on the new server
5. Make sure that NTP is running on the server and that the status is OK in Reports » Devices » Overview
6. During the trial period you don't have to register any server. Outside the trial period, you have to register the server's hardware key, which is visible in Configuration » Servers » [New Server] after starting the WANsupervisor service. Hardware registration is free by emailing sales@andrisoft.com
7. Define a new Sensor and set its Server parameter accordingly
8. Start the new Sensor from Configuration » Components
9. Watch the event log to see if there are any errors or warnings

## Server Troubleshooting

- ✓ For the server to be operational, make sure it always runs the WANsupervisor service and that its clock is synchronized with NTP. You can verify the operational status of each server and component in Reports » Devices » Overview » Servers
- ✓ The WANsupervisor service stops when the MySQL/MariaDB service running on the Console server is restarted or unavailable even for a short amount of time (e.g. during a network outage). In this case, either restart WANsupervisor manually or use automated tools such as systemd, monitd or similar
- ✓ You can discover performance-related issues by monitoring Reports » Server » [Server] » Server Graphs and Reports » Server » [Server] » Server Events
- ✓ If the DB crashes (usually due to power failures) execute `/opt/andrisoft/bin/WANmaintenance repair_db`

## Configuration » Components » Packet Sensor

In switched networks, only the packets for a specific device reach the device's network card. If the server running the **Packet Sensor** is not deployed in-line in the main data path, a network TAP, or a switch/router that offers a "monitoring port" or "mirroring port" must be used. In this case, the network device sends copies of data packets traveling through selected ports or VLANs to the monitoring port. Packet Sensor inspects every packet it receives and conducts packet-based traffic analysis. The advantages and disadvantages of packet-based traffic monitoring are listed on page 7.

For instructions on how to configure switches or routers for port mirroring, consult their documentation.

To add a Packet Sensor, click the [+] button from the title bar of the Configuration » Components panel. To modify an existing Packet Sensor, go to Configuration » Components and click its name.

- **Sensor Name** – A short name to help you identify the Packet Sensor
- **Sensor Color** – The color used in graphs for the Packet Sensor. The default color is a random one, which can be changed by clicking the drop-down menu
- **Sensor Visibility** – Toggles the listing inside Reports » Devices
- **Device Group** – Optional parameter used for grouping up components (e.g. by location or role). It can be used to restrict the access of Guest accounts
- **Sensor Server** – The server that runs the Packet Sensor. The configuration of servers is described on page 24
- **Capture Engine** – Select the preferred packet capturing engine. LibPcap requires no additional setup but it is too slow for multi-gigabit sniffing and it does not run on multiple threads. Netmap and PF\_RING are much faster but both require the installation of additional kernel modules. DPDK provides the fastest packet capturing engine and allows packet forwarding but it is very complicated to configure.
  - *Embedded LibPcap* – Select to use the built-in LibPcap 1.8.1 library
  - *System LibPcap* – Select to use the LibPcap library installed by the Linux distribution-
  - *Netmap* – Select to use the Netmap framework to speed up packet processing

- *PF\_RING* – Select to use the PF\_RING framework to speed up packet processing. Click the button on the right for PF\_RING-specific settings
- *DPDK* – Select to use the DPDK framework, then click the button on the right of the Capture Engine field to configure DPDK-specific parameters as described in Appendix 3 on page 71
- **Sniffing Interface** – The network interface(s) listened by the Packet Sensor. If the server running the Packet Sensor is deployed in-line, then this field must contain the network interface that receives the traffic entering your network. The PF\_RING framework allows listening to multiple physical interfaces simultaneously when the interfaces are entered separated by semicolon “;”
- **CPU Threads** – Packet Sensor can run multi-threaded on a given set of CPU cores. Each thread increases the RAM usage. On most systems, activating more than 6 CPU threads hurts performance
- **Link Speed IN / OUT** – Enter the speed (bandwidth, capacity) of the monitored link. The values are used for percentage-based reports
- **Sensor License** – The license used by the Packet Sensor. Wanguard provides all features; Wansight does not provide traffic anomaly detection and reaction
- **Stats Engine** – Collects traffic tops and AS graphs:
  - *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
  - *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. Permits the detection of threshold violations for external IPs
  - *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. Permits the detection of threshold violations for external IPs
- **Stats Engine Options** – When Stats Engine is set to Full you can click the button next to it. To enable Transit AS tops and graphs, enter the path to an existing **BGP Dump File** exported by BGPd in MTR format, and the IPv4 and optionally IPv6 address of the BGP router
- **IP Zone** – Packet Sensor needs an IP Zone from which to learn about your network's boundaries and to extract per-subnet settings. IP Zones are described in the “IP Zone” chapter on page 22
- **BPF Expression** – You can filter the type of traffic the Packet Sensor receives using a tcpdump-style syntax
- **IP Validation** – This option is the most frequently-used way to distinguish the direction of the packets:
  - *Off* – Packet Sensor analyzes all traffic and uses MAC Validation to identify the direction of traffic
  - *On* – Packet Sensor analyzes the traffic that has the source and/or the destination IP in the selected IP Zone
  - *Strict* – Packet Sensor analyzes the traffic that has either the source or the destination IP in the selected IP Zone
  - *Exclusive* – Packet Sensor analyzes the traffic that has the destination IP in the selected IP Zone, but not the source IP
- **MAC Validation/Options** – This option can be used to distinguish the direction of the packets or to ignore unwanted OSI Layer 2 traffic:
  - *None* – Packet Sensor analyzes all traffic and uses IP Validation to identify 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 (:)
- **Granularity** – Interval between successive updates for traffic parameters
- **Sampling (1/N)** – Contains the packet sampling rate. On most systems, the correct value is 1
- **Comments** – Comments about the Packet Sensor can be saved here. They are not visible elsewhere

To start the Packet Sensor, click the small button displayed next to its name in Configuration » Components. Ensure that the Packet Sensor starts correctly by watching the event log (details on page 40).

If the Packet Sensor starts without errors, but you can't see any data collected by it in Reports » Devices » Overview, follow the troubleshooting steps listed below.

## Packet Sensor Troubleshooting

- ✓ Look for warnings or errors produced by the Packet Sensor in the event log (details on page 40)
- ✓ Ensure that you have correctly configured the Packet Sensor. Each configuration field is described in this chapter
- ✓ The event log error "*License key not compatible with the existing server*" can be fixed by sending the string from Configuration » Servers » [Packet Sensor server] » Hardware Key to sales@andrisoft.com
- ✓ Make sure that the sniffing interface is up:  

```
ip link show <interface>
```
- ✓ Ensure that you have correctly configured the switch/TAP to send packets to the server on the configured interface
- ✓ Verify whether the server is receiving packets via the configured interface:  

```
tcpdump -i <interface_usually_eth1_or_p1p2> -n -c 100
```
- ✓ When **IP Validation** is not disabled, make sure that the selected IP Zone contains all your subnets
- ✓ If the CPU usage of the Packet Sensor is too high, set the **Stats Engine** parameter to "Basic", install PF\_RING, Netmap or DPDK to enable multi-threading, or use a network adapter that allows distributing Packet Sensors over multiple CPU cores
- ✓ To troubleshoot Sensor graph or IP graph issues, follow the Graphs Troubleshooting guide on page 19
- ✓ For PF\_RING-related issues, contact ntop.org. To increase the maximum number of PF\_RING programs from 64 to 256, increase the MAX\_NUM\_RING\_SOCKETS defined in kernel/linux/pf\_ring.h and recompile the pf\_ring kernel module
- ✓ The system process responsible for capturing packets is called WANtrafficlogger. There will be as many processes active as the number of packet dumps active in Reports » Tools » Packets
- ✓ Make sure you are running the latest version of the software by checking Help » Software Updates

## Packet Sensor Optimization Steps for Intel 82599

To distribute the packet-processing tasks of the Packet Sensor over multiple CPU cores when using an adapter with the Intel 82599 chipset (Intel X520, Intel X540, HP X560, etc.):

- ✓ Follow the documentation and optimization guides provided by the network adapter vendor
- ✓ Install PF\_RING and switch to the PF\_RING-aware ixgbe driver
- ✓ See the number of RSS queues allocated by the ixgbe driver by executing `dmesg`, or by listing `/var/log/messages` or `/var/log/syslog`. By default, the number of RSS queues is equal to the number of CPU cores when hyper-threading is off, or double the number of CPU cores when hyper-threading is on. You can set the number of RSS queues manually, by loading `ixgbe.ko` with the `RSS=<number>` option
- ✓ Enable multithreading in the Packet Sensor configuration or define multiple Packet Sensors, each listening to `ethX@queue_id` or `ethX@queue_range` and add them to a Sensor Cluster to have a unified reporting and anomaly detection domain. No additional licensing is needed because all Packet Sensors defined to listen to a single interface use a single Sensor license

On a quad-core CPU with multithreading, the ixgbe driver allocates 8 RSS queues. In this case, if you define a Packet Sensor for `ethX@0-3` and another one for `ethX@4-7`, the packet-processing task will be distributed over 2 CPU cores. PF\_RING exposes up to 32 RSS queues.

## Configuration » Components » Flow Sensor

Many routers and switches can collect IP traffic statistics and periodically export them in form of flow records to **Flow Sensor**. Since the flow protocol already performs pre-aggregation of traffic data, the flow data sent to Flow Sensor is much smaller than the monitored traffic, and this makes the Flow Sensor a good option for monitoring remote or high-traffic networks. The advantages and disadvantages of flow-based monitoring are listed on page 7.

For detailed instructions on how to enable NetFlow, sFlow or IPFIX on your network device, please consult its documentation. Appendix 2 on page 66 lists some configuration examples for a few Cisco IOS, CatOS, and Juniper devices.

To add a Flow Sensor, click the [+] button from the title bar of the Configuration » Components panel. To modify an existing Flow Sensor, go to Configuration » Components and click its name.

| Index | Interface Name          | Direction | Stats    | Speed IN | Speed OUT | Color       |
|-------|-------------------------|-----------|----------|----------|-----------|-------------|
| 1     | TenGigE_0_7_0_0         | Upstream  | Extended | 1 Gbps   | 1 Gbps    | Dark Red    |
| 2     | GigabitEthernet0_1_0... | Auto      | Extended | 10 Gbps  | 10 Gbps   | Green       |
| 3     | Bundle-Ether1           | Auto      | Extended | 40 Gbps  | 40 Gbps   | Light Green |
| 4     | Switch-Port-Channel 4   | Auto      | Extended | 100 Gbps | 100 Gbps  | Blue        |

- **Sensor Name** – A short name to help you identify the Flow Sensor
- **Reports Visibility** – Toggles the visibility inside Reports » Devices
- **Device Group** – Optional description used used to group up components (e.g. by location or role). It can be used to restrict the access of Guest accounts
- **Sensor Server** – The server that runs the Flow Sensor. The configuration of servers is described on page 24
- **Listener IP:Port** – The IP address (IPv4 or IPv6) of the network interface that receives the flow packets, and the destination port
- **Flow Exporter IP** – IP address of the flow exporter (router, switch, probe). Usually, it is the loopback address of the router. For sFlow exporters, enter the IP that sends flow packets, not the Agent IP
- **SNMP Settings** – Click the button on the right side of the Flow Exporter IP field. You must enable SNMP

on the flow exporter to allow Console to automatically extract interface information. When SNMP settings are not configured, you must manually enter the SNMP index, speed, etc. for each interface

- **Flow Exporter TZ** – Set the time offset between the time zone (TZ) of the Flow Sensor server and the time zone of the flow exporter. Running NTP on both devices to keep their clocks synchronized is a critical requirement
- **Flow Protocol** – Flow protocol used by the flow exporter: NetFlow, IPFIX or sFlow
- **Flows Timeout (s)** – Some flow exporters (e.g. Juniper MX) maintain the start time of flows that were already exported. If this is the case then you need to set here the same flow active/inactive timeout value (in seconds) as the one defined in the flow exporter's configuration
- **Sensor License** – The license type allocated to the Flow Sensor. Vanguard provides all features; Wansight does not provide traffic anomaly detection and reaction
- **Flow Collector** – When enabled, flow data is stored in a space-efficient binary format. LZO is the fastest. BZ2 offers the best compression rate but it's 30 times slower than LZO. LZ4 offers a compromise between speed and efficiency. Flow records can be queried in Reports » Tools » Flows
- **IP Zone** – Flow Sensor needs an IP Zone from which to learn the monitored network's boundaries and to extract per-subnet settings. For more information about IP Zones consult the dedicated chapter from page 22
- **Repeater IP:Port** – An embedded packet repeater can send all incoming flows to another flow collector or host. To use this optional feature enter the IP of the other flow collector and a port of your choice
- **IP Validation** – This option can be used to distinguish the direction of traffic or to ignore certain flows:
  - *Off* – Flow Sensor examines all flows and the traffic direction is established by the interface
  - *On* – Flow Sensor examines the flows that have the source IP and/or the destination IP inside the selected IP Zone. This is the recommended setting for most setups
  - *Strict* – Flow Sensor examines the flows that have either the source IP or the destination IP inside the IP Zone
  - *Exclusive* – Flow Sensor examines the flows that have the destination IP inside the IP Zone
- **IP Validation Options** – Set the **Log Invalidated Flows** field to *Periodically* if you want to see in the event log the percentage of invalidated flows and 10 flows failing validation, once every 10 ticks
- **AS Validation** – Flows from BGP-enabled routers can contain the source and destination Autonomous System number (ASN). In most configurations if the AS number is set to 0 then the IP address belongs to your network. This rarely-used option is used for establishing traffic direction. AS validation has three choices:
  - *Off* – Disables AS validation
  - *On* – Flow Sensor examines only the flows that have the source ASN and/or the destination ASN inside the local AS list (defined below)
  - *Strict* – Flow Sensor examines only the flows that have either the source ASN or the destination ASN inside the local AS list (defined below)
- **AS Validation Options** – When AS Validation is enabled, you can enter all your AS numbers (separated by space) into the **Local AS List** field. Set the **Log Invalidated Flows** field to *Periodically* if you want to see in the event log the percentage of invalidated flows and 10 flows failing validation, once every 10 ticks
- **Granularity** – Low values increase the accuracy of Sensor graphs, at the expense of increasing the RAM usage

- **Sampling (1/N)** – Enter the sampling rate configured on the flow exporter, or “1” when no sampling rate is configured. For NetFlow v9 and sFlow the value entered here is ignored because the flow protocol automatically adjusts the sampling rate. To force a particular sampling value, enter it as a negative value
- **Monitored Interfaces** – List of interfaces that should be monitored. To avoid producing duplicate flow entries, add only upstream interfaces
  - *SNMP Index* – The interfaces are internally identifiable only by their SNMP indexes. Enter the index manually, or configure SNMP settings
  - *Interface Name* – A short description used to identify the monitored interface. Descriptions longer than 10 characters clutter some reports
  - *Interface Color* – The color used in graphs for the interface. The default color is a random one, which can be changed by clicking the drop-down menu
  - *Traffic Direction* – Direction of traffic entering the interface, relative to your network:
    - “Auto” – When selected, the direction of traffic is established by IP and/or AS Validation alone. This is the recommended setting in most cases
    - “Upstream” – Set for upstream interfaces, e.g. peering interfaces, interfaces connected to the Internet
    - “Downstream” – Set for downstream interfaces, e.g. customer interfaces, interfaces connected to your network
    - “Null” – Traffic to Null interfaces is discarded by the router and should be ignored
  - *Stats Engine* – Collects various traffic tops and AS (Autonomous System) data:
    - “Basic” – Enables tops for Internal IPs, IP protocols, versions and TCP/UDP ports. It adds a very small performance penalty
    - “Extended” (recommended) – Enables all tops from “Basic” as well as tops and graphs for autonomous systems and countries, but slightly increases the CPU usage. When the router does not export AS information in flows (e.g. non-BGP router), Flow Sensor uses an internal GeoIP database to obtain AS data. Live stats for autonomous systems and countries may not be 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. Set the value to “Extended”, unless you know what you’re doing. Only this option permits the detection of threshold violations for external IPs
  - *Stats Engine Options* – When Stats Engine is set to “Extended” or “Full”, you can click the button next to it. To enable Transit AS tops and graphs, enter the path to an existing BGP Dump File exported by BGPd in MTR format, and the IPv4 and optionally IPv6 address of the BGP router
  - *Link Speed In & Link Speed Out* – Enter the speed (bandwidth, capacity) of the interface. The values are used for percentage-based reports
- **Comments** – Comments about the Flow Sensor can be saved here. These observations are not visible elsewhere

To start the Flow Sensor, click the small button displayed next to its name in Configuration » Components. Ensure that the Flow Sensor starts correctly by watching the event log (details on page 40).



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

## Flow Sensor Troubleshooting

- ✓ Look for warnings or errors produced by the Flow Sensor in the event log (details on page 40)
- ✓ Check if you have correctly configured the Flow Sensor. Each configuration parameter is described in the previous section
- ✓ Event log error "*License key not compatible with the existing server*" can be fixed by sending the string from Configuration » Servers » [Flow Sensor server] » Hardware Key to sales@andrisoft.com
- ✓ Ensure that the server is receiving flow packets on the configured **Listener IP:Port**:

```
tcpdump -i <interface_eth0_p1p1_etc> -n -c 100 host <flow_exporter_ip> and udp and port <destination_port>
```

- ✓ Make sure that the local firewall permits the Flow Sensor to receive flow packets:
 

```
iptables -L -n -v && iptables -t raw -L -n -v
```
- ✓ Ensure that the clocks of both devices are synchronized with NTP. When the devices do not reside in the same time zone, adjust the **Time Settings** parameter accordingly
- ✓ Flow Sensor may crash during spoofed attacks for not having enough RAM when a monitored interface has the *Stats Engine* parameter set to "Full". It is highly recommended to set the **Stats Engine** parameter to "Extended" instead of "Full" on systems with low amounts of RAM
- ✓ When you add interfaces with the **Traffic Direction** parameter set to "Auto", make sure that the IP Zone you have selected contains all your IP blocks because **IP Validation** and/or **AS Validation** will be used to establish traffic direction. To capture a sample of flows failing validation in the event log, set the **Log Invalidated Flows** parameter to "Periodically"
- ✓ In order to provide fast and up-to-date traffic statistics, the Flow Sensor accepts only flows describing traffic from the last 5 minutes. All flows aged and exported with a delay exceeding 300 seconds are ignored, and the event log contains the warning "*Received flow <starting/ending> <X> seconds ago*"

When the warnings refer to the starting time, make sure that the clocks between the server and the router are synchronized, the flow exporter is properly configured, and the time zone and **Flow Timeout** are correctly set.

When the warnings refer to the ending time, make sure that the clocks are synchronized, the time zone is correctly set, the flow exporter is properly configured, and the PFC PIC is not overloaded (on Juniper in particular).

You can double-check whether the time of the Flow Sensor and the start/end time of flows differ by more than 300 seconds. In Reports » Tools » Flows » Flow Records, select any interface of the Flow Sensor, set Output to Debug and generate a listing for the last 5 minutes:

- Column *Date\_flow\_received* indicates the time when the Flow Sensor received the flow packet
- Column *Date\_first\_seen* indicates the time when the flow started
- Column *Date\_last\_seen* indicates the time when the flow ended

Flow Sensor does not misinterpret the start/end time of flows. A few flow exporters are known to have bugs, limitations or inconsistencies regarding flow aging and stamping flow packets with the correct time. In this case, contact your vendor to make sure that the flow exporter is correctly configured, and it is able

to expire flows in under 5 minutes. In some cases, a router reboot fixes the issue.

In some JunOS versions there is a flow export rate limit with a default of 1k pps, which leads to flow aging errors. To raise the limit to 40k pps you need to execute:

```
set forwarding-options sampling instance NETFLOW family inet output inline-jflow
flow-export-rate 40
```

Some Cisco IOS XE devices do not export flows using NetFlow version 5, in under 5 minutes, even when configured to do so. In this case, switch to using Flexible NetFlow

- ✓ Ensure that you have correctly configured the flow exporter to send flows to the server for each of the monitored interfaces. To list all interfaces that send flows, go to Reports » Tools » Flows » Flow Tops, select any Flow Sensor interface, set Output to Debug, set Top Type to Any Interface and generate a top for the last 10 minutes. The column In/Out\_If lists the SNMP index of every interface that exports flows, even if it was not configured as a monitored interface in the Flow Sensor configuration
- ✓ If you see statistics for only one traffic direction (inbound or outbound), go to Reports » Tools » Flows » Flow Records, and generate a listing for the last 10 minutes. If all your IPs are listed in a single column, check the flow exporter's configuration and feature list. Not all devices can export flows in both directions (e.g. some Brocade equipment generates only inbound sFlow) or with the same interface SNMP index
- ✓ The traffic readings of the Flow Sensor may differ from the SNMP Sensor or from other SNMP-based monitoring tools. Flow Sensor counts In/Out traffic as traffic entering/exiting the IP Zone (when **IP Validation** is enabled), unlike SNMP tools that count In/Out traffic as traffic entering/exiting the interface. You can double-check the traffic readings of a Flow Sensor by configuring an SNMP Sensor to monitor the same flow exporter (page 35)
- ✓ If the Flow Sensor does not show the correct statistics after upgrading the router's firmware, the SNMP index of the interfaces may have changed. In this case, adjust the SNMP indexes for the monitored interfaces manually, or redefine them
- ✓ If the traffic is too low after upgrading to JunOS 15.1F2 or 16.1R1, execute:  

```
set chassis fpc inline-services flow-table-size ipv4-flow-table-size 15
```
- ✓ To troubleshoot Sensor graph or IP graph issues, follow the Graphs Troubleshooting guide on page 19
- ✓ Make sure you are running the latest version of the software by checking Help » Software Updates

## Configuration » Components » SNMP Sensor

**SNMP Sensor** monitors the bandwidth usage of routers and switches on a port-by-port basis. SNMP Sensor queries devices (e.g. routers, switches, servers) for the traffic counters of each port with small data packets. These are triggering reply packets from the device. The advantages and disadvantages of monitoring traffic by SNMP are listed on page 7. For detailed instructions on how to enable SNMP on your network device, consult its documentation.

To add an SNMP Sensor click the [+] button from the title bar of the Configuration » Components panel. To modify an existing SNMP Sensor, go to Configuration » Components and click its name.

- **Sensor Name** – A short name to help you identify the SNMP Sensor
- **Sensor Visibility** – Toggles the listing inside Reports » Devices
- **Device Group** – Optional description to group up components (e.g. by location or role). It can be used to restrict the access of Guest accounts
- **Sensor Server** – Which server runs the SNMP Sensor. It is recommended to run all SNMP Sensors on the Console server. The configuration of servers is described on page 24
- **Device IP:Port** – Enter the IP address and SNMP port (161 by default) of the networking device
- **Interface Discovery** – Activates or deactivates interface discovery:
  - *Monitor all interfaces* – Select to add all interfaces automatically to the SNMP Sensor. The interface names are based on the **Interface Name** setting available when pressing the [SNMP Tester & SNMP Object Identifier] button located next to the **Device IP:Port** field. This is not the recommended option because it clutters the Console

- *Monitor defined interfaces* – Select to monitor only the interfaces from the Monitored Interfaces grid
- **Polling Interval** – Polling is the process of sending the SNMP request periodically to the device in order to retrieve information. A low polling interval ( $\pm 1$  minute) gives you granular reports but may place an increased load on your server if you poll a large number of interfaces
- **Timeout (ms)** – The timeout value should be at least a little more than double the time it takes for a packet to travel the longest route between devices on your network. The default value is 1000 milliseconds (1 second)
- **Retries** – This value represents the number of times the SNMP Sensor retries a failed SNMP request defined as any SNMP request that does not receive a response within the Timeout (ms) defined above. The default value is 2
- **Sensor License** – License used by the SNMP Sensor. Vanguard provides all features (although severely limited by the SNMP technology); Wansight does not provide traffic anomaly detection and reaction
- **IP Zone** – When a Vanguard license is being used, the SNMP Sensor can check thresholds listed in the selected IP Zone with the following precondition (because SNMP does not provide any information about IPs or protocols):
  - Subnet must be “0.0.0.0/0”
  - Domain must be “subnet”
  - Value must be absolute, not percentage
  - Decoder must be “IP”
- **SNMP Protocol** – Select the SNMP protocol used for authentication:
  - *SNMP v1* – Easy to set up – only requires a plaintext community. Supports only 32-bit counters and it has very little security
  - *SNMP v2c* – Version 2c is identical to version 1, except it adds support for 64-bit counters. This is imperative when monitoring gigabit interfaces. Even a 1Gbps interface can wrap a 32-bit counter in 34 seconds, which means that a 32-bit counter being polled at one-minute intervals is useless. Select this option instead of v1 in most cases
  - *SNMP v3* – Adds security to the 64-bit counters. SNMP version 3 adds both encryption and authentication, which can be used together or separately. Setup is much more complex than just defining a community string
- **Community String** – SNMP v1 and v2c credentials serve as a type of password that is authenticated by confirming a match between the string provided here and the SNMP Community String stored as a MIB object on an SNMP-enabled, managed device
- **Security Level & Name** – SNMP v3-only. SNMP Sensor supports the following set of security levels as defined in the USM MIB (RFC 2574):
  - *noAuthnoPriv* – Communication without authentication and privacy
  - *authNoPriv* – Communication with authentication and without privacy
  - *authPriv* – Communication with authentication and privacy
- **Authentication Protocol & Passphrase** – SNMP v3-only. The protocols used for Authentication are *MD5* and *SHA* (Secure Hash Algorithm)
- **Privacy Protocol & Passphrase** – SNMP v3-only. An indication of whether messages sent on behalf of this user can be protected from disclosure, and if so, the type of privacy protocol which is used. This

option takes the value *DES* (CBC-DES Symmetric Encryption) or *AES* (Advanced Encryption Standard)

- **Monitored Network Interfaces** – The grid must contain the interfaces that should be monitored. To avoid mirrored graphs, add only the upstream interfaces. Settings per interface:
  - *SNMP Index* – The interfaces are identifiable by their unique indexes
  - *Interface Name* – A short description used to identify the monitored interface. Descriptions longer than 10 characters clutter some reports. By default, the auto-filled interface name is retrieved from the ifAlias OID. To change the OID used for the interface name click the [**SNMP Tester & SNMP Object Identifier**] button located next to the **Device IP:Port** field
  - *Interface Color* – Color used in graphs for the interface. The default color is a random one, which can be changed by clicking the drop-down menu
  - *Traffic Direction* – Direction of the traffic entering the interface, from the user's perspective:
    - “Unset” – Traffic entering the interface is considered “downstream”; traffic exiting the interface is considered “upstream”
    - “Upstream” – Set for upstream interfaces, e.g. peering interfaces, interfaces connected to the Internet
    - “Downstream” – Set for downstream interfaces, e.g. customer interfaces, interfaces connected to your network
    - “Null” – Traffic to Null interfaces is ignored
  - *Link Speed In & Link Speed Out* – Enter the speed (bandwidth, capacity) of the interface. The values are used for percentage-based reports
- **Comments** – Comments about the SNMP Sensor can be saved here. These observations are not visible elsewhere

To start the SNMP Sensor, click the small button displayed next to its name in Configuration » Components. Ensure that the SNMP Sensor starts correctly by watching the event log (details on page 40).

If the SNMP Sensor starts without errors, but you cannot see any data collected by it in Reports » Devices » Overview after more than 5 minutes, follow the troubleshooting guide below.

## SNMP Sensor Troubleshooting

- ✓ Look for warnings or errors produced by the SNMP Sensor in the event log (details on page 40)
- ✓ Ensure that you have correctly configured the SNMP Sensor. Each configuration field is described in depth in this chapter
- ✓ Event log error “*License key not compatible with the existing server*” can be fixed by sending the string from Configuration » Servers » [SNMP Sensor server] » Hardware Key to sales@andrisoft.com
- ✓ Verify if the Console can reach the device by clicking the [**SNMP Tester & SNMP OIDs**] button from the SNMP Sensor Configuration window, then press [**Query Device**]
- ✓ Permit the server to contact the SNMP device, by configuring its ACL
- ✓ If Sensor graphs are spiky, increase the Polling Interval value.
- ✓ Make sure you are running the latest version of the software by checking Help » Software Updates

## Configuration » Components » Sensor Cluster

**Sensor Cluster** aggregates traffic data provided by Packet Sensors and Flow Sensors into a single IP graphing domain.

To add a Sensor Cluster, click the [+] button found on the title bar of the Configuration » Components panel. To configure an existing Sensor Cluster, go to Configuration » Components, and click its name.

- **Sensor Name** – A short name to help you identify the Sensor Cluster
- **Server Color** – Color used in graphs for the Sensor Cluster. The default color is a random one, which can be changed by clicking the drop-down menu
- **Reports Visibility** – Toggles the listing inside Reports » Devices
- **Device Group** – Optional description used to group up components (e.g. by location or role). It can be used to restrict the access of Guest accounts
- **Sensor Server** – Which server runs the Sensor Cluster. It is recommended to run Sensor Clusters on the Console server. The configuration of servers is described on page 24
- **Link Speed IN / OUT** – Summed-up speeds (bandwidth, capacity) of the aggregated interfaces. These values can be used for percentage-based reports
- **Associated Sensors** – Select which Sensor interfaces are aggregated by the Sensor Cluster
- **IP Zone** – Sensor Cluster extracts from the selected IP Zone per-subnet settings about thresholds and/or IP graphing. For more information about IP Zones consult the “IP Zone” chapter on page 22
- **IP Graphing** – Select “Aggregated” to enable IP graphing by the Sensor Cluster for the summed up traffic data, and disable IP graphing by the associated Sensors. Select “Not Aggregated” to enable IP graphing by each associated Sensor and to disable IP graphing by the Sensor Cluster
- **Comments** – Comments about the Sensor Cluster can be saved here. These observations are not visible elsewhere

To start the Sensor Cluster, click the small button displayed next to its name in Configuration » Components. Ensure that the Sensor Cluster starts correctly by watching the event log (see details on page 40) and by watching Reports » Devices » Overview.

## Configuration » Schedulers » Scheduled Reports

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

If you want to receive periodic reports by email without having to log in, go to Configuration » Schedulers and click the [+] button from the title bar of the panel.

The screenshot shows the 'Scheduled Report Configuration' window. At the top, the 'Report Name' is 'Daily Report for Internal Zone' and 'Email To' is 'noc@isp.com'. Below this are four report type buttons: 'IP Group Report' (selected), 'Sensor Report', 'Filter Report', and 'Server Report'. The 'IP Group' is set to 'Internal Zone' and 'Sensors' to 'All'. There are two checked sections: 'Include Anomaly List' and 'Include Anomaly Overview'. The 'Include Anomaly List' section has dropdowns for 'Fields' (Anomaly #, Prefix, IP G), 'Class' (All), 'Decoders' (All), 'Direction' (Inbound, Outbound), and 'Count' (Top 10). The 'Include Anomaly Overview' section has dropdowns for 'Data Units' (Active Anomalies/s), 'Consolidation' (Maximum), and a checked 'Group Sensors' checkbox. The 'Scheduler' section has tabs for 'Daily' (selected), 'Weekly', 'Monthly', and 'Once'. Below the tabs are 'Reporting Time (H:M)' (08:00) and 'Report Time Frame' (Previous Day). At the bottom, there is a 'Comments' section and three buttons: 'Save', 'Save & Execute Now', and 'Delete'.

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 from the **Scheduler** tab. Only the active tab is taken into consideration.

To see how the email would look like without waiting for the preconfigured time, enter your email address, and then click the **[Save & Execute Now]** button. You should receive an email containing the report within a few seconds. If you do not, verify the settings from Configuration » General Settings » Outgoing Email.

All emails are formatted as HTML messages and include MIME attachments.

## Configuration » Schedulers » Event Reporting

Events are short text messages that describe errors, warnings or the change of an operational status. They are generated by Wansight components and logged by Console.

You can list events in Reports » Devices » [Component Name] » [Component Type] Event sub-tab. To search, sort or filter event messages, click the small down arrow that appears when hovering over the Event column header. To see additional details about an event click the [+] button from the first column.

To see a recent list of **Latest Events**, click the small bottom edge of the window to raise the South Region, or press Ctrl+E. On one side the Latest Events tab displays the latest 50 events, while on the other side it shows a 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 its importance:

- **MELTDOWN** – Meltdown events are generated in severe situations, such as hardware failures
- **CRITICAL** – Critical events are generated when significant software errors occur, such as a memory exhaustion situation
- **ERROR** – Error events are usually caused by misconfigurations, communication errors between components, or even bugs. Sensors auto-recover from errors by restarting themselves
- **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 in
- **DEBUG** – Debug events are generated only to help troubleshooting or the development process

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

**Event Reporting Configuration**

Polling Time: Every Minute Minimum Severity: WARNING

Send Events by Email  
 To: noc@isp  
 C.C.: B.C.C.:  
 Subject: Warnings and errors from Vanguard

Send Events to Syslog  
 Syslog Host: localhost Syslog Facility: LOCAL5

Send Events as SNMP Traps  
 SNMP Host: 192.168.200.71 Community: test  
 SNMP OID: 600.1.3.6.1.4.1.5089.1.0.1.1.3.6.1.4.1.5089.2.0.999

Comments

Save



## Configuration » General Settings » Outgoing Email

Console sends notification emails using the settings from Configuration » General Settings » Outgoing Email:

- **From Email** – The address you would like to appear as the sender
- **From Name** – The name as you would like it to appear as the sender
- **Mailer** – Console supports several mailing systems:
  - *PHP Mail* – Use the PHP mail() function. To use it, you may have to configure a Mail Transfer Agent (Postfix, Qmail, Sendmail) on the Console server
  - *SMTP* – Use the integrated SMTP support to send emails directly, without using a local Mail Transfer Agent
  - *Sendmail* – Send emails using the sendmail command. To use it, you may have to configure a Mail Transfer Agent (Postfix, Qmail, Sendmail) on the Console server
- **SMTP Security** – Security options:
  - *None* – No encryption
  - *SSL* – Enable SSL encryption
  - *TLS* – Enable TLS encryption
- **SMTP Host** – Specify the SMTP server(s). You can include backup SMTP server(s) separated by the “;” character
- **SMTP Port** – TCP port to connect to, usually 25 (insecure) or 587 (secured by SSL/TLS)
- **SMTP Login/Password** – Credentials used for SMTP authentication. When the fields are empty, no authentication is performed
- **Email Tester** – Send a test email to verify the settings

If you can send emails with the Email Tester, but you are not receiving emails sent via Response actions, check if there are errors when executing the command “php /opt/andrisoft/webroot/rep\_reports.php” from CLI.

## Configuration » General Settings » User Management

To add, modify or delete Console user accounts go to Configuration » General Settings » User Management.

Each Console user must be assigned to one role / access level:

- **Administrator** – Has full privileges and is allowed to manage other user accounts. Is the only role allowed to access Configuration » General Settings » License Manager
- **Operator** – Can change any configuration but is not authorized to modify other user accounts
- **Guest** – Has read-only access to Console, without access to any configuration. Can have a granular, permission-based access to specific reports, dashboards, Sensors, IP groups, tools, etc.

To add a Console account, press [**Add User**] and then select the desired role. You can modify an account by double-clicking it, or by selecting it and by pressing the [**Modify User**] button.

The **Enabled** checkbox enables or disables the selected account.

There are two **Authentication** options:

- *Local Password* – The user will be authenticated with the password entered in the **Password** field. All passwords are hashed and only the hash is stored in the database, so it's not possible to decrypt them
- *Remote Authentication* – The user will be authenticated by remote LDAP or RADIUS servers configured in Configuration » General Settings » User Authentication (details on page 44)

The **Full Name**, **Company**, **Position**, **Email**, **Phone** and **Comments** fields are optional. These details are not visible anywhere else.

**Landing Tab** shows the tab that opens immediately after logging in. The list is dynamic and expands as you

add Sensors, dashboards, IP groups, etc. You should probably set the Landing Tab to a relevant dashboard or report.

**Minimum Severity** shows the minimum severity level of events displayed in Console.

**Reports Region** lets you switch the position of the Reports Region (described on page 14).

**Configuration Region** lets you switch the position of the Configuration Region (described on page 14).

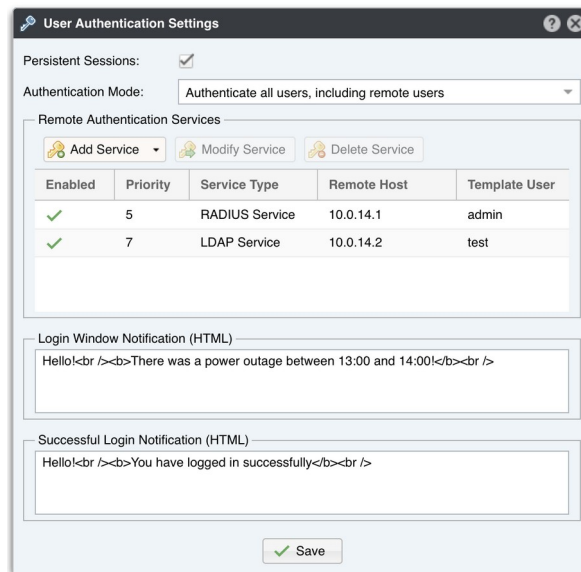
**Console Theme** allows you to change how Console looks after re-login.

**Console Notifications** manages the visual and audio notifications sent by Responses.

**REST API Access** controls whether the user has access to the REST API using his credentials. The REST API is documented on [http://<console\\_ip>/wanguard-api-ui](http://<console_ip>/wanguard-api-ui) or [https://<console\\_ip>/wanguard-api-ui](https://<console_ip>/wanguard-api-ui).

## Configuration » General Settings » User Authentication

To configure remote authentication mechanisms and login window settings click Configuration » General Settings » User Authentication.



**Persistent Sessions** enables cookie-based authentication for Console users that select the *Remember* option in the login screen. Subsequent sessions skip the login screen for the next 30 days or until the user logs out.

**Authentication Mode** enables or disables the authentication of Console users that are not defined in Configuration » General Settings » User Management but are defined in LDAP or Radius.

Console permits the use of external Radius and LDAP servers for end user authentication.

**LDAP** server settings:

- **Priority** – You can set the order in which Console connects to multiple authentication services. The authentication process stops after the first successful authentication
- **Template User** – Remotely-authenticated users without a Console account have the privileges of the Template User
- **LDAP Host** – IP or hostname of the LDAP server. To connect to a LDAP server by SSL, set this parameter as `ldaps://<IP>/`
- **Login Attribute** – Enter the LDAP attribute that contains the username. For Active Directory it usually is `mailNickname` or `sAMAccountName`, for OpenLDAP or IBM Directory Server it could be `uid`
- **LDAP Base DN** – Specify the location in the LDAP hierarchy where Console should begin searching for usernames for authorization requests. The base DN may be something equivalent to the organization, group, or domain name (AD) of the external directory: `dc=domain,dc=com`

- **Bind User DN/Password** – Distinguished name and password for a LDAP user permitted to search within the defined Base DN
- **Search Filter** – Can contain rules that restrict which users are authenticated using the current configuration. For example, the string "|(department=\*NOC\*)(department=ISP)" only allows users to authenticate in Console from departments containing the string "NOC" or (|) from the "ISP" department

**RADIUS** server settings:

- **Priority** – You can set the order in which Console connects to multiple authentication services. The authentication process stops after the first successful authentication
- **Template User** – Remotely authenticated users without a Console account have the privileges of the Template User
- **RADIUS Host** – IP or hostname of the Radius server
- **RADIUS Port** – Port through which the Radius server is listening for authentication requests
- **RADIUS Protocol** – Protocol used for authentication purposes:
  - **PAP** (Password Authentication Protocol) – provides a simple method for the peer to establish its identity using a 2-way handshake
  - **CHAP** (Challenge-Handshake Authentication Protocol) – authenticates a user or network host to an authentication entity
  - **MSCHAP** – is the Microsoft version of the Challenge-handshake authentication protocol, CHAP
  - **MSCHAP2** – is another version of Microsoft version of the Challenge-handshake authentication protocol, CHAP
- **RADIUS Secret** – Enter the credentials for connecting to the Radius server.

The contents of the **Login Window Notification** field is shown inside the Console login window.

The contents of the **Successful Window Notification** field is shown inside the Console window after log in.

## Reports » Tools » Flows

**Reports » Tools** contains a link to the **Flows** item only when there is at least one Flow Sensor defined. Here you can list, aggregate, filter and sort flow records, or generate traffic tops and statistics.

There are 2 sub-tabs, located at the left lower side of the window:

### Flow Records

Here you can list and filter flow data.

- **Sensor Interfaces** – Select the interfaces you are interested in. Administrators can restrict the interfaces visible by guest accounts
- **Flow Filtering Expression** – Here you can enter a filtering expression for flows. Click the star icon on the right to open a window that shows you the correct syntax. Frequently-used filtering expressions can be saved there and reused at any later time
- **Export** – If the output is not very large, it can be viewed, converted to PDF, emailed or printed. If you need to list huge amounts of flow data, doing it solely from the web browser may not be a good idea. In this case, select the “CLI” option to view the shell command used for flow listing. You can then execute the command from the shell and forward the output to a file
- **Time Range** – Select a predefined time range, or select “Custom...” to enter a specific time interval, in order to list flows that started or ended inside that interval. Time zone differences between the Console server and remote Flow Sensor servers are not adjusted automatically
- **Limit Flows** – Show only the first N flows
- **Output** – You can select several output formats, or you can type in your own format (the syntax is listed at the bottom of the Flow Filtering Expression help)
- **Aggregation** – By default, flows are not aggregated. By checking the appropriate options, you can select how to aggregate flows. You can aggregate entire subnets by selecting src(dst)IPv4(IPv6)/<subnet bits>
- **Sorting** – When listing flows sent from different interfaces, you can sort them after to the start time of the flows. Otherwise, flows are listed in the sequence of the selected interfaces

### Flow Tops

Here you can generate tops from flow data.

- **Sensor Interfaces** – Select the interfaces you are interested in. Administrators can restrict the interfaces visible by guest accounts
- **Top Type** – Select the desired top type from the drop-down menu
- **Flow Filtering Expression** – Here you can enter a filtering expression for flows. Click the star icon on the

right to open a window that shows you the correct syntax. Frequently-used filtering expressions can be saved there and reused at any later time

- **Export** – If the output is not very large, it can be viewed, converted to PDF, emailed or printed. If you need to list huge amounts of top data, doing it solely from the web browser may not be a good idea. In this case, select the “CLI” option to view the shell command used for flow listing. You can then execute the command from the shell and forward the output to a file
- **Time Range** – Select a predefined time range, or select “Custom...” to enter a specific time interval, in order to count only flows that started or ended inside the interval. Time zone differences between the Console server and remote Flow Sensor servers are not adjusted automatically
- **Order By** – Select the sorting unit
- **Limit** – Limit the output to only those records whose packets or bytes match the specified condition
- **Top** – Limit the top listing to the first N records
- **Output** – You can select several output formats, or you can type in your own format (the syntax is listed at the bottom of the Flow Filtering Expression help)
- **Aggregation** – By default, flows are not aggregated. By checking the appropriate options, you can select how to aggregate flows. You can aggregate entire subnets by selecting src(dst)IPv4(IPv6)/<subnet bits>

## Reports » Tools » Packets

**Reports » Tools** contains a link to the **Packets** item only when there is at least one Packet Sensor in use. The number of packet dumps (network trace, packet trace) that are active is displayed within the panel.

Here you can easily capture packets from Packet Sensors deployed on various parts of your network. You can view the content of packets directly from the user interface. Console provides an integrated packet analyzer that resembles the popular WireShark software.

There are 2 sub-tabs located at the lower left side of the window:

### Active Packet Dumps

Administrators, operators, and guests with packet capturing privileges can generate packet dumps by clicking the **[Capture Packets]** button.

- **Description** – An optional short description to help you identify the packet dump
- **Packet Sensor(s)** – Select the Packet Sensor that will capture the packets. Administrators can restrict the Packet Sensors that are visible to guest accounts
- **BPF Expression** – Click the star icon on the right to open a window that describes the Berkley Packet Filter (BPF) syntax. Frequently used BPF expressions can be saved there and reused at a later time. Entering a BPF expression is mandatory. To capture all IP packets, enter “ip”
- **Max. Running Time** – Maximum running time of the capturing thread / process
- **Stop Capture Time** – When Max. Running Time is set to “Unlimited”, you can set the exact time when the capturing thread stops
- **Max. File Size (MB)** – This option is used for splitting packet dumps into multiple files of <number> Mbytes. Before writing a raw packet to a file, Packet Sensor checks whether the file is currently larger than <number> and, if so, closes the current file and opens a new one



- **Max. Packets** – The capture stops after receiving <number> packets
- **Max. Files 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 up to <value> packets every 5 seconds
- **Packet Payload** – Select “Full” to capture the entire packet, “Only Layer 3” to zero-out the packet’s payload except for the IP header, or “Only Layer 4” to retain only the TCP, UDP and ICMP headers
- **Snapshot Length** – Sniff <number> bytes of data from each packet rather than the default of 65535 bytes. 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 this <number> to the smallest number that will capture the protocol information you are interested in
- **Filename Prefix** – Name of the packet dump file. If any file-rotation options are used, a number will be appended to the filename
- **Comments** – This field may contain comments about the packet dump

All active Packet Dumps are listed in a table having the following format:

- **Status** – Indicates the status of the capture process. It is green if it is running, and red if it is not
- **Packet Logger** – Packet Sensor or Packet Filter used for capturing the packets
- **Packets / Dropped** – Number of packets captured and number of packets dropped
- **Files / Size** – Number of dump files generated and the size of the latest dump file
- **Description [BPF]** – Description and the BPF expression used for the dump
- **From** – Date when the capture process started
- **Until** – Time limit or other conditions that will cause the capturing process to stop
- **Sampling** – Type of sampling being used
- **Payload** – The value of the “Packet Payload” parameter
- **Comments** – The value of the “Comments” parameter
- **Actions** – Click on the first icon to view the latest dump file in an integrated packet analyzer interface. Click on the second icon to download the latest dump file to your computer. If downloading does not work, but viewing does, increase the values of the *max\_execution\_time* and *memory\_limit* from php.ini. Click the third icon to stop the capturing process

## Packet Dump Archive

By default, the captures are sorted by time in descending order. By clicking the down arrow of any column header, you can apply row filters, change sorting direction and toggle the visibility of columns.

The [+] sign from the first column expands each row for additional information about the capture and provides access to packet dump files. The columns are explained in the previous section.

## Reports » Devices » Overview

Here you can see the most important operational parameters collected in realtime from all active Wansight components and servers.

The **Console** table displays the following data:

|                                |   |
|--------------------------------|---|
| <b>Status</b>                  | A green check mark indicates that Console is functioning properly. When a red "X" appears, enable the WANsupervisor service on the Console server |
| <b>Online Users</b>            | Active Console sessions   |
| <b>Avg. DB Bits/s (In/Out)</b> | Average number of bits/s sent and received since the start of the database  |
| <b>Avg. DB Queries/s</b>       | Average number of queries per second since the start of the database  |
| <b>DB Clients</b>              | DB clients that are currently using the database  |
| <b>DB Connections</b>          | Active connections to the database  |
| <b>DB Size</b>                 | Disk space used by the database   |
| <b>Free DB Disk</b>            | Disk space available on the partition configured to store the Console database  |
| <b>Free Graphs Disk</b>        | Disk space available on the partition configured to store IP graphs   |
| <b>Time Zone</b>               | Time zone of the Console server   |
| <b>Console Time</b>            | The clock from the Console server   |
| <b>Uptime</b>                  | Uptime of the database  |

The **Servers** table displays the following data for each server that runs software components of Wansight:

|                    |  |
|--------------------|--|
| <b>Status</b>      | A green check mark indicates that the server is connected to the database. When a red "X" is displayed, start the WANsupervisor service and make sure that the clock from the Console server is synchronized with the clock from the remote server |
| <b>Server Name</b> | Displays the name of the server and the associated color. Click to open a new tab with data specific to the server. Administrators and operators can right-click to open the Server Configuration window   |
| <b>Load</b>        | Load average reported by the Linux kernel for the last 5 minutes   |
| <b>Free RAM</b>    | Available RAM. Swap memory not counted   |
| <b>CPU% User</b>   | Percentage of CPU resources used by the user space processes. Can be >100% on multiple cores/CPU's (e.g. the maximum value for a quad-core system is 400%)   |
| <b>CPU% System</b> | Percentage of CPU resources used by the kernel. Can be >100% on multiple cores/CPU's (e.g. the maximum value for a quad-core system is 400%)   |
| <b>CPU% IOWait</b> | Percentage of CPU resources waiting for I/O operations to complete. A high number  |

|                               |   |
|-------------------------------|---|
|                               | indicates an I/O bottleneck   |
| <b>CPU% Idle</b>              | Percentage of idle CPU resources. Can be >100% on multiple cores/CPUs (e.g. the maximum value for a quad-core system is 400%) |
| <b>Free Flows Disk</b>        | Disk space available on the partition that is configured to store flows   |
| <b>Free Dumps Disk</b>        | Disk space available on the partition that is configured to store packet dumps  |
| <b>Contexts/IRQs/SoftIRQs</b> | Context switches, hardware interrupts and software interrupts per second  |
| <b>Uptime</b>                 | Uptime of the operating system  |

The **Sensor Clusters** table displays the following data when there is at least one Sensor Cluster active:

|                          |  |
|--------------------------|--|
| <b>Status</b>            | A green check mark indicates that the Sensor Cluster is connected to the database. If you see a red "X" instead, make sure that the WANsupervisor service is running and look for errors in the event log (see page 40)                                    |
| <b>Sensor Name</b>       | Displays the name of the Sensor Cluster and the associated color. Click to open a new tab with data specific to the Sensor Cluster. Administrators and operators can right-click to open the Sensor Cluster configuration window                           |
| <b>Pkts/s (In / Out)</b> | Inbound and outbound packets/second throughput   |
| <b>Inbound Bits/s</b>    | Inbound bits/second throughput and the usage percent   |
| <b>Outbound Bits/s</b>   | Outbound bits/second throughput and the usage percent  |
| <b>Received Pkts/s</b>   | Packet/s reported by the associated Sensors  |
| <b>IPs (Int./Ext.)</b>   | Number of IP addresses that sent or received traffic. The Int(ernal)/Ext(ernal) IPs are the IPs from inside/outside the IP Zone. The Stats Engine parameter from the associated Sensors' configurations enables or disables the monitoring of external IPs |
| <b>Dropped</b>           | Packets dropped by the Server Cluster  |
| <b>CPU%</b>              | Percentage of CPUs used by the process   |
| <b>RAM</b>               | Amount of memory used by the process   |
| <b>Start Time</b>        | Time when the Sensor Cluster instance started  |
| <b>Server</b>            | Which server runs the Sensor Cluster. Click to open a tab with specific data. Administrators and operators can right-click to open the Server Configuration window   |

The **Packet Sensors** table displays the following data when there is at least one Packet Sensor active:

|                    |  |
|--------------------|--|
| <b>Status</b>      | A green check mark indicates that the Packet Sensor is connected to the database. If you see a red "X" instead, make sure that the WANsupervisor service is running and look for errors in the event log (see page 40) |
| <b>Sensor Name</b> | Displays the name of the Packet Sensor and the associated color. Click to open a new tab with specific information. Administrators and operators can right-click to open the Packet Sensor Configuration window        |

|                          |  |
|--------------------------|--|
| <b>Pkts/s (In / Out)</b> | Inbound and outbound packets/second throughput after IP or MAC validation  |
| <b>Inbound Bits/s</b>    | Inbound bits/second throughput after IP / MAC Validation, and the usage percent  |
| <b>Outbound Bits/s</b>   | Outbound bits/second throughput after IP / MAC Validation, and the usage percent   |
| <b>Received Pkts/s</b>   | Rate of sniffed packets before IP / MAC Validation   |
| <b>IPs (Int / Ext)</b>   | IP addresses that sent or received traffic. The Int(ernal)/Ext(ernal) IPs are the IPs from inside/outside the IP Zone. The Stats Engine parameter from the Sensor configuration enables or disables the monitoring of external IPs |
| <b>Dropped</b>           | Packets dropped by the packet capturing engine. A high number usually indicates a sniffing performance problem   |
| <b>CPU%</b>              | Percentage of CPUs used by the process   |
| <b>RAM</b>               | Amount of memory used by the process   |
| <b>Start Time</b>        | Time when the Packet Sensor started  |
| <b>Server</b>            | Which server runs the Packet Sensor. Click to open a tab with specific data. Administrators and operators can right-click to open the Server Configuration window  |

The **Flow Sensors** table displays the following data when there is at least one Flow Sensor active:

|                          |  |
|--------------------------|--|
| <b>Status</b>            | A green check mark indicates that the Flow Sensor is connected to the database. If you see a red "X" instead, make sure that the WANsupervisor service is running and look for errors in the event log (see page 40)   |
| <b>Sensor Name</b>       | Displays the name of the Flow Sensor. Click to open a new tab with specific data. Administrators and operators can right-click to open the Flow Sensor Configuration window  |
| <b>Interface</b>         | Interface name and the associated color. If the interface names are missing for more than 5 minutes after the Flow Sensor has started, check the troubleshooting guide from page 33  |
| <b>Pkts/s (In / Out)</b> | Inbound and outbound packets/second throughput after IP / AS Validation  |
| <b>Inbound Bits/s</b>    | Inbound bits/second throughput after IP / AS Validation, and usage percent   |
| <b>Outbound Bits/s</b>   | Outbound bits/second throughput after IP / AS Validation, and usage percent  |
| <b>IPs (Int / Ext)</b>   | IP addresses that send or receive traffic. The Int(ernal)/Ext(ernal) IPs are the IPs from inside/ outside the IP Zone. The Stats Engine parameter from the Sensor configuration enables or disables the monitoring of external IPs   |
| <b>Flows/s</b>           | Flows per second received by the Flow Sensor   |
| <b>Flows Delay</b>       | Because traffic data must be aggregated, 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.<br>Flow Sensor cannot run with flow delays higher than 5 minutes |
| <b>Dropped</b>           | Unaccounted flows. A high number indicates a performance problem of the Flow Sensor or a network connectivity issue with the flow exporter   |
| <b>CPU%</b>              | Percentage of CPU resources used by the Flow Sensor process  |
| <b>RAM</b>               | Amount of RAM used by the Flow Sensor process  |

|                   |   |
|-------------------|---|
| <b>Start Time</b> | Time when the Flow Sensor started   |
| <b>Server</b>     | Which server runs the Flow Sensor. Click to open a tab with specific data. Administrators and operators can right-click to open the Server Configuration window |

The **SNMP Sensors** table displays the following data when there is at least one SNMP Sensor active:

|                              |  |
|------------------------------|--|
| <b>Status</b>                | A green check mark indicates that the SNMP Sensor is connected to the database. If you see a red "X" instead, make sure that the WANsupervisor service is running and look for errors in the event log (see page 40)   |
| <b>Sensor Name</b>           | Displays the name of the SNMP Sensor. Click to open a new tab with specific data. Administrators and operators can right-click to open the SNMP Sensor Configuration window  |
| <b>Interface</b>             | Interface name and the associated color  |
| <b>Pkts/s (In / Out)</b>     | Inbound and outbound packets/second throughput   |
| <b>Inbound Bits/s</b>        | Inbound bits/second throughput and usage percent   |
| <b>Outbound Bits/s</b>       | Outbound bits/second throughput and usage percent  |
| <b>Errors/s (In / Out)</b>   | For packet-oriented interfaces, it represents the number of inbound and outbound packets that contained errors, preventing them from being deliverable to a higher-layer protocol. For character-oriented or fixed-length interfaces, it represents the number of inbound transmission units that contained errors preventing them from being deliverable to a higher-layer protocol   |
| <b>Discards/s (In / Out)</b> | Discarded inbound and outbound packets even though no errors were detected to prevent them from being deliverable to a higher-layer protocol. One possible reason for discarding such a packet could be to free up buffer space  |
| <b>Oper. Status</b>          | Current operational state of the interface. The <i>Testing</i> state indicates that no operational packets can be passed. If Administrative Status is <i>Down</i> then Operational Status should be <i>Down</i> . If Administrative Status is changed to <i>Up</i> then Operational Status should change to <i>Up</i> if the interface is ready to transmit and receive network traffic; it should change to <i>Dormant</i> if the interface is waiting for external actions (such as a serial line waiting for an incoming connection); it should remain in the <i>Down</i> state if and only if there is a fault that prevents it from going to the <i>Up</i> state; it should remain in the <i>NotPresent</i> state if the interface has missing (typically, hardware) components |
| <b>Admin. Status</b>         | Desired state of the interface. The <i>Testing</i> state indicates that no operational packets can be passed. When a managed system initializes, all interfaces start with the Administrative Status in the <i>Down</i> state. As a result of either explicit management action or per configuration information retained by the managed system, the Administrative Status is then changed to either the <i>Up</i> or <i>Testing</i> states (or remains in the <i>Down</i> state)  |
| <b>CPU%</b>                  | Percentage of CPU resources used by the process  |
| <b>RAM</b>                   | Amount of RAM used by the process  |
| <b>Start Time</b>            | Time when the SNMP Sensor started  |
| <b>Server</b>                | Which server runs the SNMP Sensor. Click to open a tab with specific data. Administrators and operators can right-click to open the Server Configuration window  |

## Reports » Devices » Sensors

Clicking on a Sensor opens a tab with specific information. The tab includes a few sub-tabs located at the lower side of the window. All sub-tabs share the following common toolbar fields:

- **Sensor Interfaces** – Select the Sensor interfaces you are interested in, or select “All”. Administrators can restrict which Sensors are accessible by guest accounts
- **Time Range** – Select a predefined time range, or select “Custom...” to enter a specific time interval

### Sensor Dashboard

This tab allows you to view the most relevant data collected by Sensors in form of widgets. The configuration of the Sensor Dashboard does not apply to a particular Sensor, so the changes you make here are visible for other Sensor Dashboards as well. The operation of dashboards is described in the “Reports » Dashboards” chapter on page 59.

The configuration of Sensor widgets is outlined in the following paragraphs.

### Sensor Graphs

This sub-tab allows you to view a variety of Sensor-related histograms for the selected Sensor Interface(s):

- **Data Units** – Select one or more data units:
  - *Most Used* – Frequently-used data units
  - *Packets* – Inbound packets/second (+ on Y-axis) and outbound packets/second (- on Y-axis)
  - *Bits* – Inbound bits/second (+ on Y-axis) and outbound bits/second (- on Y-axis)
  - *Applications* – Sensor 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. These graphs are not generated when the Sensor’s Stats Engine parameter is set to “Disabled”
  - *Internal & External IPs* – Number of IP addresses that send or receive traffic. The “internal” and “external” IPs are the hosts located inside or outside the IP Zone. The Sensor’s “Stats Engine” parameter 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 blocks. A spike in the external IPs graph usually means that you have received a spoofed attack
  - *Received Frames* – For Packet Sensor, it represents the number of packets/s received before IP or MAC validation. For Flow Sensor, it represents the number of flows/s received before IP or AS validation
  - *Dropped Frames* – For Packet Sensor, it represents the number of packets dropped by the packet capturing engine. A high number indicates a sniffing performance problem. For Flow Sensor, it represents the number of unaccounted flows. A high number indicates a wrong configuration or a network connectivity issue with the flow exporter

- *Unknown Frames* – For Packet Sensor, it represents the rate of packets not passing IP validation. For Flow Sensor, it represents the rate of invalidated flows
- *Unknown Sources* – Source IP addresses that did not pass IP validation
- *Unknown Destinations* – Destination IP addresses that did not pass IP validation
- *Avg. Packet Size* – Average packet size in bits/packet
- *CPU%* – Percentage of CPU resources used by the Sensor process
- *RAM* – Amount of RAM used by the Sensor process
- *Load* – Load reported by the Linux kernel for 5 minute intervals
- *IP Graphs* – Number of IP graphs files updated
- *IP Accounting* – Number of IP accounting records updated
- *Profile Graphs* – Number of profiling files updated
- *IP Graphs Time* – Seconds needed to update the IP graphs files
- *Profile Graphs Time* – Seconds needed to update the profiling files
- *Processing Time* – Seconds needed to perform traffic analysis functions
- *IP Structures* – Number of Internal IP structures necessary to keep track of IPs
- *IP Structure RAM* – RAM bytes used by each IP structure
- *Dataplane* – Parameters collected from the DPDK-based Capture Engine
- *Bytes/time unit* – Bytes/<time unit> throughput value
- **Size** – Select a predefined graph dimension or enter a custom one in “<X> x <Y>” format, where <X> and <Y> are the X-axis and Y-axis pixels
- **Title** – Graphs have an automatically-generated title for the “Auto” option, no title for the “None” option, or you can enter your own text to be rendered as a title
- **Legend** – Select the level of detail for the graph’s legend
- **Consolidation** – If you are interested in spikes, choose the *MAXIMUM* aggregation type. If you are interested in average values, choose the *AVERAGE* aggregation type. If you are interested in low values, choose the *MINIMUM* aggregation type
- **Grouping**
  - **Sensor Interfaces** – Select to generate a single graph for the Sensor Interfaces selected
- **Stacking**
  - **Sensor Interfaces** – Select to view the summed up, stacked values for multiple Sensor Interfaces

## Sensor Tops

This sub-tab allows you to generate various traffic tops per Sensor Interface. The Stats Engine parameter from the Sensor configuration manages the collection of various Sensor tops.

- **Decoders** – Select the decoder that analyzes the type of traffic that interests you
- **Top Unit** – Select a top type:

- *Talkers* – Hosts from your network that sent or received the most traffic for the selected decoder. Not available when the Stats Engine parameter from the Sensor configuration is set to “Disabled”
- *IP Groups* – IP groups that sent or received the most traffic for the selected decoder. Not available when the Stats Engine parameter from the Sensor configuration is set to “Disabled”
- *External IPs* – External IPs that sent or received the most traffic for the selected decoder. Available when the Stats Engine parameter from the Sensor configuration is set to “Extended” or “Full”
- *Autonomous Systems* – Autonomous systems that sent or received the most traffic. Available only when the Stats Engine parameter from the Sensor configuration is set to “Extended” or “Full”
- *Transit Autonomous Systems* – Transit autonomous systems that sent or received the most traffic. Available only when the Sensor is configured to extract Transit AS data from a BGP dump file
- *Countries* – Countries that sent or received the most traffic. Available only when the Stats Engine parameter from the Sensor configuration is set to “Extended” or “Full”
- *TCP Ports* – Most-used TCP ports. Not available when the Stats Engine parameter from the Sensor configuration is set to “Disabled”
- *UDP Ports* – Most-used UDP ports. Not available when the Stats Engine parameter from the Sensor configuration is set to “Disabled”
- *IP Protocols* – Most-used IP protocols (the protocol used in the data portion of the IP datagram). Not available when the Stats Engine parameter from the Sensor configuration is set to “Disabled”
- *IP Versions* – Counters for IPv4 traffic and IPv6 traffic. Not available when the Stats Engine parameter from the Sensor configuration is set to “Disabled”
- **Traffic Direction** – Direction of traffic, *All*, *Inbound* or *Outbound*
- **Display Options** – Various display options
- **Grouping**
  - **Sensor Interfaces** – When unchecked, a different top is generated for each selected Sensor Interface. When checked, top data is combined

You can increase the number of top records and add new decoders in Configuration » General Settings » Graphs & Storage, see page 16.

Generating tops for many Sensor Interfaces and for long time ranges may take minutes. If the report page timeouts, increase the *max\_execution\_time* parameter from *php.ini*.

## Sensor Events

This sub-tab lists events generated by the selected Sensor(s) for the selected time range. The events are described in the “Event Reporting” chapter on page 40.

## AS Graphs

Flow Sensor and Packet Sensor can generate per-autonomous system bandwidth histograms. This feature is



enabled for Packet Sensor when the Stats Engine parameter is set to "Full", and for Flow Sensor when the Stats Engine parameter is set to "Full" or "Extended".

The inbound traffic represents the traffic received by the AS, while the outbound traffic represents the traffic sent from the AS.

- **AS Number(s)**

Select one of the following options:

- *Upstream* – Select to see the traffic sent to or coming from the Autonomous Systems from the list on the right
- *Transit* – Select to see the traffic that transited the Autonomous Systems from the list on the right
- *Peering* – Select to see traffic to/from your AS peers (PrevAdjacentAS and NextAdjacentAS in NetFlow v9) from the list on the right
- *Downstream* – Select to see the traffic to/from your downstream Autonomous Systems from the list on the right

Click the star icon on the right to open a window containing the correct syntax for the AS list. Frequently-searched AS numbers can be saved there, and used at a later time. To see a list of AS numbers owned by a particular organization, go to Help » IP & AS Information » AS Numbers List or consult <https://bgp.he.net>

- **Size** – Select a predefined graph dimension or enter a custom one in a "<X> x <Y>" format, where <X> and <Y> are the X-axis and Y-axis pixels
- **Title** – Graphs have an automatically-generated title for "Auto", no title for "None", or you can enter your own text to be rendered as a title
- **Legend** – Select the level of detail for the graph's legend
- **Grouping**
  - **Sensor Interfaces** – When unchecked, a different graph is generated for each selected Sensor Interface. When checked, the data is combined
  - **ASNs** – Select if you want to view a single graph for multiple AS numbers
- **Stacking**
  - **ASNs** – Select to stack up to 20 ASNs into a single graph

## Country Graphs

Flow Sensor and Packet Sensor can generate per-country bandwidth histograms. This feature is enabled when the Sensor's Stats Engine parameter is set to "Full" or "Extended".

- **Countries** – Select the country or countries from the drop-down list, or click the star icon on the right to open a window with saved selections for continents and world regions
- **Size** – Select a predefined graph dimension or enter a custom one in a "<X> x <Y>" format, where <X> and <Y> are the X-axis and Y-axis pixels
- **Title** – Graphs have an automatically-generated title for "Auto", no title for "None", or you can enter your own text to be rendered as a title
- **Legend** – Select the level of detail for the graph's legend

- **Grouping**
  - **Sensor Interfaces** – Select to generate a single graph for the selected Sensor Interfaces
  - **Countries** – Select to view a single graph when multiple countries are selected
- **Stacking**
  - **Countries** – Select to stack up to 20 countries into a single graph

## Flow Records

You can list and filter the flow data collected for the selected Flow Sensor Interfaces. The options are described in detail by the “Flow Collectors” chapter on page 46.

## Flow Tops

You can generate tops from the flow data collected for the selected Sensor Interfaces. The options are described in detail in the “Flow Collectors” chapter on page 46.

## Reports » Dashboards

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

Any dashboard can be configured to refresh all its widgets, flicker-free, at intervals ranging from 5 seconds to 15 minutes.

A few sample dashboards are included by default. If you are a Console administrator or operator you can **create** and configure your own dashboards by clicking Reports » Dashboards » [+] button from the panel's title. Guest accounts are not allowed to add or make modifications to dashboards.

In the dashboard **configuration**, you can edit the name of the dashboard, set permissions, layout, or choose to override the time range of widgets with the time range of the dashboard.

Each dashboard contains **widgets**. 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 configure a widget, click the second icon from its title bar. To delete a widget, click the third icon from its title bar.

Along with a few 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 described in other chapters.

## Reports » IP Addresses & Groups

This chapter describes how to generate detailed traffic reports for any IP address, block or group included in Configuration » Network & Policy » [IP Zone].

**Reports » IP Addresses** allows you to quickly generate traffic reports for IP addresses and blocks, either entered manually on the upper side of the panel, or selected from the expandable tree below.

**Reports » IP Groups** lists IP groups defined in IP Zones. Select an IP group to generate a traffic report for all IP blocks belonging to it. To search for a specific IP group, enter a sub-string contained in its name on the upper side of the panel.

The traffic report tab includes a few sub-tabs located on the lower side of the window. All sub-tabs share the following common toolbar fields:

- **Sensor Interfaces** – Select the Sensor Interfaces you are interested in. Administrators can restrict the Sensors accessible by guest accounts
- **Time Range** – Select a predefined time range, or select “Custom...” to enter a specific time interval

## IP Dashboard

Here you can group the most relevant data collected for the selected Sensor Interfaces and for the selected IP address, block or group. The configuration of this dashboard does not apply to a particular IP address, block or group, so the changes you make will be visible for other IP dashboards as well. The operation of dashboards is described in the “Reports » Dashboards” chapter on page 59.

## IP Graphs

You can generate IP graphs only for the IP addresses, blocks and groups specifically defined in your IP Zone(s), or that belong to a subnet having the IP Graphing parameter set to “Yes”.

- **Decoders & Data Unit** – Select the decoders and data unit you are interested in
- **Size** – Select a predefined graph dimension or enter a custom one in a “<X> x <Y>” format, where <X> and <Y> are the X-axis and Y-axis pixels
- **Title** – Graphs have an automatically-generated title for “Auto”, no title for “None”, or you can enter your own text to be rendered as a title
- **Legend** – Select the detail of the graph’s legend
- **Consolidation** – If you are interested in spikes, choose the *MAXIMUM* aggregation type. If you are interested in average values, choose the *AVERAGE* aggregation type. If you are interested in low values, choose the *MINIMUM* aggregation type
- **Direction** – Generates a graph for both directions, swap inbound (+ Y axis) with outbound (- Y axis), or show only inbound or outbound traffic

- **Grouping**
  - **Sensor Interfaces** – Generates a single graph for the selected Sensor Interfaces
  - **Subnet IPs** – Uncheck this option if you want a different traffic graph displayed for every IP address contained in the selected IP block or IP group. Do not uncheck this option on large subnets
- **Stacking**
  - **Decoders** – Select to view the summed up, stacked values for the selected decoders
  - **Sensor Interfaces** – Select to view the summed up, stacked values for multiple Sensor Interfaces
- **Permissions**
  - **Decoder Conflict** – If decoders can be included one within the other (e.g. IP contains TCP which 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 above, IP will be displayed as IP OTHER and TCP as TCP OTHER. However, when you select TCP, HTTP and TCP+SYN as decoders, the TCP+SYN decoder can be included in both TCP and HTTP, thus generating a decoder conflict. Check this option to stop detection of conflicting decoders, in order to generate more intuitive but potentially inaccurate traffic graphs
  - **Use Per-IP Data** – Creates a subnet graph by aggregating the IP graph data generated for every IP address contained in the selected IP block or group. This option will increase the load of the server if used frequently on large subnets. Use this option carefully, only when the IP block or group is not explicitly defined in the IP Zone but it is included in a larger subnet defined with the IP Graphing parameter set to “Yes”

The decoders, data units, and aggregation types can be modified in Configuration » General Settings » Graphs & Storage (see page 16).

## IP Accounting

You can generate IP accounting reports only for the IP addresses, blocks and groups specifically defined in your IP Zone(s), or that belong to a subnet that has the IP Accounting parameter set to “Yes”.

- **Decoders & Data Unit** – Select the decoders and data unit that you are interested in
- **Report Interval** – Select the minimum interval used to aggregate the accounting data: *Daily, Weekly, Monthly, Yearly*. The minimum accuracy of traffic accounting reports is 24 hours, therefore when you select a shorter time range you will still see the accounting data collected for the whole day
- **Direction** – Show both directions or only a single one
- **Group Sensor Interfaces** – Generates a single traffic accounting report for multiple Sensor Interfaces
- **Show IPs** – Check this option for the traffic accounting report to display each IP address contained in the selected IP block or group. Selecting this option also enables the option below
- **Use Per-IP Data** – Creates a traffic accounting report by aggregating the IP accounting data generated for every IP address contained in the selected IP block or group. This option will increase the load of the server if used frequently on large subnets. Use this option carefully, only when the selected IP block or group is not explicitly defined in the IP Zone but it is included in a larger subnet defined with the IP Accounting parameter set to “Yes”

- **Display Raw Values** – Check this option to avoid displaying values with metric prefixes

The decoders can be modified in Configuration » General Settings » Graphs & Storage (see page 16).

## Flow Records

You can list and filter the flow data collected for the selected Flow Sensor Interfaces and IP block, host or group. The options are described in detail by the “Flow Collectors” chapter on page 46.

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

## Flow Tops

You can generate tops from the flow data collected for the selected Flow Sensor Interfaces and IP block, host or group. The options are described in detail by the “Flow Collectors” chapter on page 46.

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

## Reports » Servers

Click on a server name anywhere in the user interface to open a tab containing server-specific information. This tab includes a few sub-tabs, located at the lower side of the window. All sub-tabs share the following common toolbar fields:

- **Servers** – Select the servers you are interested in, or select “All”. Administrators can restrict the servers available to guest accounts
- **Time Range** – Select a predefined time range, or select “Custom...” to enter a specific time interval

### Console / Server Dashboard

Allows you to group the most relevant server-related data. The configuration for this dashboard does not apply to a particular server, so the changes you make will be visible for other server dashboards as well. The operation of dashboards is described in the “Reports » Dashboards” chapter on page 59.

The configuration of Server and Console widgets is described in the following paragraphs.

### Console / Server Graphs

Server Graphs allows you to generate various histograms for the selected server(s):

- **Data Units** – Select one or more data units:
  - *Most Used* – Frequently-used data units
  - *System Load* – Load reported by the Linux kernel for a 5 minute period
  - *Free RAM* – Available RAM. Swap memory is not counted
  - *Database/Graphs/SSD/Flow Collector/Package Dumps Disk - Free space* – How much disk space is available for each file-system path
  - *Uptime* – Uptime of the operating system
  - *CPU% system/userspace/niced/idle* – Percentages of CPU resources used by the system, userspace processes, processes running with increased (nice) priority, and idle loop
  - *Number of processes* – Total number of processes that are running
  - *Hardware/Software CPU Interrupts* – CPU interrupts made by hardware and software events
  - *Context Switches* – Indicates how much time the system spends on multi-tasking
  - *Running Components* – Sensor instances
  - *Clock Delta* – Difference of time between the selected server and the Console server, in seconds. If the value is not zero run ntpd to keep the clock synchronized on all servers
  - *Database/Graphs/SSD/Package Dumps/Flow Collector Disk - Total* – How much disk space is allocated for the partitions that store the paths

- *Database/Graphs/SSD/Packet Dumps/Flow Collector Disk - Free Inodes* – Free inodes held by the partitions that store the paths
- *Database/Graphs/SSD/Packet Dumps/Flow Collector Disk - Ops/s* – Reads and writes made on the partitions that store the paths
- *Database/Graphs/SSD/Packet Dumps/Flow Collector Disk - Bytes/s* – Bytes/s on the partitions that store the paths
- *Server Interface(s) - Packets/Bits/Errors/Dropped* – Interface statistics collected for the network interfaces defined in the Configuration » Servers
- **Size** – Select a predefined graph dimension or enter a custom one in a “<X> x <Y>” format, where <X> and <Y> are the X-axis and Y-axis pixels
- **Title** – Graphs can have an automatically-generated title for the “Auto” option or title for the “None” option, or you can enter your own text to be rendered as a title
- **Legend** – Select the level of detail for the graph legend
- **Consolidation** – If you are interested in spikes, choose the *MAXIMUM* aggregation type. If you are interested in average values, choose the *AVERAGE* aggregation type. If you are interested in low values, choose the *MINIMUM* aggregation type
- **Grouping**
  - **Servers** – Generate a single graph for the selected servers
  - **Server Interfaces** – Generate a single graph for the network interfaces of the selected servers
- **Stacking**
  - **Servers** – Shows the summed up, stacked values for the selected servers

## Console Events

This sub-tab is visible only when opening the Console tab. It lists events generated by Console. Events are described in the “Event Reporting” chapter on page 40.

## Server Events

Lists events generated by the selected server(s). The events are described in the “Event Reporting” chapter on page 40.

## Server Commands

Console administrators can execute shell commands on the selected server(s) and see the output in the browser. The commands are executed by the WANsupervisor service with the “andrisoft” user’s privileges. To prevent the execution of CLI commands via Console, start the WANsupervisor service with the “-n” option.



## Appendix 1 – IPv4 Subnet CIDR Notation

Wansight uses extensively IP addresses and IP classes with the CIDR notation. To view details about any IPv4 subnet click Help → Subnet Calculator.

| CIDR MASK | 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 older IOS Devices

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
```

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. 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. 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
```

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. 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 IOS XE

Traditional NetFlow is being replaced with flexible NetFlow on newer IOS versions.

```
conf t
flow exporter WGFlowSensor
destination <ip_address>
source gi0/0/1
transport udp 9991
export-protocol netflow-v5
flow monitor WGFlowSensor
record netflow ipv4 original-input
exporter WGFlowSensor
cache timeout active 120 #in seconds
exit
int gi0/0/2
ip flow monitor WGFlowSensor input
exit
exit
wr mem
```

## Configuring NDE on IOS XR

A sample configuration for IOS XR:

```
flow exporter-map wanguard
version v9
options interface-table timeout 300
options vrf-table timeout 300
options sampler-table timeout 300
!
transport udp <port>
```

```
source Loopback8648
destination <ip_address>
!
flow monitor-map IPV4-FMM
record ipv4
exporter wanguard
cache entries 16384
cache timeout active 60
cache timeout inactive 30
!
flow monitor-map IPV6-FMM
record ipv6
exporter wanguard
cache entries 16384
cache timeout active 60
cache timeout inactive 30
!
sampler-map 1-of-128
random 1 out-of 128
!

interface TenGigE0/0/2/1
description Upstream Interface
...
flow ipv4 monitor IPV4-FMM sampler 1-of-128 ingress
flow ipv4 monitor IPV4-FMM sampler 1-of-128 egress
flow ipv6 monitor IPV6-FMM sampler 1-of-128 ingress
flow ipv6 monitor IPV6-FMM sampler 1-of-128 egress
!
```

## Configuring NDE on IOS-XE Everest

```
flow record v4r2
match ipv4 tos
match ipv4 protocol
match ipv4 source address
match ipv4 destination address
match transport source-port
match transport destination-port
match interface input
match flow sampler
collect routing next-hop address ipv4
collect ipv4 source mask
collect ipv4 destination mask
collect transport tcp flags
collect interface output
collect counter bytes
collect counter packets
collect timestamp sys-uptime first
collect timestamp sys-uptime last
collect routing source as 4-octet
collect routing destination as 4-octet
!

flow exporter exporter1
description "exported to external"
destination x.x.x.x
source Loopback0
transport udp 2000
```

```

flow monitor v4r2
exporter external1
cache timeout active 120
record v4r2

sampler external1
mode random 1 out-of 1000

interface GigabitEthernet0/0/0
ip flow monitor v4r2 sampler external1 input

```

## 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;
            }
        }
    }
}

```

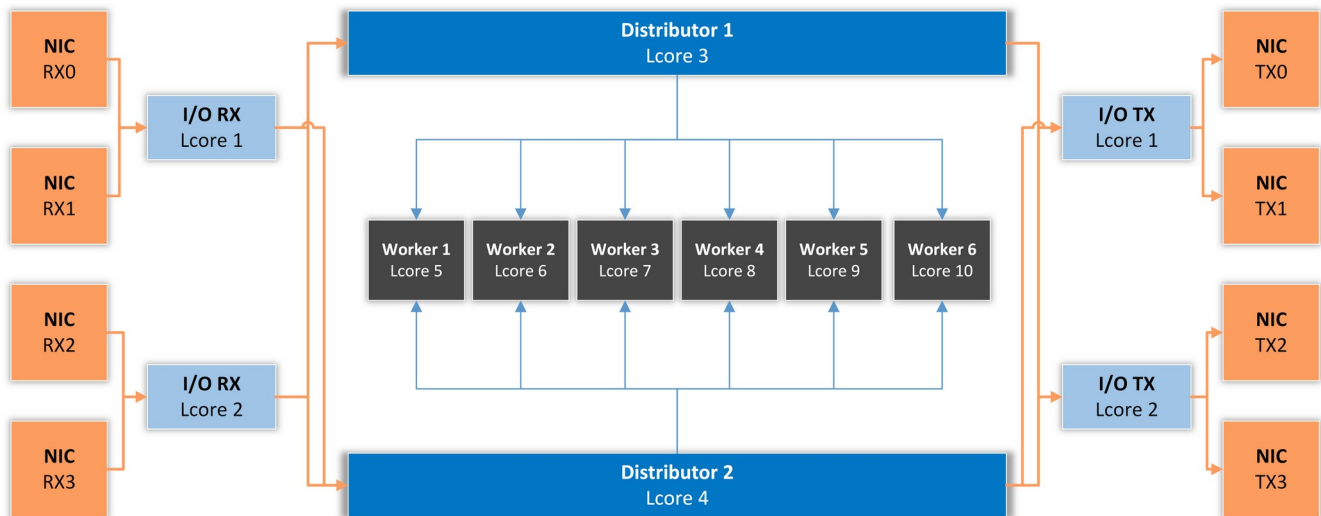
## Appendix 3 – DPDK Configuration

Wansight 8.0 supports DPDK version 19.11.2 on Intel microarchitectures starting with Sandy Bridge (Ivy Bridge, Haswell, Broadwell, Skylake, etc.). The code is currently optimized for Broadwell. NICs that need special drivers for DPDK (e.g. Mellanox) might not be supported.

To use DPDK 19.11.2, follow the installation guide from <http://www.dpdk.org> and allocate at least 8 hugepages each with 1 GB page size. There are several BIOS optimization settings required, as well as a number of kernel parameters that increase the performance of the server. It may be possible to purchase preconfigured appliances, already optimized for DPDK, from <https://www.andrisoft.com/hardware/anti-ddos-appliance>.

### Application Workflow

The architecture of the application is similar to the one presented in the following diagram which illustrates a specific case of two I/O RX and two I/O TX lcores (logical CPU cores) off-loading the packet Input/Output overhead incurred by four NIC ports, with each I/O lcore handling RX/TX for two NIC ports. The RX lcores are dispatching the packets toward two Distributor cores which are distributing them to six Worker lcores.



**I/O RX Lcore** performs packet RX from the assigned NIC RX rings and then dispatches the received packets to one or more distributor lcores using RSS or a round-robin algorithm.

**Distributor Lcore** reads packets from one or more I/O RX lcores, extracts packet metadata, performs the Dataplane firewall's functionality, and dispatches packet metadata to one or more Worker lcores.

**Worker Lcore** performs the most heavy weight and CPU-intensive tasks such as traffic analysis and attack detection.

**I/O TX Lcore** performs packet TX for a predefined set of NIC ports. The packets are forwarded in batches of minimum 4, so the latency will be very high (>50 ms!) if the application forwards just a few packets per second. On

thousands of packets/s the latency falls well under 1 millisecond.

The application needs to use one **Master Lcore** to aggregate data from the workers.

## DPDK Capture Engine Options

- **EAL Options** – See the DPDK Getting Started Guide for more information on this mandatory parameter
- **RX Parameters** – The syntax is “(PORT,QUEUE,LCORE)..” and represents a list of NIC RX ports and queues handled by the I/O RX lcores. This parameter also implicitly defines the list of I/O RX lcores. This is a mandatory parameter
- **Distributor Mode** – Specify the algorithm used to dispatch packets from the RX to the Distributor lcores:
  - *Round-robin* – The load is shared equally between the Distributor lcores. This is the best option when the packets are not forwarded
  - *RSS* – The packets with the same RSS value are always dispatched to the same Distributor lcore. This is the best option when packets are forwarded, mainly because it maintains the order of the packets
  - *Custom* – Select this option to be able to specify the Distributor lcore for each RX port. In this case, the RX Parameters syntax becomes “(PORT,QUEUE,LCORE,DISTRIBUTOR\_LCORE\_NO)..”
- **Distributor Lcores** – Enter the lcore of the Distributor thread, or a list of lcores separated by comma. This is a mandatory parameter
- **Worker Lcores** – The list of worker lcores. This is a mandatory parameter
- **Master LCORE** – Set an lcore to be used exclusively for thread management purposes. The recommended value is the hyper-thread core of CPU 0 because its performance is not important. This is a mandatory parameter
- **Forwarding Mode** – Specify the TX functionality:
  - *Disabled* – The packets are not forwarded, so the application behaves like a passive sniffer
  - *Transparent Bridge* – All Ethernet frames are forwarded without any intervention, so the application works like a transparent bridge. This is the fastest forwarding method.
  - *IP Forwarding* – The application performs several tasks for each packet. If it's an ARP packet querying for the MAC address of one of the interfaces defined below it responds to that query. On all other packets, it rewrites the source MAC address with the output interface MAC, and it rewrites the destination MAC with the MAC address defined below. The application is not performing RFC 1812 checks and is not decreasing the TTL value. This forwarding method is necessary when the server is deployed out-of-line with traffic redirected by BGP. For latency considerations see the previous option
- **TX Parameters** – The syntax is “(PORT,LCORE)..” and it defines a list of NIC TX ports handled by the I/O TX lcores. This parameter also implicitly defines the list of I/O TX lcores. This parameter is mandatory when the Forwarding Mode is not set to Disabled
- **Forwarding Table** – The syntax is “(PORT\_IN,PORT\_OUT)..” and it defines the output interface depending on the input interface
- **Interface IPs** – The syntax is “(PORT,IPV4)..” and it defines the IP of each port. This parameter is used when the Forwarding Mode is set to IP Forwarding but it does not ensure a true TCP/IP stack on the interface. The application will respond to ARP requests, but it's highly recommended to set the ARP table manually on the router because the application could respond to ARP requests with a high delay due to bulk processing

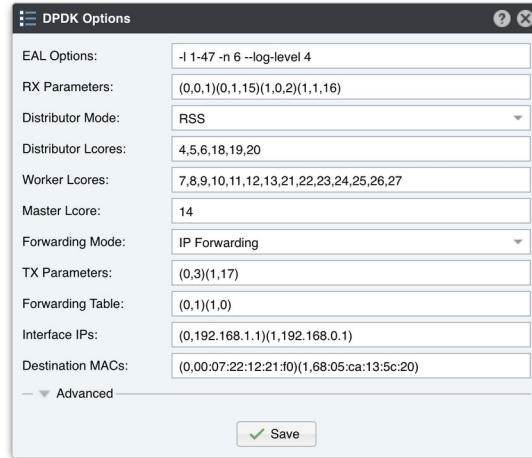


- **Destination MACs** – The syntax is “(PORT,MAC\_ADDRESS)..” and it defines the gateway MAC address for each port. This option is used when the Forwarding Mode is set to IP Forwarding
- **Maximum Frame Size** – If the network uses jumbo frames, enter the maximum frame size (usually 9000). Otherwise the default value is 1518, which captures normal Ethernet frames
- **IP Hash Table Size** – By default, the IPs are tracked using a hash table with 524288 elements for each worker lcore, IP version and traffic direction
- **Int. IP Mempool Size** – The default value is 70000 which means that each worker lcore pre-allocates RAM space to hold traffic information for up to 70000 IPs. The mempool is refreshed every 1 to 5 seconds, so to reach this limit all hosts must send or receive traffic during this period. The RAM space required per IP is listed in Sensor Graphs by selecting the Data Unit “IP Structure RAM”
- **Ext. IP Mempool Size** – This mempool is used for recording traffic information for external IP addresses. The default value is 120000 per worker lcore
- **Ring Sizes** – The accepted format is “A, B, C, D”:
  - A = The size (in number of buffer descriptors) of each of the NIC RX rings read by the I/O RX lcores
  - B = The size (in number of elements) of each of the software rings used by the I/O RX lcores to send packets to worker lcores
  - C = The size (in number of elements) of each of the software rings used by the worker lcores to send packets to I/O TX lcores
  - D = The size (in number of buffer descriptors) of each of the NIC TX rings written by I/O TX lcoresThe default values are “1024, 1024, 1024, 1024” which are optimal for the Intel ixgbe driver. Other network controllers and/or drivers might use different values
- **Burst Sizes** – The accepted format is “(A, B), (C, D), (E, F)”:
  - A = The I/O RX lcore read burst size from NIC RX
  - B = The I/O RX lcore write burst size to the output software rings
  - C = The worker lcore read burst size from the input software rings
  - D = The worker lcore write burst size to the output software rings
  - E = The I/O TX lcore read burst size from the input software rings
  - F = The I/O TX lcore write burst size to the NIC TX

The default values are “(144,144),(144,144),(144,144)” when Forwarding Mode is disabled, and “(8,8),(8,8),(8,8)” when Forwarding Mode is enabled. A burst size of 8 effectively means that the software will process at least 8 packets in parallel. So, on a traffic of 1 packet/s you will see a significant delay. It is not possible to use values less than 4 in DPDK 18.11 or 8 in DPDK 19.11

## DPDK Configuration Example

Execute the script `usertools/cpu_layout.py` from the your `dpdk` directory to see the CPU layout of your server. The following configuration assumes this CPU layout of a 14-core Xeon processor: Core 0 [0, 14], Core 1 [1, 15], Core 2 [2, 16], Core 3 [3, 17], Core 4 [4, 18], Core 5 [5, 19], Core 6 [6, 20], Core 8 [7, 21], Core 9 [8, 22], Core 10 [9, 23], Core 11 [10, 24], Core 12 [11, 25], Core 13 [12, 26], Core 14 [13, 27].



EAL Options contains the parameter “-l 1-27” which configures DPDK to use the lcores 1 to 27 (28 lcores = 14-core CPU with Hyper-threading enabled). The parameter “-n 4” configures DPDK to use 4 memory channels which is the maximum of what the reference Intel Xeon CPU (14-core Broadwell) supports.

The RX parameters configure the application to listen to the first two DPDK-enabled interfaces (0 and 1), on two NIC queues (0 and 1), and to use two CPU cores for this task (15 and 16 are the hyper-threads of cores 1 and 2).

The Distributor Mode setting ensures that the packets will be forwarded in the same order.

Three CPU cores are used for the Dataplane firewall and to distribute packets to the workers: 4, 5 and 6 (18, 19 and 20 are hyper-threads).

Seven CPU cores are used for packet analysis and attack detection: 7 to 13 (21 to 27 are hyper-threads).

The Master lcore is the hyper-thread of CPU core 0 which is used by the OS.

The TX parameters configure the application to use a single CPU core for TX. Lcore 3 sends packets over port 0, while the lcore 17 (hyper-thread of CPU core 3) sends packets over port 1.

The Forwarding Table value specifies that incoming packets on port 0 should be sent to port 1, and vice versa.

The next two parameters set the IPs and the destination MACs for both ports.