



WANSIGHT 5.4

Console + Sniffing Sensor + Flow Sensor

User Guide

Copyright ©2014 Andrisoft SRL
All rights reserved.
Revision 2.40

Copyright & trademark notices

This edition applies to version 5.x of the licensed program WANSIGHT and to all subsequent releases and modifications until otherwise indicated in new editions.

Notices

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

Copyright Acknowledgment

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

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

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

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

ANDRISOFT S.R.L.

Sales: sales@andrisoft.com

Technical Support: support@andrisoft.com

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

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

Table of Contents

1. IP Traffic Monitoring and Accounting with WANSIGHT.....	4
WANSIGHT Key Features & Benefits.....	4
WANSIGHT Components.....	4
2. A first look at the WANSIGHT Console.....	5
Side Region – used for navigation throughout the Console.....	5
Central Region – home of tabbed reports and dashboards.....	5
South Region – provides a quick look at the latest events, live statistics and graphs.....	5
Upper-right Menus – Help menu and User menu.....	5
3. Reports » Tools.....	6
Flow Collectors.....	6
Flows List.....	6
Flows Tops.....	7
Packet Analyzers.....	8
Active Captures.....	8
Captures Archive.....	10
4. Reports » Components.....	11
Overview.....	11
Console.....	11
Servers.....	12
Active Sniffing Sensors.....	12
Active Flow Sensors.....	13
Sensors.....	14
Sensor Dashboard.....	14
Sensor Tops.....	14
Sensor Graphs.....	15
Flows List.....	17
Flows Top.....	17
AS Graphs.....	17
Country Graphs.....	18
Sensor Events.....	19
5. Reports » Dashboards.....	20
6. Reports » IP Addresses & Groups.....	21
IP Dashboard.....	21
IP Graphs.....	21
IP Accounting.....	22
Flows List.....	23
Flows Tops.....	23
7. Reports » Servers.....	24
Server or Console Dashboard.....	24
Server Graphs.....	24
Server or Console Events.....	25
Server Commands.....	25
8. Installation Guide.....	26
System Requirements.....	26
Sniffing Sensor – Minimum Hardware Requirements.....	26
Flow Sensor – Minimum Hardware Requirements.....	27

Console Hardware Requirements.....	27
Software Installation & Download.....	28
Opening the Console for the First Time.....	28
Licensing Procedure.....	28
Quick Configuration Steps.....	28
9. Storage & Graphs Configuration.....	29
10.IP Zone Configuration.....	30
11.Choosing a Method of Traffic Monitoring.....	31
Comparison between Packet-Based and Flow-Based Monitoring.....	31
12.Sniffing Sensor Configuration.....	33
Troubleshooting the Sniffing Sensor.....	35
13.Flow Sensor Configuration.....	36
Troubleshooting the Flow Sensor.....	39
14.Scheduled Reports.....	40
15.Events Reporting.....	41
16.Users Management.....	42
17.Appendix 1 – Network Basics You Should Be Aware Of.....	43
IPv4 Subnet CIDR Notation.....	45
18.Appendix 2 – Configuring NetFlow Data Export.....	46
Configuring NDE on an IOS Device.....	46
Configuring NDE on a CatOS Device.....	47
Configuring NDE on a Native IOS Device.....	47
Configuring NDE on a 4000 Series Switch.....	48
Configuring NDE on a Juniper Router (non-MX).....	48

IP Traffic Monitoring and Accounting with WANSIGHT

Businesses all over the world rely on Andrisoft WANSIGHT when it comes to monitoring their network traffic. WANSIGHT includes all features of WANGUARD that don't relate to traffic anomalies.

WANSIGHT Key Features & Benefits

- **FULL NETWORK VISIBILITY** – Supports the latest IP traffic monitoring technologies: packet sniffing at 10 Gbps; NetFlow v5, v7 and v9; sFlow, IPFIX, NetStream, jFlow, cflowd.
- **ADVANCED WEB CONSOLE** – Consolidated management and reporting through a single, interactive and configurable HTML5 web portal with customizable dashboards, user roles, remote authentication, etc.
- **PACKET SNIFFER** – Includes a distributed packet sniffer that can save packet dumps from different parts of the network. View packet details in a Wireshark-like web interface.
- **FLOW COLLECTOR** – Provides a fully featured NetFlow, sFlow, and IPFIX Analyzer and Collector that saves flows in a compressed format for long term storage. Flows can easily be searched, filtered, sorted and exported.
- **COMPLEX ANALYTICS** – Generates the most complex reports with aggregated data for hosts, departments, interfaces, applications, ports, protocols, countries, autonomous systems, and more.
- **REAL-TIME REPORTING** – Bandwidth graphs are animated and have a short-term accuracy of just 5 sec.
- **HISTORICAL REPORTING** – You can view reports from the last 5 seconds to the last 10 years by selecting any custom time period. Bandwidth histograms contain 95th-percentile values for burstable billing.
- **SCHEDULED REPORTING** – Any report can be generated automatically and emailed to interested parties at preconfigured intervals of time.
- **FAST & SCALABLE** – The software was designed to run on commodity hardware. Its components can be distributed on clustered servers.

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

WANSIGHT Components

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

WANSIGHT relies on the **Sniffing Sensor** or **Flow Sensor** to provide in-depth traffic analysis, traffic accounting and bandwidth monitoring. The collected information enables you to generate complex traffic reports, graphs and tops; instantly pin down the cause of network incidents; understand patterns in application performance and make the right capacity-planning decisions.

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

A first look at the WANSIGHT Console

If you are an administrator seeking instruction on how to configure WANSIGHT, see the “Installation Guide” chapter on page 26.

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

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

Side Region – used for navigation throughout the Console

The Side Region is located on the east or west (default) edge of the window, according to the user's preference. If it is not visible, then has been either collapsed or hidden by an administrator. Clicking the edge of regions expands or collapses them.

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

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

Central Region – home of tabbed reports and dashboards

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

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

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

Upper-right Menus – Help menu and User menu

The Help menu contains the User Manual, a few select tools and the About window. Dependent on context, the User Manual will open at the chapter describing the last-opened window or tab. The Contextual Help works only with Adobe PDF Reader.

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

Reports » Tools

The **Reports » Tools** panel contains links to the **Flow Collectors** tab and to the **Packet Analyzers** tab.

Flow Collectors

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

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

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

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

Flows List

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

- **Sensors**

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

- **Time Frame**

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

- **Flows Filtering Expression**

Here you can enter a filtering expression for flows. Click the light bulb icon on the right to open a window that will show you the correct syntax. Frequently-used flows filters can be saved there and reused at a later time.

- **Output**

You can select several output formats, or you can type your own format. Click the light bulb icon on the left to open a window that will show you the correct syntax.

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

- **Export**

If the output is not very large, it can be emailed or printed.

If you need to list huge amounts of flow data, doing so solely from within the web browser may not be the best idea. In this case, select the “Dump” option to view the CLI command used to list the flows. You can execute the command locally, forward the output to a file, etc.

- **Aggregation**

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

- **Limit Flows**

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

- **Sorting**

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

Flows Tops

You can generate tops from the saved flow data by entering the fields below:

- **Sensors**

Select the Flow Sensor interfaces that capture the traffic you are interested in. Multiple selections can be made. Administrators can filter which Sensors are available to users.

- **Time Frame**

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

- **Flows Filtering Expression**

Here you can enter a filtering expression for flows. Click the light bulb icon on the right to open a window that will show you the correct syntax. Frequently-used flows filters can be saved there and reused at a later time.

- **Output**

You can select several output formats, or you can type your own format. Click the light bulb icon on the left to open a window that will show you the correct syntax.

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

- **Export**

If the output is not very large, it can be emailed or printed.

If you need to list huge amounts of flow data, doing so solely from within the web browser may not be the best idea. In this case, select the “Dump” option to view the CLI command used to list the flows. You can execute the command locally, forward the output to a file, etc.

- **Top Type**

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

- **Aggregation**

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

- **Limit**

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

- **Top**

Limit the statistics to the first top N.

Packet Analyzers

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

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

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

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

Active Captures

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

- **Description**

A short description to help you identify the traffic capture.

- **Sniffing Sensors**

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

- **BPF Expression**

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

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

- **Max Running Time**

The maximum running time.

- **Stop Capture On**

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

- **Max File Size (MB)**

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

- **Max Packets**

The capture stops after receiving <number> packets.

- **Max File Number**

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

- **Time Rotation (s)**

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

- **Sampling Type & Value**

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

- **Filename Prefix**

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

- **Snapshot (bytes/pkt)**

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

- **Comments**

This field may contain comments about the traffic capture.

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

- **Description [BPF]**

The traffic capture's description and BPF expression.

- **Sampling**

The type of sampling being used.

- **From**

The date when the Sniffing Sensor started capturing packets.

- **Until**

The time or the conditions that will cause the Sniffing Sensor to stop capturing the traffic.

- **Status**

Indicates the status of the capture. It is green if the capturing thread is active.

- **Interface**

The Sniffing Sensor or the Filter that captures packets.

- **Files / Size**

The number of dump files generated and the size of the latest dump file.

- **Packets**

The number of packets captured.

- **Actions**

Click the first icon to view the latest dump file in a Wireshark-like web interface. Click the second icon to download the latest dump file to your computer. Click the third icon to stop the capture.

Captures Archive

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

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

Reports » Components

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

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

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

Sensor tabs provide data specific to the selected Sensor.

Overview

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

Console

The Console System table has the following format:

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

Servers

The Servers table has the following format:

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

Active Sniffing Sensors

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

Status	If the active Sniffing Sensor is functioning properly then a green check mark is displayed. If the Console cannot manage or reach the Sniffing Sensor, then a red “X” is displayed. In this case, make sure that the Sniffing Sensor is configured correctly, and the WANsupervisor daemon is running, and look for errors in the events log (see page 41).
Sensor Name	Displays the name of the Sniffing Sensor and a colored box with the graph color as defined in configuration. Click it to open a new tab with data specific to the Sensor. Administrators and Operators can right-click it to open the Sensor's configuration.
Pkts/s (In / Out)	The inbound and outbound packets/second throughput after validation.

Inbound Bits/s	The inbound bits/second throughput after validation and the inbound usage percent.
Outbound Bits/s	The outbound bits/second throughput after validation and the outbound usage percent.
Received Pkts/s	The rate of sniffed packets before validation.
IPs (Int / Ext)	The number of IP addresses that sent or received traffic. The Int(ernal)/Ext(ernal) IPs are the IPs inside/outside the IP Zone. The top generator value from the Sensor configuration enables or disables the monitoring of external IPs.
Dropped	The rate of packets dropped in the capturing process. A high number indicates a sniffing performance problem.
CPU%	The percentage of CPUs used by the Sniffing Sensor process.
RAM	The amount of memory used by the Sniffing Sensor process.
Start Time	The date when the Sniffing Sensor started.
Server	The server that runs the Sniffing Sensor. Click to open a new tab with data specific to the server. Administrators and operators can right-click to open the server configuration.

Active Flow Sensors

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

Status	If the active Flow Sensor is functioning properly, then a green check mark is displayed. If the Console cannot manage or reach the Flow Sensor, then a red "X" is displayed. In this case, make sure that Flow Sensor is configured correctly and the WANsupervisor daemon is running, and look for errors in the events log (see page 41).
Sensor Name	Displays the name of the Flow Sensor. Click to open a new tab with data specific to the Sensor. Administrators and operators can right-click to open the sensor configuration.
Interface	The interface name and a colored box with the configured graph color. If the interface names are missing more than 2 minutes after the Sensor started, please check that the flow exporter's clock is synchronized with the server.
Pkts/s (In / Out)	The inbound and outbound packets/second throughput after validation.
Inbound Bits/s	The inbound bits/second throughput after validation and inbound usage percent.
Outbound Bits/s	The outbound bits/second throughput after validation and outbound usage percent.
IPs (Int / Ext)	The number of IP addresses that sent or received traffic. The Int(ernal)/Ext(ernal) IPs are the IPs inside/outside the IP Zone. The top generator value from the Sensor configuration enables or disables the monitoring of external IPs.
Flows/s	The rate of flows per second received by the Flow Sensor.

Flows Delay	Because traffic data must be aggregated first, flow devices export flows with a configured delay. Some devices export flows much later than the configured delay and this field contains the maximum flows delay detected by the Flow Sensor. The Flow Sensor cannot run with delays over 5 minutes. To minimize the RAM usage and optimize the performance of the Flow Sensor process, the flows must be exported as soon as possible.
Dropped	The number of unaccounted flows. A high number indicates a performance problem with the Sensor or a network connectivity issue with the flow exporter.
CPU%	The percentage of CPUs used by the Flow Sensor process.
RAM	The amount of memory used by the Flow Sensor process.
Start Time	The date when the Flow Sensor started.
Server	The server that runs the Flow Sensor. Click it to open a new tab with data specific to the server. Administrators and operators can right-click to open the server configuration.

Sensors

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

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

Sensor Dashboard

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

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

The operation of dashboards is documented in the Reports » Dashboards chapter (page 20).

Sensor Tops

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

- **Top Type**

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

The number of top items and decoders can be modified in the Storage & Graphs Configuration (see page 29).

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

Sensor Graphs

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

- **Data Units**
Select one or more parameters:
 - *Most Used* – Shows the most used data units, each in a different graph.
 - *Packets* – The packets/second rate.
 - *Bits* – The bits/second throughput.
 - *Applications* – Sensors can collect application-specific distribution data for: HTTP, HTTPS, SMTP, POP3, IMAP, SNMP, FTP, SSH, TELNET, SQL, NETBIOS, MS-DS, MS-RDP, DNS, ICMP and OTHERS.
 - *Bytes* – The bytes/second throughput.

- *Internal/External IPs* – The number of IP addresses that sent or received traffic. The Internal/External IPs are the IPs inside/outside the IP Zone. The top generator value from the Sensor configuration enables or disables monitoring of External IPs. A spike in the Internal IPs graph usually means that an IP class scan was performed against your IP block. A spike in the External IPs graph usually means that you received a spoofed attack.
 - *Received Frames* – For Sniffing Sensors, represents the rate of packets received before validation. For Flow Sensors, represents the rate of flows received before validation.
 - *Dropped Frames* – For Sniffing Sensors, represents the number of packets dropped in the capturing process. A high number indicates a sniffing performance problem. For Flow Sensors, represents the number of unaccounted flows. A high number indicates a performance problem with the Flow Sensor or a network connectivity issue with the flow exporter.
 - *Unknown Frames* – For Sniffing Sensors, represents the rate of invalidated packets. For Flow Sensors, represents the rate of invalidated flows.
 - *Unknown Sources* – The number of source IP addresses that did not pass validation.
 - *Unknown Destinations* – The number of destination IP addresses that did not pass validation.
 - *Avg Packet Size* – The average packet size in bits/packet.
 - *CPU%* – The percentage of CPUs used by the Sensor process.
 - *RAM* – The amount of memory used by the Sensor process.
 - *Load* – The load of the operating system for the last 5 minutes.
 - *IP Graphs* – The number of updated IP graphs files.
 - *IP Accounting* – The number of updated IP accounting records.
 - *HW Graphs* – The number of updated traffic profiling files.
 - *IP Graphs Time* – The number of seconds needed to update the IP graphs files.
 - *HW Graphs Time* – The number of seconds needed to update the traffic profiling files.
 - *Processing Time* – The number of seconds needed to perform traffic analysis functions.
 - *IP Structures* – The number of internal IP structures.
 - *IP Structure RAM* – The number of RAM bytes used by each IP structure.
- **Graphs Size**
You can select predefined sizes or you can enter your own size in the “<X pixels> x <Y pixels>” format.
 - **Graphs Title**
Graphs can have an automatically-generated title for the “Default” option or no title for the “None” option, or you can enter your own text to be rendered as a title.
 - **Graph Legend**
Select how detailed the graph's legend should be.
 - **Consolidation**
If you are interested in spikes, select the *MAXIMUM* aggregation type. If you are interested in average values, select the *AVERAGE* aggregation type. If you are interested in low values, select the *MINIMUM*

aggregation type.

- **Graph Options**

- *Stack Sensors* – If unchecked, each selected Sensor generates different a graph. If checked, all selected Sensors generate a single graph that contains the combined data.
- *Show Totals* – If multiple Sensors are used, render the data units stacked inside the graph.

Flows List

This is available only for Flow Sensors.

You can list and filter the flow data saved by the Flow Sensor. The options are listed on page 6, in the Flow Collectors chapter.

Flows Top

This is available only for Flow Sensors.

You can generate tops and process and filter the flow data saved by the Flow Sensor. The options are documented on page 6, in the Flow Collectors chapter.

AS Graphs

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

The parameters are:

- **Sensors**

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

- **Time Frame**

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

- **AS Numbers**

Click the lightbulb icon on the right to open a window containing the correct syntax. Frequently-used AS numbers can be saved there, for use at a later time.

If you don't know what AS number(s) a particular ISP has, go to Help » IP & AS Information » AS Numbers List. There you can apply different filters by clicking a table header's down arrow.

- **Export**

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

- **Refresh**

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

- **Graphs Size**

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

- **Graphs Title**

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

- **Graph Options**

- *Stack Sensors* – If unchecked, a different AS graph is generated for every Sensor. Otherwise, a single AS graph that contains summed traffic data is generated for all selected Sensors.
- *Stack ASNs* – If you entered multiple AS Numbers, then you can sum all of them in a single AS graph. This is useful with ISPs and AS owners that have more than 1 allocated AS number.

Country Graphs

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

The parameters are:

- **Sensors**

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

- **Time Frame**

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

- **Countries**

Select the country or countries from the drop-down list, or click the light bulb icon on the right to open a window with saved selections.

- **Export**

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

- **Refresh**

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

- **Graphs Size**

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

- **Graphs Title**

Graphs can have an automatically-generated title for the “Default” option or no title for the “None”

option, or you can enter your own text to be rendered as a title.

- **Graph Options**

- *Stack Sensors* – If unchecked, a different graph is generated for every Sensor. Otherwise, a single graph that contains the summed traffic data is generated for all selected Sensors.
- *Stack Countries* – If you selected multiple countries, then you can sum all of them in a single graph, or you can see a separate graph for each country.

Sensor Events

The list of events generated by the selected Sensor(s) for the selected time frame. Events are explained in the Events Reporting chapter (see page 41).

Reports » Dashboards

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

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

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

Add some **widgets** to your dashboard. To sort them, click the title bar and move them around. To collapse a widget, click the first icon on the widget title bar. To edit a widget, click the second icon from its title bar. To delete a widget, click the third icon from its title bar.

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

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

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

Reports » IP Addresses & Groups

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

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

The **Reports » IP Groups** panel lists all IP groups defined in IP Zones. You can search or filter them by entering a sub-string contained in the name of the IP group you are interested in. You can use IP groups to generate reports for customers with multiple IP blocks allocated to them. To do that, use the same IP group name for all the customer's IP blocks.

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

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

- **Sensor**

Select the Sensors you are interested in or select “All” to select all Sensors. Multiple selections can be made. Administrators can filter which Sensors are available to users.

- **Time Frame**

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

IP Dashboard

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

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

The operation of Dashboards is documented in the Reports » Dashboards chapter on page 20.

IP Graphs

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

- **Decoders & Data Unit**

Select the decoders that analyze the traffic you are interested in. Data units available: *Packets*, *Bits* and *Bytes*.

- **Graphs Size**

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

- **Graph Title**

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

- **Graph Legend**

Select how detailed the graph's legend should be.

- **Consolidation**

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

- **Graphs Stacking**

- *Stack Sensors* – If unchecked, each selected Sensor generates a different graph. If checked, all selected Sensors generate a single combined graph.
- *Stack Decoders* – If unchecked, each selected decoder generates a different graph.
- *Stack IPs* – Uncheck this option if you want a different traffic graph displayed for every IP address contained in the IP class or IP group. Use carefully, because when this option is used with a /24 CIDR, 256 traffic graphs will be displayed, one for each IP address in the “C” class.
- *Stack Conflicts* – If decoders can be included one within the other (e.g. TOTAL contains TCP that contains HTTP and HTTPS), the graph will display stacked decoders to show the most specific ones. This generates both accurate and intuitive traffic graphs. In the example, TOTAL will be displayed as TOTAL OTHER and TCP as TCP OTHER. However, when you select TCP, HTTP and TCP+SYN as decoders, then the TCP+SYN decoder can be included in both TCP and HTTP, thus generating a decoder conflict. Check this to stop detection of conflicts between decoders, but keep in mind that graphs may be less accurate.
- *Stack Recursively* – When checked, subnet graphs can be created from the IPs graphs that are contained in the subnet.

The number of decoders, Data Units and aggregation types can be modified in the Storage & Graphs Configuration, see page 29.

IP Accounting

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

- **Decoders & Data Unit**

Select the decoders that analyze the traffic you are interested in. Data units available: *Packets*, *Bits* and *Bytes*.

- **Report Type**

Select the interval you want to use to aggregate the accounting data: *Daily*, *Weekly*, *Monthly*, or *Yearly*. The minimum time interval for accounting reports is 1 day, so if you select a shorter time interval, you will still see the data collected for the whole day.

- **Sum IPs**

Uncheck this option if you want a different traffic accounting report displayed for every IP address contained in the IP class or IP group. Use carefully, because when this option is used for a large IP block like a /24 CIDR, 256 traffic accounting reports will be displayed, one for each IP address in the “C” class.

- **Sum Sensors**

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

The number of decoders can be modified in the Storage & Graphs Configuration (see page 29).

Flows List

You can list and filter the saved flow data for the IP class, host or group. The options are documented on page 6 in the Flow Collectors chapter.

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

Flows Tops

You can generate tops from the saved flow data for the IP class, host or IP Group. The options are documented on page 6 in the Flow Collectors chapter.

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

Reports » Servers

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

- **Servers**

Select the servers you are interested in, or select “All” to select all servers. Multiple selections can be made. Administrators can filter what servers are available to users.

- **Time Frame**

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

Server or Console Dashboard

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

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

The operation of dashboards is documented in the Reports » Dashboards chapter on page 20.

Server Graphs

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

- **Data Units**

Select one or more parameters:

- *Most Used* – Shows most-used data units, each in a different graph.
- *System Load* – The load reported by the operating system.
- *Free RAM* – The available RAM, not counting the swap space.
- *Database/Graphs/SSD/Flow Collector/Packet Dumps Disk - Free space* – How much disk space is available for each path.
- *Uptime* – The uptime of the operating system.
- *CPU% system/userspace/niced/idle* – The percentages of CPUs used by the system, used by the userspace processes, used by the processes with increased (nice) priorities, and idle.
- *Number of processes* – The total number of processes that are running on the system.
- *Hardware/Software CPU Interrupts* – The number of CPU interrupts made by hardware events, drivers, etc.
- *Context Switches* – Indicates how much time the system spends on multi-tasking.

- *Running Components* – The number of Sensor processes.
 - *Clock Delta* – The difference of time between the server and the Console, in seconds. Must be 0.
 - *Database/Graphs/SSD/Packet Dumps/Flow Collector Disk - Total* – How much disk space is allocated for the partitions that hold the path.
 - *Database/Graphs/SSD/Packet Dumps/Flow Collector Disk - Free Inodes* – The number of free inodes held by the partitions that hold the paths.
 - *Database/Graphs/SSD/Packet Dumps/Flow Collector Disk - Ops/s* – The number of reads and writes for the partitions that hold the paths.
 - *Database/Graphs/SSD/Packet Dumps/Flow Collector Disk - Bytes/s* – The number of bytes/s for the partitions that hold the paths.
 - *Server Interfaces - Packets/Bits/Errors/Dropped* – Various statistics collected for the interfaces defined in the Server Configuration window.
- **Graphs Size**
You can select predefined sizes or you can enter your own size in the “<X pixels> x <Y pixels>” format.
 - **Graphs Title**
Graphs can have an automatically-generated title for the “Default” option or title for the “None” option, or you can enter your own text to be rendered as a title.
 - **Graph Legend**
Select how detailed the graph's legend should be.
 - **Consolidation**
If you are interested in spikes, select the *MAXIMUM* aggregation type. If you are interested in average values, select the *AVERAGE* aggregation type. If you are interested in low values, select the *MINIMUM* aggregation type.
 - **Graph Options**
 - *Stack Servers* – If unchecked, each selected server generates a different graph. If checked, all selected servers generate single combined graph.
 - *Show Totals* – If multiple servers are used, render the total of data units in a stacked graph.

Server or Console Events

The list of events generated by the selected server(s) or by the Console for the selected time frame. Events are explained in the Events Reporting chapter (see page 41).

Server Commands

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

Installation Guide

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

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

System Requirements

WANSIGHT 5.x has been tested with the following distributions: **Red Hat Enterprise Linux 5.x or 6.x** (commercial Linux distribution), **CentOS 5.x or 6.x** (free, Red Hat Enterprise Linux-based distribution), **OpenSUSE 12.x or 13.x** (free, Novel Enterprise Linux-based distribution), **Debian Linux 6.0 or 7.0** (free, community-supported distribution), **Ubuntu 12.x**. Other distributions may work but have not yet been tested.

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

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

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

Sniffing Sensor – Minimum Hardware Requirements

Sniffing Capacity	1 Gigabit Ethernet	10 Gigabit Ethernet
Architecture	x86 (32 or 64 bit)	x86 (64 bit)
CPU	2.0 GHz dual-core Xeon	2.8 GHz quad-core Xeon
RAM	1 GB	2 GB
Network Cards	Gigabit Ethernet with NAPI support Fast Ethernet for management	10 GbE card. Intel 82599 chipset recommended Fast Ethernet for management
Operating System	RHEL 5 / CentOS 5 or 6, Debian 6 or 7, Ubuntu Server 12, OpenSUSE 12 or 13	RHEL 5 / CentOS 5 or 6, Debian 6 or 7, Ubuntu Server 12, OpenSUSE 12 or 13

Disk Space	10 GB (including operating system)	10 GB (including operating system)
-------------------	------------------------------------	------------------------------------

Flow Sensor – Minimum Hardware Requirements

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

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

Console Hardware Requirements

Capacity	< 5 Managed Components*
Architecture	x86 (32 or 64 bit)
CPU	2.4 GHz dual-core Xeon
Memory	1 GB
Network Cards	Gigabit Ethernet
Operating System	RHEL / CentOS 5 or 6, Debian 6 or 7, Ubuntu Server 12, OpenSUSE 12 or 13
Software Packages	Apache 2.x+, php 5.2+, mysql 5.x, rrdtool 1.3+, ping, whois, traceroute, telnet, wireshark, tcpdump
Disk Space	10 GB (including operating system) + additional storage to store IP graphs data

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

To access the web interface provided by the Console, one of the following web browsers is required (others should also work but have not been tested): Google Chrome, Firefox 3.5 or later, Safari 3.0 or later, Internet Explorer 8 or later. Older versions of Internet Explorer contain a very slow javascript engine and won't work well. For the best Console experience, we highly recommend Chrome or Firefox and a 1280 x 1024 pixels or higher resolution display.

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

Software Installation & Download

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

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

Opening the Console for the First Time

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

To log into the Console, open `http://<hostname>/wansight`. If the page cannot be displayed, make sure the Apache web server is running and the firewall does not block incoming traffic on port 80. You can also access the Console securely by HTTPS if the Apache web server was configured with SSL/TLS support.

Licensing Procedure

If you have not yet licensed WANSIGHT you will be asked to do so. Upload the *trial.key* file we sent you by email by clicking the key icon. The license key contains encrypted information about the licensed capabilities of the software. You can switch between WANGUARD and WANSIGHT solely by changing the license key.

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

To understand how to navigate within the Console, please read the chapter beginning on page 5.

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

Quick Configuration Steps

- Estimate storage requirements, review decoders and graphs parameters – page 29
- Add your IP address ranges and important IPs to an IP Zone – page 30
- Add and configure a Sensor, then start it – page 31
- Watch for errors in events log. Receive error notifications by email – page 41
- Generate reports and send them periodically by email – page 40
- Create your own dashboards and add useful widgets – page 20
- Configure Console accounts for your staff or customers – page 42

Storage & Graphs Configuration

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

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

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

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

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

This method creates one file for every IP address directly in the defined Graphs Disk Path. If the RRDCache daemon is being used to speed up graphs I/O, add its socket path (unix:/var/rrdtool/rrdcached/rrdcached.sock on Redhat). The Andrisoft Knowledge Base contains an article on configuring RRDCacheD.

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

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

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

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

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

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

All IP graph parameters have a direct impact on the storage space required in the Console's file system. The *Disk space required for each IP graph file* value will be updated when you click the <<Update>> button.

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

IP Zone Configuration

IP Zones are hierarchical, tree-like structures in which you should include your IP address ranges. To add an IP Zone go to Configuration » Network » + » Add IP Zone. Sensors use IP Zones to learn about your network and to extract per-subnet settings. An IP Zone may be used by multiple Sensors.

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

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

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

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

To enter IP addresses or IP blocks, you must use the CIDR notation. To enter individual hosts in IP Zones you must use the /32 CIDR for IPv4, or /128 for IPv6. For more information about the CIDR notation, see Appendix 1 from page 43.

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

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

- **IP Group.** This combo box should contain a short description of the selected prefix. Setting the same IP group for more than one subnet will allow you to easily generate combined reports.
- **IP Graphs.** If set to “Yes,” the Console will collect graph data for every IP contained in the selected prefix.
- **IP Accounting.** If set to “Yes,” the Console will save daily accounting data for each IP contained in the selected prefix.

Enabling IP graphs and IP accounting for very large prefixes (e.g. 0.0.0.0/0) is probably going to generate useless data that can potentially overload the system.

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

Choosing a Method of Traffic Monitoring

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

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

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

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

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

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

Comparison between Packet-Based and Flow-Based Monitoring

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

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

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

The table below lists the main differences between Sensors:

Sensor Type	Sniffing Sensor	Flow Sensor
Capturing Technology	- Port Mirroring (SPAN, Roving Analysis Port) - Network TAP - In-line appliance	- NetFlow version 5, 7, 9 (jFlow, NetStream, cflowd) - sFlow version 4, 5 - IPFIX
Maximum Traffic Capacity per Sensor	10 GigE >150,000 endpoints*	multiples of 10 Gbps >100,000 endpoints*
IP Graphs Accuracy	≥ 5 seconds	≥ 20 seconds
Traffic Validation Options	IP classes, MAC addresses, VLANs, BPF	IP classes, Interfaces, AS Numbers, Ingress/Egress
Packet Analyzer	Yes	No
Flow Collector	No	Yes

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

Sniffing Sensor Configuration

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

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

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

The Sniffing Sensor Configuration window contains the following fields:

- **Sensor Name**

A short name to help you identify the Sniffing Sensor.

- **Devices Group**

Optional description used within the Console to group multiple components by location, role, etc.

- **Graph Color**

The color used in graphs for this Sensor. The default color is a random one, but you can change it by entering another HTML color code or by clicking the drop-down menu.

- **Sensor License**

The license used by the Sensor. WANGUARD provides all features; WANSIGHT does not provide traffic anomaly detection and reaction.

- **Sensor Server**

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

- **Sniffing Interface**

The network interface listened by the Sniffing Sensor. The Linux network interface naming convention is eth0 for the first Ethernet interface, eth1.900 for the second Ethernet interface with 802.1Q VLAN 900, etc.

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

- **Link Speed IN / OUT**

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

- **IP Zone**

The Sensor needs an IP Zone from which to learn about your network and to extract per-subnet settings.

IP Zones are documented in the IP Zones Setup chapter on page 30.

- **IP Validation**

This option can be used to distinguish the direction of the packets or to ignore certain IPs:

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

- **MAC Validation/Address**

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

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

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

- **BPF Expression**

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

- **Use PF_RING**

Enable if you have PF_RING installed on the server. PF_RING provides high-speed packet analysis by decreasing the CPU usage of the Sniffing Sensor.

- **Top Generator**

Allows generation of traffic tops:

- *Basic* – Enables tops for Internal IPs, IP protocols, versions and TCP/UDP ports. It is the recommended value because it adds a very small performance penalty to the Sniffing Sensor.
- *Extended* – Enables all tops from *Basic* as well as tops for External IPs (IPs not included in the IP Zone). It adds performance penalty of over 20%, especially during spoofed attacks.
- *Full* – Enables all tops from *Extended* as well as tops and graphs for autonomous systems. It adds a performance penalty of over 20%, especially during spoofed attacks.

- **Comments**

Comments about the Sensor configuration can be saved here. They are not visible elsewhere.

To start the Sniffing Sensor, click the gray square button next to the Sensor's name from Configuration » Components. Check that the Sniffing Sensor starts properly by watching the events log (details on page 41).

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

Troubleshooting the Sniffing Sensor

- ✓ Look for warnings or errors produced by the Sniffing Sensor in the events log (details on page 41).
- ✓ Check that you have correctly configured the Sniffing Sensor. Each configuration field is explained in detail in the previous paragraph.
- ✓ Check that the sniffing interface is up using the “ifconfig <ethX>” command.
- ✓ Check that you have correctly configured the switch/TAP to send packets to the server on the configured interface.
- ✓ You can verify whether the server is receiving the packets through the configured interface with a tool like tcpdump. The syntax is “tcpdump -i <interface_usually_eth0> -n -c 100”.
- ✓ When IP Validation is not disabled, make sure that the IP Zone contains all your subnets.

Flow Sensor Configuration

Many routers and switches can collect IP traffic statistics on monitored interfaces, and later export those statistics as flow records to the **Flow Sensor**. Because the flow protocol already performs pre-aggregation of traffic data, the flows of data sent to the Flow Sensor are much smaller than the monitored traffic. This makes the Flow Sensor a good option for monitoring remote or high-traffic networks.

For detailed instructions on how to enable NetFlow, sFlow or IPFIX on your network device, please consult the appropriate user guide from your vendor. Appendix 2 from page 46 contains an example on how to configure NetFlow on Cisco IOS, CatOS and Juniper.

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

The Flow Sensor Configuration window contains the following fields:

- **Sensor Name**

A short name to help you identify the Flow Sensor.

- **Devices Group**

Optional description used within the Console to group multiple components by location, role, etc.

- **Sensor Server**

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

- **Listener IP:Port**

The IP address of the network interface that receives flows and the destination port.

- **Repeater IP:Port**

Send all incoming flows to another host/collector by enabling the embedded packet repeater (optional).

- **Flow Collector**

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

- **Sensor License**

The license used by the Sensor. WANGUARD provides all features, WANSIGHT doesn't provide traffic anomaly detection and reaction.

- **Flow Protocol**

The type of flows exported towards the Sensor: NetFlow, IPFIX or sFlow. For IPFIX implementations that keep the start time of the flows (e.g. Juniper MX), select the second value.

- **Flow Exporter IP**

The IP address of the router, switch, probe etc. Usually the loopback0 address of the router. For sFlow

exporters, enter the IP that sends flows, not the Agent IP.

- **Sampling (1/N)**

Must contain the sampling rate configured on the flow exporter, or 1 when no sampling rate is configured. For NeFlow v9 and sFlow the value provided here is ignored because the sampling rate is automatically adjusted by the protocol.

- **Flow Timeout (s)**

The flow-active/inactive-timeout value from the IPFIX exporter that keeps the start time of the flows.

- **Time Settings**

The time offset between the time zone of the Flow Sensor's server and the flow exporter. Run NTP on both devices to keep their clocks synchronized. This is a critical requirement for the Flow Sensor.

- **SNMP Community**

The read-only SNMP community of the flow exporter allows the Console to gather interface information. If this field is left empty, you must enter the SNMP index, speed, etc. manually for each interface.

- **Monitored Interfaces**

The list of interfaces that should be monitored. Add only upstream interfaces if possible, to avoid producing duplicate flow entries. Settings per interface:

- *Interface Name* – A short description used to identify the monitored interface. Descriptions longer than 10 characters clutter some reports.
- *Graph Color* – The color used in graphs for this interface. The default color is a random one, but you can change it by entering another HTML color code or by clicking the drop-down menu.
- *SNMP Index* – The interfaces are identifiable in flows only when their SNMP indexes is known.
- *Traffic Direction* – The direction of the traffic entering the interface:
 - Select “Inbound” for upstream interfaces, e.g. peering interfaces.
 - Select “Outbound” for downstream interfaces, e.g. customer interfaces.
 - Select “Mixed” to establish the direction by IP/AS Validation.
 - Traffic entering the “Null” interface is discarded by the router and also by the Flow Sensor.
- *Link Speed In & Link Speed Out* – The capacity of the interface. If set, it is used to generate reports based on usage percentage.
- *Top Generator* – Allows generating traffic tops:
 - “Basic” – Enables tops for Internal IPs, IP protocols, versions and TCP/UDP ports. It adds a very small performance penalty to the Flow Sensor.
 - “Extended” – Enables all tops from *Basic* as well as tops and graphs for autonomous systems and countries, but increases the CPU usage of the Flow Sensor by a few percentage points. If the router doesn't export AS information (e.g. non-BGP router), the Sensor uses an internal GeoIP database to get ASNs. Live stats for autonomous systems and countries are not very accurate.
 - “Full” – Enables all tops from *Extended* as well as tops for external IPs (IPs not included in the IP

Zone), but increases the RAM usage several times over, especially during spoofed attacks. Live stats for autonomous systems and countries are very accurate.

- **IP Zone**

The Sensor needs an IP Zone from which to learn about your network and to extract per-subnet settings. For more information about IP Zones, please consult the IP Zones Setup chapter on page 30.

- **IP Validation**

This option can be used to distinguish the direction of the traffic or to ignore certain flows:

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

- **AS Validation**

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

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

AS validation has three options:

- *Off* – Disables AS validation.
- *On* – Only flows that have the source ASN and/or the destination ASN set to 0 are analyzed.
- *Strict* – Only flows that have either the source ASN or the destination ASN set to 0 are analyzed.

- **Graphs Accuracy**

Low values increase the accuracy of Sensor graphs, at the expense of RAM usage. Setting this to under 20 seconds is not recommended.

- **Comments**

Comments about the Sensor configuration can be saved here. These are not visible elsewhere.

To start the Flow Sensor, click the gray square button next to the Sensor's name from Configuration » Components. Check that The Flow Sensor starts properly by watching the events log (details on page 41).

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

Troubleshooting the Flow Sensor

- ✓ Look for warnings or errors produced by the Flow Sensor in the events log (details on page 41).
- ✓ Check that you have correctly configured the Flow Sensor. Each configuration field is explained in detail in the previous paragraph.
- ✓ You can verify that the server is receiving the flow packets on the configured Listener IP and Port with a tool like *tcpdump*. The syntax is “`tcpdump -i <interface_usually_eth0> -n -c 100 <flow_exporter_ip> and udp and <destination_port>`”.
- ✓ You can check if the local firewall, which allows the Flow Sensor to receive the flow packets, is enabled with the *iptables* command. The syntax is “`iptables -L -n -v`”.
- ✓ The clocks of both devices are synchronized with NTP. If the devices don't reside in the same time zone, adjust the time zone offset in the Flow Sensor configuration.
- ✓ The flow exporter's active/inactive flow timeout settings are set to less than 300 seconds. Flows sent with a delay of more than 300 seconds are automatically discarded and a warning is sent to the events log.
- ✓ Check that you have correctly configured the flow exporter to send flows to the server for each of the monitored interfaces. To see the interfaces that send flows, go to Reports » Tools » Flow Collectors » Flows Top, select the Flow Sensor, set Output to Debug, set Top Type to Any interface and generate the top for the last 10 minutes. The In/Out_If column shows the SNMP index of every interface that exports flows.
- ✓ If you see statistics for only one traffic direction (inbound or outbound), go to Reports » Tools » Flow Collectors » Flows List, and generate a listing for the last 10 minutes. If all your IPs are in a single column, check the flow exporter's configuration and feature list. Not all devices can export flows in both directions (e.g. Brocade XMR) or with the same SNMP interface index.
- ✓ If you defined interfaces with the Traffic Direction parameter set to “Mixed,” then make sure that the IP Zone you have selected for the Flow Sensor contains all your IP blocks.
- ✓ If the Flow Sensor does not show the correct statistics after upgrading the router's firmware, the SNMP index of interfaces has probably changed. In this case, enter the new SNMP index for each interface.

Scheduled Reports

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

If you want to receive periodic reports by email without having to log into the Console, go to Configuration » Schedulers » Add Report.

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

You can include more than one email address in the **Email To** field by separating addresses with a comma.

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

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

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

Events Reporting

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

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

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

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

As an administrator, you should keep events with high severities under surveillance!

Configure the Console to send important events periodically by email, Syslog or SNMP in Configuration » Schedulers » Events Reporting.

Users Management

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

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

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

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

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

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

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

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

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

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

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

Appendix 1 – Network Basics You Should Be Aware Of

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

IP Addresses

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

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

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

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

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

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

IP Classes

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

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

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

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

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

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

IPv4 Subnet CIDR Notation

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

Appendix 2 – Configuring NetFlow Data Export

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

Configuring NDE on an IOS Device

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

First enable Cisco Express Forwarding:

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

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

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

For example:

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

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

Enable the exports of these flows with the global commands:

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

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

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

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

The following commands break up flows into shorter segments: 1 minute for active traffic and 30 seconds for inactive traffic. The Flow Sensor drops flows older than 5 minutes!

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

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

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

Configuring NDE on a CatOS Device

In privileged mode on the Supervisor Engine enable NDE:

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

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

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

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

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

The following commands break up flows into shorter segments: ~1 minute for active flows and ~ 30 seconds for inactive flows. The Flow Sensor drops flows older than 5 minutes!

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

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

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

And enable NDE:

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

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

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

Configuring NDE on a Native IOS Device

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


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

The following commands break up flows into shorter segments: ~1 minute for active flows and ~ 30 seconds for inactive flows. The Flow Sensor drops flows older than 5 minutes!

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

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

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

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

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

Configuring NDE on a 4000 Series Switch

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

Configuring NDE on a Juniper Router (non-MX)

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

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

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