



Doc. No. 78-4494-01

# Release Notes for NetFlow FlowAnalyzer Release 1.0

---

**September 1997**

This release note describes the NetFlow FlowAnalyzer application Release 1.0.

## Contents

NetFlow FlowAnalyzer Introduction	2
Product Features	2
Product Packaging	2
Product Documentation	2
Related Publications	3
Platform Support	3
Workstation Requirements	3
Installing the FlowAnalyzer	4
Important Notes	4
Install from CD-ROM	4
Set Up the Display User Interface	5
Set Up the DisplayServer	6
Quick Start the DisplayServer	6
Customize the DisplayServer	7
Collection Configuration	8
Memory Management	9

---

### Corporate Headquarters

Cisco Systems, Inc.  
170 West Tasman Drive  
San Jose, CA 95134-1706  
USA

Copyright © 1997  
Cisco Systems, Inc.  
All rights reserved.

Running the DisplayServer	13
Helpful Hints and FAQs	13
Cisco Connection Online	15
Documentation CD-ROM	16

## NetFlow FlowAnalyzer Introduction

The NetFlow FlowAnalyzer is a software application that reads, analyzes, and displays data collected by the NetFlow FlowCollector from NetFlow switch-enabled routers. This application provides users with information about the traffic passing through the routers.

The FlowAnalyzer application includes two programs: NFADisplayServer and NFADisplay. The NFADisplayServer program retrieves traffic statistics from the NetFlow FlowCollector application. The NFADisplay program displays the traffic statistics from the NFADisplayServer program in graphs and charts.

## Product Features

The NetFlow FlowAnalyzer application is a network management application that includes a server and a user interface to display data collected by the NetFlow FlowCollector application. The FlowAnalyzer includes two programs:

- FlowAnalyzer DisplayServer—Receives requests from the FlowAnalyzer user interface and provides data collected in flat files created by the FlowCollector to be displayed by the NetFlow FlowAnalyzer user interface.
- FlowAnalyzer User Interface—Exists within a browser such as Netscape to display the data collected by the FlowCollector in a variety of formats including tables, graphs, and spreadsheets.

## Product Packaging

The NetFlow FlowAnalyzer product suite includes the following files:

- NFA1\_0.README (this file)
- NFA1\_0.<platform>.Z (<platform> is "SOL" or "HP\_10")
- NFA1\_0.setup.sh (installation script)

---

**Note** The NetFlow FlowAnalyzer application can be downloaded from the Web and is also available on the application's CD-ROM.

---

## Product Documentation

The *NetFlow FlowAnalyzer Installation and User Guide* provides task-based information for installing and using the NetFlow FlowAnalyzer application. The manual is available in hard copy and on the Documentation CD-ROM.

You can access Cisco technical documentation on the Web at <http://www.cisco.com>. For more information, refer to the “Cisco Connection Online” section later in this document.

---

**Note** All documents mentioned are available as printed manuals or electronic documents.

---

## Related Publications

For additional information about using the NetFlow FlowCollector application to collect data from multiple routers, refer to the *NetFlow FlowCollector Installation and User Guide*.

For documentation about available Cisco IOS features, refer to the Cisco IOS, Release 11.1(2) or later, configuration guides and command references, which are located in the Cisco IOS Release 11.1(2) database.

For more information about electronic documentation, refer to the “Documentation CD-ROM” section later in this document.

---

**Note** The most up-to-date Cisco IOS documentation can be found on the latest Documentation CD-ROM and on the Web. These electronic documents contain updates and modifications made after the paper documents were printed.

---

## Platform Support

The NetFlow FlowAnalyzer application Release 1.0 supports the following platforms:

- HP-UX—Version 10.2 on a Class C machine with at least 256 MB RAM running standard Java 1.02 (HP-UX Java 1.0.3)
- Solaris—Version 2.5.1 on an ULTRA-1 workstation with at least 256 MB RAM (Java 1.02 included)

---

**Note** For HP\_10 installations, Java 1.0.2 is required. For both HP\_10 and SOL installations, a Web server is required.

---

## Workstation Requirements

To ensure a successful installation of the FlowAnalyzer application, ensure that your workstation meets the following requirements:

- A Solaris or HP-UX workstation with 400 MB of free logical memory.
- Enough memory and 6 MB disk space available to install the FlowAnalyzer application.
- A browser to run the NFAalyzer.html file. (The recommended version of Netscape browser is 3.01.)

---

**Note** In order to obtain the best performance, the FlowAnalyzer should be installed on a workstation separate from the FlowCollector.

---

### Installing the FlowAnalyzer

The FlowAnalyzer application is available on CD-ROM. This section provides instructions for installing the FlowAnalyzer. The NFADisplayServer program runs on Solaris 2.5.1 and uses Java classes and native libraries. The NFADisplay program includes a set of Java classes and an html file and should be installed on the same machine where your web server is running.

#### Important Notes

- 1 If you have installed this package before, you must remove it before you install a new one using the package remove utility:  

```
$ pkgrm NFA
```

or

```
$ pkgrm CSCOnfa
```
- 2 If you have been using a previous version of this tool and you have saved some data in the /opt/CSCOnfa/NFAServer/exported\_files directory by pressing "Export" in the Display, you should copy this data out to a safe location. The directory will be deleted when you do a pkgrm.
- 3 Because of system performance requirements, you do *not* run the FlowAnalyzer on the same machine as the FlowCollector.
- 4 This application must be installed on a machine hosting a web server.
- 5 If you install this application on an HP machine, you must also install Java 1.0.2 on that machine.
- 6 Do not alter the contents of the directories or the files in the directories of the FlowCollector. Doing so may cause the FlowAnalyzer to not perform as expected.
- 7 The AS.txt file provided in the /opt/CSCOnfa/NFAServer/imported\_files directory is meant as a template. The AS number and strings provided in this file are subject to change based on the registration activities of Internic. You are expected to supply the tools that will update this file with the latest and most accurate data.

#### Install from CD-ROM

Use the following installation procedures after installation of the FlowAnalyzer software package:

**Step 1** Log in as root.

```
$ su root
```

**Step 2** Copy the tar file from the CD and untar the file.

For Solaris

```
$ cp NFA1_0.SOL.tar  
$ tar -xvf NFA1_0.SOL.tar
```

For HP-UX

```
$ cp NFA1_0.HP_10.tar  
$ tar -xvf NFA1_0.HP_10.tar
```

**Step 3** Run the installation script and answer all questions.

```
$ chmod +x NFA1_0.setup.sh
$ ./NFA1_0.setup.sh NFA1_0.<platform>.Z
```

For the Solaris platform enter the following:

```
$ ./NFA1_0.setup.sh ./NFA1_0.SOL.Z
```

For the HP-UX platform enter the following:

```
$ ./NFA1_0.setup.sh ./NFA1_0.HPUX.Z
```

## Set Up the Display User Interface

To set up the NFADisplay program on your workstation, use the following procedure:

---

**Note** To run the AnalyzerDisplay, you must install the NFADisplayServer on the same machine where your http web server is running.

---

**Step 1** Edit the NFAAnalyzerB.html file located in /opt/CSCOnfa/NFADisplay/bin directory so that NFDISPLAYSERVERPORT (6544 is the default) uses the port value that the server is listening on.

Also, the NFDISPLAYSERVERDATAPATH must point to the directory where the data is collected.

**Step 2** Move the FlowAnalyzer files to a web server directory. For example:

```
$ cd /http-web-server-root-directory/docs
```

The http-web-server-root-directory is the directory where the http web server is installed.

**Step 3** Create a directory for the nfa files and move to that directory.

```
$ mkdir nfa
$ cd nfa
```

**Step 4** Copy the nfa files to this new directory.

```
$ cp -rf /opt/CSCOnfa/NFADisplay/bin/* .
```

**Step 5** Open the NFAAnalyzer.html file by clicking on the **Open** button or use the **Open Location** option from the File pull down menu in your browser. Enter the following:

```
http://http-web-servername/nfa/NFAAnalyzer.html
```

The http-web-servername is the hostname where the http web server is running.

---

**Note** A characteristic of browsers is that they cache applets. You should flush the browser's cache if a pre-beta version of the FlowAnalyzer exists. This can be done by the browser using the Options->Network Preferences->Cache menu item.

---

### Set Up the DisplayServer

To set up the FlowAnalyzer DisplayServer program to retrieve traffic statistics from the NetFlow FlowCollector application, you can configure the DisplayServer program using the resource file or using the default parameters. Some additional setup is required for installation on the HP-UX platform.

### Quick Start the DisplayServer

The following quick start procedures for the DisplayServer are provided here for your convenience.

#### Solaris Quick Start

**Step 1** Verify that you have at least 256 MB of physical RAM and 400 MB of available swap space. If you do not, refer to the section entitled “Memory Management.”

To reduce the MaxMB value, edit the NFADS.resources file. For example:

```
MaxMB {your_value}
```

with

```
{your_value} = (physical RAM - 32) MB
```

**Step 2** Start the DisplayServer program from the {DisplayServerBaseDir}/bin directory to begin listening to Port 6544.

```
$ start.DisplayServer
```

#### HP-UX Quick Start

**Step 1** Verify that you have at least 256 MB of physical RAM and 300 MB of available swap space. Enter

```
$ /usr/sbin/swapinfo -m
```

If your system does not meet those criteria, refer to the section entitled “Memory Management.”

To reduce the MaxMB value, edit the NFADS.resources file. For example:

```
MaxMB {your_value}
```

with

```
{your_value} = (physical RAM - 32) MB
```

**Step 2** Ensure that your execution path is set with HP-UX Java version 1.0.3. If it is not, refer to the section entitled “Customize the DisplayServer.”

**Step 3** Start the DisplayServer program from the {DisplayServerBaseDir}/bin directory to begin listening to Port 6544.

```
$ start.DisplayServer
```

## Customize the DisplayServer

The DisplayServer's resource file, NFADS.resource, contains a variety of parameters for configuring the DisplayServer to retrieve summarized traffic statistics from the NetFlow FlowCollector.

To redefine configuration parameters from their default values in the NFADS.resources file, use the guidelines described in the following paragraphs.

- Redefine one parameter per line.
- You must use a valid parameter keyword when you are starting a new line; otherwise, the line is ignored.
- If a parameter is defined more than once, the last parameter defined is used.
- Parsing stops at the first blank line of the file.

The format to use for redefining the parameters is

```
<parameter_keyword> <new_value>
```

## Define Router Aliases

The DisplayServer program includes two files used to change the information displayed by the FlowAnalyzer user interface. The files are located in the NFAServer/imported\_files directory. Edit the following files to define router aliases:

- RouterAliases.txt—Allows you to define
  - A NetFlow reporting router textual alias (used instead of an IP address)
  - An Autonomous System (AS) number for each NetFlow reporting router
- AS.txt—Allows you to define the text associated with the AS numbers in the ASMatrix aggregation schemes

## RouterAliases.txt File

The NetFlow reporting router defined in the RouterAliases.txt file has three columns.

- 1 The first column contains the router's IP address in decimal-byte format (a.b.c.d).
- 2 The second column contains the router's identifying text that is displayed in the router list of the user interface. This column can be left blank.
- 3 The third column contains the NetFlow reporting router's local AS number which is substituted for AS 0 in the ASMatrix aggregation scheme. The router's textual AS description begins with "(local)" in order to show the substitution for AS 0 has been made.

In the RouterAliases.txt file, you can leave the router's identifying text blank (second column) and still define the router's AS number in the third column. This allows you to continue to view the router and its IP address and still have the AS 0 replacement. This also implies that you cannot use a number from 0 to 65535 as a router identifying string, unless you also provide an AS number on that line. The DisplayServer's interpretation of the RouterAliases.txt file is printed near the top of the DisplayServer's log file.

## AS.txt File

The AS.txt file has two columns; the first column includes the AS number and the second column is used to define a textual description associated with the AS number.

### Collection Configuration

Run the NetFlow FlowCollector application using the GMT reference. To run the FlowCollector in this mode, you must edit the `nf.resources` file to uncomment the line `GMT_FLAG`. The `GMT_FLAG` parameter by default is turned on (yes). If you do not uncomment the `GMT_FLAG` parameter, you must use a distinct `DataSetPath` for each local time zone in your collection of data.

There can be only one time zone for each `DataSetPath` directory of your FlowCollector database. For example:

```
GMT_FLAG          yes
```

The `NFADisplayServer` can accommodate the shift to daylight saving time and will support the locally named file for a single time zone used for each `DataSetPath` (Data Location). The `DataSetPath` is defined for each thread in your `nfconfig` file and is the same as the `DataLocation`. The FlowCollector database might have anomalies if you collect different time zones from the same `DataSetPath`, or if you are not running the GMT reference when daylight saving time causes the local clock to be shifted back.

The router name should always be specified in the decimal byte format (a.b.c.d) to avoid ambiguity, which could cause the analysis module to miss requested data. For example, the format should be

```
171.69.204.5
```

Specify that the FlowCollector should use the dot-decimal format only to prevent data from the same router from being stored in only one location in the FlowCollector database. For example, the format should be

```
DEV_DOTTEDADDRESS  yes
```

Use a consistent router alias. If the FlowCollector cannot find the router alias, it will use the decimal byte format, `DEV_DOTTEDADDRESS`. To avoid inconsistencies using router identifiers, use the `DEV_DOTTEDADDRESS` format.

### Java Note for HP-UX

Run HP-UX Java 1.0.3, which is available if you click the Software button on the URL

```
http://hpcc920.external.hp.com/gsyinternet/hpjdk/productbrief.html
```

You next select "1.0.3 Release for HP-UX 10.20 or 10.10..." Problems can arise if you use inconsistent router identifiers of HP's 1.0.2 and 1.0.3 releases. Make sure you are running the correct version:

```
$ which java
/opt/java/bin/java
bee /tmp_mnt/home/jwiggins/HPUX

$ java -version
java version "1.03.01 HP-UX 10.20: 970327"
bee /tmp_mnt/home/jwiggins/HPUX
$
```

Specifically, you should delete any `/usr/bin/java` script or file, and you should not refer to the `/usr/java` directory on your disk. If you have these references, they should be deleted, and you should run a clean install of HP 1.0.3 only.

Additional setup is required if you are running the `DisplayServer` on an HP-UX platform.

**Step 1** Install the Java runtime environment Version 1.0.2 (HP-UX Version 1.0.3) on your workstation.



**Step 2** Ensure that the correct Java runtime environment is in your execution path. Use one of the following methods:

Method 1: Verify that Java is in your execution path.

```
$ which java
/opt/java/bin/java
```

Verify the Java runtime version installed.

```
$ java -version
java version "1.0.3ss:08/01/96-23:00"
```

Method 2: The shell script `set_path.JAVABIN` returns the correct location of the Java runtime environment. If you have performed the standard installation of Java, it is not necessary to edit the `set_path.JAVAVIN` file. To check the location of Java, enter

```
$ set_path.JAVABIN
/opt/java/bin/java
```

If you have installed the HP-UX version of Java in the `/opt/java/bin` default directory, the `set_path.JAVABIN` script is correct and your installation of HP-UX is complete.

If you have installed the HP-UX version of Java in different location, you must edit the `set_path.JAVABIN` script file to include the location of the java executable file which is actually a wrapper shell script named "java."

## Memory Management

You use two memory management parameters to configure the DisplayServer for your workstation and the FlowAnalyzer application:

MaxMBperCommand	280	most used on one command
MaxMB	224	size of memory pool

The MaxMBperCommand parameter is automatically truncated to MaxMB.

The memory use of DisplayServer is limited primarily by the MaxMB parameter in the `NFADS.resources` file. When tuning the memory configuration, be careful not to configure the DisplayServer to use too much memory and exhaust your workstation system swap space.

If that does happen, the operating system may terminate the DisplayServer and cause a possibly large core dump in the `/opt/CSCOnfa/NFAServer/bin` directory. You can save time and inhibit this core dump by creating an unwritable core file in the DisplayServer bin directory. Doing so creates an unwritable core file in the `NFADisplayServer` directory:

```
$ rm core
$ touch core
$ chmod 444 core
```

If Java runs out of memory, the message "java.lang.OutOfMemoryError" appears in your log file. If this happens, you should stop and then restart your DisplayServer, using the stop and start shell scripts provided.

## Recommended Memory Setup

The following memory setup procedures are required for the Solaris and HP-UX platforms. For more information about memory setup, refer to the `/opt/CSCOnfa/NFAServer/bin/README` file.

Solaris Platform

This memory setup procedure runs on the Solaris platform in normal running mode with all applications other than DisplayServer active.

**Step 1** Verify the amount of available logical memory on your workstation. You can run `vmstat` to check your available memory. (The first line of `vmstat`'s output is not valid.)

```
$ vmstat 5

procs          memory                page          ...
r b w   swap  free  re  mf pi po fr de...
0 0 0   74312 51536  0  16 64 20 46  0...
0 0 0  420056 223296  0   1  6  0  0  0...
0 0 0  420056 223280  0   0  4  0  0  0...
0 0 0  420056 223272  0   0  0  0  0  0...
0 0 0  420056 223272  0   0  0  0  0  0...
0 0 0  420056 223264  0   0  1  0  0  0...
...
```

The column "swap" shows the number of kilobytes of space available. Divide this number by 1024 to calculate the amount of memory available. For example, 420056 divided by 1024 is approximately 410 MB.

```
SWAP_AVAIL = 410 MB
```

Another method of verifying the amount of swap space available is to run the "swap -s" program. For example:

```
$ swap -s
total: 117808k bytes allocated + 60160k reserved = \
      117968k used, 550152k available
```

**Step 2** Configure the DisplayServer to allow a large amount of unused memory. The largest reasonable starting value for MaxMB is

```
$ SWAP_AVAILABLE - ALLOWANCE
```

The allowance should be at least 100 MB. For example, a reasonable value is

```
MaxMB = (410-100) = 310 => 224
```

The value 224 is selected because this is 32 MB less than the physical memory in the workstation. The DisplayServer program can use up to an additional 76 MB of memory, so a cushion of 110 MB is allowed. For example:

```
410 MB available - (224 + 76)MB = 110 MB
```

For optimum performance, keep the MaxMB below the amount of actual RAM you have in your workstation. This is particularly important if you plan to process up to the MaxMB of "Detail" disk data in a single command. Failure to allow for some "headroom" of physical RAM can cause disk thrashing while you are accessing the dispersed locations of data over the range of flow "keys." CPU utilization can drop below 2 percent, and the time to get a response to a command can be very large.

For example, on a busy Ultra-1 with 256 MB of actual RAM, "headroom" of 32 MB, and MaxMB = 224, the performance is sufficient to minimize disk thrashing.

On a moderately busy ULTRA-1, the expected performance for the DisplayServer is approximately 0.6 to 3.5 MB of disk data per second. Processing and sorting 400 MB of HostMatrix data should therefore take about 90 seconds, and 100 MB of DetailHostMatrix could require about 90 seconds. Performance varies with the amount of paging/context switching required.

If you must change the maxMB setting, run "vmstat 60" to monitor your swap space and activity when the DisplayServer is processing large volumes of "Detail\*" aggregation scheme data. The "perfmeter" program is also useful when you are tuning for the "spot" of maximal processing storage capacity with minimal disk thrashing.

## HP-UX Platform

This memory setup procedure runs on the HP-UX platform. The HP-UX memory setup procedure is based on the assumption that you have the recommended 256 MB of physical RAM and at least 350 MB of free logical memory space.

---

**Note** The HP-UX shell tools top and vmstat are not available with the NetFlow FlowAnalyzer application Version 1.0.

---

The MaxMB parameter in the NFADS.resources file limits the amount of memory used to store data during command processing. The MaxMB value should not be more than the actual amount of physical RAM in your workstation minus 32 MB. If you have 256 MB of RAM, you should keep MaxMB  $\leq 224$  in order to minimize disk thrashing and poor performance. The recommended value for the workstation is "MaxMB 224."

The workstation needs to be configured so that the DisplayServer application is allowed to use the MaxMB memory space, plus about 32 MB. The system parameter "maxdsize" is the kernel "processlimit" of the amount of memory a single application is allowed to use. The recommended value is maxdsize  $\geq (\text{MaxMB} + 32)$  MB. Consequently, the maxdsize should be configured to represent at least 256 MB on our recommended platform.

**Step 1** Monitor the amount of memory used and the amount of memory remaining. Run "swapinfo" under root access to monitor logical memory:

```
$ swapinfo
while [ 1 ]
do
/usr/sbin/swapinfo -m
sleep 3
done
```

If you run the "top" program while processing a large command, you may find that the "%CPU" is small; the process may be causing disk thrashing.

The following example shows the DisplayServer process efficiently processing several commands simultaneously:

```
$ /bin/top
System: bee                               Mon Jun 30 14:23:56 1997
Load averages: 1.03, 0.76, 0.78
101 processes: 99 sleeping, 2 running
Cpu states:
LOAD   USER   NICE   SYS   IDLE  BLOCK  SWAIT  INTR  SSYS
1.03  98.0%  0.0%  2.0%  0.0%  0.0%  0.0%  0.0%  0.0%

Memory: 53456K (50940K) real, 71836K (60848K) virtual, 5228K free Page# 1/8

  TTU  PID USERNAME PRI NI  SIZE  RES  STATE  TIME %WCPU %CPU COMMAND
  p4 26209 jwiggins 236 20 48156K 48220K run    2:36 97.90 97.72 java
  p7 26146 jwiggins 168 20   316K   312K sleep  0:02  0.38  0.38 top
  p8 26257 jwiggins 178 20   316K   328K run    0:00  0.40  0.38 top
  ?   859 root      154 20   432K   200K sleep  1:29  0.30  0.30 automount
... other processes ...
```

- Step 2** Calculate the amount of available memory. It is advisable to keep MaxMB below the amount of actual RAM you have on your workstation. This is particularly important if you plan to process up to the MaxMB of "Detail" disk data in a single command. Failure to allow for some "headroom" of physical RAM will cause disk thrashing while you are accessing the dispersed locations of data over the range of flow "keys."

### Managing Files on Solaris

The NFADisplayServer running on a Solaris host can serve data only from the host's local disks and from valid NFS mounts to that machine. An automounter may allow your NFADisplayServer's workstation to NFS mount many shared file systems that are not really valid NFS mount points.

To obtain a list of valid NFS mount points from a NetFlow database machine, log in to the database machine and type the command **showmount -e**. For example, suppose that you have data on your\_database\_WS:

```
$ /usr/sbin/showmount -e
export list for your_database_WS:
/u0                               [list_of_hosts_having_permission]
/u1                               [list_of_hosts_having_permission]
```

Any of the Solaris machines on the list\_of\_hosts\_having\_permission can run the NFADisplayServer and serve any local NetFlow data they have in addition to the your\_database\_WS. You may also obtain the same valid list without logging in to the database workstation by typing

```
$ /usr/sbin/showmount -e your_database_WS
```

You can expect better performance if you use file structures local to the NFADisplayServer's workstation.

### File Export

The files exported from the DisplayServer if you use the Export button are located in the exported\_files directory. The location for a standard installation is

```
/opt/CSCOnfa/NFAserver/exported_files
```

## Running the DisplayServer

This section describes how to start and stop the DisplayServer and check the DisplayServer's status while you are running the FlowAnalyzer on the Solaris platform.

### Start the Process

To start the DisplayServer, run the start.DisplayServer shell script. The DisplayServer command starts and generates a log file of sessions.

```
$ /opt/CSCOnfa/NFAServer/bin/start.DisplayServer [server_logfile]
```

If you do not include a <server\_logfile>, the filename **server.out** is used. If a server\_logfile already exists, the logfile is stored in the lowest-numbered server\_logfileNUM (NUM is a nonnegative integer).

### Stop the Process

To stop the DisplayServer, run the stop.DisplayServer shell script:

```
$ /opt/CSCOnfa/NFAServer/bin/stop.DisplayServer
```

### Check the Status

To check the status of the DisplayServer, run the check.DisplayServer shell script:

```
$ /opt/CSCOnfa/NFAServer/bin/check.DisplayServer
```

## Helpful Hints and FAQs

The following are helpful hints and frequently asked questions for both the Solaris and HP-UX platforms.

Q: Do I really need to run with MaxMB 224?

A: No, it does not need to run that high. We have been running the MaxMB with as little as 40 MB for the size of the dynamic memory pool.

The heap and stack could account for up to approximately 32 MB, and the Java environment itself takes approximately 10 to 12 MB. Consequently, if you run with MaxMB = 40, the whole DisplayServer could still take almost 100 MB, in the extreme case. It is not recommended that you run with MaxMB < 40.

When you reduce the MaxMB value, the DisplayServer may run out of memory on Detail\* or CallRecord commands, as well as other commands over longer time intervals, which require large storage for all of the unique keys. It is questionable whether you can service more than a few commands at one time without running out of memory. If you are tight on memory, you might want to use MaxMB = 128.

Note 1: The MaxMB should be no larger than (amount of physical memory – 32 MB) unless you are willing to live with *severe* disk thrashing, and intervals up to 50 times longer for a command to finish, when a large command requiring most of this MaxMB is serviced.

Note 2: Why so much memory? Because the DisplayServer is returning the "topN" flows, it needs to maintain storage for each flow or unique "key" in the database. For example, if you request a scan of time with files containing a 300 MB set of files, and a third of the entries in the files are unique, over 100 MB of memory is required just for storage of the keys for each unique flow.

Note 3: CSV\_FORMAT: The DisplayServer assumes that the CSV\_FORMAT is *not* selected in your nfcollector "nf.resources" file. This means that the field delimiter is assumed to be the vertical bar character ( | ). The DisplayServer will run with varying field delimiters for each file, but it will be a little faster if you keep the field delimiter as the nfcollector default ( | ).

### HP-UX Platform

Note 1: For HP-UX only. Setting maxdsize, the maximum memory a single process may consume. You may check or set the maxdsize of your HP-UX workstation by using HP's tool "sam," the system administration manager. You need root privileges to run this tool. To check the setting, follow these steps:

- Step 1**    /usr/sbin/sam.
- Step 2**    Select "Kernel Configuration."
- Step 3**    Select "Configurable Parameters."
- Step 4**    Scroll down to "maxdsize" using the arrow keys.
- Step 5**    When you are finished, use Tab key to select "File," CR to get pulldown menu, and E to exit.

The "Current Value" and "Pending Value" are in the number of octets. To obtain the megabyte values, divide these numbers by  $(1024 \times 1024) = 1,048,576$ . Details on how to set this and other parameters are on your online manual page, "man sam."

## Cisco Connection Online

Cisco Connection Online (CCO) is Cisco Systems' primary, real-time support channel. Maintenance customers and partners can self-register on CCO to obtain additional information and services.

Available 24 hours a day, 7 days a week, CCO provides a wealth of standard and value-added services to Cisco's customers and business partners. CCO services include product information, product documentation, software updates, release notes, technical tips, the Bug Navigator, configuration notes, brochures, descriptions of service offerings, and download access to public and authorized files.

CCO serves a wide variety of users through two interfaces that are updated and enhanced simultaneously: a character-based version and a multimedia version that resides on the World Wide Web (WWW). The character-based CCO supports Zmodem, Kermit, Xmodem, FTP, and Internet e-mail, and it is excellent for quick access to information over lower bandwidths. The WWW version of CCO provides richly formatted documents with photographs, figures, graphics, and video, as well as hyperlinks to related information.

You can access CCO in the following ways:

- WWW: <http://www.cisco.com>
- WWW: <http://www-europe.cisco.com>
- WWW: <http://www-china.cisco.com>
- Telnet: [cco.cisco.com](telnet://cco.cisco.com)
- Modem: From North America, 408 526-8070; from Europe, 33 1 64 46 40 82. Use the following terminal settings: VT100 emulation; databits: 8; parity: none; stop bits: 1; and connection rates up to 28.8 kbps.

For a copy of CCO's Frequently Asked Questions (FAQ), contact [cco-help@cisco.com](mailto:cco-help@cisco.com). For additional information, contact [cco-team@cisco.com](mailto:cco-team@cisco.com).

---

**Note** If you are a network administrator and need personal technical assistance with a Cisco product that is under warranty or covered by a maintenance contract, contact Cisco's Technical Assistance Center (TAC) at 800 553-2447, 408 526-7209, or [tac@cisco.com](mailto:tac@cisco.com). To obtain general information about Cisco Systems, Cisco products, or upgrades, contact 800 553-6387, 408 526-7208, or [cs-rep@cisco.com](mailto:cs-rep@cisco.com).

---

## Documentation CD-ROM

Cisco documentation and additional literature are available in a CD-ROM package, which ships with your product. The Documentation CD-ROM, a member of the Cisco Connection Family, is updated monthly. Therefore, it might be more up to date than printed documentation. To order additional copies of the Documentation CD-ROM, contact your local sales representative or call customer service. The CD-ROM package is available as a single package or through an annual subscription. You can also access Cisco documentation on the World Wide Web at <http://www.cisco.com>, <http://www-china.cisco.com>, or <http://www-europe.cisco.com>.

---

This document is to be used in conjunction with the *NetFlow FlowAnalyzer Installation and User Guide* publication.

AccessPath, AtmDirector, Cache Director System, CD-PAC, Cisco IOS, the Cisco IOS logo, *CiscoLink*, the Cisco Powered Network logo, ClickStart, ControlStream, Fast Step, FragmentFree, IGX, JumpStart, LAN<sup>2</sup>LAN Enterprise, LAN<sup>2</sup>LAN Remote Office, MICA, NetBeyond, NetFlow, Netsys Technologies, *Packet*, PIX, Point and Click Internetworking, RouteStream, SMARTnet, StrataSphere, StrataSphere BILLder, StrataSphere Connection Manager, StrataSphere Modeler, StrataSphere Optimizer, Stratm, StreamView, SwitchProbe, *The Cell*, TokenSwitch, TrafficDirector, VirtualStream, VlanDirector, Workgroup Director, Workgroup Stack, and XCI are trademarks; The Network Works. No Excuses. is a service mark; and BPX, Catalyst, Cisco, Cisco Systems, the Cisco Systems logo, EtherChannel, FastHub, FastPacket, ForeSight, IPX, LightStream, OptiClass, Phase/IP, StrataCom, and StrataView Plus are registered trademarks of Cisco Systems, Inc. in the U.S. and certain other countries. All other trademarks mentioned in this document are the property of their respective owners.

Copyright © 1997, Cisco Systems, Inc.  
All rights reserved. Printed in USA.  
978R