Managing the FlowAnalyzer

This chapter provides information for managing the FlowAnalyzer application. It contains the following sections:

- Accessing NetFlow Data Files—This section explains why data stored on nonlocal file systems can be inaccessible to the FlowAnalyzer application.
- Managing FlowAnalyzer Files and Directories—This section describes directory maintenance tasks and provides tips on how to efficiently use disk space on your FlowAnalyzer workstation.
- Optimizing FlowAnalyzer Memory Use—This section provides tips on how to efficiently use memory on your FlowAnalyzer workstation.
- Controlling FlowAnalyzer Modules—This section tells you how to start, stop, and check the status of the modules that make up the FlowAnalyzer application.
- Performance Tuning Considerations—This section explains how fast the FlowAnalyzer application can be expected to run, given the amount of NetFlow data being manipulated and the platform on which the FlowAnalyzer application is running.

Accessing NetFlow Data Files

The Display module of the FlowAnalyzer application might not be able to display desired NetFlow data if that data is not stored on an accessible file system. For example, the DisplayServer module can retrieve NetFlow data only from the following:

- Local disks attached to workstations on which the DisplayServer module is running
- File systems with valid Network file System (NFS) mounts to the workstation on which the DisplayServer module is running

In order for NFS mounts to work as expected in your FlowAnalyzer environment, such mounts must pertain to NetFlow data operations involving the local disks of HP-UX platforms exclusively, or to NetFlow data operations involving the local disks of Solaris platforms exclusively.

In other words, an HP-UX-based system can read a disk only on another HP-UX-based platform. Likewise, a Solaris-based system can read a disk only on another Solaris-based platform.

Note If the host DisplayServer workstation makes use of an automounter facility, this facility can mount many shared file systems that may not be valid NFS mounting hosts. In such cases, the DisplayServer module might not be able to retrieve NetFlow data stored on such file systems.

To obtain a list of valid NFS mounting hosts for a NetFlow database workstation in your network, you can issue the **showmount -e** command.

For example, if you have stored Netflow data in a file system called “my_database_WS,” you can issue the **showmount -e** command, as shown below, to list the available hosts in your network that have NFS mounting permission for this file system:

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

Note If you issue the **showmount -e** command from a Solaris-based platform, you may see hosts listed in the command output that pertain to file systems stored on HP-UX platforms. However, you will not be able to read NetFlow data from such platforms.

You can run the DisplayServer module and service user requests for NetFlow data stored on disk volumes anywhere in the network, provided that such volumes are associated with a platform of the same type as that being used by the “my_database_WS” file system.

To obtain valid NFS mounts to file systems not listed in the output of the **showmount** command, you will need to consult with your system administrator.

Note The performance of the DisplayServer module will be optimized if you use file systems that are local to the workstation on which the DisplayServer is running.

Managing FlowAnalyzer Files and Directories

This section describes the files and directories used by the FlowAnalyzer application. It also explains how to manage and maintain these entities.

Maintaining the UtilityServer Output Directories

To monitor and maintain the UtilityServer module, you should perform the following tasks periodically:

- Check the log files for errors
- Delete old log files

The UtilityServer module stores log files in the directory /opt/CSCOnfa/NFAUtility/logs. These log files are numbered sequentially, as shown below:

— NFAU.log

— NFAU1.log

— NFAU2.log

— Etc.

Whenever the UtilityServer module is started, a new log file is created.

In addition to the /opt/CSCOnfa/NFAUtility/logs directory, the UtilityServer module writes output to the /opt/CSCOnfa/NFAUtility/data directory. This directory contains various data files that are used by the UtilityServer module and the DisplayServer module. Hence, the directory is for internal use only by the FlowAnalyzer application.

Maintaining the Exported Files Directory

The files exported from the DisplayServer module (when you invoke the **File, Export** pull-down function in the Display module window) are stored in a file named "exported_files."

For a standard FlowAnalyzer installation, this file is stored in the following directory:

```
/opt/CSCOnfa/NFAServer/exported_files
```

Periodically, you should remove any files from this directory that are no longer needed.

Maintaining the DisplayServer Cache Directory

A NoWait feature of the DisplayServer module is used to request a background job or to service a NetFlow command that generates voluminous output. In the latter case, the NoWait option enables the command output to be stored in the DisplayServer's /opt/CSCOnfa/NFAServer/Cache directory, from which stored data can be retrieved in smaller segments.

This section explains how you can define parameters in the NFADS.resources file to limit the size and number of data files stored in the /opt/CSCOnfa/NFAServer/Cache directory.

For instructions on editing the NFADS.resources file, see the section entitled "Customizing the NFADS.resources File" in Chapter 2.

The parameters that you can define in the NFADS.resources file to manage the storage of NoWait command response files in the /opt/CSCOnfa/NFAServer/Cache directory are described briefly below:

- Max_Stored_NoWait_MB—The maximum aggregate disk space allotted for storing the NoWait command response files in the /opt/CSCOnfa/NFAServer/Cache directory.
- Max_Stored_NoWait_Files—The maximum number of NoWait command response files that can be stored in the /opt/CSCOnfa/NFAServer/Cache directory.
- Tgt_Stored_NoWait_Percent—The percentage of NoWait command response files that will be deleted from the /opt/CSCOnfa/NFAServer/Cache directory when the value of either the Max_Stored_NoWait_MB parameter or the Max_Stored_NoWait_Files parameter is exceeded.

The DisplayServer module checks the state of the /opt/CSCOnfa/NFAServer/Cache directory upon completion of each NoWait command. If the value of either the Max_Stored_NoWait_MB parameter or the Max_Stored_NoWait_Files parameter is exceeded, the DisplayServer module deletes files in chronological order until both the target storage value (in MB) and the target file count are reached.

The DisplayServer module multiplies the value of the Tgt_Stored_NoWait_Percent parameter by the value of both the Max_Stored_NoWait_MB parameter and the Max_Stored_NoWait_Files parameter to determine how much file storage space and how many files, respectively, must be deleted.

Note, however, that the file relating to the most recently completed NoWait command is never deleted, even if its size exceeds the value of the Max_Stored_NoWait_MB parameter. Hence, there will always be at least one file in the /opt/CSCOnfa/NFAServer/Cache directory.

Optimizing FlowAnalyzer Memory Use

This section presents procedures that will enable the FlowAnalyzer workstation to use available memory resources efficiently.

Calculating Available Memory and Adjusting the MaxMB Value

This section explains how to calculate the logical memory capacity of the workstation on which the Solaris or HP-UX platform is running and how to fine-tune the system for efficient use of memory.

If your FlowAnalyzer workstation meets the recommended memory guidelines (256 MB of physical memory and 400 MB of free logical memory), you might not need to perform the procedures in this section.

However, if your system does not meet the recommended memory guidelines, or if you are experiencing FlowAnalyzer performance problems, you should perform the appropriate procedure below for your operating platform.

Calculating Available Memory for the Solaris Platform

Use the procedure in this section to calculate the available memory for use by the Solaris platform when all the FlowAnalyzer modules are running (except for the DisplayServer module).

For instructions on starting and stopping the DisplayServer module, see the later section in this chapter entitled “Controlling FlowAnalyzer Modules.”

To calculate the amount of logical memory (swap space) available on your FlowAnalyzer workstation, perform the following procedure:

Step 1 Run the vmstat program with “5” as an argument (to refresh command output every five seconds), as shown below:

```
$ 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... (this line does not yield valid output)
 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...
...
```

Note that the first line of output from the vmstat program is not valid.

The column labeled “swap” shows the number of kilobytes of swap space available on your FlowAnalyzer workstation.

To calculate the amount of memory available in MB, divide this number by 1024.

For example, divide 420056 by 1024, which equals approximately 410 MB of swap space.

Another way to determine the amount of logical memory available for swap space on your FlowAnalyzer workstation is to issue the **swap -s** command, as shown below:

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

Step 2 By means of the MaxMB parameter in the NFADS.resources file, you can configure the DisplayServer module to use a large amount of memory (MaxMB).

The MaxMB parameter in the NFADS.resources file limits the amount of memory that will be used by the FlowAnalyzer to store data when a command is being processed.

See the section entitled "Customizing the NFADS.resources File" in Chapter 2 for instructions on changing the parameters in the NFADS.resources file.

You should observe the following rules in calculating the largest reasonable starting value for the MaxMB parameter:

Rule 1—The MaxMB value must be less than or equal to the following:

Workstation's_Physical_Memory - 32

Rule 2—The MaxMB value must be less than or equal to the following:

SWAP_AVAILABLE - 100

Thus, the value of the MaxMB parameter should be no larger than either of the following:

(a) The workstation's actual physical RAM (256 MB in this case), minus 32 MB (see Rule 1)

(b) The workstation's available swap space (410 MB in this case), minus 100 MB (see Rule 2)

The formula for calculating physical memory is:

MaxMB—Less than or equal to 256 – 32 = 224 MB

The formula for calculating swap space is:

MaxMB—Less than or equal to 410 – 100 = 310 MB

Therefore, the value of the MaxMB parameter should be no larger than 224 MB, which is the smaller of the two results above.

If you violate Rule 1 in setting the value of the MaxMB parameter, severe performance degradation can result, including disk thrashing.

Similarly, if you violate Rule 2 in setting the value of the MaxMB parameter, the system software might run out of swap space, in which case, active processes will be killed.

For information about the expected performance of the FlowAnalyzer, see the later section in this chapter entitled "Performance Tuning Considerations."

Calculating Available Memory for the HP-UX Platform

Use the procedure in this section to calculate the available memory for use by the HP-UX platform.

For purposes of this procedure, it is assumed that your FlowAnalyzer workstation has the recommended 256 MB of physical RAM and at least 350 MB of free logical memory.

The value of the MaxMB parameter in the NFADS.resources file limits the amount of memory that can be used to store NetFlow data during command processing.

The value of the MaxMB parameter should not exceed the actual amount of physical RAM available in your FlowAnalyzer workstation, minus 32 MB.

If your FlowAnalyzer workstation has 256 MB of RAM, you should ensure that the value of the MaxMB parameter is less than or equal to 224 to minimize the possibility of disk thrashing and poor system performance. The recommended value for the MaxMB parameter is 224.

You should configure the FlowAnalyzer workstation so that the DisplayServer is allowed to use the amount of memory defined by the value of the MaxMB parameter, plus approximately 32 MB.

The system parameter “maxdsize” defines the kernel process memory limit that a single application is allowed to use. The recommended value of this parameter is:

MaxMB + 32 MB

The value of the maxdsize parameter should be set to at least 256 MB for use by the HP-UX platform.

To monitor the amount of FlowAnalyzer workstation memory being used and the amount of memory remaining when the DisplayServer module is stressed with a heavy load, perform the following steps:

Step 1 Enter the following commands and observe the resulting output for a few minutes:

```
$ su root
password: <enter the password>
# while [ 1 ]
do
/usr/sbin/swapinfo -m
sleep 5
echo
done
```

The output is refreshed every 5 seconds and takes the following form:

```
             Mb      Mb      Mb   PCT  START/      Mb
TYPE      AVAIL    USED    FREE  USED   LIMIT RESERVE  PRI  NAME
dev         300       0     298    0% 1789952       -    1  /dev/dsk/c0t6d0
reserve       -      61     -61
memory      198     126      72   64%
```

Step 2 From the output, you can determine the amount of logical memory available for use as swap space.

In the output, unused memory is shown in the “Mb FREE” column; the percentage of memory in use is shown in the “PCT USED” column.

Step 3 When you want to stop monitoring memory use, issue the **^C** command.

Step 4 Edit the NFADS.resources file to configure the DisplayServer module to use a large amount of unused memory (MaxMB).

See the section entitled “Customizing the NFADS.resources File” in Chapter 2 for instructions on changing the value of parameters in the NFADS.resources file.

If you determined in Step 2 that the system is running out of memory resources, reduce the value of the MaxMB parameter in the NFADS.resources file.

Conversely, if the FlowAnalyzer workstation has memory to spare, increase the value of the MaxMB parameter accordingly.

Note If you set the value of the MaxMB parameter in the NFADS.resources file too high, disk thrashing may occur.

In general, you should keep the value of the MaxMB parameter at a lower level than the amount of physical RAM in your FlowAnalyzer workstation. Maintaining the value of the MaxMB parameter at a lower level is particularly important if you plan to issue a network command whose memory consumption in processing NetFlow data approaches the value of the MaxMB parameter. Failure to maintain the value of the MaxMB parameter at a lower level will result in disk thrashing.

Controlling FlowAnalyzer Modules

This section tells you how to start, stop, and check the status of FlowAnalyzer modules individually.

If you want to start all the FlowAnalyzer modules at once, see the section entitled “Starting the FlowAnalyzer” in Chapter 2.

If you want to start the FlowAnalyzer modules individually, start them in the following order:

1. TrawhoisServer module
2. UtilityServer module
3. DisplayServer module
4. Display module

Running the TrawhoisServer Module

This section tells you how to start, check the status of, and stop the TrawhoisServer module.

Starting the TrawhoisServer Module

To start the TrawhoisServer module, perform the following steps:

Step 1 Log in as root:

```
$ su root
password: <enter the password>
```

Step 2 Run the start.TrawhoisServer shell script, as shown below:

```
# /opt/CSCOnfa/NFATrawhois/trawho-2.03/bin/start.TrawhoisServer
```

Checking the Status of the TrawhoisServer Module

To check the status of an active TrawhoisServer process, run the check.TrawhoisServer shell script, as shown below:

```
$ /opt/CSCOnfa/NFATrawhois/trawho-2.03/bin/check.TrawhoisServer
```

Stopping the TrawhoisServer Module

To stop an active TrawhoisServer process, perform the following steps:

Step 1 Log in as root:

```
$ su root
password: <enter the password>
```

Step 2 Run the stop.TrawhoisServer shell script, as shown below:

```
# /opt/CSCOnfa/NFATrawhois/trawho-2.03/bin/stop.TrawhoisServer
```

Running the UtilityServer Module

This section tells you how to start, check the status of, and stop the UtilityServer module.

Starting the UtilityServer Module

To start the UtilityServer module, perform the following steps:

Step 1 Log in as root:

```
$ su root
password: <enter the password>
```

Step 2 Run the start.UtilityServer shell script, as shown below:

```
# /opt/CSCOnfa/NFAUtility/bin/start.UtilityServer
```

Checking the Status of the UtilityServer Module

To check the status of an active UtilityServer process, run the check.UtilityServer shell script, as shown below:

```
$ /opt/CSCOnfa/NFAUtility/bin/check.UtilityServer
```

Stopping the UtilityServer Module

To stop an active UtilityServer process, perform the following steps:

Step 1 Log in as root:

```
$ su root
password: <enter the password>
```

Step 2 Run the stop.UtilityServer shell script, as shown below:

```
# /opt/CSCOnfa/NFAUtility/bin/stop.UtilityServer
```

Running the DisplayServer Module

This section tells you how to start, check the status of, and stop the DisplayServer module.

Starting the DisplayServer Module

To start the DisplayServer module, perform the following steps:

Step 1 Log in as root:

```
$ su root
password: <enter the password>
```

Step 2 To start the DisplayServer module, run the start.DisplayServer shell script, as shown below:

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

This command starts the DisplayServer module and generates a log file of DisplayServer sessions. If you do not specify a <server_logfile>, the system uses the file name “server.out” by default.

If a server_logfile already exists, the log file output is stored in the lowest-numbered server_logfileNUM file (with “NUM” as a non-negative integer).

Note The startup script, start.dsa, in the /opt/CSCOnfa/NFAServer/bin/ directory is for exclusive use by Cisco engineering personnel in doing advanced troubleshooting. Do not run this script.

Checking the Status of the DisplayServer Module

To check the status of an active DisplayServer process, run the check.DisplayServer shell script, as shown below:

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

Stopping the DisplayServer Module

To stop an active DisplayServer process, perform the following steps:

Step 1 Log in as root:

```
$ su root
password: <enter the password>
```

Step 2 Run the stop.DisplayServer shell script, as show below:

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

Running the Display Module

This section tells you how to start and stop the Display module.

Starting the Display Module

To start the Display module, enter the following command:

```
$ /opt/CSCOnfa/NFADisplay/bin/start.Display
```

Stopping the Display Module

You can stop the Display module from within the module itself.

To exit from the Display module, perform the following steps:

Step 1 If you wish to save any changes made to the data tree structure in the Netflow Data area of the Display module window, select **File, Save** from the pull-down menu of the Display module window.

Step 2 Select **File, Quit** from the pull-down menu.

Performance Tuning Considerations

The FlowAnalyzer may run slowly due to the amount of NetFlow data it is processing.

In a busy network, traffic data can be amassed at a rate of hundreds of megabytes per hour. Hence, searching or sorting through what you may think is a limited set of NetFlow data can, in fact, turn out to be a huge, time-consuming task.

Table 4-1 shows the performance of the FlowAnalyzer application doing search operations on two different workstations. In each case, FlowAnalyzer performance was evaluated during search operations on 100 MB of NetFlow data.

Table 4-1 Comparative Performance of FlowAnalyzer Workstations

Platform	Processing Speed	Search Time for 100 MB of Data
ULTRA-1 workstation	1 MB per second	96 seconds
SPARC-20 workstation	0.5 MB per second	206 seconds

FlowAnalyzer performance varies with the amount of paging/context switching required in processing Netflow data. If you need to change the value of the MaxMB parameter to improve system performance (see the earlier section entitled “Optimizing FlowAnalyzer Memory Use”), it is useful to run the vmstat 60 program (refreshes output every 60 seconds) to monitor swap space and disk activity when the DisplayServer module is processing large volumes of NetFlow data for any one of the “Detail” aggregation schemes.

It is also useful to run the perfmeter program when you are tuning the NetFlow system to ensure maximum storage capacity for processed NetFlow data. Maximizing storage capacity for NetFlow data reduces the likelihood of disk thrashing.

You should also note that better system performance always results when you access desired NetFlow data on a workstation that is also the host for the DisplayServer module, as opposed to mounting a non-local file system for obtaining desired Netflow data.