NetFlow Tracker

User’s Guide Version 3.1

May 2007

CISCO SYSTEMS
Technology
Developer
Partner

Crannog Software NetFlow Tracker

Software License Agreement

This is a legal agreement between you (“You”/ “the End User””), and Fluke Corporation, a Washington corporation, its subsidiaries and affiliates, including Fluke Networks (“Fluke”), with offices at 6920 Seaway Boulevard, Everett, Washington, 98203, USA.

BY DOWNLOADING OR OTHERWISE ELECTRONICALLY RECEIVING THIS SOFTWARE PRODUCT (“PRODUCT”) IN ACCORDANCE WITH OUR SOFTWARE DELIVERY PROCEDURES OR BY OPENING THE SEALED DISK PACKAGE WHICH CONTAINS THE PRODUCT, YOU ARE AGREEING TO BE BOUND BY THE TERMS OF THIS AGREEMENT. IF YOU DO NOT AGREE TO THE TERMS OF THIS AGREEMENT, PROMPTLY DELETE THE DOWNLOADED OR ELECTRONICALLY RECEIVED SOFTWARE FROM YOUR COMPUTER SYSTEM AND NOTIFY US OF SAME IN ORDER TO CLAIM AND, IF YOU HAVE RECEIVED A SEALED CD-ROM PACKAGE, RETURN THE UNOPENED DISK PACKAGE AND THE ACCOMPANYING ITEMS (INCLUDING MANUALS) TO A FLUKE REPRESENTATIVE, FOR REFUND OF THE PRICE PAID.

1. GRANT OF LICENCE AND PAYMENT OF FEES

Provided that You have paid the applicable licence fee, Fluke grants You a non-exclusive and non-transferable, revocable licence to use one copy of the Product on the maximum number of servers and the maximum number of devices specified in your purchase order, or if not so specified, on a single server and a single device by a single user, and only for the purpose of carrying out your business in the country specified in your order. This Product is licensed for internal use by You, the end user only. The Product is not licensed for provision of a public service by You or for the provision of any fee generating service by You to a third party.

In the event that at any time You wish to extend the permitted number of servers or devices above the permitted amount, You must contact Fluke or the reseller from whom you purchased the Product (“the Reseller”) and an additional licence fee may be agreed and a new licence issued for the requested additional number of servers/devices.

Fluke or your Reseller may require that You provide written certification showing the geographical locations, type and serial number of all computer hardware on which the Software is being used, together with confirmation that the Product is being used in accordance with the conditions of this Agreement. You shall permit Fluke or your Reseller, and/or their respective agents to inspect and have access to any premises, and to the computer equipment located there, at or on which the Software is being kept or used, and any records kept pursuant to this Agreement, for the purposes of ensuring that the Customer is complying with the terms of this licence, provided that Fluke/your Reseller provides reasonable advance notice to the Customer of such inspections, which shall take place at reasonable times.

2. EVALUATION, UPDATES, UPGRADES AND SUPPORT AND MAINTENANCE

EVALUATION. If a provided licence key is labelled “Evaluation”, Fluke grants You the right to use the Product enabled by that key solely for the purpose of evaluation, and the Product will cease to function seven (7) days from enabling (or after such longer period as may be agreed by Fluke and confirmed by Fluke or your Reseller in writing), at which time the licence grant for that Product also ends. After the evaluation period, You may either purchase a full licence to use the Product from your Reseller or directly from Fluke, or You must promptly return to Fluke or cease to use the Evaluation Product and all associated documentation. The warranty set out in Clause 5 shall not apply in respect of Product downloaded for evaluation purposes.

UPDATES. Please refer to the release notes accompanying any new versions, updates or upgrades (“Updates”) prior to installation. Fluke will inform You or your Reseller of any Updates which it may develop from time to time and may licence any such Updates to You for a reasonable charge. To the extent that Fluke issues any Updates to You under the terms of this Agreement, any reference to the Product herein shall be deemed to include such Updates.

If You have purchased the maintenance and support services from Fluke then subject to payment of the support fees, Fluke shall provide such services in respect of the Product in accordance with the provisions of the Support and Maintenance Agreement contained in Appendix 1.

3. COPYRIGHT

All intellectual property rights in the Product belong to Fluke and You acknowledge that You have no ownership claims or rights whatsoever in the Product. You may (a) make one copy of the Product solely for backup or archival purposes and keep this securely, or (b) transfer the software to a secure single hard disk provided that You keep the original solely and securely for backup or archival purpose. You may not copy the written materials accompanying the Product. You shall not remove or alter Fluke’s copyright or other intellectual property rights notices included in the Product or in and any associated documentation. You must notify Fluke forthwith if You become aware of any unauthorized use of the Product by any third party.

4. OTHER RESTRICTIONS

You shall not sub-licence, distribute, market, lease, sell, commercially exploit, loan or give away the Product or any associated documentation. For the avoidance of doubt, this licence does not grant any rights in the Product to, and may not be assigned, sub-licenced or otherwise transferred to, any connected person, where the term connected person includes but is not limited to the End User’s subsidiaries, affiliates or any other persons in any way connected with the End User, whether present or future. The Product and accompanying written materials may not be used on more than the permitted number of servers at any one time or for in excess of the permitted number of devices. Subject always to any rights which You may enjoy under applicable law (provided that such rights are exercised strictly in accordance with applicable law) and except as expressly provided in this Agreement, You may not reproduce, modify, adapt, translate, decompile, disassemble or reverse engineer the Product in any manner. You shall not merge or integrate the Product into any other computer program or work, and You shall not create derivative works of the Product. Fluke reserves all rights not expressly granted under this Agreement.

5. LIMITED WARRANTY

Fluke warrants that during the warranty period (a) the Product will perform substantially in accordance with its accompanying written materials, and (b) the media on which the Product is furnished shall be free from defects in materials and workmanship. The warranty period applicable to the Product shall be ninety (90) days from the date of delivery of the Product or, if longer, the shortest warranty period permitted in respect of the Product under applicable law (“Warranty Period”). The warranty for any hardware accompanying the Product shall be as stated on the warranty card shipped with the hardware.

If, within the Warranty Period, You notify Fluke of any defect or fault in the Product in consequence of which the Product fails to perform substantially in accordance with its accompanying written materials, and such defect or fault does not result from You, or anyone acting with your authority, having amended, modified or used the Product for a purpose or in a context other than the purpose or context for which it was designed or licensed according to this Agreement, or as a result of accident, power failure or surge or other hazards, Fluke shall, at Fluke’s sole option and absolute discretion, do one of the following:

(i) repair the Product; or

(ii) replace the Product; or

(iii) repay to You all license fees which You have paid to Fluke under this Agreement.

Fluke does not warrant that the operation of the Product will be uninterrupted or error or interruption free.

6. CUSTOMER REMEDIES

You must call your Fluke representative for an authorization to return any item during the 90 day warranty period referred to in clause 5 above, and You will be supplied with a return authorisation number and an address for returning the item together with a copy of your receipt. You acknowledge that your sole remedy for any defect in the Product will be Your rights under clause 5.

7. NO OTHER WARRANTIES

FLUKE AND/OR ITS SUPPLIERS, DISCLAIM ALL OTHER WARRANTIES, EITHER EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WITH RESPECT TO THE PRODUCT, THE ACCOMPANYING WRITTEN MATERIALS AND ANY ACCOMPANYING HARDWARE AND YOU AGREE THAT THIS IS FAIR AND REASONABLE. THE EXPRESS TERMS OF THIS AGREEMENT ARE IN LIEU OF ALL WARRANTIES, CONDITIONS, UNDERTAKINGS, TERMS OF OBLIGATIONS IMPLIED BY STATUTE, COMMON LAW, TRADE USAGE, COURSE OF DEALING OR OTHERWISE, ALL OF WHICH ARE HEREBY EXCLUDED TO THE FULLEST EXTENT PERMITTED BY LAW.

8. NO LIABILITY FOR CONSEQUENTIAL DAMAGES

IN NO EVENT SHALL FLUKE AND/OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, CONSEQUENTIAL OR ECONOMIC LOSS OR DAMAGES WHATSOEVER OR FOR ANY LOSS OF PROFITS, REVENUE, BUSINESS, SAVINGS, GOODWILL, CAPITAL, ADDITIONAL ADMINISTRATIVE TIME OR DATA ARISING OUT A DEFECT IN THE PRODUCT OR THE USE OF OR INABILITY TO USE THE PRODUCT, EVEN IF FLUKE HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

9. TERMINATION

Either party shall be entitled forthwith to terminate this Agreement by written notice if the other Party commits any material breach of any of the provisions of this Agreement and, fails to remedy the same within sixty (60) days after receipt of a written notice from the non-breaching Party giving full particulars of the breach and requiring it to be remedied.

You shall be obliged to notify Fluke in writing of any change in the control or ownership of the End User and Fluke shall be entitled forthwith to terminate this Agreement by written notice.

This Agreement shall automatically terminate if replaced at any time with a new licence agreement.

The right to terminate this Agreement given by this clause 9 will be without prejudice to any other accrued right or remedy of either Party including accrued rights or remedies in respect of the breach concerned (if any) or any other breach, or which the Parties have accrued prior to termination.

10. INDEMNIFICATION

You shall indemnify Fluke in full and hold Fluke harmless in respect of any loss, damages, proceedings, suits, third party claims, judgements, awards, expenses and costs (including legal costs) incurred by or taken against Fluke as a result of the negligence, fault, error, omission, act or breach of You or of your employees, staff, contractors, agents or representatives or for any breach of this Agreement whatsoever by You.

Notwithstanding any other provision of this Agreement, the aggregate liability of Fluke for or in respect of all breaches of its contractual obligations under this Agreement and for all representations, statements and tortious acts or omissions (including negligence but excluding negligence causing loss of life or personal injury) arising under or in connection with this Agreement shall in no event exceed the licence fee paid by You pursuant to this Agreement prior to the date of the breach.

11. CONFIDENTIAL INFORMATION AND SECURITY

During and after this Agreement, the Parties will keep in confidence and use only for the purposes of this Agreement all Confidential Information. Confidential Information means information belonging or relating to the Parties, their business or affairs, including without limitation, information relating to research, development, Product, processes, analyses, data, algorithms, diagrams, graphs, methods of manufacture, trade secrets, business plans, customers, finances, personnel data, and other material or information considered confidential and proprietary by the Parties or which either Party is otherwise informed is confidential or might or ought reasonably expect that the other Party would regard as confidential or which is marked "Confidential". For the avoidance of doubt, You shall treat the Product and any accompanying documentation as Confidential Information. Confidential Information does not include any information (i) which one Party lawfully knew before the other Party disclosed it to that Party; (ii) which has become publicly known through no wrongful act of either Party, or either Parties' employees or agents; or (iii) which either Party developed independently, as evidenced by appropriate documentation; or (iv) which is required to be disclosed by law.

The Parties will procure and ensure that each of its employees, agents, servants, sub-contractors and advisers will comply with the provisions contained in this clause. If either Party becomes aware of any breach of confidence by any of its employees, officers, representatives, servants, agents or sub-contractors it shall promptly notify the other Party and give the other Party all reasonable assistance in connection with any proceedings which the other Party may institute against any such person. This clause 11 shall survive the termination of this Agreement.

notwithstanding the above confidentiality provisions, in accepting this licence agreement, You agree that, subject to any applicable data protection laws, Fluke may use your business name and logo for the purposes of marketing and promotion of the product and its business and You hereby grant Fluke a limited licence to use your business name and logo for these purposes.

12. EXPORT CONTROL

You shall be responsible for and agree to comply with all laws and regulations of the United States and other countries ("Export Laws") to ensure that the Product is not exported directly, or indirectly in violation of Export Laws or used for any purpose prohibited by Export laws.

13. GOVERNING LAW AND JURISDICTION

13.1 This Agreement and all relationships created hereby will in all respects be governed by and construed in accordance with the laws of the state of washington, united states of america, in respect of all matters arising out of or in connection with this agreement. The Parties hereby submit to the exclusive jurisdiction of the washington Courts. NOTHING IN THIS CLAUSE SHALL PREVENT FLUKE FROM TAKING AN ACTION FOR PROTECTIVE OR PROVISIONAL RELIEF IN THE COURTS OF ANY OTHER STATE.

14. MISCELLANEOUS

14.1 The provisions of clauses 3, 7, 8, 10, 11, 12, 13 and 14 and the obligation on you to pay the licence fee shall survive the termination or expiry of this Agreement.

14.2 This Agreement is personal to You and You shall not assign, sub-licence or otherwise transfer this Agreement or any part of your rights or obligations hereunder whether in whole or in part save in accordance with this Agreement and with the prior written consent of Fluke and You shall not allow the Product to become the subject of any charge, lien or encumbrance of whatever nature. Nothing in this Agreement shall preclude the Licensor from assigning the Product or any related documentation or its rights and obligations under this Agreement to a third party and You hereby consent to any such future assignment.

14.3 This Agreement and the Support and Maintenance Agreement supersede all prior representations, arrangements, understandings and agreements between the Parties herein relating to the subject matter hereof, and sets out the entire and complete agreement and understanding between the Parties relating to the subject matter hereof.

14.4 If any provisions of the Agreement are held to be unenforceable, illegal or void in whole or in part the remaining portions of the Agreement shall remain in full force and effect.

14.5 No party shall be liable to the other for any delay or non-performance of its obligations under this Agreement (save for your obligation to pay the fees in accordance with clause 1) arising from any cause or causes beyond its reasonable control including, without limitation, any of the following: act of God, governmental act, tempest, war, fire, flood, explosion, civil commotion, industrial unrest of whatever nature or lack of or inability to obtain power, supplies or resources.

14.6 A waiver by either party to this Agreement of any breach by the other party of any of the terms of this Agreement or the acquiescence of such party in any act which but for such acquiescence would be a breach as aforesaid, will not operate as a waiver of any rights or the exercise thereof.

14.7 No alterations to these terms and conditions shall be effective unless contained in a written document made subsequent to the date of the terms and conditions signed by the parties which are expressly stated to amend the terms and conditions of this Agreement.

Appendix 1 to End User Licence

Terms and Conditions for Fluke Support and Maintenance Service

1. Definitions

1.1 In this Agreement and in the Schedules hereto, save where the context so admits or requires, the following definitions shall have the following meanings:

"Intellectual Property Rights" includes, without limitation, copyrights, discoveries, concepts, domain names, patents, secret processes, database rights, technologies, know how, inventions, ideas, improvements, information, trade secrets, all copyright works, business methods, logos, designs, trademarks, service marks, topography and semi-conductor chip rights, business names, literary, dramatic, musical and artistic works anywhere in the world (whether any of the foregoing is registered or unregistered and including any application in relation to any of the aforesaid).

"Licence Agreement" means the product licence agreement under the terms of which the Product is licensed to You and which is entered into simultaneously with this Agreement.

"Retail Prices Index" means that the Consumer Price Index as published monthly by the Central Statistics Office or Ireland or any of its successors.

"Support Charges" shall mean the applicable annual support fee as published in Fluke's price list.

"Support Hours" means those hours specified in the Schedule during which Fluke shall provide the Support Services described in this Agreement.

"Support Services" shall mean the maintenance and support services provided by Fluke under the terms of this Agreement as detailed in the Schedule.

"Working Day" means any day other than Saturday or Sunday or a bank or a public holiday in Ireland.

Capitalised terms which are not defined herein shall have same meaning as under the Licence Agreement.

1.2 In the event of any inconsistency between the Schedule and any terms or provisions of any clause contained in this Agreement, the terms or provisions in the clause of the Agreement shall prevail.

1.3 Words in the singular shall include the plural and vice versa where the context so admits or requires and words importing one gender include every other gender.

1.4 The headings in this Agreement are for ease of reference only and do not form part of the contents of this Agreement and shall not affect its interpretation.

1.5 Save as provided for elsewhere in this Agreement, this Agreement including the Schedule represents the entire of the understanding of the parties concerning the subject matter hereof, viz, the provision of support and maintenance services by Fluke to You, and overrides and supersedes all prior promises, representations, understandings, arrangements, agreements, letters of intent or heads of agreement concerning the same which are hereby revoked by mutual consent of the parties.

1.6 No alterations to these terms and conditions shall be effective unless contained in a written document made subsequent to the date of the terms and conditions signed by the parties which are expressly stated to amend the terms and conditions of this Agreement.

1.7 The contents of the Schedule form an integral part of this Agreement and shall have as full effect as if it were incorporated in the body of this Agreement and the expressions "this Agreement" and "the Agreement" used in the Schedule shall mean this Agreement and any reference to "this Agreement" shall be deemed to include the Schedule.

2. Support Services

2.1 In consideration of the payment by You of the Support Charges, Fluke agrees to provide the Support Services.

3. Support Charges

3.1 You shall pay the Support Charges to Fluke annually in advance. The Support Charges shall be paid within 30 days after receipt of Fluke's invoice thereof. No Support Services will be provided until payment in full has been received by Fluke. In the event of late payment, interest shall be charged at the rate of interest referred to in the European Communities (Late Payment in Commercial Transactions) Regulations 2002, from the date of invoice until the date of actual payment, such interest to accrue daily and both before and after judgement.

3.2 The Support Charges (including the charges for support outside of the Support Hours) may at Fluke's sole discretion be increased annually in accordance with the annual increase in the Retail Prices Index.

3.3 The Support Charges payable under the terms of this Agreement are related to the Support Services specified in Schedule 2. Additional support is subject to Fluke's then standard rates.

3.4 All Support Charges referred to in this Agreement are exclusive and net of any taxes, duties or such other additional sums which shall be paid by You including, but without prejudice to the generality of the foregoing, VAT, excise tax, tax on sales, property or use, import or other duties whether levied in respect of this Agreement, the Support Services or otherwise.

4. Undertakings by You

You undertake:

4.1 To maintain accurate and up to date records of the number and location of all copies of the Product supplied to You under the terms of the Licence Agreement and in relation to the numbers of users of such.

4.2 To co-operate with Fluke's personnel in the diagnosis of any error or defect in the Product or Updates reported by You.

4.3 To make available to Fluke, all reasonable information, facilities, services and access required by Fluke in order to perform the Support Services.

5. Supplier's Undertakings

5.1 Fluke shall use its reasonable commercial endeavours to ensure that the Support Services will be performed in such a way as to cause only minimal interruptions to your business processes (other than any pre-agreed unavoidable interruption which in Fluke's sole discretion is required in order to perform the Support Services in a proper and efficient manner).

5.2 Fluke shall use its reasonable commercial endeavours to ensure that the Support Services are performed with reasonable skill and care.

5.3 The express terms of this Agreement are in lieu of all warranties, conditions, undertakings, terms of obligations implied by statute, common law, trade usage, course of dealing or otherwise, all of which are hereby excluded to the fullest extent permitted by law.

5.4 Without prejudice to the generality of clause 5.3 and for the avoidance of doubt, to the fullest extent permitted by law all terms implied by Sections 13, 14 and 15 of the Sale of Goods Act, 1893 are hereby excluded and all terms implied by the Sale of Goods and Supply of Service Act, 1980 including, without prejudice to the generality of the foregoing, Section 39, are hereby excluded and the parties agree that this is fair and reasonable.

6. Limitation of Liability and indemnity

6.1 You shall indemnify Fluke in full and hold Fluke harmless in respect of any loss, damages, proceedings, suits, third party claims, judgements, awards, expenses and costs (including legal costs) incurred by or taken against Fluke as a result of the negligence, fault, error, omission, act or breach of You or of your employees, staff, contractors, agents or representatives or for any breach of this Agreement whatsoever by You.

6.2 In no event will Fluke be liable to You for any special, incidental, indirect, punitive or consequential loss or damages, any loss of business, revenue or profits, loss of use, loss of data, loss of savings or anticipated savings, loss of investments, loss of goodwill, capital costs or loss of extra administrative cost, whether occasioned by the negligence, fault, error, omission, act or breach of the Fluke, its employees, contractors or sub-contractors whether or not foreseeable, arising out of or in connection with this Agreement, whether in an action based on contract, equity or tort including negligence or other legal theory.

6.3 Notwithstanding any other provision of this Agreement, the aggregate liability of Fluke for or in respect of all breaches of its contractual obligations under this Agreement and for all representations, statements and tortious acts or omissions (including negligence but excluding negligence causing loss of life or personal injury) arising under or in connection with this Agreement shall in no event exceed the Support Charges paid by You pursuant to this Agreement prior to the date of the breach.

7. Intellectual Property Rights

7.1 Ownership of all Intellectual Property Rights in the Product and any accompanying documentation is governed by the provisions of the Licence Agreement.

8. Termination

8.1 You can terminate this Agreement at any time after the first anniversary of this Agreement by giving to Fluke not less than 90 days' written notice.

8.2 Either Party may terminate this Agreement by written notice to the other Party where:

8.2.1 the other party has committed a material breach of the terms or conditions of this Agreement including the terms, conditions and provisions of the Schedule and where the breaching party has failed to remedy such breach within sixty (60) days after receiving written notice from the non-breaching party requiring it so to do; and

8.2.2 the other party makes any arrangement or composition with its creditors or pass a resolution or where a Court shall make an order that the defaulting party shall be wound up (save and excepting only a member's winding up for the purposes of reconstruction or amalgamation to which the other party has been approved in writing prior to such) or where an examiner or a receiver or a liquidator is appointed over the other a Party's business.

8.3 On termination of this Support Agreement all rights and obligations of the parties under this Support Agreement shall automatically terminate except for any rights of action which may have accrued prior to termination and any obligations which expressly or by implication are intended to commence or continue in effect on or after termination.

9. Confidential Information and Security

9.1 During and after this Agreement, the Parties will keep in confidence and use only for the purposes of this Agreement all Confidential Information. Confidential Information means information belonging or relating to the Parties, their business or affairs, including without limitation, information relating to research, development, Product, processes, analyses, data, algorithms, diagrams, graphs, methods of manufacture, trade secrets, business plans, customers, finances, personnel data, and other material or information considered confidential and proprietary by the parties or which either party is otherwise informed is confidential or might or ought reasonably expect that the other party would regard as confidential or which is marked "Confidential". Confidential Information does not include any information (i) which one party knew before the other party disclosed it to that party; (ii) which has become publicly known through no wrongful act of either party, or either parties' employees or agents; or (iii) which either party developed independently, as evidenced by appropriate documentation; or (iv) which is required to be disclosed by law.

9.2 The Parties will procure and ensure that each of its employees, agents, servants, sub-contractors and advisers will comply with the provisions contained in this clause.

9.3 If either Party becomes aware of any breach of confidence by any of its employees, officers, representatives, servants, agents or sub-contractors it shall promptly notify the other Party and give the other Party all reasonable assistance in connection with any proceedings which the other Party may institute against any such person.

9.4 This clause shall survive the termination of this Agreement.

10 Miscellaneous

10.1 This Agreement is personal to You and You shall not assign, sub-licence or otherwise transfer this Agreement or any part of its right or obligations hereunder whether in whole or in part without the prior written consent of Fluke. Nothing in this Agreement shall preclude Fluke from assigning or sublicensing its rights and obligations under this Agreement.

10.2 If any provisions of the Agreement are held to be unenforceable, illegal or void in whole or in part the remaining portions of the Agreement shall remain in full force and effect.

10.3 No Party shall be liable to the other for any delay or non-performance of its obligations under this Agreement (save for the obligation of You to pay the Support Charges in accordance with clause 3) arising from any cause or causes beyond its reasonable control including, without limitation, any of the following: act of God, governmental act, tempest, war, fire, flood, explosion, civil commotion, industrial unrest of whatever nature or lack of or inability to obtain power, supplies or resources.

10.4 A waiver by either party to this Agreement of any breach by the other party of any of the terms of this Agreement or the acquiescence of such party in any act which but for such acquiescence would be a breach as aforesaid, will not operate as a waiver of any rights or the exercise thereof.

10.5 No alterations to these terms and conditions shall be effective unless contained in a written document made subsequent to the date of the terms and conditions signed by the parties which are expressly stated to amend the terms and conditions of this Agreement.

10.6 This Agreement and all relationships created hereby will in all respects be governed by and construed in accordance with the laws of Ireland in respect of all matters arising out of or in connection with this agreement. The Parties hereby submit to the exclusive jurisdiction of the Irish Courts. NOTHING IN THIS CLAUSE SHALL PREVENT FLUKE FROM TAKING AN ACTION FOR PROTECTIVE OR PROVISIONAL RELIEF IN THE COURTS OF ANY OTHER STATE.

Schedule

Support Services

1. Support Hours

The Support Hours during which Fluke shall supply the Support Services shall be between 9.30am and 5pm on Working Days.

2. Support Services

Fluke shall provide You during the Support Hours with:

2.1. technical advice and assistance by telephone, facsimile, e-mail or other electronic means as shall be necessary to resolve your difficulties and queries in relation to the Product and the Updates which You may require;

2.2. an error correction and problem solving service as follows:

if You shall discover that the then current supported version of Product fails to conform with any part of the description of the Product provided to you by Fluke then Fluke, on receiving notification of the error, shall use its reasonable endeavours to:

2.2.1 diagnose and resolve the reported error or problem; and

provide the required solution to remedy or correct the error or problem; and

2.2.3 provide You with all assistance reasonably required by You to enable You to implement the error correction supplied as soon as possible; and

2.2.4 correct errors by "fix" where Fluke, in its sole discretion, considers such to be appropriate.

2.3 Response times to technical advice and assistance queries and reported errors and problems are set out in clause 3 below.

2.4 Remote connection support shall only be provided by Fluke in the event that telephone, fax or email support does not resolve a problem.

3. Response Times

3.1 In the event of any problem arising in relation to the Product's installation and functioning, Fluke shall respond within 8 Support Hours after the logging of such an incident by You provided that the incident was logged by You during normal Support Hours. Fluke shall in turn endeavour to resolve the problem as soon as possible.

4. Exceptions to Support Services

4.1 The Support Services described in clause 2 of this Schedule shall not include service in respect of:

4.1.1 defects or errors resulting from any modifications of the Product or Updates made by any person other than Fluke;

4.1.2 incorrect use of the Product or Updates or operator error;

4.1.3 any fault in Your hardware, computer equipment or in any programs used in conjunction with the Product or Updates; or

4.1.4 defects or errors caused by the use of the Product or Updates on or with equipment or programs not approved by Fluke.

Contents

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.

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.
- 1GB 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.

Operating System Support

- Microsoft Windows™ 2000 or above; server versions will provide better performance due to more advanced disk caching and memory management.
- Solaris™ 8 or above (Intel™-compatible or Sparc™ processors).
- Any modern Linux distribution capable of running Java 1.4.2 and MySQL™ 5.0 (Intel-compatible processors).

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 Fluke Networks' 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 Fluke Networks' 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 Fluke Networks. 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 Fluke Networks' web site. Please contact support@flukenetworks.com for more details.

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 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 summarized 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:

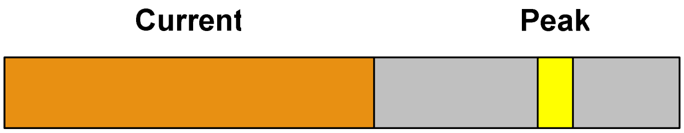

Each chart is scaled relative to the busiest device; this ensures that a high value on a chart indicates a relatively high traffic or packet rate.

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 utilization 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:

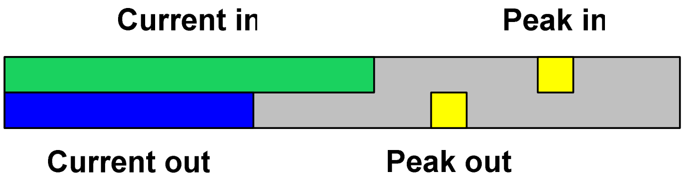

The scale of the chart depends on which column it is in; the “% Utilisation” column scales each row of each chart according to the configured speed of the interface in that direction whereas the “Relative Traffic” and “Relative Packets” are scaled relative to the busiest direction of the busiest interface. This ensures that a high value on a chart indicates either high utilization 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:

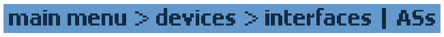

main menu > devices > interfaces | ASs

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

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:

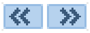

« »

Note that when you open a device or interface chart from the device or interface lists it will automatically to 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:

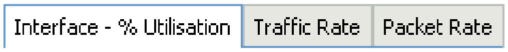

Interface - % Utilisation | Traffic Rate | Packet Rate

In this case, the utilization 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 center 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 center 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 utilization

If the source data for a report is filtered by interface, the total utilization 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:

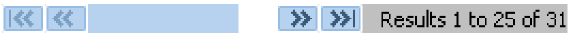

The buttons to the left of the scrollbar move to the first page of the report and back one page, respectively. Since the first page of the report is shown already, these buttons are unavailable. The buttons to the right of the scrollbar move forward one page and to the last page respectively. Clicking 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 rows 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…”:

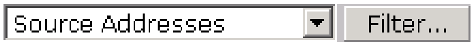

Thus if you are looking at a report of source applications, you can select an application and view a report of source addresses using that application.

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.
- **Bi-directional Address Pairs** – adds extra columns showing the traffic and packets sent from destination to source for each address pair.
- **Source Address Dissemination** – shows 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** – shows 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** – shows the pairs of connected endpoints that exchanged most traffic or packets. A single conversation represents, for example, a web browser downloading a single image.
- **Bi-directional Conversations** – adds extra columns showing the traffic and packets sent from destination to source for each conversation.
- **Source Endpoints** – shows 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** – shows the IP addresses and corresponding applications that were the destination of most traffic or packets.

- **Server-Client Sessions** – shows 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** – shows 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** – shows the ToS levels with most traffic or packets.
- **Differentiated Services** – shows the DiffServ code points with most traffic or packets.

Network Reports

- **Source ASs** – shows the autonomous systems that were the source of most traffic or packets. Note that a switch does not know anything about ASs.
- **Destination ASs** – shows the autonomous systems that were the destination of most traffic or packets.
- **AS Pairs** – shows the pairs of connected ASs that exchanged most traffic or packets.
- **Bi-directional AS Pairs** – adds extra columns showing the traffic and packets sent from destination to source for each AS pair.
- **Source Networks** – shows 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** – shows the IP subnets that were the destination of most traffic or packets.
- **Network Pairs** – shows the pairs of connected IP subnets that exchanged most traffic or packets.
- **Bi-directional Network Pairs** – adds extra columns showing the traffic and packets sent from destination to source for each network pair.

Interface Reports

- **In Interfaces** – shows 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** – shows 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.
- **VPNs** – shows the VPNs with most traffic or packets. Interfaces must be associated with VPNs in Device Settings for this report to function.
- **Next Hops** – shows the next-hop addresses that received most traffic or packets. Note that only a router can supply a next-hop address.

Traffic Identification Reports

- **Identified Applications** – shows the identified applications with most traffic or packets. See Device Settings for more information.
- **Traffic Classes** – shows the traffic classes that with most traffic or packets; see Device Settings for more.

Other Reports

- **Total** – shows just the total traffic and packets passing the filter.

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 hidden to save space and bandwidth; you can add a filter to the page by selecting it and clicking "Add".

Each filter 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. You can choose to include or exclude the items you select.

There are several ways to select items, depending on the type of the filter. Named items can be selected by highlighting them in the left-hand "Available" box and using the ">" button to move them to the right-hand "Selected" box. To deselect items highlight them in the right-hand box and click "<".

Some filters additionally allow you to manually enter an item in a box above the selected box and click "Add"; to deselect an item added in this way click "<" as before. If a filter does not have any named items there is a single selected box and a "Remove" button.

Some filters allow a range of items to be added; in this case enter the start and end of the range in the boxes provided. To select a single item, leave the right-hand box empty,

If you are logged in as an administrator or not logged in you can save a filter by clicking "Save…" at the bottom of the page. This allows you to assign a name to a set of filters for re-use later. If you have any saved filters defined they are available in the same dropdown list used to add filters to the page. Saved Filters are managed in Report Settings.

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.

Sample Size

NetFlow Tracker picks an optimal sample size for a real-time chart based upon the amount of time covered; you can override this by selecting a number of units. For example, you can create a report covering a day with each sample being an hour long.

Source Data

Long-term data is stored in samples of various sizes that are optimal for different lengths of chart; you can override the automatic selection of the source data to create charts showing, for example, a month in day-long blocks.

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.

Length

Instead of specifying a start and end time you can specify a length in units; the report will cover that number of units and end at the last full unit before the time it is opened.

Reload Interval

If you have selected a unit length or a time range that extends into the future you may want the report to refresh automatically to show new data; if so, enter the number of seconds between automatic refreshes here.

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":

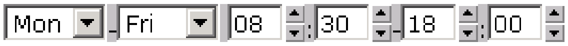

Mon - Fri 08 : 30 - 18 : 00

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.

In VPN

The in VPN filter restricts a report to just traffic where the inbound interface is part of the selected VPN(s). Interfaces must be associated with VPNs in Device Settings for this filter to function.

Out VPN

The out VPN filter selected traffic where the outbound interface is part of the selected VPN(s).

VPN

The VPN filter selects traffic where either interface is part of the selected VPN(s).

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 summarized in advance so a long-term report over several days or weeks can often be much faster than an equivalent real-rime 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:

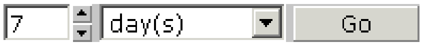

You can select any number of full minutes, hours, days, weeks, months, quarters, half-years or years. Note that if you zoom in to or out of a long-term chart, or drill down into a selection (other than one selected using the Select All button), the time range selector will not be available on the resulting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

`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. This parameter should not be used in conjunction with **`id`** or **`cid`**.

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
0029	Bi-directional Address Pairs
0030	Bi-directional Conversations
0031	Bi-directional AS Pairs
0032	Bi-directional Network Pairs
0033	Total
0034	VPNs
_flows	Full flows

id – specifies the long-term report to open. 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. This parameter should not be used in conjunction with **templid** or **cid**.

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
<id>	A custom long-term report ID

cid – specifies the executive report to open. The ID for an executive report is available in Report Settings. This parameter should not be used in conjunction with **templid** or **id**.

<id>	An executive report ID

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

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

<number>	The number of rows per page
-1	Show all rows

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.

`<heading>`	The URL-encoded column heading; note that % is URL-encoded as %25
`–<heading>`	A column to make invisible; parameters specifying invisible columns cannot be mixed with those specifying visible columns

`nelements` – specifies the number of elements to chart.

`<number>`	The number of elements to chart

`chartTitle` – specifies the chart to show.

`<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.

<table>
<tr><td rowspan="7"><b><code>&lt;sections&gt;</code></b></td><td colspan="2">The sections, formed by summing the values for each section</td></tr>
<tr><td><b>1</b></td><td>Title</td></tr>
<tr><td><b>2</b></td><td>Time range & filter description</td></tr>
<tr><td><b>4</b></td><td>Main report or chart body</td></tr>
<tr><td><b>8</b></td><td>Chart title, if applicable</td></tr>
<tr><td><b>16</b></td><td>Chart legend, if applicable</td></tr>
<tr><td><b>32</b></td><td>Result information, if applicable</td></tr>
<tr><td><b><code>–&lt;sections&gt;</code></b></td><td colspan="2">The sections that are not displayed</td></tr>
</table>

features – specifies the available 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
-<features>	The features that are not displayed	

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 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
`mon`	Weeks starting on a Monday
`tue`	Weeks starting on a Tuesday
`wed`	Weeks starting on a Wednesday
`thu`	Weeks starting on a Thursday
`fri`	Weeks starting on a Friday
`sat`	Weeks starting on a Saturday
`sun`	Weeks starting on a Sunday
`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 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 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
`mon`	Weeks starting on a Monday
`tue`	Weeks starting on a Tuesday
`wed`	Weeks starting on a Wednesday
`thu`	Weeks starting on a Thursday
`fri`	Weeks starting on a Friday
`sat`	Weeks starting on a Saturday
`sun`	Weeks starting on a Sunday
`month`	Months
`quarter`	Quarters
`halfyear`	Half-years
`year`	Years

`sdate_unit` – (optional) specifies the unit to measure how long before the report is generated the time range starts. Format as for **`date_unit`** above.

`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_unit` – (optional) specifies the unit to measure how long before the report is generated the time range end. Format as for **`date_unit`** above.

`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 the first day of the unit following the current unit at the time of report generation
`1`	The last day of the time range will be the first day of the current unit at 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.

<day1>-<day2>/<time1>-<time2>	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**

timemask_exclude=true – specifies that the supplied time masks are excluded from the time range rather than included in it.

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.

`sf` – specifies a saved filter to apply to the report. The ID for a saved filter is available in Report Settings.

`<id>`	A saved filter ID

`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.

`invpn` – specifies a VPN that the input interface must be part of.

`<name>`	The VPN name; see Device Settings for more information
`<id>`	The VPN identifier

`outvpn` – specifies a VPN that the output interface must be part of. Format as for **`invpn`** above.

`vpn` – specifies a VPN that either interface must be part of. Format as for **`invpn`** above.

`srcaddr` – specifies an acceptable source address.

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

`srcaddr_exclude=true` – specifies that the supplied source addresses are excluded rather than included.

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

dstaddr_exclude=true – specifies that the supplied destination addresses are excluded rather than included.

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

addr_exclude=true – specifies that the supplied source or destination addresses are excluded rather than included.

proto – specifies an acceptable IP protocol.

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

proto_exclude=true – specifies that the supplied protocols are excluded rather than included.

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**
<name>	The name of a grouped application

srcappl_exclude=true – specifies that the supplied source applications are excluded rather than included.

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

dstappl_exclude=true – specifies that the supplied destination applications are excluded rather than included.

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

appl_exclude=true – specifies that the supplied source or destination applications are excluded rather than included.

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

recappl_exclude=true – specifies that the supplied recognised applications are excluded rather than included.

`applid` – specifies an acceptable identified application.

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

`applid_exclude=true` – specifies that the supplied identified applications are excluded rather than included.

`tos` – specifies an acceptable Type-of-Service value.

`<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

`tos_exclude=true` – specifies that the supplied Type-of-Service values are excluded rather than included.

`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`

`ds_exclude=true` – specifies that the supplied differentiated service codepoints are excluded rather than included.

`class` – specifies an acceptable traffic class.

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

`class_exclude=true` – specifies that the supplied traffic classes are excluded rather than included.

`srcas` – specifies an acceptable source autonomous system number.

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

`srcas_exclude=true` – specifies that the supplied source autonomous system numbers are excluded rather than included.

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

dstas_exclude=true – specifies that the supplied destination autonomous system numbers are excluded rather than included.

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

as_exclude=true – specifies that the supplied source or destination autonomous system numbers are excluded rather than included.

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**

srcnet_exclude=true – specifies that the supplied source subnets are excluded rather than included.

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

dstnet_exclude=true – specifies that the supplied destination subnets are excluded rather than included.

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

net_exclude=true – specifies that the supplied source or destination subnets are excluded rather than included.

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

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

srcmask_exclude=true – specifies that the supplied source subnet masks are excluded rather than included.

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

dstmask_exclude=true – specifies that the supplied destination subnet masks are excluded rather than included.

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

mask_exclude=true – specifies that the supplied source or destination subnet masks are excluded rather than included.

nexthop – specifies a next-hop address.

<addr>	The address in dotted-decimal format

nexthop_exclude=true – specifies that the supplied next-hop addresse are excluded rather than included.

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.

aclvpn – specifies a permitted VPN. Format as for **invpn** 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.

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.

`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
`29`	VPN
`30`	In VPN
`31`	Out VPN

`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.

A BGP device may be configurable to send the BGP next-hop address in its NetFlow exports; if this is the case you will have the option to store this value in place of the IP next-hop for the device.

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 the device offers enough information via SNMP or other means to automatically detect the name of each traffic class the "Automatically map traffic classes" option will be available and checked; it is recommended that you leave this setting as it is. If you uncheck this option or it is not available for a device, you must add each traffic class to NetFlow Tracker if it is not already added 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 mappings from the number the device uses to identify a traffic class to 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. Similar to traffic classes, you can choose to disable automatic mapping of identified applications; this is not recommended.

VPNs

NetFlow Tracker can associate an interface on a device with a VPN for reporting and filtering. Any number of interfaces on any number of devices can be associated with a single VPN, and their traffic will be grouped together in the VPNs report and by the VPN filters. NetFlow Tracker will assign the customer-facing interfaces of an MPLS PE router using MPLS VPN and supporting the standard SNMP MIB automatically; you can override this or assign interfaces manually by first clicking "add/delete" in the heading of the VPN column of the interfaces box for any device. Each VPN must have a unique id and name; a description is optional. To set the VPN for an interface, simply click the VPN name and choose another in the dropdown box that appears. You can set the VPN to "none" if the interface is not part of a VPN; the P interface(s) on an MPLS PE router should have their VPN set to "none" as they carry traffic from multiple VPNs.

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 an 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.
- **Work around "click to activate"** enables or disables the work around for the "click to activate and use this control" message that appears over the chart applets in Internet Explorer and Opera. Some combinations of operating system, browser and Java plug-in do not work correctly when this is enabled; if you notice that the applets do not show up or drilling down does not work you should try disabling this work around,
- **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.

Saved Filters

Saved filters can be defined that can be added wherever a filter editor appears in the software. A saved filter allows you to attach a name to, for example, a time-of-day mask or a filter that selects traffic related to a particular multi-port application or group of servers.

To create a saved filter, type a name in the box and click “New…”, then use the provided Filter Editor to define the filter. You can copy an existing filter by clicking the 🗐 icon, and you can change the order in which saved filters appear by clicking the ↑ and ↓ icons. To edit or delete a filter click its name.

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. To add a sub-report, give it a name that will identify it when you layout the executive report and select whether it is a real-time or long-term report. You can then use the provided filter editor to define the report. You can add custom parameters to alter anything about the report not configurable using the filter editor; see Report Format Parameters for more.

If you select “Default/Custom” as the time range of the report and do not add custom time range parameters the time range used will be whatever is passed to the executive report itself, or the default real-time or long-term time range according to the report.

Also note that any filters passed to the executive report are applied to the sub-reports in conjunction with whatever filters they have themselves. 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 sub-reports 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: sub-report cells and HTML cells.

A sub-report cell shows content from one of the sub-reports; the sub-report must be selected from the list provided. If the sub-report is a chart over time you can choose to output a pie chart instead; this is used in the example report below.

You should then select which sections of the sub-report you would like the cell to display, and which user-interface controls should be enabled for it. You can also select which columns to show, and if the sub-report is a chart or pie chart you can select which chart to show. Custom parameters can be supplied; see Report Format Parameters for more about which parameters are acceptable here.

If you have allowed either drilling down or the open in new window button for a report cell you must also specify how the URL is modified to create the new window. You can choose to show all sections and columns and allow all controls; this is usually the case for a complicated layout, You can also specify custom parameters. Note that you can remove a parameter from the new window's URL by giving it a blank value.

A HTML cell allows you to add your own HTML content, such as explanatory text or a company logo, to an executive report. 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. Three standard styles are offered – "Report Title" produces a cell that looks exactly like a report title, "Report Description" one with the blue background of a report's time range and filter description and "Content Cell" one with a simple white background. If you use "Report Description" as the cell style you will probably need to enclose the text in HTML tags as follows:

```
<span class="repdesctext">Test</span>
```

You can control the layout of the report by moving rows up and down and cells left and right within their rows. Complex layouts can be created by making cells span multiple columns; the |↔| and →||← buttons make a cell one column wider and narrower respectively. Finally, a cell or row can be deleted by clicking the ✕ button.

An Example Executive Report – Top Applications Today and This 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

Add a real-time sub-report with a tag of "Today". Use the filter editor to make it a Recognised Applications chart, and select the device(s) and any other filters you like. Make the length of the report 24 hours. Finally, add two custom parameters: **`nelements=5`** and **`chartWidth=400`**.

Add a long-term sub-report with a tag of "This Week". Select the correct long-term chart; you may need to define a custom long-term report to incorporate your desired filter. Make the length of the report 7 days, and add **`nelements=5`** and **`chartWidth=400`** as custom parameters.

Note that the chart width is set in the sub-report and not as a custom output parameter in a sub-report cell; 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.

Content

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

Click "Add Row" to add a row, then select "HTML" and click "Add Cell". Choose "Report Description" as the CSS class, and enter the following as the HTML:

```
<span class="repdesctext">Top applications on our Internet router
over the last 24 hours and last seven days</span>
```

You should change the text to reflect the filter applied to your sub-reports. After clicking "Ok", click |◂▸| to make the cell cover two columns.

The second row consists of a single HTML cell containing a title for the first sub-report. This time, choose "Report Title" as the CSS Class, and enter "Last 24 Hours" as the HTML. Again, make the cell cover two columns.

The third row consists of two sub-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.

For the each cell, select "Today" as the sub-report, "Results/Chart" as the only section, and "Open in a New Window" and "Drilldown" as the controls. Add **`nelements`** and **`chartWidth`** as custom new window/drilldown parameters, both with no values, so the reports resulting from drilling down or opening a cell in a new window are the default size and show the default number of elements.

To make the first cell display as a pie chart, check "Output as a Pie Chart". Also, add **`chartWidth=300`** as a custom output parameter to make the pie chart look better.

The fourth row consists of a single sub-report cell containing the chart legend for the first sub-report. No interactive controls are supported. Simply select "Today" as the report, "Legend" as the only section, and deselect all controls. Don't forget to make the cell cover two columns.

Finally, the fifth, sixth and seventh rows are the same as the second, third and fourth; however, the title HTML should be "Last 7 Days" and the sub-report "This Week" for all three sub-report cells. The seventh row consists of a single report cell containing the chart legend as above.

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.

If an application uses multiple ports or a range of ports you can define it as a grouped application. Grouped applications appear as one entry in the application reports, regardless of context. You may find that a saved filter is more useful as a means of defining, for example, the traffic relating to a networked application running on a cluster of servers.

To define a grouped application you must first give it a unique identifier and a name; you can then add application ports and ranges of ports to it.

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 clearing “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.
- **Hold back real-time data for** determines the number of seconds after its end each one-minute sample of real-time data is held in RAM before being committed to disk. You may need to increase this to avoid ignored flows.
- **MySQL can not access temporary files** should be unchecked to improve the performance of inserts database. However, it is possible that on Unix the user the NetFlow Tracker user runs as has a umask that creates temporary files that MySQL cannot read; in this case you must check this setting.
- **Number of threads to use to generate a report** controls the number of threads used to generate real-time charts over time and pie charts. You should not set this to more than the number of CPU cores in your system and are unlikely to see any benefit beyond 4.
- **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.

Backup

NetFlow Tracker can back up its configuration, and optionally its long-term and real-time databases, to a nominated folder on demand or on a schedule. The contents of the folder are erased before the backup, so ensure that you move scheduled backups to long-term storage if required. It may be advisable to schedule a backup to different locations on alternate days.

Backing up the real-time database takes a long time and it is advisable to omit it on a busy system unless it is essential.

To restore a backup you must first install exactly the same version as you had previously – you may need to contact support@flukenetworks.com to obtain this. Then, open a command prompt and issue the following commands on Windows, replacing paths as appropriate; <enter> means to press the enter or carriage return key:

```
c: <enter>

cd \nftracker <enter>

runany c:\nftracker c:\progra~1\java\j2re14~1.2_0
com.crannogsoftware.ulysses.CRestore -sourcefolder c:\backup <enter>
```

On Unix, issue the following commands in a terminal, again replacing paths as appropriate:

```
cd /usr/local/nftracker <enter>

./runany com.crannogsoftware.ulysses.CRestore -sourcefolder c:\backup
<enter>
```

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.

Average sample storage duration

This keeps track of how long it takes the system to store a one-minute sample of real-time data. If this takes more than about fifteen seconds it is a sign that the system is overloaded.

Last long-term database maintenance duration

This is how long it took to perform the last update of the long-term database; if this takes more than two to three hours you may have to reduce the number of long-term reports you have, reduce the number of devices they cover or set some of the long-term sample sizes to zero.

Last real-time database maintenance duration

This is how long it took to perform the last reorganisation of the real-time database; this should not take longer than thirty minutes.

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. Some devices do not have a configurable active flow timeout (e.g., Packeteer) and some high-end Cisco routers expose a design flaw in NetFlow that prevents the active flow timeout from being honoured; in these cases you can configure NetFlow Tracker to hold data in RAM for longer to prevent ignored flows - see Database Settings for more information.

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: Device Configuration

This is a brief guide to setting up NetFlow on various types of device. Note that if your device isn’t listed here it does not mean it is not supported by NetFlow Tracker; please ask your device vendor for a guide to enabling NetFlow.

Enabling NetFlow Export/NDE on a Cisco Router or Layer 3 Switch

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 layer 3 switch you must configure IOS on the MSFC and CatOS on the Supervisor Engine. Native IOS also requires extra commands; these are documented below.

Enabling Netflow Export on an IOS Device

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

```
ip cef
```

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

```
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.

Note that enabling or disabling NetFlow version 5 or version 9 (not version 1) on a 12000 series router causes packet forwarding to stop for a few seconds while the route processor and line card CEF tables are reloaded. To avoid interruption of service to a live network, apply this command during a change window, or include it in the startup-config file to be executed during a router reboot.

```
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` or **`ip flow ingress`** or **`ip route-cache cef`**

`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. Note that there are several commands to enable NetFlow on an interface and you must use the same command for every interface. **`ip route-cache flow`** and **`ip flow ingress`** enable NetFlow for inbound traffic on the interface; the only difference between the two is that the latter can be applied to individual sub-interfaces whereas the former must be applied to the physical interface. Be careful not to enable NetFlow for both a physical interface and one or more of its sub-interfaces.

`ip flow egress` enables NetFlow for outbound traffic on the interface and is required if you are using input filters. You may enable NetFlow for bout inbound and outbound traffic on a single interface if you are interested only in its traffic; in this case ensure that no other interface has NetFlow enabled.

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. Note that a Catalyst 4000 series switch does not support any of the commands to enable NetFlow for an interface; instead NetFlow is enabled for all interfaces using a special command documented below.

`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 summarize the active flows and give an indication of how much NetFlow data the router is exporting.

Enabling 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), or a Supervisor V, 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.

Enabling NDE on a Native IOS Device

The following commands are required in addition to the commands required to configure an IOS device above to get NetFlow information on route-switched traffic from a Catalyst 6000 or above; they are not required for a Catalyst 4000 series.

`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)E03, 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)E03, 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 Fluke Networks 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)E03, 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 NDE on a CatOS Device

A layer 3 switch running CatOS appears as two devices; the MSFC can be configured to export NetFlow information on all the packets it routes by following the instructions for configuring an IOS device above.

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 NetFlow Input Filters for Traffic Class Reporting

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.

Enabling Flow Detail Records on a Packeteer Device

A Packeteer 1200, 1550, 2500, 4500, 6500, 8500, 9500, or 10000 series running PacketWise v7.0.0 or above and having 256MB or more of memory can be configured to send either NetFlow records or a similar proprietary format to NetFlow Tracker. For more information visit http://support.packeteer.com/documentation/packetguide/rc3.1/overviews/flowdetail.htm

To enable Flow Detail Records, first log in to the PacketShaper in touch mode, then open the "flow detail records" page on the "setup" tab. In one of the collector rows, enter the IP address of the NetFlow Tracker server and one of the ports configured in the Listener Ports settings page (2055 is monitored by default). Packeteer-1 is the recommended record type for use with NetFlow Tracker; Packeteer-2 is also supported but NetFlow Tracker does not use any of the extra information and thus it is wasteful of network bandwidth between the PacketShaper and the NetFlow Tracker server. You can also choose to export NetFlow v5 records; this will prevent the Traffic Classes and Identified Applications reports and filters from functioning for the device. Finally, set the value under "Enabled" to "on" and click "apply changes…".

To ensure that NetFlow Tracker receives enough information from the device you must ensure that the "Look Community String" configured in the "SNMP" page is one of those set up in SNMP Settings, and you must set "Packeteer-0 Packets" to "on" in the "system variables" page.

If you have a recent version of PacketWise, you may have extra settings on the "system variables" page that should be changed. If available, "Intermediate FDR" should be set to "on", "Intermediate FDR Timeout " to 30000 milliseconds, and "Reset Packeteer 1/2 counters" to "on". If these settings are not available then the PacketShaper will describe all of the traffic for a long-lived flow in one record, and NetFlow Tracker will account for it all in the minute during which the flow ended. This will lead to large spikes in charts for the device.

Enabling NetFlow on an Enterasys Device

NetFlow Tracker supports Enterasys devices capable of exporting NetFlow version 9 exports. To enable NetFlow, enter the following commands while logged in to the router with read/write access:

```
set netflow cache enable
```

This enables NetFlow.

```
set netflow 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.

```
set netflow export-interval 1
```

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

```
set netflow port <port-string> enable
```

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.

```
set netflow export-version 9
```

This sets the export version. Version 9 is required for NetFlow Tracker to be able to associate NetFlow information with the interfaces it relates to.

Using sflowtool to Convert sFlow Records to NetFlow

NetFlow Tracker does not directly support devices which export sFlow records; however, the developer of sFlow provides a tool to convert sFlow records to NetFlow records, available at http://www.inmon.com/technology/sflowTools.php. This is a simple command-line utility which can be run as a daemon on Unix or a service on Windows by using one of the many free service installers available. The required command line options are:

`-p <port>`

This sets the incoming port number; the device should be configured to send sFlow records to this port on the address of the server running sflowtool.

`-c <address>`

This sets the address of the NetFlow Tracker server.

`-d <port>`

This sets the port on the NetFlow Tracker server that NetFlow records are sent to; this must be one of the ports configured in the Listener Ports settings page (2055 is monitored by default).

`-S`

This asks the tool to create NetFlow packets with the same source address as the incoming sFlow records, thus tricking NetFlow Tracker into believing that the NetFlow packets came directly from the device. Note that the tool will need to be run as root on Unix systems or as an administrator on Windows for this to work. If you use a service installer on Windows to run the tool it will be run under the built-in system account which is similar to an administrator account.

Note that support for this feature depends upon how the tool was compiled from source code and on operating system support – Windows XP does not support IP address spoofing, for example, and as a result recent Windows versions of the tool do not offer the feature on any version of Windows.

`-e`

This includes the peer AS numbers in the generated NetFlow records rather than the default origin AS numbers.

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 utilization 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://logging.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://tomcat.apache.org/. 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://opennms.svn.sourceforge.net/viewvc/opennms/opennms/branches/OPENNMS/src/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 which is no longer available. This is distributed under the Advantys Freeware license contract, a copy of which is available at http://web.archive.org/web/20031209160524/http://www.jspsmart.com/liblocal/docs/legal.htm.

IE5.5+ PNG Alpha Fix

NetFlow Tracker includes the IE5.5+ PNG Alpha Fix v1.0RC4, available at http://www.twinhelix.com/css/iepngfix/demo/. This is distributed under the CC-GNU Lesser GNU Public License, a copy of which is available at http://creativecommons.org/licenses/LGPL/2.1/deed.en.

Apache Xerces Java

NetFlow Tracker includes Apache Xerces Java v2.9.0, available at http://xerces.apache.org/xerces2-j/. This is distributed under the Apache Software License, a copy of which is available at http://www.apache.org/LICENSE.