



CRANNOGSOFTWARE
making networks assets, not overheads

NetFlow Tracker User's Guide

Version 2.1.1

www.crannog-software.com

Contents

INTRODUCTION	5
What is NetFlow?	5
What is NetFlow Tracker?	5
Features and Benefits	5
INSTALLATION	6
Minimum System Requirements	6
Pre-installation Checks	6
Installation on Microsoft Windows™	7
Installation on Solaris and Linux	8
Post-installation Tasks	8
USING NETFLOW TRACKER	10
Real-time Data	10
Long-term Data	10
Executive Reports	10
Network Overview	10
Devices	10
Per-AS data	12
Working with Charts	12
Working with Pie Charts	14
Working with Tabular Reports	14
Report Templates	15
Creating Filtered Reports	17
LONG-TERM REPORTS	21
Devices and Interfaces	21

Per-device and Per-interface Long-term Reports	21
Filter Editor	21
EXECUTIVE REPORTS	22
REPORT URL FORMAT	23
General Form	23
Report Format Parameters	23
Time Range Parameters	26
Filter Parameters	32
Security Parameters	34
Management Portal Access Control Parameters	35
PERFORMANCE TUNING	38
Disk Speed	38
Query Size	38
Database Server Settings	38
CONFIGURATION GUIDE	39
Licensing	39
Listener Ports	39
SNMP Settings	39
Device Settings	40
Security Settings	42
Management Portal Settings	42
Report Settings	43
An Example Executive Report – Top Applications Today and Last Week	45
IP Application Names	47
DiffServ Names	47
Hostname Resolution Settings	47

AS Names	47
Subnet Names	47
Database Settings	48
Archiving	48
Memory Settings	49
Performance Counters	49
APPENDIX 1: CONFIGURING NETFLOW DATA EXPORT	51
Configuring NetFlow Export on an IOS device	51
Configuring NetFlow Input Filters	52
Configuring NDE on a CatOS device	52
Configuring NDE on a Native IOS device	53
Configuring NetFlow Export on a 4000 series switch	54
APPENDIX 2: CSV FILE FORMAT	55
Chart CSV format	55
Tabular report CSV format	55
APPENDIX 3: THIRD PARTY SOFTWARE COMPONENTS	56

Introduction

This document is the user manual for NetFlow Tracker, a software product designed to collect NetFlow information from Cisco equipment and present it in a meaningful way. This document does not provide any assistance with Cisco equipment itself. Please consult your Cisco documentation for any queries you have relating to the equipment itself. For more information on NetFlow from the Cisco website, go to <http://www.cisco.com/go/netflow>.

This manual is regularly updated. Visit

<http://www.crannog-software.com/netflowtracker.cfm> to download the latest version.

What is NetFlow?

A network flow is a sequence of packets between a given source and destination in one direction only. Cisco routers store and export information about the network flows they handle for network management purposes; high-end routers and switches use network flows to accelerate security processing. In order to distinguish flows from one another, the source and destination addresses and application (TCP/UDP) port numbers are used. The IP Type of Service byte, protocol type and the ifIndex of the input interface are also used to uniquely identify the flow to which a packet belongs.

What is NetFlow Tracker?

NetFlow Tracker provides a powerful but easy-to-use set of dynamic charts and reports to help the network administrator make sense of the NetFlow information provided by his routers. The focus is on troubleshooting and diagnostics; long-term analysis is not catered for.

Features and Benefits

- Highly detailed view of network traffic without the need for costly probes.
- Web-based front end allows users anywhere on the network to use the system.
- Straightforward installation and configuration.
- Can be installed on Windows, Linux and Solaris based servers.
- Per-minute resolution.
- Traffic statistics visible just minutes after the event.
- Allows rapid diagnosis of network congestion and failure.
- Useful when configuring QoS to examine the effect of a change in policy.
- Stores one week of full information by default.
- All real-time reports and charts can be filtered on any field.
- Every real-time report and chart allows drilldown on each row or area.
- Every real-time chart allows zooming in and drilling down on a selected time range.
- Custom long-term reports and charts can be created.
- Custom executive reports can be defined and easily accessed.
- Every report and chart can be formatted as CSV for further processing.
- Straightforward URL format for linking current, automatically updated charts into other applications.
- Optimized database structure ensures fast report generation under heavy load.

Installation

Minimum System Requirements

The type of system required to run NetFlow Tracker depends on the number of devices sending NetFlow information to it and the amount and nature of traffic handled by those devices. The following requirements are a guideline; the only way to determine your requirements is by testing the software's performance in your network environment.

- Single processor of Pentium III, Pentium IV or Xeon class, although multiple processors will provide a modest performance increase.
- 512MB RAM, although performance will increase with the amount of RAM available for the disk cache and database buffers.
- High performance disk subsystem with substantial free space – the exact nature of this is dependent on system load. For all but the lightest of loads, a server RAID card running RAID 5 over at least three high-performance disks is recommended. NetFlow Tracker stores and queries full information for a week; a busy enterprise router can generate in the order of 20GB of NetFlow information in this time.
- For running on Microsoft Windows™, an NT4™-compatible operating system is required. A server version will provide better performance due to its more advanced disk caching and memory management.
- For running on Solaris a Sparc based server running Solaris 8 or above is required.
- Any modern Linux distribution capable of running Java 1.4.2 and MySQL 4.1 is supported.

Pre-installation Checks

Before installing, there are a few things you need to check:

- NetFlow Tracker puts a heavy load on the system. It is strongly recommended that you install it on a dedicated server.
- You must be logged in as an administrator in order to install the software.
- NetFlow Tracker uses MySQL to provide database services. Due to the large database size and optimised structure, MySQL must be configured in a way that would seriously degrade the performance of many other types of software that use MySQL. Thus it is recommended that no other MySQL-dependent software be installed on the server running NetFlow Tracker.
- The version of MySQL used by NetFlow Tracker is significantly different to that used by Crannog Software's products NetFlow Monitor, NetWatch and ResponseWatch. If NetFlow Tracker is installed on a server running one of these products it will not function correctly. Likewise, if one of these products is installed on a server running NetFlow Tracker, both products are likely not to function correctly.
- NetFlow Tracker contains an embedded web server. Web servers normally run on port 80, but this may be in use by another web server on your system. You can choose a different port during installation or disable other web servers prior to installation if you wish.
- If you have previously configured a router for NetFlow Monitor, note that NetFlow Tracker requires a different active flow timeout or long aging timer be configured. See [Appendix 1](#) for more information.

Installation on Microsoft Windows™

Installation is straightforward and should take no more than a few minutes. If you received NetFlow Tracker on CD the setup program should start automatically. If not, simply open the CD drive in My Computer and double-click “setup.exe”. If you downloaded the software simply double-click the file you downloaded. Installation involves several steps. At each step, you can click the “Next >” button to accept the default choices and continue.

Unsupported MySQL detection

If MySQL is installed on the server already, you will see a message informing you of this and asking if you wish to continue. While it is not recommended that you do continue, it is possible. Note however that NetFlow Tracker was tested with the version of MySQL it ships with and may not function correctly with a different version. The installation program will fail if the installed version of MySQL uses a root password.

Java Runtime Environment installation

If the server does not have the required version of the Java Runtime Environment installed, you will be prompted to press Ok to install it. It will take several seconds to launch the Java installer, after which you must accept Sun's licence agreement. You will then be given the choice of Typical or Custom installation; if you wish not to have your web browser configured to use Sun's Java Plug-in you must choose Custom installation.

Welcome & Licence Agreement

Once the Java Runtime Environment is installed, you can press the “Next >” button to view Crannog Software's licence agreement, which you must agree to before pressing “Next >” again.

Customer Information

You will be asked to provide your name and company name, and whether to install the software just for yourself or for every user that logs in to the system. If you choose to install the software just for yourself, only you will see the shortcut to the web front-end and only you will be able to uninstall the software.

Setup Type

If you choose “Complete” NetFlow Tracker will be installed to the folder “nftracker” on your system drive, MySQL to the folder “MySQL” on the same drive, and the internal web server will run on port 80 if available. If port 80 is unavailable you will be prompted to choose another. If you want to change the install folders or choose a different port even if 80 is available you must choose “Custom”.

Custom Setup

You will only see this dialog if you chose custom setup above. You should see options for NetFlow Tracker and MySQL, unless an unsupported version of MySQL was detected. To change the install folder for either NetFlow Tracker or MySQL, click on the feature and then on “Change...”.

Select HTTP Port

You will only see this dialog if you chose custom setup or if port 80 is in use. You can choose a port and press “Test” to check if it is available, or simply press “Next >” which will not allow you to proceed if the port is unavailable.

Ready to Install

Click "Install" to start. Installation should take no more than a few minutes; if it appears to have stopped for a long time you should contact Crannog Software. When installation is complete you can click "Finish" to close the install program.

Accessing the web front-end

The install program will have placed a shortcut to the web front-end in a folder called "NetFlow Tracker" in the Programs section of your start menu.

Installation on Solaris and Linux

Instructions for a fresh install or an upgrade are available with the program files from Crannog Software's FTP site:

<ftp://ftp.crannog-software.com/nftracker/>

Post-installation Tasks

Access the web front-end

You can access the web front-end from any workstation on the network by opening the following address in a web browser:

<http://address:port>

Where "address" is the address of the server and "port" is the http port you chose, or 80 if you didn't choose a port.

Note that the web browser must support Java applets; when you installed the Java Runtime Environment it will have set up any browsers on the server with this capability, but you may find that other machines on your network do not display applets correctly, especially those running Windows XP. You can easily download the Java Plug-in from <http://www.javasoft.com> if you find a browser that does not support Java applets.

Open the settings page

The first thing you'll see when you access the web front-end is a splash screen displaying the product version and your licence details. This will disappear after a few seconds, or you can click anywhere on the page to dismiss it. You can then click on "Settings".

Install your licence

If you have a full or trial licence you should install it using the [Licensing](#) settings page.

Set up SNMP community strings

If any of the devices you intend monitoring do not use a read-only SNMP community of "public" you will need to add their communities to the list in [SNMP Settings](#).

Add listener ports

If you intend monitoring more than one device it is recommended that you set up one listener port per device rather than use the default port 2055 for all of them. You can add ports in the [Listener Ports](#) settings page.

Set up web front-end security

If you wish to set passwords to protect access to the web front-end and the settings pages you can do so in [Security Settings](#).

Configure your routers and switches

You must configure your devices to send NetFlow exports to the server running NetFlow Tracker, and to allow the server read-only SNMP access. Even if you have set up NetFlow before, please read the configuration guide in [Appendix 1](#).

Verify that data is being received

You can check that data is being received from a device by looking for it in the [Performance Counters](#) settings page. You should also check the [Device Settings](#) to ensure that SNMP access was successful.

Using NetFlow Tracker

Once you have installed NetFlow Tracker and configured your devices, data will be available within a few minutes. There are many ways to access this data.

Real-time Data

NetFlow Tracker stores up to fourteen days full NetFlow data with one minute resolution. This data is can be reported upon once it is several minutes old. There are several ways to view reports on this data – from the [Network Overview](#) page, from the [Devices](#) page or from the [Filter Editor](#).

Long-term Data

In addition to automatically storing full data for up to fourteen days, NetFlow Tracker can be configured to store summarised data for any length of time. Long-term data is not stored automatically; long-term reports must first be set up using the [Report Settings](#) page. See the [Long-term Reports](#) chapter for more about setting up and viewing long-term reports.

Executive Reports

You can configure custom reports using the [Report Settings](#) page that contain sections from multiple real-time or long-term reports. See the [Executive Reports](#) chapter for more about setting up and viewing executive reports.

Network Overview

The Network Overview page is accessible from the home page of the software; if you do not have user security set up (see [Security Settings](#)) it is also the default page you see when you access the software.

The page gives you a simple overview of the devices and interfaces currently carrying the most traffic on your network. You can click on a device in the pie chart or on its name to see its top applications and busiest interfaces; you can also click on an interface name to see its recent traffic and top applications. It is also possible to drilldown from any of the charts to examine the data in more detail; see [Working with Charts](#) for more about this.

Devices

While the [Network Overview](#) page is useful for quickly identifying the busiest devices and interfaces on your network, the Devices page lists all devices regardless of how busy they are. You can sort the devices by name, address, recent peak traffic rate and recent peak packet rate by clicking the appropriate column header. By default, each peak rate is the highest two-minute rate in the last six hours, but this will be different if the default time range is altered (see [Report Settings](#)). Note that the report is refreshed regularly to ensure it is always up-to-date.

Device traffic meters

In addition to the orderable columns there are two graphical meter columns that allow you to instantly see which devices are currently busy. Each chart shows you the recent peak and the current rate:




Each chart is scaled relative to the busiest device; this ensures that a high value on a chart indicates a relatively high traffic or packet rate.

If you click on either of the meters, you will open a chart of the device's recent activity in terms of traffic or packets over time. By default the last six hours will be shown.

There are various controls you can use to manipulate the chart and examine areas of it more detail; for more see [Working with Charts](#) below.

Interfaces

If you click on a device's name, you will open a page listing all of that device's interfaces. The interfaces can be sorted by name, recent peak utilisation in either direction, recent peak traffic rate in either direction and recent peak packet rate in either direction. The graphical columns on the interface status report show the recent peak and current rates in each direction on each interface:



The scale of the chart depends on which column it is in; the "% Utilisation" column scales each row of each chart according to the configured speed of the interface in that direction whereas the "Relative Traffic" and "Relative Packets" are scaled relative to the busiest direction of the busiest interface. This ensures that a high value on a chart indicates either high utilisation or a relatively high traffic or packet rate. Note that you can change the speed of an interface in [Device Settings](#); you will certainly need to do this for an asynchronous interface. You can also use the Device Settings page to hide interfaces that never export any NetFlow data.

To examine an interface in more detail you can click on its name or any of its meters. If you are unsure about which interface you want to examine, hover the mouse pointer over the interface's name to see its speed, type and extended description if available. When you click on an interface, you will open a chart showing the interface's recent bi-directional utilisation, traffic rate or packet rate over time; see [Working with Charts](#) below for more on the various controls.

Per-AS data

If your router uses BGP to route traffic it will provide source and destination origin or peer AS numbers in its NetFlow data. NetFlow Tracker creates optimised bi-directional charts for each AS just as it does for each interface. An AS chart is only available for a single device as otherwise there is a high chance that some or all traffic will be accounted for multiple times by multiple routers. You can use the [Filter Editor](#) to create a report or chart based upon an AS and data from multiple routers.

To view the ASs routed by a given router, click the ASs link in the navigation menu at the top of the interface report:

[main menu](#) > [devices](#) > [interfaces](#) | [ASs](#)

The AS list is similar to the interfaces list, but does not show percentage utilisation.

Working with Charts

Charts are one of the most useful ways of working with data in NetFlow Tracker. A chart lets you quickly pick out an area of interest to examine in further detail. A chart displays the elements that contributed most to the overall total traffic or packet rate over the charted time range. By default, at most ten elements are charted but this can be configured in the [Report Settings](#) page.

Viewing earlier or later data

You can easily look at earlier or later data by using the forward and back buttons above the chart:



Note that when you open a device or interface chart from the device or interface lists it will automatically keep up to date, but using the forward or back buttons will prevent this from happening.

Changing the displayed chart

All charts have several views, only one of which is displayed at a time. You can change which one is displayed using the tabs above the chart:

Interface - % Utilisation Traffic Rate Packet Rate

In this case, the utilisation chart is displayed and the corresponding tab is raised.

The chart legend

Each charted element has a corresponding row in the legend below the chart. The legend may also have a row for other elements that were not big enough to be charted separately. Depending on the type of chart, some elements in the legend may be underlined; this indicates that more information is available by hovering the mouse over the text.

Zooming in

You can zoom in to the chart by clicking the zoom in button on the toolbar:



This will zoom in on the centre of the chart. If you want to zoom in on a particular selection, see [Selecting a time range](#) below. Note that zooming in will stop the chart from automatically refreshing.

Zooming out

You can zoom out from the chart by clicking the zoom out button on the toolbar:



This will zoom out from the centre of the chart and will again stop the chart from automatically refreshing.

Selecting a time range

If you wish to zoom in on a particular time range you can do so by clicking and dragging the mouse across the chart. You can then zoom in on the selection using the [zoom in](#) toolbar button.

Selecting the entire time range

You can select the entire visible time range using the select all toolbar button:



Examine selected data

Once you have [selected a time range](#) as above, you can “drill down” into it by clicking the right mouse button on the selection. A context menu will pop up, allowing you to create another chart based upon any one of or all of the charted elements during the selected time range. If the chart is automatically refreshing and you used the [select all](#) button to select the time range the new chart will also automatically refresh. The types of chart you can create are described in [Report Templates](#) below.

View a standard chart as a pie chart

Most charts allow you to open a pie chart of the entire charted time range by clicking the pie chart toolbar button:



See [Working with Pie Charts](#) below for more about tabular reports.

View a standard chart as a tabular report

Most charts allow you to open a tabular report of the entire charted time range by clicking the report toolbar button:



See [Working with Tabular Reports](#) below for more about tabular reports.

Alter the filter applied to a standard chart

Most charts allow you to change the applied filter by click the filter editor toolbar button:



See [Creating Filtered Reports](#) for more about the filter editor.

View resolved domain names

If a chart shows IP addresses several of them may be underlined; this indicates that you can see the resolved domain name by hovering the mouse over the address. You can attempt to resolve more of the addresses by clicking the refresh toolbar button:



You can also reload the chart with all resolvable domain names shown in full by clicking the resolve all button:



If all resolvable domain names are displayed you can revert to the normal display of just addresses by clicking the resolve available toolbar button:



Export a chart to another application

You can convert a chart to a comma-separated value (CSV) file by clicking the CSV toolbar button:



You will be prompted to open or save the file; most databases and spreadsheets should be able to understand the format, described in [Appendix 2](#).

Print the chart

You can open a version of the currently displayed chart that is designed for printing or archiving by clicking the print button:



Open the chart in a new window

You can open the chart in its own window using the new window toolbar button:



Working with Pie Charts

Most [charts](#) can be displayed instead as a pie chart. Rather than breaking the selected time range into small chunks and charting each one, a pie chart shows each of the top element's proportion of the total octets or packets during the entire time range.

Most of the toolbar buttons used for working with a chart are also used for working with a tabular report; however there are some differences.

View a pie chart as a standard chart

You can view a pie chart as a chart over time by clicking the chart toolbar button:



Working with Tabular Reports

Most [charts](#) can be displayed instead as a tabular report. Rather than breaking the selected time range into small chunks and charting each one, a tabular report shows the entire time range in one table. A tabular report also shows every contributing element rather than just the largest ones.

Many of the toolbar buttons used for working with a chart are also used for working with a tabular report; however there are some differences.

Filtered utilisation

If the source data for a report is filtered by interface, the total utilisation of all the traffic displayed in the report as a percentage of the interface bandwidth is shown under the interface name. This can help you judge whether an element's traffic is significant or not.

View a tabular report as a chart

You can view a report as a chart by clicking the chart toolbar button:



View more rows of a tabular report

If there are more than twenty-five rows in a report it will be displayed in multiple pages to avoid long download times. The row above the column headings shows where you are in the report and allows you to page through it:



The buttons to the left of the scrollbar move to the first page of the report and back one page, respectively. Since the first page of the report is shown already, these buttons are greyed out. The buttons to the right of the scrollbar move forward one page and to the last page respectively. Clicking anywhere in the scrollbar will move to the corresponding position in the report; i.e., if you click one-third of the way along the scrollbar the page one-third of the way into the report will be shown. A blue line or box on the scrollbar indicates what page is shown and how much of the report the page represents.

Sort a tabular report


A report can be sorted on any of the columns describing the reported elements, or can be sorted by traffic or packet rate. Simply click the column heading – if you click a column heading twice it will be sorted in the opposite order.

Examine a single row

Every row in a tabular report has a radio button to its left:



You can click one of these radio buttons to select a row to drill down into. Note that only one row can be selected. To examine the data contributing to that row's figures, select the type of sub-report you'd like to open from the drop down list at the bottom of the report and click on "Filter...":



Thus if you are looking at a report of source applications, you can select an application and view a report of source addresses using that application.

Report Templates

Whenever you create a new tabular report or chart you can choose any of the standard report templates depending on what you want to examine:

Address Reports

- **Source Addresses** – shows the IP addresses that were the source of most traffic or packets.
- **Destination Addresses** – shows the destination IP addresses that were the destination of most traffic or packets.
- **Address Pairs** – shows the pairs of connected IP addresses that exchanged most traffic or packets.
- **Source Address Dissemination** – charts the source addresses that conversed with the most distinct destination addresses and that were involved in the most distinct endpoint-to-endpoint conversations. This can help detect file sharing or virus infected hosts.
- **Destination Address Popularity** – charts the destination addresses that conversed with the most distinct source addresses and that were involved in the most distinct conversations.

Session Reports

- **Protocols** – shows the IP protocols, such as TCP or UDP, used by most traffic or packets.
- **Source Applications** – shows the IP applications that were the source of most traffic or packets. An IP application is a combination of an application port and protocol; common examples are HTTP or FTP. You can assign names to applications using the [IP Application Names](#) settings page. Examining the source applications inwards on an interface can show you what applications are using your Internet bandwidth.
- **Destination Applications** – shows the IP applications that were the destination of most traffic or packets. The destination applications outwards can show the most requested applications on a link.
- **Recognised Applications** – shows the IP applications that were the source or destination of most traffic or packets. Whether the application was the source or destination depends on whether it has a name defined in the [IP Application Names](#) settings page, or if both or neither have names, whichever has the lower port number.
- **Conversations** – charts the pairs of connected endpoints that exchanged most traffic or packets. A single conversation represents, for example, a web browser downloading a single image.
- **Source Endpoints** – charts the IP addresses and corresponding applications that were the source of most traffic or packets. The top source endpoints inwards on a link are the remote services using your bandwidth.
- **Destination Endpoints** – charts the IP addresses and corresponding applications that were the destination of most traffic or packets.
- **Server-Client Sessions** – charts the pairs of connected source endpoints and destination addresses that exchanged most traffic or packets. A session might represent, for example, a web browser downloading several web pages with images from a web server.
- **Client-Server Sessions** – charts the pairs of connected source addresses and destination endpoints that exchanged the most traffic or packets. A session could represent a client's requests to a web server for several pages and images.

QoS Reports

- **Types of Service** – charts the ToS levels with most traffic or packets.
- **Differentiated Services** – charts the DiffServ code points with most traffic or packets.

Network Reports

- **Source ASs** – charts the autonomous systems that were the source of most traffic or packets. Note that a switch does not know anything about ASs.
- **Destination ASs** – charts the autonomous systems that were the destination of most traffic or packets.
- **AS pairs** – charts the pairs of connected ASs that exchanged most traffic or packets.
- **Source Networks** – charts the IP subnets that were the source of most traffic or packets. Note that a router may not know the subnet of a particular address, and a switch never knows it.
- **Destination Networks** – charts the IP subnets that were the destination of most traffic or packets.
- **Network Pairs** – charts the pairs of connected IP subnets that exchanged most traffic or packets.

Interface Reports

- **In Interfaces** – charts the router interfaces or switch ports that were the arrival point of most traffic or packets. Note that this is only meaningful for the outwards direction.
- **Out Interfaces** – charts the router interfaces or switch ports that were the departure point of most traffic or packets. Note that this is only meaningful for the inwards direction.
- **Next Hops** – charts the next-hop addresses that received most traffic or packets. Note that only a router can supply a next-hop address.

Traffic Identification

- **Identified Applications** – shows the identified applications with most traffic or packets. In order for applications to be identified the NetFlow device must support the functionality and its identified application mapping must be configured in [Device Settings](#).
- **Traffic Classes** – shows the traffic classes that with most traffic or packets. In order for traffic classes to be identified the NetFlow device must support the functionality and its traffic class mapping must be configured in [Device Settings](#).

Creating Filtered Reports

NetFlow Tracker allows any chart or tabular report to be created using a powerful dialog called the filter editor. To create a filtered report, click on “Filter Editor” on the main page.

Most of the options in the filter editor are initially collapsed to save space and bandwidth; you can expand an item by clicking the plus sign to its left. Each option allows you to specify a restriction on the source data considered for the report; if a filter is not specified it will not impose any restriction.

To edit a filter option you must type or pick the element you would like to include in the filter, and then click the “Add” button. To remove an element, select it and click “Remove”.

Report Template

You can choose the type of element you wish to report here, and specify whether you want to create a tabular report, chart or pie chart. For more about the different types of report, see [Report Templates](#) above.

Start Time

Pick the date and time of the earliest data to consider. The default value is six hours before you opened the filter editor.

End Time

Pick the date and time of the last data to consider.

Time Mask

You can use the time mask filter to select only certain times of day within the time range. For example, you can choose to only consider data between 8:30 and 18:00 on a weekday. To do this, select Monday, Friday, 8:30 and 18:00 and press “Add”:

Mon ▾ - Fri ▾ 08 ▾:30 ▾-18 ▾:00 ▾

You can add as many masks as you wish; only data within one or more masked areas is considered. If no masks are selected then all data between the start time and end time is considered.

Time Zone

You can change the time zone used to interpret the start and end times and time masks from the default of the time zone used by the NetFlow Tracker server.

Source Device

You must select which router or switch you want to consider. If you need to consider more than one device, click “Multiple...”, but be aware that if you select multiple devices there is a chance that some or all traffic may be accounted for multiple times.

In Interface

You can report on inbound traffic for an interface or set of interfaces by adding them to the in interface filter. The interfaces you can pick depend on the filtered source devices.

Out Interface

The out interface filter restricts a report to just outbound traffic from a set of interfaces. Used in combination with an in interface filter it will report on traffic that took a particular path through a router.

In/Out Interface

The in/out interface filter restricts the report to bi-directional traffic for the selected interfaces.

Source Address

You can restrict the report to traffic with a given source IP address or one of a set of source IP addresses. Type the address or domain in the box and click “Add”. If you type a domain name, all addresses resolved for that domain are added to the filter.

Destination Address

The destination address filter will report on data with one of a set of destination IP addresses.

Source/Destination Address

This filter will consider traffic either originating from or destined for the given addresses.

Protocol

You can restrict the set of IP protocols considered. For example, you may want to consider only UDP or ICMP traffic while investigating a denial-of-service attack.

Source Application

The source application filter restricts the IP protocol and source application port number. You can enter a port number and protocol manually or you can select from the configured in the [IP Application Names](#) settings page.

Destination Application

This restricts the protocol and destination application port, selectable by name.

Source/Destination Application

This filter considers traffic using the given application as either the source or destination.

Recognised Application

This filter selects traffic with the given source or destination application. Whether the source or destination application is considered depends on whether it has a name defined in the [IP Application Names](#) settings page, or if both or neither have names, whichever has the lower port number.

Identified Application

This filter selects traffic with the given identified application. In order for applications to be identified the NetFlow device must support the functionality and its identified application mapping must be configured in [Device Settings](#).

ToS

You can report only on traffic bearing any one of a set of type-of-service byte values. You build the ToS byte value by picking the priority and the minimize delay (D), maximise throughput (T), maximise reliability (R) and minimise monetary cost (M) flags. If you leave the priority or any of the flags empty then only the fields you supplied a value for are considered. Thus you can match traffic of a given priority with any flags, or with particular flags set or unset but any priority and any values for the other flags.

DiffServ

This will select only traffic bearing one of the selected differentiated service code points. Since DiffServ and ToS use the same field in the IP header you should not use both filters at the same time. You can assign a name to a code point using the [DiffServ Names](#) settings page.

Traffic Class

This filter selects traffic with the given traffic class. In order for traffic classes to be identified the NetFlow device must support the functionality and its traffic class mapping must be configured in [Device Settings](#).

Source AS

You can select traffic bearing one of a set of source AS numbers. Whether this is the origin or peer AS depends on the configuration of the router (see [Appendix 1](#)). You can enter an AS number manually or select from the set of private-use ASs configured in the [AS Names](#) settings page; note that you cannot select public ASs by name to avoid the filter page being excessively large.

Destination AS

This restricts the source data to traffic bearing the given destination origin or peer ASs.

Source/Destination AS

This filter considers traffic to or from the given origin or peer ASs.

Source Subnet

This will select traffic with the given source subnet. You can enter the network address and mask length manually or select from the subnets configured in the [Subnet Names](#) settings page. Note that the subnet mask used by the router to route the traffic is ignored when applying this filter.

Destination Subnet

This filter selects traffic with the given destination subnets. Note that a destination subnet filter of 224.0.0.0/4 will select multicast traffic.

Source/Destination Subnet

This filter selects traffic to or from the given subnets.

Source Mask

This will select traffic routed using the given source network mask.

Destination Mask

This filter selects traffic with the given destination network mask.

Source/Destination Mask

This filter selects traffic with the given source or destination network mask.

Next Hop

This will filter traffic according to the next hop used by the router in routing the traffic.

Long-term Reports

Long-term reports allow you to look at data over much longer time ranges than is possible with the standard real-time database. The data for long-term reports is summarised in advance so a long-term report over several days or weeks can often be much faster than an equivalent real-time one.

Long-term reports are not created automatically – you must first identify which reports you would like to see over the long-term and set them up in [Report Settings](#).

To access your long-term reports, click on “Long-term Reports” on the software’s homepage. You can then access your long-term reports in two ways: the [Devices](#) page or the [Long-term Filter Editor](#).

Devices and Interfaces

The long-term device and interface pages are very similar to the [real-time](#) versions, but there are several differences. Most noticeable is the time range selector at the bottom of the page. The default time range for a long-term report is the last seven full days according to the time zone of the NetFlow Tracker server; this can be changed in [Report Settings](#). The time range selector will change the time range of the current report or chart, and of any reports or charts opened by interacting with it:



You can select any number of full minutes, hours, days, weeks, months, quarters, half-years or years. Note that if you zoom in to or out of a long-term chart, or drill down into a selection (other than one selected using the [Select All](#) button), the time range selector will not be available on the resulting chart.

Another major difference is that while the real-time device and interface pages show the peak and most recent traffic and packet rates over the displayed time range, the long-term versions show the peak and average rates. You can also sort the pages by the average rates.

Per-device and Per-interface Long-term Reports

When you select a range of time on a long-term device or interface chart and right-click to drill down you will either find that no charts are available or the set is limited. The only reports that you can access in this way are ones that are created as per-device, per-inbound interface or per-outbound interface in [Report Settings](#).

Filter Editor

You can access any long-term report through the long-term filter editor. It is the only way you can access custom long-term reports that are created as basic reports.

The long-term filter editor is a much simplified version of its [real-time](#) counterpart. You must select the report and time range to view. If the report did not have a time mask applied to it when it was created you will be able to apply one using the [Time Mask](#) and [Time Zone](#) editors. The time range and time mask editors behave exactly like their counterparts in the real-time [Filter Editor](#).

If you select a per-device, per-inbound interface or per-outbound interface report you must also specify what device or interface to report upon. The editors for selecting a device or interface are slightly different to their counterparts in the real-time [Filter Editor](#) in that they allow only one item to be selected.

Executive Reports

An executive report is a pre-defined template that contains one or more charts or tabular reports. Executive reports can be created to show related information on one page and to allow quick access to commonly-used reports. Executive reports are defined in [Report Settings](#) and accessed by clicking on “Executive Reports” on the software’s home page.

Report URL Format

You can easily generate your own URLs or modify automatically created ones for use in network management portals favourites lists.

General Form

`http://<server>:<port>/report.jsp?prm=value&prm=value...`

with:

server	The domain name or IP address of the NetFlow Tracker server
port	The HTTP port of the NetFlow Tracker server
prm, value	A named parameter and its value; supply as many parameters as necessary in any order with each prm=value pair separated by an ampersand.

Report Format Parameters

templid – specifies the [report template](#) to use.

0000	Source Addresses
0001	Destination Addresses
0002	Address Pairs
0003	Protocols
0006	Source Applications
0007	Destination Applications
0008	Source Endpoints
0009	Destination Endpoints
0010	Server-Client Sessions
0011	Client-Server Sessions
0012	Conversations
0013	Types of Service
0014	Differentiated Services
0015	Source ASs
0016	Destination ASs
0017	AS Pairs
0018	Source Networks
0019	Destination Networks
0020	Network Pairs
0021	In Interfaces
0022	Out Interfaces
0023	Next Hops
0024	Source Address Dissemination
0025	Destination Address Popularity
0026	Recognised Applications
0027	Traffic Classes
0028	Identified Applications
_flows	Full flows

id – specifies the [long-term report](#) to use. It is possible to enable several standard long-term reports in [Report Settings](#); the ids for these reports are given below. The id for a custom report is available in [Report Settings](#).

0000	Source Addresses per inbound interface
0001	Source Addresses per outbound interface
0002	Destination Addresses per inbound interface
0003	Destination Addresses per outbound interface
0004	Recognised Applications per inbound interface
0005	Recognised Applications per outbound interface
0100	Source Addresses per source device
0101	Destination Addresses per source device
0102	Recognised Applications per source device

cid – specifies the [executive report](#) to use. The id for an executive report is available in [Report Settings](#).

output – specifies if a tabular report or chart will be generated.

table	A tabular report will be generated (default)
chart	A chart over time will be generated
pie	A pie chart will be generated; this value can be used as an output parameter in an executive report to display a pie chart if the sub-report is a chart over time

nrecords – specifies the number of rows to show per page of a tabular report.

<number>	The number of rows per page
-----------------------	-----------------------------

others – specifies that a tabular report shows an “others” row instead of a page navigator. Note that long-term tabular reports always show an “others” row.

true	An “others” row is shown instead of a page navigator
false	No “others” row is shown (default)

visible – specifies a visible column of a report or chart; this parameter should be specified as many times as is necessary to include all desired columns. By default, all columns are visible. This parameter can be used as an output parameter in an executive report.

<heading>	The column heading
------------------------	--------------------

chartTitle – specifies the chart to show. This parameter can be used as an output parameter in an executive report.

<title>	The chart title
----------------------	-----------------

chartWidth – specifies the width of the chart. This parameter can be used as an output parameter in an executive report.

<width>	The chart width in pixels
----------------------	---------------------------

chartHeight – specifies the height of the chart. This parameter can be used as an output parameter in an executive report.

<height>	The chart height in pixels
-----------------------	----------------------------

sections – specifies the report sections to output. This parameter can be used as an output parameter in an executive report.

<sections>	The sections, formed by summing the values for each section.	
	1	Title
	2	Time range & filter description
	4	Main report or chart body
	8	Chart title, if applicable
	16	Chart legend, if applicable
	32	Result information, if applicable

features – specifies the available interactive report features. Some values for this parameter can be used as an output parameter in an executive report; these are indicated with an asterisk (*) below.

<features>	The features, formed by summing the values for each feature.	
	1	Navigation Menu
	2	Select All button, if applicable
	4	Zoom In button, if applicable
	8	Zoom Out button, if applicable
	48	Open as Tabular Report, Chart or Pie buttons as applicable
	64	Filter Editor button, if applicable
	128	Refresh and Resolve All buttons, if applicable
	256	Print and CSV buttons, if applicable
	512	Open in New Window button*
	1024	Drilldown controls*
	2048	Direct drilldown links (found in navigation reports)*
	4096	Page navigator
	8192	Sortable column headers
16384	Chart scrollbar	
32768	Chart selection headers	
65536	Time range editor, if specified	

resolve – specifies how domain names will be handled in a report with an IP address column.

all	All domain names will be resolved and shown in full
available	Only already resolved domain names will be shown, as tooltips (default)

format – specifies the output format of the report or chart.

html	Fully interactive HTML (default)
print	Printable/saveable HTML
csv	Comma separated values

reload – specifies the number of seconds between automatic refreshes of the report. This is best used in conjunction with one of the dynamic time ranges, below. Only the interactive HTML format supports this parameter.

-1	The report will not reload automatically (default)
<seconds>	Number of seconds between refreshes

Time Range Parameters

The time range can be specified in one of several ways. If no time range is specified a default will be used.

Start and end time

An fixed start and end time can be specified in UTC, which is the number of milliseconds since 1 Jan 1970, or in plain text.

stime – specifies the start of the required time range.

<time>	The time in milliseconds UTC
<dd>/<MM>/<yyyy>%20<HH>:<mm>	The time, with <dd> being the date, <MM> the month, <yyyy> the year, %20 a URL-encoded space character, <HH> being the hour in the 24-hour clock and <mm> being the minutes

etime – specifies the end of the required time range.

<time>	The time in milliseconds UTC
<dd>/<MM>/<yyyy>%20<HH>:<mm>	The time, with <dd> being the date, <MM> the month, <yyyy> the year, %20 a URL-encoded space character, <HH> being the hour in the 24-hour clock and <mm> being the minutes

Fixed length

If you would like to create a URL that will always show a current time range, you can specify a certain number of milliseconds ending at the time the report is generated.

length – specifies the length of the required time range.

<millis>	The length in milliseconds
----------	----------------------------

Calendar-based (simple)

A simple calendar-based time range is a given number of units ending either when the report is generated or at the end of the last full unit before the report is generated.

unit – specifies the unit to measure the time range in.

hour	Hours
day	Days
week	Weeks
month	Months
quarter	Quarters
halfyear	Half-years
year	Years

nunitsago – specifies the number of units before the time of report generation the time range should end.

0	The time range will end at the time of report generation; an incomplete unit will be counted as full
1	The time range will extend to the end of the last full unit before the time of report generation (default)
<number>	The time range will extend to the end of this number of full units before the time of report generation

nunits – specifies the number of units required. Note that this may include a partial unit.

1	The time range will extend for a single unit (default)
<number>	The time range will extend for this number of units

Calendar-based (advanced)

An advanced calendar-based time range has an optional start date specified as a given number of units before the time of report generation, defaulting to the day of report generation. The start time is specified in plain text. The optional end date is specified in the same manner as the start date, defaulting to the same day as the start date. Finally, the end time is specified in plain text.

date_unit – (optional) specifies the unit to measure how long before the report is generated the time range starts and ends.

day	Days
week	Weeks
month	Months
quarter	Quarters
halfyear	Half-years
year	Years

sdate_nunitsago – (optional) specifies the number of units before the time of report generation of the first day of the time range.

1	The first day of the time range will be the first day of the current unit at the time of report generation (default)
<number>	The first day of the time range will be at the start of this number of full units before the time of report generation

edate_nunitsago – (optional) specifies the number of units before the time of report generation of the last day of the time range.

0	The last day of the time range will be at the end of the current unit at the time of report generation; this is likely to be later than the time of report generation
1	The last day of the time range will be at the end of the last full unit before the time of report generation (default)
<number>	The time range will extend to the end of this number of full units before the time of report generation

stime – specifies the time of day at which the time range starts.

<HH>:<mm>	The time, with <HH> being the hour in the 24-hour clock and <mm> being the minutes
------------------------------	--

etime – specifies the time of day at which the time range ends.

<HH>:<mm>	The time, with <HH> being the hour in the 24-hour clock and <mm> being the minutes
------------------------------	--

Applying a time-of-day mask to the time range

If the time range is longer than a day, you may wish to restrict it to just certain times on each day. You can select only working hours or only non-working hours, for example.

Note that if a long-term report has a configured time zone or mask, this parameter will have no effect.

timemask – specifies an inclusive mask to apply the to time range. To specify multiple inclusive masks, include a parameter name and value in the URL for each mask.

<code><day1>-<day2>/<time1>-<time2></code>	The range of weekdays and the times on those weekdays to include in the mask with a weekday being one of SUN, MON, TUE, WED, THU, FRI or SAT , day2 coming on or after day1 in the list above, a time being in the 24-hour form hh:mm , and time2 being after time1
--	--

Specifying a time zone

By default the time zone used to interpret calendar-based time ranges and time-of-day masks is the time zone of the NetFlow Tracker server. You can specify a non-default time zone if you wish.

Note that if a long-term report has a configured time zone or mask, this parameter will have no effect.

timezone – specifies the time zone of the report.

0	(GMT-12:00) International Date Line West
1	(GMT-11:00) Midway Island, Samoa
2	(GMT-10:00) Hawaii
3	(GMT-09:00) Alaska
4	(GMT-08:00) Pacific Time (US & Canada); Tijuana
15	(GMT-07:00) Arizona
10	(GMT-07:00) Mountain Time (US & Canada)
13	(GMT-07:00) Chihuahua, La Paz, Mazatlan
33	(GMT-06:00) Central America
20	(GMT-06:00) Central Time (US & Canada)
30	(GMT-06:00) Guadalajara, Mexico City, Monterrey
25	(GMT-06:00) Saskatchewan
45	(GMT-05:00) Bogota, Lima, Quito
35	(GMT-05:00) Eastern Time (US & Canada)
40	(GMT-05:00) Indiana (East)
50	(GMT-04:00) Atlantic Time (Canada)
55	(GMT-04:00) Caracas, La Paz
56	(GMT-04:00) Santiago
60	(GMT-03:30) Newfoundland
65	(GMT-03:00) Brasilia
70	(GMT-03:00) Buenos Aires, Georgetown

73	(GMT-03:00) Greenland
75	(GMT-02:00) Mid-Atlantic
80	(GMT-01:00) Azores
83	(GMT-01:00) Cape Verde Is.
90	(GMT) Casablanca, Monrovia
85	(GMT) Greenwich Mean Time : Dublin, Edinburgh, Lisbon, London
110	(GMT+01:00) Amsterdam, Berlin, Bern, Rome, Stockholm, Vienna
95	(GMT+01:00) Belgrade, Bratislava, Budapest, Ljubljana, Prague
105	(GMT+01:00) Brussels, Copenhagen, Madrid, Paris
100	(GMT+01:00) Sarajevo, Skopje, Warsaw, Zagreb
113	(GMT+01:00) West Central Africa
130	(GMT+02:00) Athens, Beirut, Istanbul, Minsk
115	(GMT+02:00) Bucharest
120	(GMT+02:00) Cairo
140	(GMT+02:00) Harare, Pretoria
125	(GMT+02:00) Helsinki, Kyiv, Riga, Sofia, Tallinn, Vilnius
135	(GMT+02:00) Jerusalem
158	(GMT+03:00) Baghdad
150	(GMT+03:00) Kuwait, Riyadh
145	(GMT+03:00) Moscow, St. Petersburg, Volgograd
155	(GMT+03:00) Nairobi
160	(GMT+03:30) Tehran
165	(GMT+04:00) Abu Dhabi, Muscat
170	(GMT+04:00) Baku, Tbilisi, Yerevan
175	(GMT+04:30) Kabul
180	(GMT+05:00) Ekaterinburg
185	(GMT+05:00) Islamabad, Karachi, Tashkent
190	(GMT+05:30) Chennai, Kolkata, Mumbai, New Delhi
193	(GMT+05:45) Kathmandu
201	(GMT+06:00) Almaty, Novosibirsk"
195	(GMT+06:00) Astana, Dhaka
200	(GMT+06:00) Sri Jayawardenepura
203	(GMT+06:30) Rangoon
205	(GMT+07:00) Bangkok, Hanoi, Jakarta
207	(GMT+07:00) Krasnoyarsk"
210	(GMT+08:00) Beijing, Chongqing, Hong Kong, Urumqi
227	(GMT+08:00) Irkutsk, Ulaan Bataar
215	(GMT+08:00) Kuala Lumpur, Singapore
225	(GMT+08:00) Perth
220	(GMT+08:00) Taipei
235	(GMT+09:00) Osaka, Sapporo, Tokyo
230	(GMT+09:00) Seoul
240	(GMT+09:00) Yakutsk
250	(GMT+09:30) Adelaide
245	(GMT+09:30) Darwin
260	(GMT+10:00) Brisbane
255	(GMT+10:00) Canberra, Melbourne, Sydney
275	(GMT+10:00) Guam, Port Moresby
265	(GMT+10:00) Hobart
270	(GMT+10:00) Vladivostok
280	(GMT+11:00) Magadan, Solomon Is., New Caledonia

290	(GMT+12:00) Auckland, Wellington
285	(GMT+12:00) Fiji, Kamchatka, Marshall Is.
300	(GMT+13:00) Nuku'alofa

Specifying the chart sample size

When you create a real-time chart the system chooses a sample size that will create as close to 150 samples over the full width of the chart as possible. If you want to you can specify a different sample size to show, for example, a day in hour-long samples or a month in day-long samples.

sample_unit – specifies the unit to measure the sample size in.

minute	Minutes
hour	Hours
day	Days
week	Weeks
month	Months
quarter	Quarters
halfyear	Half-years
year	Years

sample_nunits – specifies the number of units in each sample

1	Each sample will be one unit long (default)
<number>	Each sample will be this number of units long

Specifying the source long-term data

When you create a long-term chart or tabular report, the source data is chosen so the time range will be in as close to 150 samples as possible. You can override this if you wish.

range – specifies the source long-term data to use

daily	Daily data (ten minute samples) will be used
weekly	Weekly data (one hour samples) will be used
monthly	Monthly data (six hour samples) will be used
quarterly	Quarterly data (twelve hour samples) will be used
halfyearly	Half-yearly data (one-day samples) will be used
yearly	Yearly data (two-day samples) will be used

sample – specifies the source long-term data to use

10minute	Daily data (ten minute samples) will be used
1hour	Weekly data (one hour samples) will be used
6hour	Monthly data (six hour samples) will be used
12hour	Quarterly data (twelve hour samples) will be used
1day	Half-yearly data (one-day samples) will be used
2day	Yearly data (two-day samples) will be used

Filter Parameters

Any number of filters can be applied to a report. Each filter is a set of acceptable values for a certain aspect of the source data. If a filter is not specified then all values for that aspect are accepted.

To specify multiple acceptable values for a filter, include the parameter name and value in the URL once for each value.

Note that the filters that can be applied to a long-term report depend upon its type.

device – specifies the address of an acceptable NetFlow-exporting device.

<addr>	The address in dotted-decimal format (a.b.c.d)
---------------------	---

inif – specifies an acceptable input interface, thus selecting inbound traffic on the interface.

<addr>/<id>	The interface with addr being the address of the NetFlow-exporting device in dotted-decimal format and id being the NetFlow Tracker-specific interface identifier
<addr>/-<ifindex>	The interface with addr being the address of the NetFlow-exporting device in dotted-decimal format and ifindex being the current SNMP interface index assigned to the interface

outif – specifies an acceptable output interface, thus selecting outbound traffic on the interface. Format as for **inif** above.

if – specifies an acceptable input or output interface of the flow, thus selecting traffic passed in both directions across the interface. Format as for **inif** above.

srcaddr – specifies an acceptable source address.

<addr>	The address in dotted-decimal format
---------------------	--------------------------------------

dstaddr – specifies an acceptable destination address. Format as for **srcaddr** above.

addr – specifies an acceptable source or destination address. Format as for **srcaddr** above.

proto – specifies an acceptable IP protocol.

<name>	The protocol name, such as TCP or UDP
<number>	The protocol number, in the range 0-255

srcappl – specifies an acceptable source IP application.

<port>/<name>	The application, with port being the application port number in the range 0-65535 and name being the protocol name, such as TCP or UDP
<port>/<number>	The application, with port being the application port number in the range 0-65535 and num being the protocol number in the range 0-255

dstappl – specifies an acceptable destination IP application. Format as for **srcappl** above.

appl – specifies an acceptable source or destination IP application port. Format as for **srcappl** above.

recappl – specifies an acceptable recognised IP application port. Format as for **srcappl** above.

applid – specifies an acceptable identified application.

<id>	The identified application identifier; see Device Settings for more information
-------------------	---

tos – specifies an acceptable Type-of-Service byte.

<prec>	The precedence, in the range 0-7
<tos>	A string of letters indicating which ToS bits must be set or unset, each letter being one of D , T , R or M for low delay, high throughput, high reliability and minimise monetary cost respectively, or d , t , r or m for normal delay, normal throughput, normal reliability and normal monetary cost; any bits not specified as set or unset will be disregarded
<prec>%20<tos>	The precedence and ToS as above; %20 being a URL-encoded space character

ds – specifies an acceptable differentiated service codepoint.

<name>	The assigned name of the codepoint
<code>	The six-digit binary representation of the codepoint
<byte>	The value of the entire Type-of-Service byte, in the range 0-255

class – specifies an acceptable traffic class.

<id>	The traffic class identifier; see Device Settings for more information
-------------------	--

srcas – specifies an acceptable source autonomous system number.

<as>	The AS number, in the range 0-65535
-------------------	-------------------------------------

dstas – specifies an acceptable destination autonomous system number. Format as for **srcas** above.

as – specifies an acceptable source or destination autonomous system number. Format as for **srcas** above.

srcnet – specifies an acceptable source subnet. Note that the subnet mask supplied by the router is ignored.

<addr>/<mask>	The subnet, with addr being the network address in dotted-decimal format and mask being the mask length, in the range 0-32
----------------------------------	--

dstnet – specifies an acceptable destination subnet. Format as for **srcnet** above.

net – specifies an acceptable source or destination subnet. Format as for **srcnet** above.

srcmask – specifies an acceptable source subnet mask, as supplied by the router.

<mask>	The mask length, in the range 0-32
---------------------	------------------------------------

dstmask – specifies an acceptable destination subnet mask. Format as for **srcmask** above.

mask – specifies an acceptable source or destination subnet mask. Format as for **srcmask** above.

nexthop – specifies a next-hop address.

<addr>	The address in dotted-decimal format
---------------------	--------------------------------------

Security Parameters

If a username and password is required to access a report it can be specified in the URL.

j_username – specifies the username.

<username>	The username
-------------------------	--------------

j_password – specifies the password.

<password>	The password
-------------------------	--------------

Management Portal Access Control Parameters

NetFlow Tracker allows management portals to set up restricted access to the system for multiple users. So long as it is possible to conceal the initial URL sent to NetFlow Tracker it is possible for the user to fully interact with the resulting report while being prevented from accessing certain data.

Portal access requires that the restricted users can only access NetFlow Tracker via the portal's proxy server. You can use your firewall to hide the NetFlow Tracker server from the Internet, or you can simply configure password protection. The management portal must also be registered with NetFlow Tracker using the [Management Portal Settings](#) page.

Access restrictions are set up by including the management portal's secret value in the URL along with a set of allowed devices, interfaces, reports, filters and interactive features. If no restrictions of a particular type are set, then all elements of that type are allowed, with the exception that if no device restrictions are set they are implied from the interface restrictions. Since this URL contains the management portal's secret value, it is important that it is not visible to the user; most management portals have a way to provide access through their proxy while concealing the actual URL being sent to the underlying server.

Note that requests from a management portal are authenticated automatically so a username and password does not need to be included in the URL.

When NetFlow Tracker creates a report in response to a request from a management portal, any interaction with that report will cause a cryptographically secure identifier to be included in the URL sent to the server. If a request from a management portal contains neither the correct secret value nor a valid identifier, or attempts to access a resource forbidden by the access restrictions originally supplied by the management portal, it will be rejected.

portalsecret – specifies the secret value assigned to the management portal in [Management Portal Settings](#).

<secret>	The secret value
-----------------------	------------------

acldevice – specifies the address of a permitted NetFlow-exporting device. Format as for **device** above.

aclif – specifies a permitted interface. Format as for **inif** above.

acltemplid – specifies a permitted report template.

null	No report templates are permitted
<id>	A permitted report template; see templid in Report Format Parameters above for permitted values

aclid – specifies a permitted long-term report. Format as for **id** above.

null	No long-term reports are permitted
<id>	A permitted long-term report; see id in Report Format Parameters above for permitted values

aclcid – specifies a permitted executive report. Format as for **cid** above.

null	No executive reports are permitted
<id>	A permitted executive report; see cid in Report Format Parameters above for permitted values

aclfiltereditor – specifies a filter that will appear in the [Filter Editor](#). Note that it will be possible for the user to create reports with other filters by drilling down or manually editing a URL.

null	No filter editors are permitted
0	Source Device
1	Source Address
2	Dest Address
3	Src/Dest Address
4	Next Hop
5	In Interface
6	Out Interface
7	In/Out Interface
8	Protocol
12	Source Application
13	Dest Application
14	Src/Dest Application
15	ToS
16	DiffServ
17	Source AS
18	Dest AS
19	Src/Dest AS
20	Source Subnet
21	Dest Subnet
22	Src/Dest Subnet
23	Source Mask
24	Dest Mask
25	Src/Dest Mask
26	Recognised Application
27	Traffic Class
28	Identified Application

aclfeatures – specifies the permitted interactive report features.

<features>	The features, formed by summing the values for each feature.	
	1	Navigation Menu
	2	Select All button, if applicable
	4	Zoom In button, if applicable
	8	Zoom Out button, if applicable
	48	Open as Tabular Report, Chart or Pie buttons as applicable
	64	Filter Editor button, if applicable
	128	Refresh and Resolve All buttons, if applicable
	256	Print and CSV buttons, if applicable
	512	Open in New Window button
	1024	Drilldown controls
	2048	Direct drilldown links (found in navigation reports)
	4096	Page navigator
	8192	Sortable column headers
16384	Chart scrollbar	
32768	Chart selection headers	
65536	Time range editor, if specified	

Performance Tuning

There are several factors that influence how quickly a given report is generated:

Disk Speed

The first step in creating a report is reading the raw data from disk; increasing the speed of the disk subsystem will make reporting faster. A high-quality server RAID card running a striped pattern such as RAID 5 over fast disks is recommended; more disks will make the array faster. In addition, extra RAM can be used by the operating system for a disk cache.

Query Size

The amount of raw data that needs to be read from disk is dependent on the number of source devices selected, the data load of those devices and the amount of time selected. Indexes are not used due to the increase in database size they would cause, so any other filters have no impact on the amount of raw data read from the disk. If possible, avoid reporting over multiple devices and over long periods of time. It is likely that a report over multiple devices will account for some traffic multiple times.

Database Server Settings

The database server used by NetFlow Tracker can be tuned to improve query speed if you have a fast disk subsystem or lots of RAM, or both. See [Database Settings](#) for details.

Configuration Guide

To open any of the settings pages, click “Settings” on the main page. If you have password protection enabled you may have to login as an administrative user to see the link. Each settings page controls a single aspect of the software; if you make any changes you must click “Ok” on the page before they will be applied and changed. “Cancel” will return to the main settings page without altering anything. It is recommended that you do not use the “Back” button in your web browser as it can cause changes to be lost.

Licensing

You can check the status of your licence or apply a new one using this page. If you received a licence file, load it by clicking “Browse” to locate the file, then click “Load”. If you received your licence in text form, paste it into the large box and press “Decode”. Either way, the licence details will be updated to reflect the new licence. You must click “Ok” to use the new licence.

Listener Ports

NetFlow Tracker listens for NetFlow packets sent to it by any number of routers. When you set up NetFlow exporting on a router, you are asked to provide a port number on the server to send exports to. This is normally 2055, and this is the default used by NetFlow Tracker. However, if you are sending NetFlow exports to NetFlow Tracker from more than one router it is recommended that you use a different port for each one.

To do this, simply add the port numbers you wish to use to the list. You can also choose to listen on all local IP addresses or only one if the server running NetFlow Tracker has more than one IP address and you wish to listen for NetFlow exports on a specific address rather than on all of them.

When you have added all the ports you wish to listen for NetFlow exports on, click “Ok”. If you get an error message, it is probably because one or more of the ports are in use already. They will be marked with an asterisk (*). Remove these ports and add others until there are no errors.

Under very heavy load you may need to increase the size of the buffer used for each listener; see [missed flows](#) under [Performance Counters](#) below for more.

SNMP Settings

Whenever NetFlow Tracker receives exports from a previously unknown device it attempts to scan the device using SNMP to discover its name and the properties of its interfaces. A password called a community is required to use SNMP, and in many cases a default community of “public” is set up on a device. If your devices do not have a read-only community of “public” set up you should add the communities they so use to this list. NetFlow Tracker attempts each one in turn when a new device is detected, so you should put the most frequently used communities first in the list.

You can also set the timeout and number of retries used for SNMP requests; it is unlikely you will need to alter these.

Device Settings

Device List

This page allows you to check the status of a known device and override the interface descriptions and speeds obtained from it.

The name and address of each known device is listed, along with an icon indicating its status; an exclamation (!) indicates that the device could not be contacted using SNMP or it is being ignored due to a license violation and an hourglass (⌚) indicates that the device is currently being scanned and cannot be edited. You can update the list to see if a scan has finished by clicking “Refresh”. If no icon is displayed the device is working correctly.

Clicking the name of the device you wish to edit will open a new page. It is important to remember that any changes you make to any device are only applied when you click “Ok” in the main device settings page.

Device Settings

The settings page for a single device allows you to set its SNMP properties, override the name and local AS number detected using SNMP and override the default “Show interface descriptions” [Report Settings](#) value for the device. The local AS number is required to get correct AS numbers for traffic routed to or from the local AS in a BGP environment; if you do not use BGP this value should be left blank.

SNMP

If the device does not support SNMP you can change the SNMP mode to “Don't use SNMP”. This will assign default properties to each interface encountered in NetFlow exports from the device. It is also possible to freeze a device's configuration by changing the mode to “Keep current configuration” – this will cause any new interface encountered to be ignored, so should be used with caution. If possible you should allow NetFlow Tracker to use SNMP to scan a device as the numbers used to identify the inbound and outbound interfaces in NetFlow exports are not constant and SNMP is the only way NetFlow Tracker can work out a correct correlation between an identifiers and physical interface or port.

You can request an immediate rescan of an SNMP device by clicking “Rescan”. This will scan the device using the SNMP version and community specified in the page but **will not** save them; you must click “Ok” on the main device settings page before any changes are applied. Note that NetFlow Tracker rescans a device when it is restarted, if a new interface is encountered or if it appears the device was rebooted, so you will not normally have to manually rescan a device.

If you are unable to change the configuration of the router or switch, or if an interface is asynchronous, you can override the description or inwards and outwards speed used in reports here. You can also supply interface descriptions and speeds for a non-SNMP compatible device. You should note that if the speed or description supplied by the device changes between SNMP scans NetFlow Tracker uses that speed or description, even if you have previously overridden it. Thus the most recently set description or speed is used, whether it was set on the device or within NetFlow Tracker.

If you wish to prevent interfaces that never report any NetFlow data from appearing in the [interface status report](#) and [Filter Editor](#) check the box corresponding to the interface in the “inactive” column. If the configuration of the device has changed there may be some unused interfaces listed separately; it is likely you will want to mark these as inactive.

Archiving

You can choose to archive old real-time data for the device rather than delete it by checking “Archive real-time data”. See [Archiving](#) for more information.

Traffic Classes

Some types of device can export information about the traffic class used to help route the traffic involved in each flow. Currently some Cisco devices and Packeteer devices support this feature; see [Appendix 1](#) for required configuration. If you have one or more devices that export traffic class information, you must add each traffic class to NetFlow Tracker and configure a map from the device's class ID to the NetFlow Tracker traffic class for each class on each device. To add traffic classes, click on “add/delete” in the heading of the traffic class box for any device. You will then be able to add traffic classes; you must give each one a unique identifier that will be used if you create a URL with a traffic class filter (see [Filter Parameters](#)). Note that this identifier does not need to be the same as the identifier exported by any of your devices for the traffic class.

Once you have added the traffic classes your devices use you must configure a map between the number each device uses to identify a traffic class and the actual traffic class you added. To do this, enter the device's class ID, select the relevant traffic class and click “Add” for each class exported by the device.

Identified Applications

Identified applications are very like traffic classes and are configured in the same way. Unlike a traffic class, which is used by the device to block or apply QoS settings to traffic, an identified application is an accounting tool. Currently only Packeteer devices support this feature; see [Appendix 1](#) for required configuration.

Deleting a Device

Finally, you can delete a device by clicking “Delete”; although the device will only be deleted when you click “Ok” in the main device settings page there is no way to cancel deleting a device except by pressing “Cancel” in the main device settings page and thus losing any other changes. You should also note that if the device is still sending exports to the software it will reappear.

Security Settings

You can set up password protection of the web front end to NetFlow Tracker by adding user accounts here. To add an account, type a login and the same password twice, and tick the administrator box if you wish the user to be able to configure the system. Click "Add" to add the user. To delete an existing user, tick the box above the "Delete" button corresponding to the user and click "Delete". You can also reset a user's password and whether or not the account is an administrator.

You must also choose what level of protection you desire. You can choose not to protect access at all; to protect only access to the settings pages or to protect both configuration and normal access. If you protect access of any sort you will need to add at least one administrator account.

You can also change the page that users see when they access the server without specifying a page (i.e., `http://server/`). You can specify a custom homepage that applies to all users, including the default one when logging in is not required. You can also specify a custom homepage for any user account.

Ensure that the URL of any custom homepage is relative to the server's root; for example, the standard homepage would be specified as "index.jsp" and the [Network Overview](#) would be specified as "report.jsp?cid=_topdevices". Note that since version 2.1, new installs of NetFlow Tracker have the Network Overview pre-configured as a custom homepage.

You can use your own html page if you wish by putting it in the "customweb" folder under the NetFlow Tracker install folder; it is then available from the NetFlow Tracker server as, for example, `http://server/customweb/file.html`, so the homepage would be simply `customweb/file.html`.

Management Portal Settings

If you wish to use a management portal to set up restricted access to NetFlow Tracker for multiple users you must first register it with NetFlow Tracker. Please see [Management Portal Access Control Parameters](#) under [Report URL Format](#) for more details of this feature.

To register a management portal, enter the IP address NetFlow Tracker will see as the source of HTTP requests and a secure secret value that will be included in requests made by the portal and click "Add". To remove a registered management portal, tick the box above the "Delete" button corresponding the portal and click "Delete".

Report Settings

This page lets you configure various values affecting the way reports and charts appear in NetFlow Tracker.

- **Rows per tabular report page** is the number of rows shown on each page of a tabular report. Note that the device and interface status reports show all rows on a single page.
- **Elements considered per chart/long-term block** determines the accuracy of a real-time or long-term chart, and of a long-term tabular report. When a chart is generated only the largest elements are considered from each block when determining the elements to chart. Since it is possible that the highest elements overall may not be the highest elements in each block of the chart, it is important that more elements are considered from each block than the eventual number of charted elements.
- **Charted elements** is the maximum number of elements displayed on a chart, excluding the “Others” element.
- **Long-term tabular report rows** is the maximum number of rows displayed on a long-term tabular report. Note that setting this value higher than the number of rows per tabular report page has no effect. Also note that the accuracy of a long-term tabular report depends upon the number of elements considered per chart block.
- **Default real-time report time range** is the time span used for any real-time report or chart where one is not specified – thus it is the time range of the device, interface and AS status reports and charts and the default time range selected in the filter editor.
- **Reload interval** is the number of minutes between automatic refreshes of the device, interface and AS status reports and charts.
- **Show hostnames in reports** controls whether reports and charts are opened with all resolvable hostnames resolved and shown by default.
- **Show chart legends in descending order** controls whether the rows of a chart legend are shown in the same order as the corresponding tabular report, or in the same order as the areas are drawn on the chart.
- **Show interface descriptions** controls whether the description of an interface is used, when available, in filter descriptions instead of the name.
- **Standard long-term reports are disabled** controls whether the standard set of per-device and per-interface long-term reports are disabled.
- **Default long-term report time range** is the time span used for any long-term report where one is not specified.

Long-term Reports

NetFlow Tracker allows any report that can be created using the [Filter Editor](#) to be set up as a long-term report. A custom long-term report has a name, a report template and a type. It can also have its own storage settings overriding those in [Database Settings](#), a time mask and a filter.

The report type determines how it is accessed. A basic report is created across the entire system, and thus it is strongly recommended that it has a filter on at least [source device](#). A basic report can only be accessed from the [long-term filter editor](#).

A long-term report can also be created for each device in the system, or for each interface inbound or outbound. These reports can still have a filter or time mask applied if desired. A per-device, inbound interface or outbound interface report can be accessed from the [long-term filter editor](#) or by drilling down from the long-term [device or interface charts](#).

To create a custom long-term report, enter a name and select a report template and type and click “New...”. A new page for the report will be opened, allowing you to give the report non-default storage settings, a time mask and a filter. Click “Ok” to go back the main Report Settings page or “Delete” to cancel.

You can delete a long-term report or edit its name, storage settings and filter by clicking its name. It is not possible to change the report template, type or time mask of an existing report due to the way long-term data is stored.

Executive Reports

An executive report is a pre-configured template that contains one or more reports or charts and user-defined HTML content. They can be used to provide easy access to often-used reports or to group related reports together on one page.

To create an executive report, enter a name and click “New...”. You can edit an existing report by clicking its name.

The first part of defining an executive report is specifying the sub-reports that you would like to embed within it. These are specified by URL, and can be any type of real-time or long-term report you like. You will need to specify the report template or long-term report id and the output format (chart, pie, table). You can, and probably should, include [filter](#) and [relative time range](#) in the URLs, but note that if you omit them, they are taken from the URL passed to the executive report itself. If no time range is specified in the sub-report URL or the URL passed to the executive report, each sub-report will use the default real-time or long-term time range according to its type. Please be careful about using unfiltered sub-reports as they will be accessible from the [Executive Reports](#) homepage without a means of supplying a filter, and this could cause problems. Thus it is recommended that they are used only in conjunction with a portal system.

Once you have added the sub-report URLs to the executive report, you must then specify the report content. The executive report is made up of rows, and each row contains one or more cells. A cell can be configured to span a number of columns, allowing complex layouts. To add a row, click the “Add Row” button; you can then add cells to the row. There are two types of cells: report cells and HTML cells.

A report cell shows content from one of the sub-reports; the sub-report must be selected from the list provided. The output parameters are used to control what is displayed in the cell. The most important parameters used here are **sections** and **features**. See [Report Format Parameters](#) for more details.

If you have allowed either drilling down or the open in new window button for a report cell you must also specify new window parameters that are added to the URL created by combining the sub-report parameters and the output parameters for the cell. In most cases these will be “sections=&features=”. Note that not supplying a value for a parameter here will remove that parameter from the URL.

You can also include your own HTML content such as explanatory text or a company logo in an executive report using a HTML cell. You can include any HTML content you like, including links and images. You can include images stored in the “customweb” folder under NetFlow Tracker’s install folder; they are accessible as “customweb/<filename>.<ext>”.

A HTML cell has a CSS style that is used to control its appearance. You are likely to use “title”, which will produce a cell that looks like a report title; “repdesc” which will produce a cell that looks like a report time range and filter description; or “content” which will produce a cell that has a simple off-white background. You can use HTML constructs like “span” to apply styles to the text within the cell; if you use “repdesc” as the cell style you will probably need to enclose the text in a span with the style “repdesctext”.

An Example Executive Report – Top Applications Today and Last Week

This report contains two sub reports, one showing top applications for a device over the last 24 hours and the other over 7 days. The reports are shown as pie and time charts, and HTML cells are used to annotate the report.

Sub-reports

```
templid=0026&output=chart&nelements=5&chartWidth=400
&device=10.100.50.250&length=86400000
```

and

```
id=0102&output=chart&nelements=5&chartWidth=400&device=10.100.50.250
&unit=day&nunits=7
```

Note that the chart width is set in the URL to the sub-report; this is because the chart width is used to determine the sample size or source long-term data and if we were to simply control the size of the chart using the output parameters the samples may be an inappropriate size. Of course, it is possible to specify the [sample size](#) or [source long-term data](#) in the URL if desired.

Content

The first row consists of a single HTML cell containing a short description of the report.

Column Span	2
CSS Class	repdesc
HTML	Top applications on our Internet router over the last 24 hours and last seven days

The second row consists of a single HTML cell containing a title for the first sub-report.

Column Span	2
CSS Class	title
HTML	Last 24 Hours

The third row consists of two report cells, one containing a pie chart of the first sub-report and one containing a chart over time for the same sub-report. Each chart allows drilling down and opening in a new window.

Column Span	1
Sub-report	<code>templid=0026&...&length=86400000</code>
Output Parameters	<code>output=pie&chartWidth=300&sections=4&features=1536</code>
New Window Parameters	<code>output=pie&nelements=&chartWidth=&sections=&features=</code>

Column Span	1
Sub-report	<code>templid=0026&...&length=86400000</code>
Output Parameters	<code>sections=4&features=1536</code>
New Window Parameters	<code>nelements=&chartWidth=&sections=&features=</code>

The fourth row consists of a single report cell containing the chart legend for the first sub-report. No interactive features are supported.

Column Span	2
Sub-report	<code>templid=0026&...&length=86400000</code>
Output Parameters	<code>sections=16&features=0</code>
New Window Parameters	

The fifth row consists of a single HTML cell containing a title for the second sub-report.

Column Span	2
CSS Class	<code>title</code>
HTML	<code>Last 7 Days</code>

The sixth row consists of two report cells containing charts as above.

Column Span	1
Sub-report	<code>id=0102&...&unit=day&nunits=7</code>
Output Parameters	<code>output=pie&chartWidth=300&sections=4&features=1536</code>
New Window Parameters	<code>output=pie&nelements=&chartWidth=&sections=&features=</code>

Column Span	1
Sub-report	<code>id=0102&...&unit=day&nunits=7</code>
Output Parameters	<code>sections=4&features=1536</code>
New Window Parameters	<code>nelements=&chartWidth=&sections=&features=</code>

The seventh row consists of a single report cell containing the chart legend as above.

Column Span	2
Sub-report	<code>id=0102&...&unit=day&nunits=7</code>
Output Parameters	<code>sections=16&features=0</code>
New Window Parameters	

IP Application Names

NetFlow Tracker receives application information in the form of a protocol number and port number. These correspond directly to specific network applications. Many are predefined (well-known ports) while others (registered ports) are defined by the software manufacturer. NetFlow Tracker comes configured with the well-known ports as well as many others. You can edit this list yourself with this page. By default, ports below 1024 are not shown on this page as they normally don't need to be changed but, if required, these can be shown by clicking (more...) in the title of the Port column. A comprehensive list of all the well-known and registered ports is available at <http://www.iana.org/assignments/port-numbers>.

DiffServ Names

NetFlow Tracker can filter and report by differentiated service code point; you can assign names to each of the 64 code points here. The standard code point names are already configured.

Hostname Resolution Settings

This page lets you configure aspects of the resolution of hostnames for addresses encountered on reports. These are cached to increase reporting speed and reduce the amount of network traffic generated by the NetFlow Tracker when generating a report. You can change how long a resolved hostname is cached for, the default being 30 minutes, and how long a failure to resolve a hostname for a given address is remembered, the default being 10 seconds. You can also control the size of the cache and the number of threads used to resolve hostnames. If you find that hostname resolution is not working, click "Defaults" to put the settings back to useful default values. Click "Ok" to accept your changes or "Cancel" to abort.

Should you wish to clear the cache of resolved hostnames, disable resolution by unchecking "Enable hostname resolution" and clicking "Ok", then go back into the configuration page and enable resolution again by checking "Enable hostname resolution" and clicking "Ok".

AS Names

This page lets you assign names to AS numbers appearing in reports. AS numbers below 34816 are assigned by several agencies; NetFlow Tracker comes with many of these ASs already named. Numbers between 34816 and 64511 are held by the IANA and should not be used. Numbers above 64511 are for private use and can be named using this page. You can assign or edit the name for a public or reserved AS by clicking "(more...)" in the title of the AS column.

Subnet Names

This page lets you assign names to the IP subnets that appear in reports. The network mask length appearing in a network report is the one used by the router to route the traffic described, so you may need to configure names for subnets that overlap.

Database Settings

This page lets you improve the performance of reports and charts, and change the number of days for which data is retained.

- **Expect large result sets** controls the method by which the database server manipulates raw data. If you have a fast disk subsystem you should set this to “Always” to ensure reports over large amounts of data perform well. If you have a slower disk subsystem, lots of RAM and a relatively small amount of data, you might consider setting this to “Never”, but bear in mind that reports over large amounts of data may take considerably longer to run.
- **Maximum in-memory temporary table size** is the maximum amount of memory the database server will use during a query when it has been told not to expect a large result set. Increasing this will increase the amount of data that can be reported on with “Expect large result sets” set to “Never” before there is a significant drop in performance.
- **Sort buffer size** is the size of the buffer used to reduce the amount of disk seeks when sorting rows for grouping or final display. Increasing this will improve reporting speed, but you are unlikely to see much improvement for sizes above 128MB.
- **Store real-time data for** allows you to change the number of days full real-time data is stored for. You can reduce this to save disk space, or increase it if you are sure you have enough free space.
- **Store 10 minute, 1hour, etc. long-term data for** allows you to change how long the different types of long-term data are stored for. Each type of data allows a long-term chart to display blocks of that size; if the block size is not specified when opening a long-term report the closest available size to the ideal for the selected time range is chosen.
- **Use compression** to reduce the amount of disk space used, but note that it is likely to slow down your reports.

Archiving

NetFlow Tracker can be configured to archive real-time data older than the age configured in [Database Settings](#) to a nominated location rather than delete it. Archiving is enabled for a device in [Device Settings](#); the archiving settings page allows you to set the archive location and mount archived data back into the system for reporting using the [Filter Editor](#).

You can choose to have all archives stored in the archive folder, or you can choose to store in sub folders for each device and/or day, Please note that NetFlow Tracker does not delete archive files so you must ensure that they are moved from the archive directory to permanent storage.

To mount an archive, enter the directory containing it in the box under “Mount Archives” and press “List”; you can then select archives and press “Mount”. When there are archives mounted they appear under “Currently Mounted Archives” and can be unmounted by selecting and pressing “Unmount”. Note that mounting and unmounting archives does not affect the archive file itself.

Mounting an archive from a device that was deleted or was never present on the server is not supported.

Memory Settings

NetFlow Tracker uses a small amount of memory during its normal operation. You can control this amount by changing the values here, but it is not likely to be necessary. Note that it is possible to prevent the software from working by setting inappropriate values. Note also that this page is not available on Unix installations; to change the memory settings on Unix the “start” script must be edited.

Performance Counters

The performance counters can help diagnose problems setting up NetFlow Tracker. Counters are stored for each device the software has received data from. The counters are kept from when the system is started; you can reset them at any time.

NetFlow Data Received

This counter shows the number of exports and the amount of NetFlow data received by the software from each device. Note that this is not the amount of traffic described by the exports but the LAN traffic generated by the exports themselves.

Traffic Described

This counter keeps track of the total amount of network traffic across all interfaces in each direction described by NetFlow exports received from each device.

Ignored Flows

Flows are ignored if they arrive too late to be processed. If you see a large number of ignored flows you should ensure the inactive timeout or short aging time are correctly set as described in [Appendix 1](#).

Unprocessed Flowsets

NetFlow version 9 flows are encoded in a flexible manner using templates that are exported by the router every few seconds. For a period after starting NetFlow Tracker or after a router reboot, flows may be received without NetFlow Tracker knowing how to decode them.

Interface Scans

The software must scan the interface list of each device exporting to it whenever the device or the software is restarted. A large number of rescans, particularly failed ones, indicates a problem.

Missed Flows

NetFlow version 5 and 7 exports contain a sequence number to allow a NetFlow collector to detect when exports are missed. Exports can be missed due to network congestion or a busy router. If a switch or router is reordering the UDP packets containing NetFlow exports you will see missed flows being registered. Note that each export normally contains information on about 30 flows.

If the NetFlow Tracker server is under very heavy load it may drop packets itself. If you suspect this is happening, try increasing the receive buffer size in [Listener Ports](#).

Missed Exports

NetFlow version 9 exports contain a sequence number to allow a NetFlow collector to detect when exports are missed. Unlike the version 5 or 7 sequence number, this only allows the number of missed exports to be counted rather than the number of missed flows.

No Out Interface

The router sends flows with no out interface whenever an access control list lookup fails or whenever multicast traffic is routed. A high number of flows without out interfaces is normal.

No In Interface

If flows arrive with no in interface it may indicate a configuration problem on a Catalyst switch. Please contact technical support.

Appendix 1: Configuring NetFlow Data Export

This is a brief guide to setting up NetFlow on a Cisco routing or route-switching device. For more information on this subject, visit <http://www.cisco.com/go/netflow>. We recommend that only people with experience in configuring Cisco devices follow these steps. If in doubt, contact your network administrator or Cisco consultant. Note that if you are running hybrid mode on a Supervisor Engine you must configure both CatOS on the Supervisor Engine and IOS on the MSFC. If you are running Native IOS the commands are slightly different.

Configuring NetFlow Export on an IOS device

In configure mode on the router or MSFC, issue the following to enable NetFlow Export:

```
ip flow-export destination <address> 2055
```

Use the address of your NetFlow Tracker machine and one of the ports configured in the [Listener Ports](#) settings page. Port 2055 is monitored by default.

```
ip flow-export source loopback 0
```

The source interface is used to set the source IP address of the NetFlow exports sent by the router. NetFlow Tracker will make SNMP requests of the router on this address. If you experience problems you can set the source interface to an Ethernet or WAN interface instead of the loopback.

```
ip flow-export version 5 [peer-as | origin-as]  
or
```

```
ip flow-export version 9 [peer-as | origin-as]
```

This sets the export version. Version 5 and Version 9 both support all of the features NetFlow Tracker is capable of using; if you have a Native IOS switch you may need to use version 9 to work around a bug – this is described below. If your router uses BGP, you can specify that either the origin or peer ASs are included in exports – it is not possible to include both.

```
ip flow-cache timeout active 1
```

This breaks up long-lived flows into one-minute segments.

```
ip flow-cache timeout inactive 15
```

This ensures that flows that have finished are exported in a timely manner.

```
interface <interface>  
ip route-cache flow  
bandwidth <kbps>  
exit
```

You need to enable NetFlow on each interface through which traffic you are interested in will flow. This will normally be the Ethernet and WAN interfaces. You may also need to set the speed of the interface in kilobits per second. It is especially important to set the speed for frame relay or ATM virtual circuits.

```
ip cef
```

This enables Cisco Express Forwarding, which is required for NetFlow in most recent IOS releases.

show ip flow export

This will show the current NetFlow configuration. Issue this in normal (not configuration) mode.

show ip cache flow**show ip cache verbose flow**

These commands issued in normal mode summarise the active flows and give an indication of how much NetFlow data the router is exporting.

Configuring NetFlow Input Filters

IOS versions 12.2(25)S, 12.2(27)SBC and 12.3(4)T and greater support the NetFlow Input Filters feature, which can be used by NetFlow Tracker to report upon the traffic class used to route each flow.

flow-sampler-map allflows**mode random one-out-of 1****exit**

Create a flow sampler that exports every flow record.

policy-map netflowpolicymap**class <class>****netflow-sampler allflows****exit****exit**

Create a policy map containing NetFlow sampling actions; you must include each class that you would like information on.

interface <interface>**service-policy input netflowpolicymap****exit**

Associate the policy map with an interface; you must associate the policy map with each NetFlow-enabled interface that you would like traffic class information from.

Configuring NDE on a CatOS device

In privileged mode on the Supervisor Engine, issue the following to enable NDE:

set system name <name>

Set the name of your switch. Note that even if the prompt has been set to the name of the switch you still need this command.

set mls nde <address> 2055

Use the address of your NetFlow Tracker machine and one of the ports configured in the [Listener Ports](#) settings page. Port 2055 is monitored by default.

set mls nde version 7

This sets the export version. Version 7 is the most recent full export version supported by switches.

set mls agingtime long 64

This breaks up long-lived flows into (roughly) one-minute segments.

```
set mls agingtime 32
```

This ensures that flows that have finished are exported in a timely manner.

```
set mls flow full
```

This sets the flow mask to full flows. This is required to get useful information from the switch.

```
set mls bridged-flow-statistics enable <vlanlist>
```

CatOS 7.(2) or higher is required for this command, which enables NDE for all traffic within the specified VLANs rather than just inter-VLAN traffic.

```
set mls nde enable
```

This enables NDE.

```
show mls nde
```

```
show mls debug
```

These commands can help debug your NDE configuration.

Configuring NDE on a Native IOS device

In configure mode on the Supervisor Engine, follow the instructions for an IOS device above, and then issue the following to enable NDE:

```
mls netflow
```

This enables NetFlow on the supervisor.

```
mls nde sender version 5
```

or

```
mls nde sender version 7
```

This sets the export version. Due to several IOS bugs, the export version you must use on the supervisor is dependent on your hardware configuration and IOS version:

- Distributed Forwarding Cards and 12.1(13)EO3, 12.1(18.1)E, 12.2(13.6)S, 12.2(15.1)S, 12.2(17a)SX or above: use version 5. Note that this configuration will cause the [Performance Counters](#) to report missed flows that are not actually missed; this is the result of an IOS bug fixed in the SXF strains.
- Distributed Forwarding Cards and older than 12.1(13)EO3, 12.1(18.1)E, 12.2(13.6)S, 12.2(15.1)S or 12.2(17a)SX: this configuration will cause serious problems, so please contact Crannog Software if your device matches this description.
- No Distributed Forwarding Cards and 12.0(24)S, 12.2(18)S, 12.3(1) or above: use version 5 and configure the MSFC to export version 9 as described above.
- No Distributed Forwarding Cards and 12.1(13)EO3, 12.1(18.1)E, 12.2(13.6)S, 12.2(15.1)S, 12.2(17a)SX or above: use version 5.
- Anything else: use version 7. Note that version 7 may not include AS or subnet mask information.

```
mls aging long 64
```

This breaks up long-lived flows into (roughly) one-minute segments.

```
mls aging normal 32
```

This ensures that flows that have finished are exported in a timely manner.

```
mls flow ip interface-full
```

```
mls nde interface
```

or

```
mls flow ip full
```

If you have a Supervisor Engine 2 or 720 running IOS version 12.1.13(E) or higher the first two commands are required to put interface and routing information into the NetFlow Exports. This information is unavailable with any earlier IOS version on the Supervisor Engine 2 or 720.

If you have a Supervisor Engine 1 the third command is required to put full information into the NetFlow Exports.

```
ip flow ingress layer2-switched vlan <vlanlist>
```

```
ip flow export layer2-switched vlan <vlanlist>
```

A PFC3B or PFC3BXL running 12.2(18)SXE or higher is required for this command, which enables NDE for all traffic within the specified VLANs rather than just inter-VLAN traffic.

Configuring NetFlow Export on a 4000 series switch

The 4000 and 4500 series switches require a Supervisor IV with a NetFlow Services daughter card (WS-F4531) and IOS version 12.1(19)EW or above to support NetFlow. First configure the device as for an IOS device above, omitting the command **ip route-cache flow** on each interface, and then issue the following:

```
ip route-cache flow infer-fields
```

This ensures routing information is included in the flows.

Appendix 2: CSV File Format

Every standard chart and tabular report can be converted to comma-separated-value format for importing into a database server or spreadsheet.

Chart CSV format

Each section is separated by a row of “=” signs. The first section is the chart title; the second is the time range and filter.

Each subsequent section represents a single chart, equivalent to the tabs above the chart in an interactive chart. If a utilisation chart is present it will be included in the CSV file but with identical data to the traffic rate chart. The first line of the section is the name of the chart. The next two rows contain the start and end time of each sample in milliseconds UTC. Each has an empty column at the start to accommodate the description of each data row below. Each data row consists of a description followed by an octet or packet count for each sample.

Tabular report CSV format

Each section is separated by a row of “=” signs. The first section is the report title; the second is the time range and filter.

The third section starts with the title of each column, separated by a comma. Each subsequent line in the section is a row with each value separated by a comma, and text values contained within double quotes. There are several differences between a report viewed in a browser and one converted to CSV; in CSV format all rows are included, information normally available by hovering the mouse over a label is unavailable, and traffic and packets passed are output as simple counts rather than rates.

The fourth section contains column totals, again separated by commas. There will usually be empty values in the total row corresponding to non-numeric columns.

Appendix 3: Third Party Software Components

This product includes software developed by the Apache Software Foundation (<http://www.apache.org/>).

This product includes software developed by Advantys (<http://www.advantys.com>).

Jakarta Log4j

NetFlow Tracker includes Jakarta Log4j v1.1.3, available at <http://jakarta.apache.org/log4j/>. This is distributed under the Apache Software License, a copy of which is available at <http://www.apache.org/LICENSE>.

Jakarta Tomcat

NetFlow Tracker includes Jakarta Tomcat v3.3.2, available at <http://jakarta.apache.org/tomcat/>. This is distributed under the Apache Software License, a copy of which is available at <http://www.apache.org/LICENSE>.

joeSNMP

NetFlow Tracker includes joeSNMP v0.2.6, available at <http://www.opennms.org/files/releases/joeSNMP/>. This is distributed under the Lesser GNU Public License, a copy of which is available at <http://www.gnu.org/licenses/lgpl.html>.

jspSmartUpload

NetFlow Tracker includes jspSmartUpload v2.1, available at <http://www.jspsmart.com/>. This is distributed under the Advantys Freeware license contract, a copy of which is available at <http://www.jspsmart.com/liblocal/docs/legal.htm#free>.