

HP OpenView Performance Insight

NetFlow Preprocessor User Guide

Software Version: 3.0

Reporting and Network Solutions 7.0



November 2004

© Copyright 2004 Hewlett-Packard Development Company, L.P.

Legal Notices

Warranty

Hewlett-Packard makes no warranty of any kind with regard to this document, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose. Hewlett-Packard shall not be held liable for errors contained herein or direct, indirect, special, incidental or consequential damages in connection with the furnishing, performance, or use of this material.

A copy of the specific warranty terms applicable to your Hewlett-Packard product can be obtained from your local Sales and Service Office.

Restricted Rights Legend

Use, duplication, or disclosure by the U.S. Government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause in DFARS 252.227-7013.

Hewlett-Packard Company
United States of America

Rights for non-DOD U.S. Government Departments and Agencies are as set forth in FAR 52.227-19(c)(1,2).

Copyright Notices

© Copyright 2002–2004 Hewlett-Packard Development Company, L.P., all rights reserved.

No part of this document may be copied, reproduced, or translated into another language without the prior written consent of Hewlett-Packard Company. The information contained in this material is subject to change without notice.

Trademark Notices

OpenView is a U.S. registered trademark of Hewlett-Packard Development Company, L.P.

Java™ is a U.S. trademark of Sun Microsystems, Inc.

Oracle® is a registered U.S. trademark of Oracle Corporation, Redwood City, California.

UNIX® is a registered trademark of The Open Group.

Windows® and Windows NT® are U.S. registered trademarks of Microsoft Corporation.

Support

Please visit the HP OpenView web site at:

<http://www.managementsoftware.hp.com/>

This site provides contact information and details about the products, services, and support that HP OpenView offers. To reach the Support site, click **Support**. The Support site provides a fast and efficient way to access interactive technical support tools needed to manage your business. As a valuable support customer, you can benefit by using the support site to:

- Search for documents of interest
- Submit and track progress on support cases
- Manage a support contract
- Look up HP support contacts
- Review information about available services
- Join online discussions with other software customers
- Research and register for software training

Most of the support areas require you to register as an HP Passport user and log in. You may also require a support contract.

To find more information about access levels, go to the following URL:

http://support.openview.hp.com/access_level.jsp

To register for an HP Passport ID, go to:

<http://passport.hp.com/hpp2/newuser.do>

Chapter 1	Overview	7
	Collecting and Processing Flow Data	7
	NetFlow Preprocessor Functions	8
	Sources for Additional Information	8
Chapter 2	Installing the Preprocessor	11
	Installation Prerequisites	11
	Installing the Preprocessor	12
	Testing for Correct Installation	13
	Package Contents	13
	Removing the NetFlow Preprocessor	13
Chapter 3	Preprocessor Configuration	15
	Master Configuration File	15
	Domain Lookup File	21
	Protocol Lookup File	22
	Application Lookup File	22
	Configuring Flow Collector Applications	23
Chapter 4	Troubleshooting	25
	Identifying the Cisco NetFlow FlowCollector User	25
	Error Messages and Warnings	26
	Perl Not Installed Correctly at /usr/local/bin	29
	Files in the Bin Directory Will Not Run	29
	Output File Is Empty	30
Index		31

Overview

Tracing congestion to specific applications, servers, and clients is now easier and faster thanks to a suite of NetFlow reporting packages available from HP OpenView. This suite consists of the following products:

- NetFlow Preprocessor
- NetFlow Interface Report Pack/NetFlow Interface Datapipe
- NetFlow Global View Report Pack/NetFlow Global View Datapipe

The NetFlow Preprocessor is a prerequisite for both report packs. You will install the NetFlow Preprocessor on the same system where the flow collector application resides. You will install the report packs and datapipes on your OVPI server.

Collecting and Processing Flow Data

A flow is a group of packets moving between source and destination devices. The packets in this group are moving in the same direction, they share the same protocol, and they use the same transport-layer information. The traffic generated by a browser on a PC, using HTTP to request information from a website, is one flow; the response from the website to the PC is a second flow.

Devices that support NetFlow are able to record data about flows and send UDP datagrams to a configured destination. The destination must be running a collector application, such as Cisco's NetFlow FlowCollector or HP's Internet Usage Manager (IUM). The collector application accepts the datagram, performs decoding and aggregation, and then writes a new file in one of two formats, CallRecord format or DetailCallRecord format.

The NetFlow FlowCollector application performs these tasks:

- Receives flow statistics from Cisco devices
- Aggregates and stores data
- Produces an output file in a format (CallRecord or DetailCallRecord) defined by Cisco
- Calls the NetFlow Preprocessor when the delimited ASCII file is ready for processing



The call from the flow collector application to the NetFlow Preprocessor is possible only when the flow collector application and the NetFlow Preprocessor reside on the same system.

The NetFlow Preprocessor supports DetailCallRecord format and CallRecord format. If you are running the NetFlow Interface Report Pack, you want the collector application to output records in DetailCallRecord format. If you are running the NetFlow Global View Report Pack, you want the flow collector application to output records in CallRecord format.



Although the NetFlow Global View Datapipe will accept records in DetailCallRecord format, this format is not suitable for the NetFlow Global View Report Pack and may produce undesirable results.

NetFlow Preprocessor Functions

The NetFlow Preprocessor processes data from the collector application and then creates a new file in a format suitable for reading by an OVPI datapipe. The datapipes that collect data for the NetFlow Interface Report Pack and the NetFlow Global View Report Pack read the contents of the new file and populate tables in the OVPI database.

In NetFlow Interface reports, and in NetFlow Global View reports, you will see hourly, daily, and monthly trends for specific types of traffic. Will your response to congestion start with these reports? Probably not. Your response to congestion will likely start with the Interface Reporting Report Pack, where you will find out which interfaces have high utilization. As soon as you know which interfaces are affected, you can use NetFlow Interface and NetFlow Global View reports to find out which traffic types are experiencing congestion and where each traffic type originated.

The NetFlow Preprocessor performs these tasks:

- Filters the data in the file created by NetFlow FlowCollector
- Performs various groupings and aggregations
- Matches pairs of uni-directional records to create bi-directional data
- Allows further filtering on aggregated data
- Creates an output file in a format required by:
 - NetFlow Interface Datapipe
 - NetFlow Global View Datapipe
- Stores the output file on a local or remote file system

Filtering and aggregation are configurable, depending on protocol, port, and address information. Output can be restricted to a certain number of records, a percentage of the total traffic, or minimum transfer rates. For details, see [Chapter 3, Preprocessor Configuration](#).

Sources for Additional Information

The following documents are related to this manual:

- *NetFlow Preprocessor 3.0 Release Statement*
- *NetFlow Interface Report Pack 3.0 User Guide*

- *NetFlow Global View Report Pack 2.0 User Guide*
- *Interface Reporting Report Pack 4.6 User Guide*
- *NetFlow FlowCollector Installation and User Guide* [Cisco]



The last document in this list contains information about two record formats, CallRecord and DetailCallRecord.

Manuals for the core product, OVPI, and manuals for the reporting solutions that run on OVPI can be downloaded from the following site:

<http://www.hp.com/managementsoftware>

Select **Support > Product Manuals** to open the **Product Manuals Search** page. The user guides for OVPI are listed under **Performance Insight**. The user guides for reporting solutions, NNM SPIs, and NNM correlators are listed under **Reporting and Network Solutions**.

The manuals listed under **Reporting and Network Solutions** indicate the month and year of publication. When a manual is revised, the date of publication will change even if the software version number does not change. Since we post revised user guides on a regular basis, you should search this site for updates before using an older PDF that may not be the latest PDF available.

Installing the Preprocessor

This chapter covers the following topics:

- Installation prerequisites
- Installing the preprocessor
- Testing for correct installation
- Package contents
- Removing the preprocessor

Installation Prerequisites

Install the NetFlow Preprocessor on the system where the flow collector application is running. When the NetFlow Preprocessor and the flow collector application are running on the same system, the flow collector application can call the preprocessor automatically. The automatic call is not possible when the preprocessor and NetFlow FlowCollector are running on different systems. If you do not use an automatic call from the collection software, you are responsible for launching the preprocessor using some other means.

Router Configuration

The devices you want to monitor must be configured to use NetFlow. They must export NetFlow datagrams to the address and port that your collector software is configured to listen to. Refer to your hardware vendor's documentation for more information about configuring devices to export NetFlow datagrams to a particular address and port.

DetailCallRecord Format

Your flow collector application must be configured to export records in DetailCallRecord format or CallRecord format. If records are exported in any other format, the NetFlow Preprocessor will be unable to process the file, and reports will not contain any data.

Collector Configuration

It is important that you do not enable any additional processes available in the collector application that map or aggregate data. These processes will conceal data that the preprocessor needs to see. If this data is concealed, your reports will be incomplete or misleading.

Detailed information about configuring NetFlow FlowCollector can be found in the *NetFlow FlowCollector Installation and User Guide*, published by Cisco. For information about configuring IUM to create DetailCallRecord files, contact your HP IUM representative.

Perl Installed and Running

Perl 5.x is a prerequisite and must be installed before running the preprocessor. Perl is shipped with some software packages (for example, OVPI) and some operating systems. For more information on obtaining and installing Perl, please go to <http://www.perl.com> or the website maintained by the vendor of your operating system. For MS Windows, go to:

www.activestate.com.

Perl should be available via the user's PATH environment variable. No additional Perl libraries or modules are required. For UNIX systems, the Perl executable (or a symbolic link to it) should exist as /usr/bin/perl.

Installing the Preprocessor

The NetFlow Preprocessor is supplied as an archive file in standard compressed-file formats ("tarball" and zip). The archive file is included with every OVPI package that requires the NetFlow Preprocessor. The archive file can be found in the Preprocessor subdirectory beneath the datapipe in the Packages directory on your OVPI system. For example:

```
{DPIPE_HOME}/packages/Netflow_Interfaces_Datapipe/Preprocessor
```

After reviewing the current *NetFlow Preprocessor Release Statement* for any information about known problems and issues, follow these steps to install the preprocessor:

- 1 If you have a previous installation of the preprocessor and have set up automatic calls to it, disable these calls.
- 2 Remove or rename the preprocessor files and/or directory.
- 3 To ensure that the correct ownership and permissions are applied to NetFlow Preprocessor files on UNIX systems, log in as the same user who will run the NetFlow FlowCollector application.
- 4 Create a destination directory for installing the package.
- 5 Using uncompress and tar or an unzip utility, unpack the appropriate archive file (zip or tar.Z) to the destination directory; make sure that the option to preserve directory/path names is selected.

If you need to identify the user running Cisco's NetFlow FlowCollector, see [Chapter 4, Troubleshooting](#).

Testing for Correct Installation

To verify correct installation of the NetFlow Preprocessor, run the following command from the directory where you installed the package:

```
perl bin/Netflow_PP.pl -h
```

If the NetFlow Preprocessor installed correctly, you will see the following message:

Usage is:

```
bin/NetflowPP.pl [-c <cfg_file>] -f <file>
```

where:

<cfg_file> is a configuration file

<file> is a NetFlow file

The default configuration file is `./cfg/netflow.cfg`.

If you do not see this message, see [Chapter 4, Troubleshooting](#).

Package Contents

Installing the archive file creates a bin directory and a cfg directory beneath the install directory. The bin directory contains the preprocessor executable, a shell script, and a batch file. The preprocessor executable, shell script, and batch file must be located in the same directory, so do not move them to different directories. The cfg directory contains the following configuration files:

- Master configuration file
- Default domain lookup file
- Default application lookup file
- Default protocol lookup file

Although the archive file imposes the proper directory structure, you may if necessary alter the default file locations by editing the master configuration file. In addition, by using a command line option, you can alter the location of the master configuration file itself.

The master configuration file may contain a password in clear text. To ensure that this password cannot be read by unauthorized users, only the owner of the configuration should be able to read it or write to it. On UNIX systems, you must own the configuration files used. When files are installed, the required permissions are set by default.

Removing the NetFlow Preprocessor

To uninstall the NetFlow Preprocessor follow these steps.

- 1 Save any configuration files that you have modified. Any configuration files that you have modified will be lost if you remove the whole directory tree.

- 2 If you configured the flow collector application to call the preprocessor automatically, remove this option from the config files. (For details, see [Chapter 3, Preprocessor Configuration](#).)
- 3 Remove or rename the files that were created by unpacking the zip file.

Preprocessor Configuration

The NetFlow Preprocessor includes a master configuration file. Although the parameters in this file are set to defaults, we strongly recommend that you enter information appropriate to your environment and your needs. Doing so will greatly enhance the value of the output from the preprocessor. In addition to the master configuration file, this chapter covers:

- Domain lookup files
- Protocol lookup file
- Application lookup file
- Configuring the flow collector application to call the preprocessor

Master Configuration File

The master configuration file is named `netflow.cfg`. By default, it is looked for in a directory named `cfg`, below the directory where the preprocessor is installed. This file contains a list of parameter/value pairs separated by an equals sign (=). The rules are:

- Comments are supported.
- Anything following a hash mark (#) on a line is ignored.
- White space is ignored unless it is embedded in a parameter.

The default configuration contains no information about domains. If a domain is unresolvable, the value for `DEFAULT` will be used and all addresses will resolve to `OTHER_DOMAINS`. In addition, application and interface will be the only differentiating factors in `DetailCallRecord` files, and the only way that the data will roll up is by application and interface.

The following stipulations apply to default settings:

- At most, only the top 100 records (by total traffic) will be output.
- Only records that contribute a maximum of 90% of the total traffic will be output.
- Any aggregated flows with less than 1000 bytes per second will be ignored.

Master Configuration File Defaults

The following table provides a list of parameters in the master configuration file and the default values, where applicable.

Parameter	Default	Description
UNKNOWN_APP	DEFAULT	Action to take on finding a combination of protocol, source, and destination ports that cannot be resolved to an application. DEFAULT = use value in DEFAULT_APP CREATE = create a name consisting of "lowerport:higherport:protocol" IGNORE = ignore the flow record
DEFAULT_APP	OTHER_APPS	Value to use for unknown applications when UNKNOWN_APP is set to DEFAULT; otherwise this parameter is not required.
UNKNOWN_DOM	DEFAULT	Action to take on finding an IP address that cannot be resolved to a domain. DEFAULT = use value in DEFAULT_DOM CREATE = create domain from the IP address IGNORE = ignore the flow record
DEFAULT_DOM	OTHER_DOMAINS	Value to use for unknown domains when UNKNOWN_DOM is set to DEFAULT; otherwise this parameter is not required.
PROTOCOLS	./cfg/protocols.cfg	Lookup file for protocols.
APPLICATIONS	./cfg/ports.cfg	Lookup file for applications.
DOMAINS	./cfg/domains.cfg	Lookup file for IP address domains.
LOG	./netflow.log	File in which to record errors and warnings. Errors encountered before this file can be opened will be directed to STDERR.
AUDIT		File in which to write audit log messages. If this parameter is not specified (the default), no audit log messages are produced.
WORK	./	Work directory for spooling output. If the output directory is located on the local system, the work directory should also be located on the same file system.
SAVE	NULL	Save directory for input data. If set to /dev/null, input files are removed; if set to a null value (that is, SAVE=), no action is taken for input files.

Parameter	Default	Description
OUT	./	Output directory. May also specify an FTP URL for the directory (for example, OUT=ftp://myserver/outdir). If the directory is on the local system, it should reside on the same file system as the work directory (see the entry for the WORK parameter). If the directory specifies using an FTP URL, the USER and PASS parameters are used to log on to the remote system. If these are not supplied, a .netrc file may be used. If this is not present, anonymous FTP is used. Obviously, the destination system must support FTP to use an FTP URL. You may also specify a directory that exists on another server by using the tools offered by your operating system for sharing disks.
USER	Anonymous	Username for logging on to the FTP server specified in OUT (required only if OUT is an FTP URL).
PASS	<email address>	Password for logging on to the FTP server specified in OUT (required only if OUT is an FTP URL).
OUT-PREFIX	NETFLOW-PP	Name of the output file will be prefixed by this value.
MIN_INTERVAL	300	Flow files with PERIOD less than this value will be rejected. (Note that PERIOD is expressed in minutes, while this parameter is in seconds.)
PCT_INCLUDE	90	The percentage of the file to output, by traffic. Output is ordered by traffic (in bytes); only records constituting the top X% of the traffic will be written, where: $X = \text{PCT_INCLUDE}$
TOP_X	100	Only the top X records will be written. If set to NULL, all records should be output.
MIN_BPS	1000	The total traffic for an aggregated output record must be equivalent to a rate of at least this value when expressed as a number of bytes per second over the whole period. For example, if an input file covers 15 minutes and this value is set to 5, each output record must include total traffic of at least 4,500 bytes: $(5 \text{ bytes/sec} * 15 \text{ minutes} * 60 \text{ seconds}) = 4,500 \text{ bytes}$.

Parameters in the Master Configuration File

This section provides comments about the parameters listed in the preceding table.

1. UNKNOWN_APP

When a combination of ports and protocols cannot be resolved to an application, the action defined by UNKNOWN_APP is taken. This can be:

DEFAULT	Assign a defined default value for the application name
CREATE	Create an application name using a concatenation of source port, destination port, and protocol, separated by colons (:)
IGNORE	Ignore the input flow record

The files that ship with the preprocessor contain “well-known” and registered applications. If you have configured any application to use particular ports and protocols without these, they will be unresolvable until you enter information in the application’s lookup file.

Make sure that the application lookup file contains any additional applications you installed.

2. DEFAULT_APP

If the UNKNOWN_APP parameter is set to DEFAULT, this value will be assigned to any unresolvable applications.

3. UNKNOWN_DOM

When an IP address cannot be resolved to a domain, the action defined by UNKNOWN_DOM is taken. The options are:

DEFAULT	Assign a defined default value for the domain name
CREATE	Create a domain name from the IP address
IGNORE	Ignore the input flow record

The files that ship with the preprocessor contain no domain information and, as such, all addresses are unresolvable. Make sure that you add domain information to the domain lookup file before changing this parameter.

4. DEFAULT_DOM

If the UNKNOWN_DOM parameter is set to DEFAULT, this value will be assigned to any irresolvable domains.

5. PROTOCOLS

Contains the name and path of the lookup file for protocols. Relative paths are relative to the directory in which the preprocessor is invoked. (See “Protocol Lookup File” later in this chapter for details about the contents of this file.) Since the numbering of protocols is much more strictly governed than ports, needing to modify protocol mappings is not likely.

6. APPLICATIONS

Contains the name and path of the lookup file for applications. Relative paths are relative to the directory in which the preprocessor is invoked. (See “Application Lookup File” later in this chapter for details about the contents of this file.) You are strongly advised to add any additional applications you have installed to the application lookup file.

7. DOMAINS

Contains the name and path of the lookup file for protocols. Relative paths are relative to the directory in which the preprocessor is invoked. (See “Domain Lookup File” later in this chapter for details about the contents of this file.) Make sure that you add domain information appropriate to your environment to the source-domain lookup file.

8. LOG

Defines the name of the file to which warning and error messages will be written. The format of error messages is consistent with OVPI standards; errors can be written to the standard `trend.log` file.

9. AUDIT

Defines the name of the file to which standard OVPI audit messages will be written. If not specified (the default), no audit records will be produced. The format of audit records is consistent with OVPI standards; records can be written to the standard `audit.log` file.

10. WORK

Defines the directory to which temporary files will be written. If output files are being saved locally (an FTP URL is not being used for OUT), this directory should be on the same file system as the output directory to avoid problems associated with spooling output.

11. SAVE

Defines what to do with the input data after it is processed. The options are:

SAVE=	A null value; no action is taken on input files
save=/dev/null	Input files are deleted
SAVE=<dir>	Input files are moved to the directory indicated by <dir>

The preprocessor does not manage processed input data. If you do not remove the data after processing, you must do this using some other mechanism, for example, OVPI's `age_files` program.

12. OUT

Defines the location where output data will be written. The location can be a directory name on the local system or the location can be an FTP-style URL. If FTP is used to move the output data to another system, `USER` and `PASS` can be used to define the username and password for logging on to the remote system. If `USER` and `PASS` are not defined, a `.netrc` file is used. If a `.netrc` is not present, using anonymous FTP will be attempted. Whether the location is a directory name or a URL, data will be spooled to a temporary file before being moved to its final destination.

The path to the output directory will depend on which NetFlow report pack is in use (NetFlow Interface or NetFlow Global View) and whether or not the `addr2name` mapping utility is being used. The default directories for the OVPI system are:

NetFlow Global View `addr2name` input directory:

```
{DPIPE_HOME}/data/ImportData/NetFlowGVDP_addr2name
```

NetFlow Global View Teel `SourceDirectory`:

```
{DPIPE_HOME}/data/ImportData/NetFlowGVDP
```

NetFlow Interface `addr2name` input directory

```
{DPIPE_HOME}/data/ImportData/NetFlowIFDP_addr2name
```

NetFlow Interface Teel `SourceDirectory`

```
{DPIPE_HOME}/data/ImportData/NetFlowIFDP
```

If you choose to put the data in a different directory on the OVPI system, you must modify the input directory for the datapipe. For details about changing the datapipe's input directory, see the user guides for NetFlow Interface and NetFlow Global View.

13. USER

Defines the username to use for FTPing output to another server if an FTP-style URL is given in `OUT`.

14. PASS

Defines the password to use for FTPing output to another server if an FTP-style URL is given in `OUT`. Since this password is available in clear text in the configuration file, certain restrictions are imposed to avoid possible security breaches. Since the configuration file must be owned by the user running the preprocessor, only the owner should have read or write access to it.

15. OUT-PREFIX

Defines the prefix that is used for output file names. Output file names are created from a concatenation of the prefix, the SOURCE address from the input file header, and the STARTTIME from the input file header, separated by periods (.). Only the prefix can be specified.

16. MIN_INTERVAL

The value of PERIOD in the input file header must be greater than or equal to the MIN_INTERVAL value. PERIOD is expressed in minutes while MIN_INTERVAL is expressed in seconds. If PERIOD is PARTIAL, the difference between ENDTIME and STARTTIME must be greater than or equal to MIN_INTERVAL.

17. PCT_INCLUDE

The maximum percentage of the file to output, by traffic. Output is ordered by traffic (in bytes); only records constituting the top X% of the traffic will be written (where X = PCT_INCLUDE).

18.TOP_X

The maximum number of records to be written. If set to NULL, all records may be output.

19. MIN_BPS

The total traffic for an aggregated output record must be equivalent to a rate of at least this value when expressed as a number of bytes per second over the whole period. For example, if an input file covers 15 minutes and this value is set to 5, each output record must include total traffic of at least 4,500 bytes, calculated as:

$$(5 \text{ bytes/sec} * 15 \text{ minutes} * 60 \text{ seconds}) = 4,500 \text{ bytes}$$

Domain Lookup File

A domain-name lookup file contains a definition of the IP addresses that make up each domain. A domain may contain one or more IP address. The rules are as follows:

- Comments are supported.
- Anything following a hash mark (#) on a line is ignored.
- White space before a domain name or trailing a domain name is ignored.
- Spaces in domain names are allowed.

For IP domains, each line consists of an IP address or range of IP addresses followed by a domain name separated by white space. IP address ranges can be defined two ways:

- Start address, followed by dash, followed by end address, for example:

```
192.168.1.2-192.168.1.254
```

- CIDR blocks, or classless IP domain range, for example:

`192.168.1.2/24`

You are strongly advised to add domain information appropriate to your environment to the source-domain lookup file.

Protocol Lookup File

Protocols are defined in a similar way to BSD style `/etc/protocols`. The content of the supplied default file is derived from the Internet Authority for Number Assignments (IANA) file “protocol-numbers.” This contains all registered protocol numbers. The rules are:

- Comments are supported.
- Anything following a hash mark (#) on a line is ignored.
- White space is ignored.
- Each line consists of a protocol name followed by a protocol number.
- The protocol number may be followed by several other fields; these fields are ignored.
- One protocol number maps to one protocol name; it is a one-to-one relationship.
- When there are multiple definitions for a particular protocol number, the first definition in the file is used; subsequent definitions are ignored.

Protocols are not likely to change. The defaults are adequate for most circumstances.

Application Lookup File

Applications are defined in a similar way to BSD style `/etc/services`. The content of the supplied default file is derived from the Internet Authority for Number Assignments (IANA) file “port-numbers.” This contains all well-known and registered applications. The rules are:

- Comments are supported.
- Anything following a hash mark (#) on a line is ignored.
- White space is ignored.
- Each line consists of an application name followed by a port/protocol pair separated by white space.
- The port/protocol pair is separated by a forward slash (/).
- Unlike BSD format, a wildcard (*) can be used for the port number.

If you install additional applications, make sure you add these applications to the application lookup file. If you use non-standard ports for any applications, these ports must be defined; otherwise their visibility in reports will be lost.

Configuring Flow Collector Applications

This section provides assistance with configuring your flow collector. For more detailed information, refer to the documentation provided by the software vendor for your collection application. Refer to the documentation provided by your hardware vendor for information on configuring NetFlow enabled devices to export NetFlow datagrams to your collection system.

Configuring Cisco's FlowCollector

You may need to modify several parameters to produce the data required for the preprocessor.

- 1 Locate the NetFlow configuration directory.
- 2 Use a text editor to modify the `nf.resources` file.
- 3 Make sure that `OUTPUT_DOTTEDADDRESS` is set to yes.
- 4 If you are collecting data from devices in multiple time zones, make sure that `GMT_FLAG` is set to yes.
- 5 Make sure that `DEVICE_DOTTEDADDRESS` is set to yes.
- 6 Make sure that `ACCEPT_PACKETS_FROM` block is commented out unless you wish to filter data depending upon its source.
- 7 If you wish to automatically call the preprocessor whenever data is created (this is the recommended approach), change the `USER_SCRIPT_LOCATION` to the fully qualified name and path for `nf2ovpi.ksh` (which can be found in the bin directory created beneath the directory in which the preprocessor is installed).
- 8 Save the file.
- 9 Locate the `NFC_CONFIGFILE` entry in the `nf.resources` file and use a text editor to modify it.
- 10 Make sure there is a section that includes `Aggregation DetailCallRecord` (or `Aggregation CallRecord` if you are not running NetFlow Interface) and insert values for `Period`, `Port`, `DataSetPath`, and `MaxUsage`. Ensure that `Compression` is set to `No` and `Binary` is also set to `No`.
- 11 Save the file.
- 12 Locate the `NFC_KNOWNPROTOCOLS` entry in the `nf.resources` file and use a text editor to modify it.
- 13 Remove or comment out the content of this file.
- 14 Save the file.
- 15 Locate the `NFC_KNOWNSRCPORTS` entry in the `nf.resources` file and use a text editor to modify it.
- 16 Remove or comment out the content of this file.
- 17 Save the file.
- 18 Stop and restart the NetFlow Collector to activate the changes.

Configuring HP's Internet Usage Manager

Add each router to your IUM collector. Make sure that the NotifyCommand feature is used to call the preprocessor. The file used to call the preprocessor varies. For UNIX, it is a shell script; for MS Windows, it is a batch file. Contact your HP IUM representative for more information about setting up a configuration suitable for producing DetailCallRecord data using IUM.

Troubleshooting

This chapter discusses:

- Identifying the Cisco NetFlow FlowCollector user
- Error messages and warnings
- Corrective actions
- Perl not installed correctly
- Files in the bin directory will not run
- Output file is empty

Identifying the Cisco NetFlow FlowCollector User

To identify the user running NetFlow FlowCollector, run the following command:

```
ps -deaf awk '/NFCollector/ {print $0}' -
```

You should see output similar to the following:

```
bin 498 493 0 Sep 12 ? 3:37 NFCollector
```

The user “bin” is running NetFlow FlowCollector.

Error Messages and Warnings

The following table contains a list of messages generated by the NetFlow Preprocessor. Causes and corrective actions are included where appropriate.

Message	Type	Recommended Action
Application resolution file (<file>) does not exist	FATAL	The application lookup file does not exist. Relative paths are relative to the directory from which the preprocessor is called. If the default shell script is used, paths are relative to the directory above the shell script itself. Check the setting for APPLICATIONS in the master configuration file.
Can't open config file (<file>): <reason>	FATAL	The master configuration file could not be opened for the reason given.
Can't open log file (<file>): <reason>	FATAL	The log file could not be opened for the reason given.
Default value must be set before it can be assigned (DEFAULT_APP)	FATAL	The action to take when an irresolvable application is found is to assign a default. This default has not been defined using the DEFAULT_APP parameter. Check the settings in the master configuration file.
Default value must be set before it can be assigned (DEFAULT_DOM)	FATAL	The action to take when an irresolvable domain is found is to assign a default. This default has not been defined using the DEFAULT_DOM parameter. Check the settings in the master configuration file.
Destination directory invalid (<dir>)	FATAL	The destination directory on the destination system is not valid. Ensure that it exists and the user has write access to it.
Domain resolution file (<file>) does not exist	FATAL	The domain lookup file does not exist. Relative paths are relative to the directory from which the preprocessor is called. If the default shell script is used, paths are relative to the directory above the shell script itself. Check the setting for DOMAIN in the master configuration file.
Duplicate application: <port> <protocol> <application>	WARNING	A duplicate definition for an application was found. The first definition found will be used. This one will be ignored.
Duplicate domain: <IP> <domain>	WARNING	A duplicate definition for a domain was found. The first definition found will be used. This one will be ignored.
Duplicate protocol: <protocol name> <protocol number>	WARNING	A duplicate definition for a protocol was found. The first definition found will be used. This one will be ignored.
Failed opening applications file, <file> <reason>	FATAL	The application lookup file could not be opened for the reason given.

Message	Type	Recommended Action
Failed opening domains file, <file>: <reason>	FATAL	The domain lookup file could not be opened for the reason given.
Failed opening ftp session <reason>	FATAL	An F'FTP session could not be initiated for the reason given.
Failed opening input file (<file>): <reason>	FATAL	The master configuration file could not be opened for the reason given.
Failed opening protocols file, <file>: <reason>	FATAL	The protocol lookup file could not be opened for the reason given.
Failed removing existing version of file on target host	FATAL	A version of the output file was found on the remote server when using F'FTP. This file could not be deleted. Check the following: 1) Why was the same output file produced? The preprocessor should not be used to process the same input file multiple times. 2) Why could it not be removed?
Failed to move <workfile> to <output>	FATAL	The work file could not be moved to the final destination directory. Check permissions on the file and directory.
Failed to move input data to save (<file> <destination>): <reason>	FATAL	The input file could not be moved to the save directory (as configured). Check permissions on the file and directory.
Failed to open audit log (<file>) <reason>	FATAL	The audit log file could not be opened for the reason given.
Failed to open work file (<file>): <reason>	FATAL	The master configuration file could not be opened for the reason given.
Failed to remove input data (<file>): <reason>	FATAL	The input file could not be removed after processing (as configured). Check permissions on the file.
FTP login failed (<ftp server>)	FATAL	Logging on to the FTP server failed. This is most likely because an invalid username and password are being used. Check the setting for USER and PASS in the master configuration file. If these settings are not used, check the values in any .netrc file used or the availability of anonymous login on the server.
FTP put failed (<reason>)	FATAL	FTP put failed for the reason given.
FTP rename failed	FATAL	FTP rename failed. The output file is transferred first to a temporary file then renamed to the final destination. Check permissions on the destination system.
Input file does not exist (<file>)	FATAL	An input file was passed to the preprocessor but it does not exist. Does the file exist? If not why not? Check the method used to call the preprocessor. Are other programs accessing the input files?

Message	Type	Recommended Action
Interval too small (<time>, <file>)	FATAL	The interval is smaller than that specified by MIN_INTERVAL, and the file has been rejected.
Invalid header field (<field>)	FATAL	A header field has no value. Check that NetFlow FlowCollector produced the file in DetailCallRecord format.
Invalid header record (<file>)	FATAL	The header record in the file is invalid. Check that NetFlow FlowCollector produced the file in DetailCallRecord format.
Invalid owner for config file (<file>)	FATAL	The user running the preprocessor must own the master configuration file. This is a security measure since the file may contain a clear text password.
Invalid permissions on config file (<file>)	FATAL	Only the owner of the master configuration file should be able to read or write the configuration file. This is a security measure, because the file may contain a clear text password.
Invalid value for UNKNOWN_APP (<application>)	FATAL	This parameter must be DEFAULT, CREATE, or IGNORE. Check the settings for UNKNOWN_APP and DEFAULT_DOM in the master configuration file.
Invalid value for UNKNOWN_DOM (<domain>)	FATAL	This parameter must be DEFAULT, CREATE, or IGNORE. Check the settings for UNKNOWN_DOM and DEFAULT_DOM in the master configuration file.
No input file specified	FATAL	No input file was specified. An input file must be passed to the preprocessor using the "-f" option.
Output directory (<file>) does not exist	FATAL	The output directory does not exist. Relative paths are relative to the directory from which the preprocessor is called. If the default shell script is used, paths are relative to the directory above the shell script itself. Check the setting for OUT in the master configuration file.
Protocol resolution file (<file>) does not exist	FATAL	The protocol lookup file does not exist. Relative paths are relative to the directory from which the preprocessor is called. If the default shell script is used, paths are relative to the directory above the shell script itself. Check the setting for PROTOCOLS in the master configuration file.
Record type must be 1 or 2 for audit log entries	FATAL	This internal coding error should not occur. Contact HP Technical Support if it does.
Required header fields are missing (<file>)	FATAL	Required header fields are not present. Ensure that NetFlow FlowCollector produced the file in DetailCallRecord format.

Message	Type	Recommended Action
Save directory (<file>) does not exist	FATAL	The save directory does not exist. Relative paths are relative to the directory from which the preprocessor is called. If the default shell script is used, paths are relative to the directory above the shell script itself. Check the setting for SAVE in the master configuration file.
Work directory (<file>) does not exist	FATAL	The work directory does not exist. Relative paths are relative to the directory from which the preprocessor is called. If the default shell script is used, paths are relative to the directory above the shell script itself. Check the setting for WORK in the master configuration file.

Perl Not Installed Correctly at /usr/local/bin

This issue applies to UNIX systems only. Ensure that you can run Perl from /usr/local/bin/perl. Run the command:

```
/usr/local/bin/perl -v
```

You will see a message similar to the following text:

```
This is Perl, v5.6.0 built for sun4-solaris
```

```
Copyright 1987-2000, Larry Wall
```

```
Perl may be copied only under the terms of either the Artistic License or
the GNU General Public License, which may be found in the Perl 5.0 source
kit.
```

```
Complete documentation for Perl, including FAQ lists, should be found on
this system using `man perl' or `perldoc perl'. If you have access to the
Internet, point your browser at http://www.perl.com/, the Perl Home Page.
```

If you do not see this message, ask your systems administrator to install Perl and create a symbolic link to the Perl executable in the /usr/local/bin directory.

Files in the Bin Directory Will Not Run

This issue applies to UNIX systems only. Ensure that execute permission has been granted on the files in the bin directory located beneath the directory where the preprocessor was installed. Running the command:

```
ls -l bin
```

should produce output similar to the following:

```
-rwx-----  1 bin      staff      18459 Jun  1 09:30 netflow_pp.pl
-rwx-----  1 bin      staff         604 Jun  1 09:30 trend_nfc.ksh
```

If you are trying to run the preprocessor manually, make sure you have the appropriate permissions. If you want to call the preprocessor using the `USER_SCRIPT_LOCATION` parameter provided by NetFlow FlowCollector, the user running NetFlow FlowCollector must have execute permission.

Output File Is Empty

The resolution action parameters are `UNKNOWN_APP` and `UNKNOWN_DOM`. If either of the resolution action parameters is set to `IGNORE`, it is possible to ignore a whole data file, especially if the domain lookup files have not been updated to reflect the environment. Use the `IGNORE` option **only** if you have configured the applications and domains for your environment.

A

- addr2name utility, **20**
- application lookup file, **22**
- APPLICATIONS, **16, 19**
- AUDIT, **16, 19**
- audit.log file, **19**
- automatic call of preprocessor, **7, 23**

B

- bin directory
 - contents, **13**
 - files will not run, **29**

C

- call preprocessor automatically, **7, 23**
- cfg directory, **13, 15**
- configuration files, list of, **13**

D

- data flow, defined, **7**
- DEFAULT_APP, **16, 18**
- DEFAULT_DOM, **16, 18**
- defaults, master configuration file, **16**
- DetailCallRecord format, **7**
 - differentiating factors in files, **15**
- directories
 - bin, **13**
 - cfg, **13, 15**
- domain, unresolvable, **15**
- domain lookup file, **21**
- DOMAINS, **16, 19**

E

- empty output file, **30**
- error messages, **26**

F

- flow, defined, **7**

I

- IANA file, **22**
- IGNORE option, problems with, **30**
- installation
 - prerequisites, **11**
 - procedure, **12**
 - verifying, **13**
- Interface Reporting Report Pack, **8**
- Internet Usage Manager, **7**
 - configuring to call preprocessor, **24**

L

- LOG, **16, 19**
- lookup files
 - application, **22**
 - domain, **21**
 - protocol, **22**

M

- master configuration file, **15**
 - editing or moving, **13**
 - parameters
 - details about, **18**
 - listed with defaults, **16**
 - read/write access, **13**
- messages, **26**
- MIN_BPS, **17, 21**
- MIN_INTERVAL, **17, 21**

N

- netflow.cfg file, **15**

NetFlow FlowCollector
 automatically calls preprocessor, **7**
 functions of, **7**
 identifying user, **25**
NetFlow Global View Report Pack, **8**
NetFlow Interface Report Pack, **8**
netrc file, **20**

O

OUT, **17, 20**
output, default number of records, **15**
output file, empty, **30**

P

package contents, **13**
PASS, **17, 20**
PCT_INCLUDE, **17, 21**
Perl
 installed incorrectly, **29**
 prerequisite for preprocessor, **12**
port numbers, in IANA file, **22**
PREFIX, **17, 21**
preprocessor
 automatically called, **7, 23**
 contents of package, **13**
 functions of, **8**
 installing, **12**
 location of archive file, **12**
 uninstalling, **13**
prerequisites for installation, **11**
Product Manuals Search page, **9**
protocol lookup file, **22**
protocol numbers, **22**
PROTOCOLS, **16, 19**

R

removing the preprocessor, **13**
report packs
 Interface Reporting, **8**
resolution action parameters, **30**
rules
 application lookup file, **22**
 domain lookup files, **21**
 protocol lookup file, **22**

S

SAVE, **16, 19**
shell script, **13**

T

temporary files, **19**
TOP_X, **17, 21**
trend.log file, **19**

U

uninstalling the preprocessor, **13**
UNKNOWN_APP, **16, 18, 30**
UNKNOWN_DOM, **16, 18, 30**
USER, **17, 20**

V

verifying installation, **13**

W

warning messages, **26**
WORK, **16, 19**