Cisco Multi NetFlow Collector Installation and Configuration Guide

Release 6.0

Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
http://www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 527-0883

Customer Order Number:
Text Part Number: OL-12883-01

Cisco Multi NetFlow Collector Installation and Configuration Guide

CONTENTS

About This Guide

Objective

The *Cisco Multi NetFlow Collector Installation and Configuration Guide* describes the Cisco Multi NetFlow Collector application, which is used with the NetFlow services data export feature on Cisco routers and Catalyst switches. This document also describes the system requirements that must be met to install the Cisco Multi NetFlow Collector product, as well as, how to install, start, and configure Cisco Multi NetFlow Collector.

NetFlow services consist of high-performance IP switching features that capture a rich set of traffic statistics exported from routers and switches while they perform their switching function. Cisco NetFlow Collector provides fast, scalable, and economical data collection from multiple export devices exporting NetFlow data records.

Cisco NetFlow Collector, Release 6.0 introduces a tiered netflow collection architecture that provides increased scalability and performance. The role of the first tier (Tier 1) maps to the NFC functionality of Cisco NetFlow Collector 5.0.3 with the addition of new features described in *Release Notes for Cisco NetFlow Collector, Release 6.0*.

Cisco NetFlow Collector, Release 6.0 supports new Cisco NetFlow Collector Tier 2 functionality, also referred to as Multi NetFlow Collector. The Multi NetFlow Collector runs on separate server hardware and provides an aggregation layer that correlates data from several Tier 1 instances.

Prior to reading this manual, you should read the *Release Notes for Cisco Multi NetFlow Collector, Release 6.0* document. These release notes provide information about known software and documentation problems and any last minute information about the Multi NetFlow Collector software not available when this guide was produced.

Audience

This guide is intended primarily for individuals with network and system administration skills. You should have a basic understanding of network design, operation, and terminology, as well as familiarity with your own network configurations. You also must have a basic familiarity with Web browsers, Red Hat Enterprise Linux, or Sun Microsystem's Solaris Operating System.

How This Guide Is Organized

This guide is organized as follows:

Chapter 1, "Overview,"describes the Cisco Multi NetFlow Collector application.

Chapter 2, "Installing the Cisco Multi NetFlow Collector,"describes how to install the Cisco Multi NetFlow Collector application.

Chapter 3, "Starting and Stopping Cisco Multi NetFlow Collector,"describes the Cisco Multi NetFlow Collector components and how to start and stop MNFC.

Appendix A, "Troubleshooting the Multi NetFlow Collector," contains troubleshooting information in case you encounter problems while using the Cisco Multi NetFlow Collector.

An Index is also provided.

Command Syntax Conventions

Table 1 describes the syntax used with the commands in this document.

Table 1 Command Syntax Guide

Convention	Description
boldface	Commands and keywords.
italic	Command input that is supplied by you.
[ ]	Keywords or arguments that appear within square brackets are optional.
{ x | x | x }	A choice of keywords (represented by x) appears in braces separated by vertical bars. You must select one.
^ or Ctrl	Represent the key labeled *Control*. For example, when you read ^*D* or *Ctrl-D*, you should hold down the Control key while you press the D key.
`screen font`	Examples of information displayed on the screen.
boldface screen font	Examples of information that you must enter.
< >	Nonprinting characters, such as passwords, appear in angled brackets.
[ ]	Default responses to system prompts appear in square brackets.

Obtaining Documentation, Obtaining Support, and Security Guidelines

For information on obtaining documentation, obtaining support, providing documentation feedback, security guidelines, and also recommended aliases and general Cisco documents, see the monthly *What's New in Cisco Product Documentation*, which also lists all new and revised Cisco technical documentation, at:

http://www.cisco.com/en/US/docs/general/whatsnew/whatsnew.html

CHAPTER 1

Overview

This chapter describes the Cisco NetFlow Collector (NFC) and Multi NetFlow Collector (MNFC) applications, that are used with the NetFlow services data export feature on Cisco routers and Catalyst switches.

This chapter includes the following sections:

What Are NetFlow Services?

NetFlow services consist of high-performance IP switching features that capture a rich set of traffic statistics exported from routers and switches while they perform their switching functions. The exported NetFlow data consists of traffic flows, which are unidirectional sequences of packets between a particular source device and destination device that share the same protocol and transport-layer information. The captured traffic statistics can be used for a wide variety of purposes, such as network analysis and planning, network management, accounting, billing, and data mining.

Because of their unidirectional nature, flows from a client to a server are differentiated from flows from the server to the client. Flows are also differentiated on the basis of protocol. For example, Hypertext Transfer Protocol (HTTP) Web packets from a particular source host to a particular destination host constitute a separate flow from File Transfer Protocol (FTP) file transfer packets between the same pair of hosts.

Routers and switches identify flows by looking for the following fields within IP packets:

- Source IP address
- Destination IP address
- Source port number
- Destination port number
- Protocol type
- Type of service (ToS)
- Input interface

Catalyst 5000 series switches can identify flows by looking at a subset of these fields. For example, they can identify flows by source and destination address only.

Note For Catalyst 5000 series switches, the analog to NetFlow services is integrated Multilayer Switching (MLS) management. Included are products, utilities, and partner applications designed to gather flow statistics, export the statistics, and collect and perform data reduction on the exported statistics. MLS management then forwards them to consumer applications for traffic monitoring, planning, and accounting.

NetFlow Services Device and IOS Release Support

You can find the most up-to-date information available to help you determine the compatibility among different Cisco hardware platforms, Cisco IOS software releases, and supported NetFlow data export versions at the following URL:

http://tools.cisco.com/ITDIT/CFN/Dispatch?SearchText=Netflow&act=featSelect&rnFeatId=null&featStartsWith=&task=TextSearch&altrole=

Note Except for descriptions requiring references to specific router or switch platforms, the remainder of this chapter and the remaining chapters of this guide use the term export device instead of the terms router and switch.

NetFlow Data Export

NetFlow data export makes NetFlow traffic statistics available for purposes of network planning, billing, and so on. An export device configured for NetFlow data export maintains a flow cache used to capture flow-based traffic statistics. Traffic statistics for each active flow are maintained in the cache and are updated when packets within each flow are switched. Periodically, summary traffic statistics for all expired flows are exported from the export device by means of User Datagram Protocol (UDP) datagrams, which NetFlow Collector receives and processes.

How and When Flow Statistics Are Exported

NetFlow data exported from the export device contains NetFlow statistics for the flow cache entries that have expired since the last export. Flow cache entries expire and are flushed from the cache when one of the following conditions occurs:

- The transport protocol indicates that the connection is completed (TCP FIN) plus a small delay to allow for the completion of the FIN acknowledgment handshaking.
- Traffic inactivity exceeds 15 seconds.

For flows that remain continuously active, flow cache entries expire after a specified period of time, for example every 30 minutes, to ensure periodic reporting of active flows.

NetFlow data export packets are sent to a user-specified destination, such as the workstation running NetFlow Collector, either when the number of recently expired flows reaches a predetermined maximum, or every second-whichever occurs first. For:

- Version 1 datagrams, up to 24 flows can be sent in a single UDP datagram of approximately 1200 bytes.

- Version 5 datagrams, up to 30 flows can be sent in a single UDP datagram of approximately 1500 bytes.
- Version 7 datagrams, up to 27 flows can be sent in a single UDP datagram of approximately 1500 bytes.
- Version 8 datagrams, the number of flows sent in a single UDP datagram varies by aggregation scheme.
- Version 9 datagrams, the number of flows is variable, and depends on the number and size of fields defined in one or more templates.

See Appendix B, "NetFlow Export Datagram Formats," in the *Cisco NetFlow Collector User Guide* for details on all versions of the NetFlow data export format.

NetFlow Data Export Formats

NetFlow exports flow information in UDP datagrams in one of five formats: Version 1 (V1), Version 5 (V5), Version 7 (V7), Version 8 (V8), or Version 9 (V9).

Version 1 is the original format supported in the initial NetFlow releases. Version 5 is an enhancement that adds Border Gateway Protocol (BGP) autonomous system information and flow sequence numbers. Version 7 is an enhancement that exclusively supports Cisco Catalyst 5000 series switches equipped with a NetFlow feature card (NFFC). V7 is not compatible with Cisco routers. Version 8 is an enhancement that adds router-based aggregation schemes. Version 9 is an enhancement to support different technologies such as Multicast, Internet Protocol Security (IPSec), and Multi Protocol Label Switching (MPLS). NetFlow Collector Release 5.0 can collect, filter, and aggregate Version 9 data in the same way it does for NetFlow Data Export Versions 1 through 8.

Versions 2, 3, 4, and 6 are not supported by NetFlow Collector. For more information on the distinctions among the NetFlow data export formats, see Appendix B, "NetFlow Export Datagram Formats," in the *Cisco NetFlow Collector User Guide.*

The following types of information are part of the detailed traffic statistics:

- Source and destination IP addresses
- Next hop address
- Input and output interface numbers
- Number of packets in the flow
- Total bytes (octets) in the flow
- First and last time stamps of packets that were switched as part of this flow
- Source and destination port numbers
- Protocol
- Type of service (ToS)
- Source and destination autonomous system (AS) numbers, either origin or peer (present in V5 and select V8 datagrams)
- Source and destination prefix mask bits (present in V5, V7, and V8 datagrams)
- Shortcut router IP address (present in V7 on Cisco Catalyst 5000 series switches only).

Caution Throughout this publication there are numerous examples of NetFlow Collector input commands and output results. Included are examples of IP addresses. Be aware that IP address examples are not usable IP addresses. The examples do not represent real-life configurations.

What Is Cisco NetFlow Collector?

The Cisco NetFlow Collector application provides fast, scalable, and economical data collection from multiple export devices exporting NetFlow data records. Figure 1-1 shows an example of a typical NetFlow data export scheme. In it, various export devices send export data to user-specified NetFlow Collector UDP and SCTP ports.

Figure 1-1 NetFlow Collector Overview

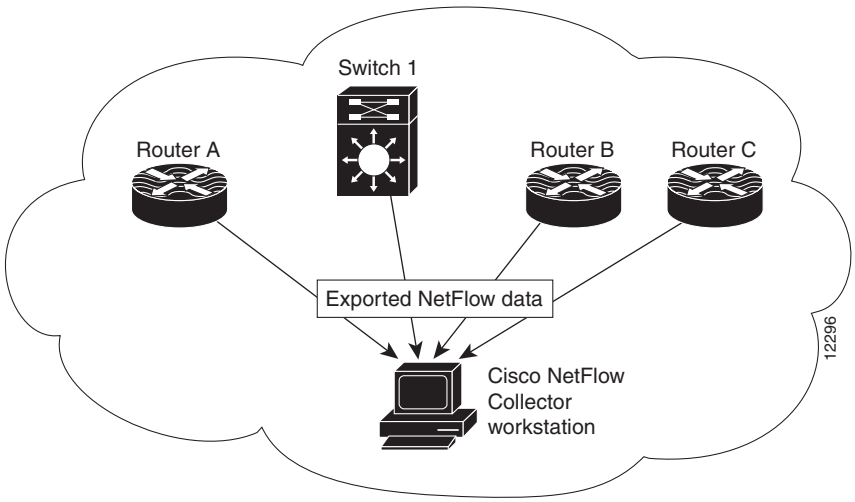

Each of the export devices in this example is configured for NetFlow data export. Part of the configuration information for each export device includes the IP address and the UDP or SCTP port number (a logical port designator) that identify NetFlow Collector as the receiver of flows from this export device. The port number is a user-configurable designator: you can configure NetFlow Collector to listen for flows on a number of different ports, and then configure your export devices so that each device exports flows to a dedicated port, or have a number of devices export flows to the same, shared port.

After you configure and start Cisco NetFlow Collector, it listens to the user-specified UDP and SCTP ports for exported flows from the export devices you have configured for NetFlow data export.

Cisco NetFlow Collector performs the following functions:

- NetFlow data collection from multiple export devices
- Reduction in data volume through filtering and aggregation
- Hierarchical data storage (helps client applications retrieve data)
- File system space management

Cisco NetFlow Collector collects and summarizes (aggregates) data into data files based on user-defined criteria specified in a NetFlow Collector *aggregator*. An *aggregator* is an aggregation task defined by a set of user-configurable attributes that specify how NetFlow Collector summarizes the traffic flows that are received. Two important aggregator attributes are:

- Aggregation schemes – defines the subset of data of interest in a traffic flow, as well as which statistics are kept
- Filter – criteria for accepting or rejecting flows that are aggregated or summarized

Cisco NetFlow Collector provides a set of predefined aggregation schemes to help you collect NetFlow export data and summarize the data (that is, aggregate the flows). You can choose one or more of these aggregation schemes to customize NetFlow Collector for your operating context. Moreover, starting in Release 5.0 you can modify any of the predefined aggregation schemes or define your own aggregation schemes based on them. You can also use filters with aggregation schemes to include or exclude certain types of NetFlow data.

For more information about threads, aggregation schemes, and filters, see Chapter 4, "Customizing the CNS NetFlow Collection Engine," in the *Cisco NetFlow Collector User Guide*.

What Is Cisco Multi NetFlow Collector?

The Cisco Multi NetFlow Collector is a the second-tier application of the NetFlow architecture. MNFC imports the data files resident in multiple NFCs and performs network-level correlation and provides a central view for all distributed Cisco NFC implementations in the network. Figure 1-2 shows an example of a typical NetFlow data export scheme. .

Figure 1-2 Multi NetFlow Collector Overview

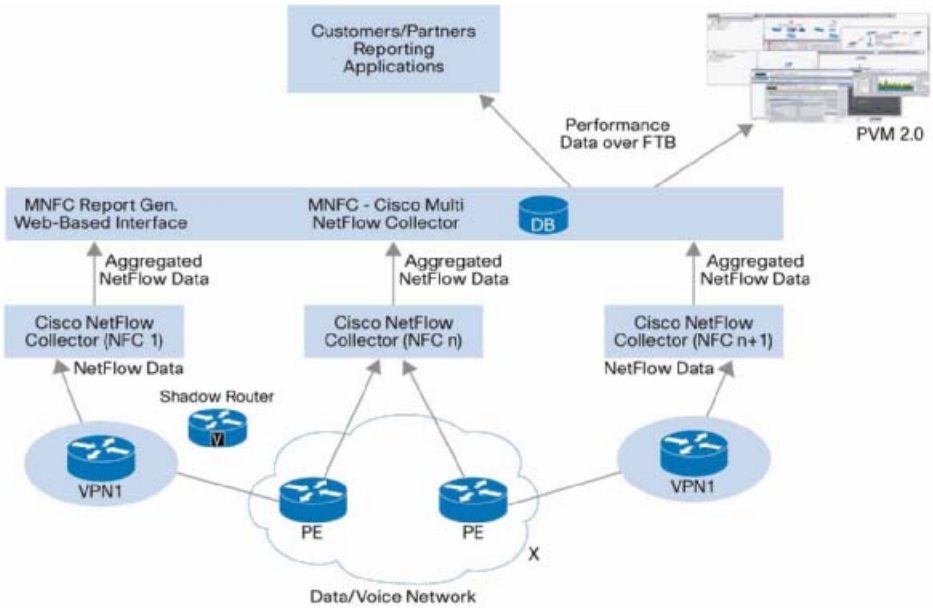

Note Cisco MNFC supports only Cisco NFC Release 6. It does not support previous NFC releases. Cisco MNFC and NFC must run on separate servers.

Table 1-1describes for MNFC features.

Table 1-1 Cisco Multi NetFlow Collector Features

Feature	Benefit
NF-Egress Packets lost and site in-out traffic summary	Monitors packets lost from IP-IP flows. You can use this feature to monitor the point o failure of each link in the network.
PE-PE, PE-CE, CE-PE, and CE-CE data collection	Provides traffic statistics between two IP networks.
Correlation traffic summary for VPN/VRF and VPN/non-VPN	Provides a view of traffic statistics for each VPN based on each VRF. You can classify and report site-to-site and non-VPN/VPN traffic summaries.
Embedded data	Provides centralized storage of all data files from multiple distributed Cisco NFC implementations for the longer period of the trending report.
Report generator	Imports the data files resident in multiple Cisco NFCs to its server to perform network-level correlation, a central view of end-to-end traffic summaries, and classification information.

CHAPTER 2

Installing the Cisco Multi NetFlow Collector

This chapter describes how to install the Cisco Multi NetFlow Collector (MNFC) application. It includes the following sections:

Requirements

Cisco MNFC supports only Cisco NFC Release 6. It does not support previous NFC releases. Cisco MNFC and NFC must run on separate servers.The following sections describe requirements for Cisco Multi NetFlow Collector, Release 6.0.

Note The CPU, RAM, and disk space recommendations listed are minimum requirements. Your actual requirements are determined by your configuration and by the volume and uniqueness of NetFlow data that is received. Actual resource usage can vary greatly depending on these factors.

Supported Operating Systems and Platforms

Multi NetFlow Collector, Release 6.0 supports the following operating systems and platforms:

- Solaris 8, 9, or 10 on a midrange server, such as the Sun Fire V490 with Quad UltraSPARC IV 1.5 GHz processors.
- Red Hat Enterprise Linux 3.0, or 4.0 (ES and AS) on a midrange server, such as an IBM x346 with a single dual-core Intel Xeon 3.8 GHz processor.

Note To serve as the concentrator in a scalable NetFlow solution, the workstation should be dedicated to the Multi NetFlow Collector and should not be running other applications.

Hardware Requirements

The Multi NetFlow Collector, Release 6.0 has the following hardware requirements:

- Minimum of 16 GB RAM, 10K SCSI, dual 70 GB disk, and dual processor on an midrange-level server.

Browser Requirements

The Multi NetFlow Collector, Release 6.0 web-based user interface is compatible with Microsoft Internet Explorer 6 and Mozilla Firefox 1.5 or greater on Windows or UNIX. The web-based UI requires that the browser support a Java virtual machine (JVM) to run applets.

Note The Sun JVM must be used; the JVM version must be 1.5 or higher. You can download Sun JVM 1.5 from the website http://java.sun.com/javase/downloads/index.jsp.

Preparing to Install Multi NetFlow Collector

To install the Multi NetFlow Collector, Release 6.0, you must complete the following:

- Set up nfcuser and informix user accounts
- Prepare raw disk partitions for MNFC's database
- Install the Cisco NetFlow Network Concentrator (CSCOmnfc) package
- Install Informix Dynamic Server (IDS)
- Configure the Informix Dynamic Server
- Set up the MNFC license

These tasks are described in the following sections.

Setting Up User Accounts

To create accounts for the MNFC user and for the IDS administrator, do the following:

Step 1 Verify that the group **informix** exists in **/etc/group**. Enter:

niscat group.org_dir

If the group exists, note its **group id**. If it does not exist, create it using the command **groupadd**.

Step 2 Verify that the user name **informix** exists in **/etc/passwd**. Run **id informix** to verify that the informix user exists.

If the user name **informix** does not exist, create both the user and group using the command **useradd**. Enter:

useradd -u 744 -d /opt/informix_home -s /bin/sh informix

If the informix user exists but is not found in **/etc/group**, enter the following:

useradd -u 744 -g informix -d /opt/informx_home -s /bin/sh informix

See the following example:

useradd -u 744 -d /opt/informix_home -s /bin/sh informix

mkdir /opt/informix_home

chown informix:informix /opt/informix_home

passwd informix

Setting Up Raw Disk Partitions for Informix

For optimal work with database we have to give it raw disk partitions so that RDBMS server will work using its own IO stacks bypassing the file system.

Following is outline of the procedure to prepare raw disk partitions.

Allocating Disk Space

Cisco recommend that the minimal number of hard disk drives required for running MNFC is two, the boot, OS, and file system partitions should be allocated on one disk drive, while other disk drives should be used as raw disks for the database.

Solaris Platform

On a Solaris platform, use the command **format** to view disk space allocation. Below is an example:

```
nfc-ott1# format
Searching for disks...done

AVAILABLE DISK SELECTIONS:
       0. c0t0d0 <SUN18G cyl 7506 alt 2 hd 19 sec 248>
          /pci@1f,4000/scsi@3/sd@0,0
       1. c0t1d0 <SUN36G cyl 24620 alt 2 hd 27 sec 107>
          /pci@1f,4000/scsi@3/sd@1,0
Specify disk (enter its number): 1
selecting c0t1d0
[disk formatted]
...............................

partition> p
Current partition table (original):
Total disk cylinders available: 24620 + 2 (reserved cylinders)

Part      Tag    Flag     Cylinders         Size            Blocks
  0       root    wm       0 -     0        1.41MB    (1/0/0)         2889
```

```
  1 unassigned    wu       1 - 12309       16.96GB    (12309/0/0) 35560701
  2 unassigned    wu   12310 - 24618       16.96GB    (12309/0/0) 35560701
  3     backup    wm       0 - 24619       33.92GB    (24620/0/0) 71127180
  4 unassigned    wm       0                0         (0/0/0)            0
  5 unassigned    wm       0                0         (0/0/0)            0
  6 unassigned    wm       0                0         (0/0/0)            0
  7 unassigned    wm       0                0         (0/0/0)            0
```

The following is an example of the raw partitions created for disk **c0t1d0**:

```
# ls -l /dev/rdsk/c0t1d0s[12]
lrwxrwxrwx   1 root     root          45 Apr 21 12:11 /dev/rdsk/c0t1d0s1 ->
../../devices/pci@1f,4000/scsi@3/sd@1,0:b,raw
lrwxrwxrwx   1 root     root          45 Apr 21 12:11 /dev/rdsk/c0t1d0s2 ->
../../devices/pci@1f,4000/scsi@3/sd@1,0:c,raw
```

To create symbolic links in **/dev** to **/dev/rdsk**, for the example above enter the following:

ln -s /dev/rdsk/c0t1d0s1 /dev/raw_dev1

ln -s /dev/rdsk/c0t1d0s2 /dev/raw_dev2

Step 3 Verify the symbolic links by entering the following:

 # ls -l /dev/raw_dev?

For this example, the following is displayed:

```
lrwxrwxrwx 1 root oldother 18 Oct 12 16:02 /dev/raw_dev1 -> /dev/rdsk/c0t1d0s1
lrwxrwxrwx 1 root oldother 18 Oct 12 16:02 /dev/raw_dev2 -> /dev/rdsk/c0t1d0s2
```

Warning **For MNFC database installation, the symbolic link in /dev must contain raw.**

Set ownership of block devices to **Informix** and chmod to **66**. See the following example:

```
# ls -l /devices/pci@1f,4000/scsi@3/sd@1,0:?,raw
crw-r-----   1 root     sys       32,496 Apr 21 12:11
/devices/pci@1f,4000/scsi@3/sd@1,0:a,raw
crw-rw----   1 informix informix  32,497 Apr 21 12:11
/devices/pci@1f,4000/scsi@3/sd@1,0:b,raw
crw-rw----   1 informix informix  32,498 Apr 21 12:11
/devices/pci@1f,4000/scsi@3/sd@1,0:c,raw
crw-r-----   1 root     sys       32,499 Apr 21 12:11
/devices/pci@1f,4000/scsi@3/sd@1,0:d,raw
crw-r-----   1 root     sys       32,500 Apr 21 12:11
/devices/pci@1f,4000/scsi@3/sd@1,0:e,raw
crw-r-----   1 root     sys       32,501 Apr 21 12:11
/devices/pci@1f,4000/scsi@3/sd@1,0:f,raw
crw-r-----   1 root     sys       32,502 Apr 21 12:11
/devices/pci@1f,4000/scsi@3/sd@1,0:g,raw
crw-r-----   1 root     sys       32,503 Apr 21 12:11
/devices/pci@1f,4000/scsi@3/sd@1,0:h,raw
```

Linux Platform

On a Linux platform, use the command **fdisk** to view disk space allocation. Below is an example for disk **sdb** and partitions **sdb[1,3]**:

```
:# fdisk /dev/sdb
...........

Disk /dev/sdb: 73.4 GB, 73407820800 bytes
255 heads, 63 sectors/track, 8924 cylinders
Units = cylinders of 16065 * 512 = 8225280 bytes

   Device Boot    Start       End    Blocks   Id  System
/dev/sdb1             2      1724  13839997+  83  Linux
/dev/sdb2             1         1      8032   83  Linux
/dev/sdb3          1725      8924  57834000   83  Linux
```

To define raw devices, use the **raw** utility. Refer to **man raw** if you need details on using the utility. For example:

raw /dev/raw/raw1 /dev/sdb1

raw /dev/raw/raw3 /dev/sdb3

To verify the creation of the raw device, enter:

more /etc/sysconfig/rawdevices

Below is an example:

```
# more /etc/sysconfig/rawdevices
```

To list the raw device, enter:

ls -l /dev/raw/raw?

Below is an example:

```
# raw device bindings
# format:  <rawdev> <major> <minor>
#          <rawdev> <blockdev>
# example: /dev/raw/raw1 /dev/sda1
#          /dev/raw/raw2 8 5
/dev/raw/raw1 /dev/sdb1
/dev/raw/raw3 /dev/sdb3
```

Warning **In /etc/sysconfig/rawdevices, use <rawdev> <blockdev> format because it contains word raw. For example, /dev/raw/raw1 /dev/sdb1.**

Set ownership of block devices to **Informix** and chmod to **66**. See the following example:

```
# ls -l /dev/raw/raw?
crw-rw----    1 informix informix 162,   1 Jun 24  2004 /dev/raw/raw1
crw-rw----    1 root     disk     162,   2 Jun 24  2004 /dev/raw/raw2
crw-rw----    1 informix informix 162,   3 Jun 24  2004 /dev/raw/raw3
crw-rw----    1 root     disk     162,   4 Jun 24  2004 /dev/raw/raw4
crw-rw----    1 root     disk     162,   5 Jun 24  2004 /dev/raw/raw5
crw-rw----    1 root     disk     162,   6 Jun 24  2004 /dev/raw/raw6
crw-rw----    1 root     disk     162,   7 Jun 24  2004 /dev/raw/raw7
crw-rw----    1 root     disk     162,   8 Jun 24  2004 /dev/raw/raw8
```

Installing Multi NetFlow Collector

The Multi NetFlow Collector, Release 6.0 is distributed on DVD. When installing from the DVD do the following:

Step 1 Log into the host as root.

Step 2 Insert the DVD in to the drive.

Step 3 Navigate to the **CSCOmnfc-6.0.0-31-standard** directory using **cd**.

Step 4 Run **MNFC_setup.sh**. The image is installed to **/opt/CSCOmnfc**.

Step 5 When prompted, specify the following:

- The OS user id to own the $MNFC_DIR subdirectory and to run the MNFC server.
- The version of JRE to be used for MNFC server if more than one version is available.
- Whether you want the MNFC application to start up when the server host is rebooted.

The following is a sample transcript:

```
# ./MNFC_setup.sh

Using software package /opt/packages/CSCOmnfc-6.0.0-31-standard/./CSCOmnfc.zip.

*********************************************************************

Cisco Multi NetFlow Collector 6.0.0 [standard image, build 31]
Copyright (c) 2006-2007 by Cisco Systems, Inc.
All rights reserved.

This product contains cryptographic features and is subject to
United States and local country laws governing import, export,
transfer and use. Delivery of Cisco cryptographic products does
not imply third-party authority to import, export, distribute
or use encryption. Importers, exporters, distributors and users
are responsible for compliance with U.S. and local country laws.

By using this product you agree to comply with applicable laws
and regulations. If you are unable to comply with U.S. and local
laws, return this product immediately.

A summary of U.S. laws governing Cisco cryptographic products may be
found at: http://www.cisco.com/wwl/export/crypto/tool/stqrg.html

If you require further assistance please contact us by sending email
to export@cisco.com.

*********************************************************************

Press Return to continue...

An existing userid must be selected as the owner of files and processes.

Enter userid: nfcuser

This will be installed and run with user id nfcuser; group id is nfcuser.

Installing new package...

Installing Java from jdk-1_5_0_10-linux-i586.bin...
```

```
Web UI login name: [nfcuser]
Web UI password for nfcuser:
Enter value again to confirm:
Setting file ownership...

Updating platform config...

Would you like to start this when the system initializes? (y/n) y

CSCOmnfc will be started automatically when the system initializes.

Installation of Cisco Multi NetFlow Collector 6.0.0 [standard image, build 31]
is complete.  Informix should now be installed and configured
as described in the product User's Guide.

Note: In order to use this product, you must first obtain a license
from Cisco at http://www.cisco.com/go/license and copy
it to /opt/CSCOmnfc/config/mnfc.lic

#
```

Step 6 Update time zone data in the Java Runtime Environment (JRE) packaged with NFC with the latest time zone data available. See Appendix B, "Updating JRE Time Zone Data" for details.

Installing Informix

You must install the Informix Dynamic Serve if the MNFC server does not have IDS 9.40. Do the following:

Step 1 Log into the host as root.

Step 2 Obtain Kernel parameters and operating system patches.

Obtain the MachineNote from IBM for the IDS 9.40 pack that will be used. Verify and update if needed the patches and kernel parameters as recommended by the MachineNote.

Step 3 Insert Disk 2 in to the CDROM drive.

Step 4 Navigate to the Linux or SunOS directory using **cd**.

Step 5 From the current directory, run the script **ids_install**.

The script **ids_install** can be executed interactively; for MNFC the only needed component is IDS itself, so the script can be executed in straight silent mode. For example:

```
INFORMIXDIR=/opt/informix
export INFORMIXDIR
./ids_install IDS
```

Configuring the Database

To configure the database, do the following.

Note The script **bin/configuredb.sh** is designed to be executed at install time. If you execute the script after MNFC is in use, then configuration entities such as Aggregators, Summarizations, and Correlators must be discarded and recreated and the MNFC server restarted.

Step 1 As root, run **/opt/CSCOmnfc/bin/configuredb.sh**.

For efficiency and performance reasons, MNFC works with application data (NetFlow records and derived or correlated tables) stored in Informix tables fragmented into 24, or 48, or 96, etc. extents depending on available space. See the supported disk fragmentation profiles in Appendix B and the Records Retention and Data Latency in Primary Table section of the *Cisco Multi NetFlow Collector User Guide*.

The availability of disk fragmentation options for a particular server configuration is determined by the allocated disk size, that is the total space estimated as available on Informix raw disk partitions. If multiple choices exist you can pick the fragmentation option manually.

In addition, you must:

- Specify which valid user account MNFC can use internally in its built-in scheduler
- Confirm the choices

The following is a sample transcript:

```
[root@IBM346 bin]$ ./configuredb.sh
Acquiring info on disk partitions...
Planning allocation of dbspaces...

The following disk fragmentation profiles are available
 based on available raw partitions space total of 214957669
 (less slices allow for greater capacity, more slices allow for longer retention.
  Please refer to MNFC Performance and Sizing Guide for details).

     SMALL profile: number of slices =  24
    MEDIUM profile: number of slices =  48
     LARGE profile: number of slices =  96

Number of slices: [96] 48
Planning allocation of dbspaces...
System account for quartz DB access: nfcuser
Account password for nfcuser:
Enter value again to confirm:

Configuring the following values:

Storage for Rootdbs device/path:     /dev/raw/raw2
Rootdbs size (kb):                   23884000
System account for quartz DB access: nfcuser
Password for quartz DB access:       *******
Number of DB slices:                 48

Proceed? [yes]

Configuring database using MEDIUM profile...

Updating /opt/informix/etc/onconfig.nfc...
Updating /opt/informix/etc/sqlhosts...
Updating /opt/CSCOmnfc/config/quartz.properties...
Updating /opt/CSCOmnfc/config/nfcdb.xml...
```

```
Finished updating files, starting Informix ...
Checking status ...
Executing DBSpace creation...
Creating dbspace 1 on /dev/raw/raw1
Verifying physical disk space, please wait ...
Space successfully added.

** WARNING **  A level 0 archive of Root DBSpace will need to be done.
Creating dbspace 2 on /dev/raw/raw1
Verifying physical disk space, please wait ...
Space successfully added.

.....................................................................................

Creating dbspace 48 on /dev/raw/raw1
Verifying physical disk space, please wait ...
Space successfully added.

** WARNING **  A level 0 archive of Root DBSpace will need to be done.
Checking DBSpace count...
Creating NFC_DB instance...
Creating NFC DB and tables.

Database created.

.....................................................................................

Database closed.

Creating metadata.

Shutting down Informix...

DB Configuration is complete.


[root@IBM346 bin]$
```

For more information on database configuration profiles and the number of allocated dbspaces, refer to Appendix B of the *Cisco Multi NetFlow Collector User Guide.*

Uninstalling Multi NetFlow Collector

To uninstall the Multi NetFlow Collector, do the following:

Step 1 Shut down the MNFC system by entering **cscomnfc stop**.

Step 2 Execute the command **uninstall.sh** outside of the **$MNFC_DIR** subtree. The subtree **$MNFC_DIR** will be removed.

Licensing

A license file is required for each host running Multi NetFlow Collector Release 6. The license is specific to the IP address of the host. You can obtaina permanent license at http://www.cisco.com/go/license.

You must have the IP address of the host on which Multi NetFlow Collector will run. To obtain a permanent license, you must also have the PAK you received after purchasing NetFlow Collector. After you enter the information, a license file is emailed to you. Copy the license file or its contents with no alterations to **/opt/CSCOmnfc/config/mnfc.lic**.

The first line of the license file contains either the demo expiration date or the word **permanent** for a permanent license, and the IP address of the host to which Multi NetFlow Collector is licensed.

If the host running Multi NetFlow Collector has more than one network card and IP address, specify the IP address associated with the hostname when licensing the product.

Note By default, Red Hat Enterprise Linux associates the system hostname with the loopback address **127.0.0.1** in **/etc/hosts**. However, for licensing to work, the hostname must be associated with the host's IP address. Edit **/etc/hosts**, remove the hostname from the loopback address entry, and add an entry for the licensed IP address.

The file **/etc/nsswitch.conf** is normally configured so that hostname lookups are first obtained from files (**/etc/hosts**).

CHAPTER 3

Starting and Stopping Cisco Multi NetFlow Collector

This chapter describes how to start and stop the Cisco Multi NetFlow Collector application and then validate that it is operating properly. It includes the following sections:

The Multi NetFlow Collector server contains number of processes which are expected to remain up and running without interruption 24/7 unless there is a need to perform a system maintenance. It also needs its database server (Informix IDS 9.40 relational database) to remain operational.

The Main MNFC Components

The Multi NetFlow Collector is implemented in Java and includes the following processes executed in separate instances of JVM:

- Storage Manager (or **nfcdb**) executed with userid **informix**
- Process Watcher
- Concentrator
- Web GUI process (as Tomcat by Apache with web application **mnfc**)
- Report Daemon
- VPN MIB Collector
- CLI Collector

Storage Manager is the only process executed with user id **informix**.All other processes are executed with user id **nfcuser**. See Table 1 for information about MNFC component processes.

Table 1 MNFC Component Processes

Short Name	Full Name	Purpose	UNIX User ID*	Default Presence**	Log Properties file in config directory
nfcdb	Storage Manager	Run-time maintenance of **nfc_db** database in IDS	informix	Mandatory	mnfcdb-log4j
pw	Process Watcher	Running of all processes except Storage Manager	nfcuser	Mandatory	mnfcpw-log4j
concentrator	Concentrator	MNFC back-end server	nfcuser	Mandatory	mnfc-log4j
web	Web GUI	Tomcat with **mnfc** webapp, MNFC front-end server	mfcuser	Mandatory	mnfcweb-log4j
rd	Report Daemon	Execution of reports	nfcuser	Mandatory	mnfcrd-log4j
vpnmibcltr	VPN MIB Collector	SNMP-based collection of CE-CE data	nfcuser	Optional	vpnmibcltr-log4j
clicollector	CLI Collector	CLI Collector	nfcuser	Optional	clic-log4j

* Instead of nfcuser, another id can be designated by the operator executing MNFC installation.

** With exception of Storage Manager and Process Watcher the presence flag is configurable via the Process Watcher's controls.

The cscomnfc Script

The script **cscomnfc** is intended as a single invocation point for starting up or shutting down all MNFC processes, including those executed using user ids **informix** and **nfcuser**. Run this script as the user root. When running this script, you must use either the **start** or **stop** parameter.

The location of the **cscomnfc** script is under **/opt/CSCOmnfc/bin** or **$MNFC_DIR/bin**.

The mnfc Script

The script **mnfc** is intended for starting and stopping of MNFC processes executed by Process Watcher using user if **nfcuser**. The script accepts mandatory verb parameter **start**, **stop**, **shutdown**, and **status**. If you use **start** or **stop**, then a second parameter is needed. Use the short name for the MNFC process. For a listing of MNFC processes, see Table 1 on page 3-2.

When starting and stopping all MNFC processes for the id **nfcuser**, you must use the following:

- For starting: **mnfc start all**
- For stopping: **mnfc shutdown**

The location of the **mnfc** script is under **/opt/CSCOmnfc/bin** or **$MNFC_DIR/bin**.

Starting and Stopping the MNFC Storage Manager

To start the MNFC storage manager, as user informix enter the shell command **startnfcdb.sh**.

To stop the MNFC storage manager, as user informix enter the shell command **stopProcess.sh nfcdb**.

The location of the **startnfcdb.sh** and **stopProcess.sh** script is under **/opt/CSCOmnfc/bin** or **$MNFC_DIR/bin**.

Starting and Stopping the MNFC Database

Before the Informix database is shutdown or restarted, the Multi NetFlow Collector application must be shutdown.

For information on operating the Informix server, refer to IDS 9.40 documentation by IBM.

APPENDIX A

Troubleshooting the Multi NetFlow Collector

This appendix provides helpful information and procedures in case you encounter problems while using Cisco Multi NetFlow Collector (MNFC).

Solving Multi NetFlow Collector Problems

This section discusses some basic problems that you might encounter while attempting to run Multi NetFlow Collector.

Symptom Licensing errors in **/opt/CSCOmnfc/logs/mnfc.log**.

Possible Cause Either the license file is missing or invalid; or system hostname and address information is not configured properly on the host.

Recommended Action

1. Determine the following:
   - The host IP address
   - The license file contents in the **/opt/CSCOmnfc/config/mnfc.lic** file
   - The **/etc/hosts** file content
   - The result of running the **hostname** command.
2. Verify that the IP address matches the **mnfc.lic** line **HOSTID=INTERNET=ipaddress**. If not, the license file is not valid for this host.
3. Verify that the hostname is set, determined from the results from the **hostname** command, and is some value other than **localhost**, localhost.**localdomain**, etc.
4. Verify that the **/etc/hosts** file contains an entry for this hostname, and that the IP address in the entry is the same IP address in the **mnfc.lic** file. Also verify that the loopback address **127.0.0.1** in the **/etc/hosts** file does not contain this hostname, which is the default for Red Hat Enterprise Linux.

Symptom Unable to login to the web UI with a general error message on the login screen about checking for errors in collector log files.

Possible Cause A licensing or configuration error has prevented the collection process from starting, and the concentrator process must be running for the user to log in.

Recommended Action

1. Check the log files for errors
   - On the Linux platform: **grep -e ERR -e WARN -e FATAL /opt/CSCOmnfc/logs/mnfc.log**
   - On the Solaris platform: enter **/usr/xpg4/bin/grep** to specify multiple -e options
2. If you find a licensing error, see the licensing symptom and Recommended Action above.
3. If you find a configuration error, you need to fix the error by editing the XML file, and then restarting the concentrator process.

Symptom The script **configuredb.sh** fails to start Informix; the file **/opt/informix/online.log** has information on shared memory failure similar to the following:

11:38:42 IBM Informix Dynamic Server Started.

11:38:42 Contiguous shared memory segment allocation failed at 0x44000000.

Allocation successful at 0x19658000.

Check SHMBASE is consistent with the value in $INFORMIXDIR/etc/onconfig.std.

If you are using the correct SHMBASE value in your ONCONFIG file, then consider this message informational only.

11:38:46 OS couldn't allocate first resident segment at SHMBASE (0x8539ed0).

Possible Cause On certain Linux platforms the default value of **SHMBASE** set to **0x44000000** is taken by other libraries or utilities.

Recommended Action Set the **SHMBASE** to a value of **0xB000000L** in the file **$MNFC/db/config/onconfig-Linux.nfc** and then re-run the **configuredb.sh** script.

APPENDIX B

Updating JRE Time Zone Data

Cisco Multi NetFlow Collector requires up-to-date time zone data in the Java Runtime Environment (JRE) bundled with MNFC so that files and reports that contain the time-of-day accurately reflect current Daylight Saving Time rules that are in effect. Time zone data in MNFC's JRE should be updated using the TZupdater tool provided by Sun Microsystems when MNFC is installed and whenever local Daylight Saving Time rules change.

Note JRE time zone data that MNFC depends on is separate from system time zone data. Although it is recommended that you should also update your operating system time zone data when local Daylight Saving Time rules change, MNFC will report times correctly if only the JRE is updated.

Sun Microsystems has provided general information about this topic in the article *Timezone Data and the Java Runtime Environment* at http://java.sun.com/javase/timezones/. To see if a change announced recently for your location is reflected in the latest data provided by Sun, see *Timezone Data Versions in the JRE Software* at http://java.sun.com/javase/timezones/tzdata_versions.html.

To update the JRE timezone information for MNFC, do the following:

Step 1 Login as the system user who owns MNFC files and processes; this account was specified when MNFC was installed.

Step 2 Run **umask** to verify that the current umask is not set to a restrictive value higher than 022. If set to a higher value, set to 022 by running **umask 022**.

Step 3 Download the latest TZupdater tool from Sun Microsystems. If you don't have a Sun Online Account, you will be prompted to create one. Information about the tool including a link to the download location of the most recent version is available at http://java.sun.com/javase/tzupdater_README.html.

Step 4 In a temporary directory, extract the zipped contents of the download file. This creates the subdirectory **tzupdater-*version*** containing the jar file **tzupdater.jar**.

Step 5 Shut down MNFC by running **/opt/CSCOnfc/bin/nfcollector shutdown**. Verify that no java processes by running **ps -ef | grep java**.

Step 6 Run the TZupdater tool to update JRE timezone data as follows:

/opt/CSCOmnfc/java/bin/java -jar *path-to-tzupdater-files*/tzupdater.jar -u

Note: On the Solaris platform, ignore package-related warning messages such as:

> ***path-to-java-bin*/java not directly found in contents file, no package resolution performed. (May not be in PKG form, not an absolute path, or is a symlink.)**

Step 7 Restart Cisco Multi NetFlow Collector. See Chapter 3, “Starting and Stopping Cisco Multi NetFlow Collector” for details.

INDEX

N

O

P

R

S

T

U

V