

IPTraf-ng User's Manual

IPTraf-ng User's Manual

Copyright © 1997, 2003 Gerard Paul Java

Copyright © 2016 Phil Cameron

This manual is released under the terms of the GNU Free Documentation License of March, 2000 as published by the Free Software Foundation, reproduced in this manual as Appendix B.

IPTraf-ng is open-source software released under the terms of the GNU General Public License version 2 or any later version as published by the Free Software Foundation, reproduced in the LICENSE file in the distribution's top-level directory.

The accompanying software and the information contained in this document are provided "AS IS" without warranty of any kind, express or implied, including, without limitation, the implied warranties of mercantability or fitness for any particular purpose.

In no event shall the author be liable for any indirect, special, consequential, or incidental damages arising from the use of this manual or the accompanying software even if the author has been advised of the possibility of such damages.

Linux is a registered trademark of Linus Torvalds. Pentium is a registered trademark of Intel Corporation. All other trademarks are property of their respective owners.

Some structure declarations were based on code copyrighted by the Regents of the University of California.

Table of Contents

| | |
|--|-----------|
| About This Document | v |
| For Additional Information | v |
| Document Conventions..... | v |
| 1. Getting Started | 1 |
| About IPTraf-ng..... | 1 |
| Installation..... | 1 |
| System Requirements..... | 1 |
| Availability..... | 2 |
| Starting and Stopping IPTraf-ng | 2 |
| Command-line Options..... | 2 |
| Using the Menus | 4 |
| Exiting IPTraf-ng | 5 |
| 2. Preparing to Use IPTraf-ng..... | 7 |
| Number Display Notations | 7 |
| Instances and Logging..... | 7 |
| Screen Update Delays | 8 |
| Supported Network Interfaces..... | 8 |
| 3. The IP Traffic Monitor..... | 11 |
| The Upper Window | 11 |
| Closed/Idle/Timed Out Connections..... | 15 |
| Sorting TCP Entries | 15 |
| Lower Window | 16 |
| Entry Details..... | 18 |
| Additional Information | 21 |
| 4. Network Interface Statistics | 23 |
| General Interface Statistics..... | 23 |
| Detailed Interface Statistics..... | 24 |
| 5. Statistical Breakdowns | 27 |
| Packet Sizes | 27 |
| TCP and UDP Traffic Statistics..... | 28 |
| Sorting TCP/UDP Entries | 29 |
| Additional Information..... | 30 |
| 6. LAN Station Statistics..... | 31 |
| Sorting the LAN Station Monitor Entries | 32 |
| Additional Information | 33 |
| 7. Filters..... | 35 |
| IP Filters..... | 35 |
| Defining a New Filter..... | 36 |
| Applying a Filter..... | 44 |
| Editing a Defined Filter..... | 44 |
| Deleting a Defined Filter..... | 45 |
| Detaching a Filter..... | 45 |
| ARP, RARP, and other Non-IP Packet Filters | 45 |
| 8. Configuring IPTraf-ng..... | 47 |
| Toggles | 47 |
| Reverse DNS Lookups | 47 |
| TCP/UDP Service Names | 47 |
| Force promiscuous..... | 48 |
| Color | 48 |
| Logging | 48 |
| Activity mode..... | 49 |
| Source MAC addrs in traffic monitor | 49 |
| Timers..... | 50 |
| TCP Timeout..... | 51 |

| | |
|---|-----------|
| Log Interval | 51 |
| Screen Update Interval | 52 |
| TCP closed/idle persistence | 52 |
| Custom Information | 52 |
| Additional ports..... | 52 |
| Delete port/range..... | 53 |
| LAN Station Identifiers..... | 53 |
| 9. Background Operation | 55 |
| A. Messages..... | 57 |
| IPTraf-ng Messages | 57 |
| Resolving Process Messages | 63 |
| B. GNU Free Documentation License | 65 |
| PREAMBLE | 65 |
| APPLICABILITY AND DEFINITIONS..... | 65 |
| VERBATIM COPYING | 66 |
| COPYING IN QUANTITY..... | 66 |
| MODIFICATIONS..... | 67 |
| COMBINING DOCUMENTS..... | 68 |
| COLLECTIONS OF DOCUMENTS..... | 68 |
| AGGREGATION WITH INDEPENDENT WORKS..... | 68 |
| TRANSLATION..... | 69 |
| TERMINATION..... | 69 |
| FUTURE REVISIONS OF THIS LICENSE..... | 69 |
| How to use this License for your documents..... | 69 |

About This Document

This document contains the instructions on how to use the IPTraf-ng network monitoring software version 1.2. This manual details the different statistical facilities, the user interface, and the important features of the software.

For Additional Information

See the included README file for summarized and late-breaking information. The CHANGES file contains a record of the changes made to the software since 1.0.0. README.resolving contains information on the reverse resolution function. See the other README files for support and development information.

Document Conventions

The following symbols and typefaces are used throughout this manual:

[]

items in brackets are optional. Brackets also denote items that may or may not be displayed onscreen depending on settings or conditions.

{ }

curly braces enclose items you choose from

|

the vertical bar separates choices in curly braces

`normal monospace`

normal monospace text in syntax specifications should be typed in exactly as presented. Because UNIX and variants are case-sensitive, case must be preserved. Monospace is also used in presenting items that appear on the screen.

monospace italics

italics in syntax specifications indicate items that are to be replaced with an actual item (e.g. *interface* should be replaced with an actual interface name, like `eth0`).

Additional information appears distinctively set apart from the main text. This information includes Notes, Tips, or Technical Notes.

Notes are additional pieces of information that may be useful or may clarify the preceding paragraphs of the manual.

Tips provide shortcuts, clarify tasks that may not be immediately obvious, or provide references to additional sources of information.

Technical notes are explanations of a more technical nature and may be of more use to programmers and advanced users.

About This Document

Chapter 1. Getting Started

About IPTraf-ng

IPTraf-ng is a network monitoring utility and traffic analyzer for IP networks. It intercepts packets and returns data about captured the network traffic in various statistical facilities.

IPTraf-ng includes these major features:

- An IP traffic monitor that shows TCP connection information (hosts, packet/byte counts, flags, window sizes), and color-coded information about other IP packets
- Statistics (counts and load rates) for network interfaces in general and detailed views
- Statistics per TCP/UDP port
- Statistical breakdown according to packet sizes
- A LAN host monitor that returns counts and loads per detected MAC address
- A powerful filtering system for users to view only interesting traffic
- Logging
- An asynchronous DNS resolver for the IP traffic monitor
- A text-based, full-color, menu-driven user interface suitable for use on all Linux systems with terminals, especially Linux consoles and color xterms
- Easy configuration
- Fully software-based. No additional hardware required

Basic knowledge of the important TCP/IP protocols (IP, TCP, UDP, ICMP, etc.) is necessary for you to best understand the information generated by the program.

Installation

System Requirements

IPTraf-ng requires:

Hardware Requirements

- 16 megabytes of physical RAM (more recommended, at least 64 MB for very busy networks)
- 2 megabytes of free disk space for installation (more will be needed if you log high amounts of traffic over time)
- Pentium-class processor or higher (Pentium-II 200 MHz or higher recommended) or equivalent.
- One or more of the supported network interfaces.

Operating System Requirements

- Linux kernel 2.2.0 or higher
- GNU C Library 2.1 or later
- ncurses 4.2 or later with the complete terminfo database in `/usr/share/terminfo`. Support for `linux`, `vt100`, `xterm`, `xterm-color` recommended.

Compilation Requirements

The following components are required when compiling IPTraf-ng from the source code.

- gcc 2.7.2.3 or later
- GNU C (glibc) development library 2.1 or later
- ncurses development libraries 4.2 or later
- git

Availability

IPTraf-ng is available in binary form from many Linux Distributions including Red Hat RHEL, Centos, Fedora, and Ubuntu.

IPTraf-ng git source repository can be cloned from from:

<https://git.fedorahosted.org/git/iptraf-ng.git>

Starting and Stopping IPTraf-ng

After installation, you can start the program, as root, by simply entering

```
iptraf-ng
```

Entering the IPTraf-ng command without any command-line parameters brings up the program's main menu. From there, you can select the facilities you want.

IPTraf-ng determines and makes use of the maximum number of lines and columns on the terminal.

Note: IPTraf-ng does not have a SIGWINCH handler; it does not adjust itself when an xterm or some other X terminal is resized.

Technical note: IPTraf-ng needs to refer to the terminfo database in `/usr/share/terminfo`. If the supplied executable program fails with `Error opening terminal`, your terminfo database may be located somewhere else. You can control the terminfo search path by using the `TERMINFO` environment variable. For example, if you're using the **sh** or **bash** shell, and your terminfo database is in `/usr/lib/terminfo` (typical for Slackware distributions), you can use the commands:

```
TERMINFO=/usr/lib/terminfo
export TERMINFO
```

You can place these commands in your `~/.profile` or the systemwide `/etc/profile` startup files.

You can also create a symbolic link named `/usr/share/terminfo` to let it point to your existing terminfo (assuming again your terminfo is in `/usr/lib/terminfo`):

```
ln -s /usr/lib/terminfo /usr/share/terminfo
```

Or you can recompile your program to use your existing ncurses library installation. If you do this, make sure you have ncurses 4.2 or later.

Command-line Options

IPTraf-ng has a few optional command-line parameters. As with most UNIX commands, IPTraf-ng command-line parameters are case-sensitive (`-l` is NOT the same as `-L`).

The following command-line parameters can be supplied to the **iptraf-ng** command:

`-i iface`

causes the IP traffic monitor to start immediately on the specified interface. If `-i all` is specified, all interfaces are monitored.

`-g`

starts the general interface statistics

`-d iface`

shows detailed statistics for the specified interface

`-s iface`

starts the TCP/UDP traffic monitor for the specified interface

`-z iface`

starts the packet size breakdown for the specified interface

`-l iface`

starts the LAN station monitor on the specified interface. If `-l all` is specified, all LAN interfaces are monitored.

`-t timeout`

The `-t` parameter, when used with one of the other parameters that specify a facility to start, tells IPTraf-ng to run the indicated facility for only timeout minutes, after which the facility exits. The `-t` parameter is ignored in menu mode.

If this parameter is not specified, the facility runs until the exit keystroke is pressed.

`-B`

Redirects all terminal output to the "bit bucket" `/dev/null`, closes standard input, and places the program in the background. This parameter can be used only with one of the `-i`, `-g`, `-d`, `-s`, `-z`, or `-l` parameters. See Background Operation in Chapter 9. `-B` is ignored in menu mode.

`-L filename`

Allows you to specify an alternate log file name when the any facility is directly started from the command line, whether in foreground or background mode. If specified in foreground mode, the log filename prompt is bypassed, even when logging is turned on in the *Configure...* menu. If this parameter is omitted in background mode, the default log filename is used.

This parameter always turns on logging.

If an absolute path is not specified, the log file will be created in the default log file directory

`-f`

Forces IPTraf-ng to clear all lock files and reset all instance counters to zero before running any facilities. IPTraf-ng will then think it's the first instance of itself.

The `-f` parameter overrides the existing locks and counters imposed by the IPTraf-ng process and by the various facilities, causing this instance to think it is the first and that there are no other facilities running. Use this parameter with great caution. A common use for this parameter is to recover from abrupt or abnormal terminations which may leave stale locks and counters still lying around.

The `-f` parameter may be used together with the others.

`-h`

displays a short help screen

While the command-line options are case-sensitive, interactive keystroke at the IPTraf-ng full-screen interface are not.

Using the Menus

Menu items with a trailing ellipsis (. . .) either pop up a submenu with further items, or require additional information before it can complete the task and return to the menu. Menu items without an ellipsis execute immediately.

Use the Up and Down arrow keys on your keyboard to move the selection bar. Press Enter to execute the selected item. Alternatively, you can also directly press the highlighted letter of the item you want. This will immediately execute the option.




Figure 1-1. The IPTraf-ng Main Menu

Exiting IPTraf-ng

You can exit IPTraf-ng with the Exit command in the main menu.

When started with one of the command-line options to directly start a statistical facility, pressing X or Q will exit the facility directly, without any confirmation. The `-t` command-line parameter will automatically exit the facility after the specified length of time without any confirmation as well. Daemon facilities started with the `-B` parameter will immediately terminate after being sent a USR2 signal. See background operation in chapter 9 for more information.

Chapter 2. Preparing to Use IPTraf-ng

This chapter provides information applicable to all of IPTraf-ng's statistical monitors.

Number Display Notations

IPTraf-ng initially returns exact counts of bytes and packets. However, as they grow larger, IPTraf-ng begins displaying them in increasingly higher denominations.

A number standing alone with no suffix represents an exact count. A number with a K following is a kilo (thousand) figure. An M, G, and T suffix represents mega (million), giga (billion), and tera (trillion) respectively. The following table shows examples.

Table 2-1. Numeric Display Notations

| | |
|---------|--------------------------------|
| 1024067 | exactly 1024067 |
| 1024K | approximately 1024000 |
| 1024M | approximately 1024000000 |
| 1024G | approximately 1024000000000 |
| 1024T | approximately 1024000000000000 |

These notations apply to both packet and byte counts.

Instances and Logging

IPTraf-ng allows multiple instances of the facilities at the same time in different processes (for example, you can run two or more IP Traffic Monitors at the same time). However only one can listen on a specific interface or all interfaces at once. The only exception is the general interface statistics, which is still restricted to only one instance at a time.

Because of this, each instance now generates log files with unique names for instances, depending on either their instance or the interface they're listening on. If the *Logging* option is turned on (see the Configuration chapter), IPTraf-ng will prompt you for a log file name while presenting a default. You may accept this default or change it. Press Enter to accept, or Ctrl+X to cancel. Canceling will turn logging off for that particular session.

If you don't specify an absolute path, the log file will be placed in:

`/var/log/iptraf-ng.`




Figure 2-1. The logfile prompt dialog

See the Logging section in the Configuration chapter for detailed information on logging. See also the documentation on each statistical facility for the default log file names.

The default log file names will also be used if the `-B` parameter is used to run IPTraf-ng in the background. You can override the defaults with the `-L` parameter. See Background Operation in Chapter 9.

Screen Update Delays

A configuration option is available to control screen update speed.

See the *Screen update interval...* configuration option under the Configuration chapter of this manual.

Supported Network Interfaces

IPTraf-ng currently supports the following network interface types and names.

`lo`

The loopback interface. Every machine has one, and has an IP address of 127.0.0.1. `lo` is also indicated if data is detected on the `dummyn` interface(s).

`ethn`

An Ethernet interface. *n* starts from 0. Therefore, `eth0` refers to the first Ethernet interface, `eth1` to the second, and so on. Most machines only have one.

`fddin`

An FDDI interface. *n* starts from 0.

`pppn`

A PPP interface. *n* starts from 0.

`sln`

A SLIP interface. *n* starts from 0.

`plipn`

PLIP interfaces. These are point-to-point IP connections using the PC parallel port.

`ipsecn`

This refers to Free s/WAN (and possibly other) logical VPN interfaces.

`sbnn`

SBNI long-range modem interfaces

`dvbn`, `sm200`, `sm300`

DVB satellite-receive interfaces

`wlann`, `wvlann`

Wireless LAN interfaces

`tunn`

general logical tunnel interfaces

`brgn`

general logical bridge interfaces

`hdlcn`

Frame Relay base (FRAD) interfaces (non-PVC)

`pvcn`

Frame Relay Permanent Virtual Circuit interfaces

Your system's network interfaces must be named according to the schemes specified above.

Chapter 3. The IP Traffic Monitor

Executing the first menu item or specifying `-i` to the `iptraf-ng` command takes you to the IP traffic monitor. The traffic monitor is a real-time monitoring system that intercepts all packets on all detected network interfaces, decodes the IP information on all IP packets and displays the appropriate information, most notably the source and destination addresses. It also determines the encapsulated protocol within the IP packet, and displays some important information about that as well.

There are two windows in the traffic monitor, both of which can be scrolled with the Up and Down cursor keys. Just press `W` to move the `Active` indicator to the window you want to control.

```
IPTraf
TCP Connections (Source Host:Port) ----- Packets -----
[ 61.9.80.40:3812 > 757
[ 61.9.4.185:http > 1001
[ 61.9.80.38:1624 > 528
[ CPE3439373939323531.cpe.net.cabl:1214 > 832
[ boh141zoy4111.bc.hsia.telus.net:1214 > 1139
[ 61.9.80.38:1334 > 652
[ 64.94.89.245:http > 533
[ 61.9.82.125:62620 > 346
[ 61.9.82.125:63612 > 277
[ 128.167.58.181:http > 467
[ 61.9.82.122:64399 > 231
[ h24-80-94-122.vn.shawcable.net:1214 > 332
TCP: 6276 entries -----

ARP request for 61.9.108.253 (46 bytes) from 00d0b7b7ea8
ARP reply from 61.9.108.253 (46 bytes) from 00000c4340a0
ICMP echo req (84 bytes) from riker.mozcom.com to w4.dcx
ICMP echo rply (84 bytes) from w4.dcx.yahoo.com to riker
Non-IP (0x4) (46 bytes) from 00d0bacceb43 to 0180c200000
Non-IP (0x4) (46 bytes) from 00d0bacceb44 to 0180c200000
Bottom ----- Elapsed time: 0:03 -----
Pkts captured (all interfaces): 208029 | TCP flow ra
Up/Dn/PgUp/PgDn-scroll M-more TCP info W-chg actu win
```

Figure 3-1. The IP traffic monitor

The Upper Window

The upper window of the traffic monitor displays the currently detected TCP connections. Information about TCP packets are displayed here. The window contains these pieces of information:

- Source address and port
- Packet count
- Byte count
- Source MAC address
- Packet Size
- Window Size
- TCP flag statuses
- Interface

The Up and Down cursor keys move an indicator bar between entries in the TCP monitor, scrolling the window if necessary. The PgUp and PgDn keys display the previous and next screenfuls of entries respectively.

The IP traffic monitor computes the data flow rate of the currently highlighted TCP flow and displays it on the lower-right corner of the screen. The flow rate is in kilobits or kilobytes per second depending on the *Activity mode* switch in the *Configure...* menu.

Because this monitoring system relies solely on packet information, it does not determine which endpoint initiated the connection. In other words, it does not know which endpoints are the client and server. This is necessary because it can operate in promiscuous mode, and as such cannot determine the socket statuses for other machines on the LAN. However, a little knowledge of the well-known TCP port numbers can give a good idea about which address is that of the server.

The system therefore displays two entries for each connection, one for each direction of the TCP connection. To make it easier to determine the direction pairs of each connection, a bracket is used to "join" both together. This bracket appears at the leftmost part of each entry.

Just because a host entry appears at the upper end of a connection bracket doesn't mean it was the initiator of the connection.

Each entry in the window contains these fields:

Source address and port

The source address and port indicator is in *address:port* format. This indicates the source machine and TCP port on that machine from which this data is coming.

The destination is the host:port at the other end of the bracket.

Packet count

The number of packets received for this direction of the TCP connection

Byte count

The number of bytes received for this direction of the TCP connection. These bytes include total IP and TCP header information, in addition to the actual data. Data link header (e.g. Ethernet and FDDI) data are not included.

Source MAC address

The address of the host on your local LAN that delivered this packet. This can be viewed by pressing M once if *Source MAC addrs* in traffic monitor is enabled in the *Configure...* menu.

Packet Size

The size of the most recently received packet. This item is visible if you press M for more TCP information. This is the size of the IP datagram only, not including the data link header.

Window Size

The advertised window size of the most recently received packet. This item is visible if you press M for more TCP information.

Flag statuses

The flags of the most recently received packet.

S

SYN. A synchronization is taking place in preparation for connection establishment. If only an S is present (S---) the source is trying to initiate a connection. If an A is also present (S-A-), this is an acknowledgment of a previous connection request, and is responding.

A

ACK. This is an acknowledgment of a previously received packet

P

PSH. A request to push all data to the top of the receiving queue

U

URG. This packet contains urgent data

RESET

RST. The source machine indicated in this direction reset the entire connection. The direction entries for reset connections become available for new connections.

DONE

The connection is done sending data in this direction, and has sent a FIN (finished) packet, but has not yet been acknowledged by the other host.

CLOSED

The FIN has been acknowledged by the other host. When both directions of a connection are marked CLOSED, the entries they occupy become available for new connection entries.

-

The flag is not set

Some other pieces of information can be viewed as well. The M key displays more TCP information. Pressing M once displays the MAC addresses of the LAN hosts that delivered the packets (if the *Source MAC addr in traffic monitor* option is enabled in the *Configure...* menu). N/A is displayed if no packets have been received from the source yet, or if the interface doesn't support MAC addresses (such as PPP interfaces).

If the *Source MAC addr in traffic monitor* option is not enabled, pressing M simply toggles between the counts and the packet and window sizes.

By default, only IP addresses are displayed, but if you have access to a name server or host table, you may enable reverse lookup for the IP addresses. Just enable reverse lookup in the *Configure...* menu.

The Asynchronous Resolving Process

The IP traffic monitor starts a process to help speed up reverse lookups without sacrificing too much keyboard control and accuracy of the counts. While reverse lookup is being conducted in the background, IP addresses will be used until the resolution is complete.

If for some reason the resolving process cannot start, and you are on the Internet, and you enable reverse lookup, your keyboard control can become very slow. This is because the standard lookup functions do not return until they have completed their tasks, and it can take several seconds for a name resolution in the foreground to complete.

The resolving process will spawn up to 200 children to process reverse DNS queries.

Tip: If you notice unusual SYN activity (too many initial (S---) but frozen SYN entries, or rapidly increasing initial SYN packets for a single connection), you may be under a SYN flooding attack or TCP port scan. Apply appropriate measures, or the targeted machines may begin denying network services.

Entries not updated within a user-configurable amount of time may get replaced with new connections. The default time is 15 minutes. This is regardless of whether the connection is closed or not. (Some unclosed connections may be due to extremely slow links or crashes at either end of the connection.) This figure can be changed at the *Configure...* menu.

Some early entries may have a > symbol in front of its packet count. This means the connection was already established when the monitor started. In other words, the figures indicated do not reflect the counts since the start of the TCP connection, but rather, since the start of the traffic monitor. Eventually, these > entries will close (or time out) and disappear. TCP entries without the > were initiated after the traffic monitor started, and the counts indicate the totals of the connection itself. Just consider entries with > partial.

Some > entries may go idle if the traffic monitor was started when these connections were already half-closed (FIN sent by one host, but data still being sent by the other). This is because the traffic monitor cannot determine if a connection was already half-closed when it started. These entries will eventually time out. (To minimize these entries, an entry is not added by the monitor until a packet with data or a SYN packet is received.)

Direction entries also become available for reuse if an ICMP Destination Unreachable message is received for the connection.

The lower part of the screen contains a summary line showing the IP, TCP, UDP, ICMP, and non-IP byte counts since the start of the monitor. The IP, TCP, UDP, and ICMP counts include only the IP datagram header and data, not the data-link headers. The non-IP count includes the data-link headers.

Technical note: IP Forwarding and Masquerading: Previous versions of IPTraf-ng issued a warning if the kernel had IP masquerading enabled due to the way the kernel masqueraded and translated the IP addresses. The new kernels no longer do it as before and IPTraf-ng now gives output properly on masquerading machines. The `-q` parameter is no longer required to suppress the warning screen.

On forwarding (non-masquerading) machines packets and TCP connections simply appear twice, one each for the incoming and outgoing interfaces if all interfaces are being monitored.

On masquerading machines, packets and connections from the internal network to the external network also appear twice, one for the internal and external interface. Packets coming from the internal network will be indicated as coming from the internal IP address that sourced them, and also as coming from the IP address of the external interface on your masquerading machine. In much the same way, packets coming in from the external network will look like they're destined for the external interface's IP address, and again as destined for the final host on the internal network.

Closed/Idle/Timed Out Connections

A TCP connection entry that closes, gets reset, or stays idle too long normally gets replaced with new connections. However, if there are too many of these, active connections may become interspersed among closed, reset, or idle entries.

IPTraf-ng can be set to automatically remove all closed, reset, and idle entries with the *TCP closed/idle persistence...* configuration option. You can also press the F key to immediately clear them at any time.

Note: The *TCP timeout...* option only tells IPTraf-ng how long it should take before a connection should be considered idle and open to replacement by new connections. This does not determine how long it remains on-screen. The *TCP closed/idle persistence...* parameter flushes entries that have been idle for the number of minutes defined by the *TCP timeout...* option.

Sorting TCP Entries

The TCP connection entries can be sorted by pressing the S key, then by selecting a sort criterion. Pressing S will display a box showing the available sort criteria. Press P to sort by packet count, B to sort by byte count. Pressing any other key cancels the sort.

The sort operation compares the larger values in each connection entry pair and sorts the counts in descending order.

Over time, the entries will go out of order as counts proceed at varying rates. Sorting is not done automatically so as not to degrade performance and accuracy.




Figure 3-2. The IP traffic monitor sort criteria

Lower Window

The lower window displays information about the other types of traffic on your network. The following protocols are detected internally:

- User Datagram Protocol (UDP)
- Internet Control Message Protocol (ICMP)
- Open Shortest-Path First (OSPF)
- Interior Gateway Routing Protocol (IGRP)
- Interior Gateway Protocol (IGP)
- Internet Group Management Protocol (IGMP)
- General Routing Encapsulation (GRE)
- Layer 2 Tunneling Protocol (L2TP)
- IPSec AH and ESP protocols (IPSec AH and IPSec ESP)
- Address Resolution Protocol (ARP)

- Reverse Address Resolution Protocol (RARP)

Other IP protocols are looked up from the `/etc/services` file. If `/etc/services` doesn't contain information about that protocol, the protocol number is indicated.

Non-IP packets are indicated as `Non-IP` in the lower window.

Note: The source and destination addresses for ARP and RARP entries are MAC addresses.

Strictly speaking, ARP and RARP packets aren't IP packets, since they are not encapsulated in an IP datagram. They're just indicated because they are integral to proper IP operation on LANs.

For all packets in the lower window, only the first IP fragment is indicated (since that contains the header of the IP-encapsulated protocol) but with no further information from the encapsulated protocol.

UDP packets are also displayed in `address:port` format while ICMP entries also contain the ICMP message type. For easier location, each type of protocol is color-coded (only on color terminals such as the Linux console).

UDP

Red on White

ICMP

Yellow on Blue

OSPF

Black on Cyan

IGRP

Bright white on Cyan

IGP

Red on Cyan

IGMP

Bright green on Blue

GRE

Blue on white

ARP

Bright white on Red

RARP

Bright white on Red

Other IP

Yellow on red

Non-IP

Yellow on Red

The lower window can hold up to 512 entries. You can scroll the lower window by using the W key to move the Active indicator to it, and by using the Up and Down cursor keys. The lower window automatically scrolls every time a new entry is added, and either the first entry or last entry is visible. Upon reaching 512 entries, old entries are thrown out as new entries are added.

Some entries may be too long to completely fit in a screen line. You can use the Left and Right cursor keys to vertically scroll the lower window when it is marked Active. If your terminal can be resized (e.g. xterm), you may do so before starting IPTraf-ng.

Entries for packets received on LAN interfaces also include the source MAC address of the LAN host which delivered it. This behavior is enabled by turning on the Source MAC addrs in traffic monitor toggle in the *Configure...* menu.

Entry Details

In general, the entries in the lower window indicate the protocol, the IP datagram size (full frame size for non-IP, including ARP and RARP), the source address, the destination address, and the network interface the packet was detected on. However, some protocols have a little more information.

ICMP

ICMP entries are displayed in this format:

```
ICMP type [(subtype)] (size bytes) from source to destination  
[(src HWaddr srcMACaddress)] on interface
```

where type could be any of the following:

echo req, echo rply

ICMP echo request and reply. Usually used by the ping program and other network monitoring and diagnostic program.

dest unrch

ICMP destination unreachable. Something failed to reach its target. The dest unreach type is supplemented with a further indicator of the problem. Destination unreachable messages for TCP traffic causes the corresponding TCP entry in the upper window to be made available for reuse by new connections.

redirect

ICMP redirect. Usually generated by a router to tell a host that a better gateway is available.

src qnch

The ICMP source quench is used to stop a host from transmitting. It's a flow control mechanism for IP.

time excd

Indicates a packet's time-to-live value expired before it got to its destination. Mostly happens if a destination is too far away. Also used by the traceroute program.

router adv

ICMP router advertisement

router sol
 ICMP router solicitation

timestamp req
 ICMP timestamp request

timestamp rep
 ICMP timestamp reply

info req
 ICMP information request

info rep
 ICMP information reply

addr mask req
 ICMP address mask request

addr mask rep
 ICMP address mask reply

param prob
 ICMP parameter problem

bad/unknown

An unrecognized ICMP packet was received, or the packet is corrupted.
The destination unreachable message also includes information on the type of error encountered. Here are the destination unreachable codes:

ntwk
 network unreachable

host
 host unreachable

proto
 protocol unreachable

port
 port unreachable

pkt fltrd
 packet filtered (normally by an access rule on a router or firewall)

DF set
 the packet has to be fragmented somewhere, but its don't fragment (DF) bit is set.

src rte fail
 source route failed

```
src isltd
    source isolated (obsolete)

net comm denied
    network communication denied

host comm denied
    host communication denied

net unrch for TOS
    network unreachable for specified IP type-of-service

host unrch for TOS
    host unreachable for specified IP type-of-service

prec violtn
    precedence violation

prec cutoff
    precedence cutoff

dest net unkn
    destination network unknown

dest host unkn
    destination network unknown
```

For more information on ICMP, see RFC 792.

OSPF

OSPF messages also include a little more information. The format of an OSPF message in the window is:

```
OSPF type (a=area r=router) (sizebytes) from source to destination
[(src HWaddr srcMACaddress)] on interface
```

The type can be one of the following:

```
hlo
    OSPF hello. Hello messages establish OSPF communications and keep routers
    informed of each other's presence.

DB desc
    OSPF Database Description

LSR
    OSPF Link State Request

LSU
    OSPF Link State Update. Messages indicating the states of the OSPF network
    links
```

LSA

OSPF Link State Acknowledgment

The entries in parentheses:

`a=area`

The area number of the OSPF message

`r=router`

The IP address of the router that generated the message. It is not necessarily the same as the source address of the encapsulating IP packet.

Many times, the destination addresses for OSPF packets are class D multicast addresses in standard dotted decimal notation or (if reverse lookup is enabled), hosts under the `MCAST.NET` domain. Such multicast addresses are defined as follows:

`224.0.0.5 (OSPF-ALL.MCAST.NET)`

OSPF all routers

`224.0.0.6 (OSPF-DSIG.MCAST.NET)`

OSPF all designated routers

See RFC 1247 for details on the OSPF protocol.

Additional Information

When started from the main menu and logging is enabled, the IP traffic monitor prompts you for a log file name. The default name is `ip_traffic-n.log` (where `n` is what instance of the traffic monitor this is (1, 2, 3, and so on)). (e.g. if this is the first instance, the default file name will be `ip_traffic-1.log`.)

When started with the `-i` parameter, the log filename can be specified with the `-L` parameter. See the Command-line Parameters section above for more information.

On busy networks, the display may become cluttered with traffic you're not interested in. To control the traffic monitor's output, you can apply a *filter*. See Chapter 7, Filters for more information on IPTraff-ng's filters.

At any time, you can press X or Q to return to the main menu (or back to the shell if the monitor was started with **iptraf-ng -i**).

Chapter 4. Network Interface Statistics

There are two network interface statistics facilities: the general interface statistics, which displays a statistical summary of all attached interfaces, and the detailed interface statistics, which shows more statistical and load information about a single selected interface.

General Interface Statistics

The second menu option displays a list of attached network interfaces, and some general packet counts. Specifically, it displays counts of IP, non-IP, and bad IP packets (packets with IP checksum errors). It also includes an activity indicator, which shows the number of kilobits and packets the interface sees per second. All figures are for incoming and outgoing packets. (Again, considering promiscuous mode for LAN interfaces, which simply causes the machine to intercept all packets). This is useful for general monitoring of all attached interfaces. If byte counts and additional information are needed for a specific interface, the *Detailed interface statistics* option is also available.

The activity indicators can be toggled between kbits/s and kbytes/s with the *Activity mode* configuration option.

The general statistics window will dynamically add new entries as packets from newly-created interfaces (e.g. new PPP interfaces) are intercepted. Long lists can be scrolled with the Up, Down, PgUp, and PgDn keys.

This monitor is affected by IPTraf-ng's filters as described in Chapter 7.

Copies of the statistics are written to the log file `iface_stats_general.log` at regular intervals if logging is enabled. See the *Logging* option in the Configuration chapter.

This facility can be started directly from the command line with the **-g** option to the **iptraf-ng** command. When started from the command line, the log filename and log interval can be specified with the **-L** and **-I** parameters respectively. See the Command-line Parameters section above for more information.

IPTraf

| Iface | Total | IP | NonIP | BadIP | Activity |
|-------|-------|-----|-------|-------|----------------|
| lo | 0 | 0 | 0 | 0 | 0.00 kbits/sec |
| eth0 | 142 | 139 | 3 | 0 | 7.00 kbits/sec |

Elapsed time: 0:00 ——— Total, IP, NonIP, and BadIP are packet counts
 Up/Down/PgUp/PgDn—scroll window X/Ctrl+X—Exit

Figure 4-1. The general interface statistics screen

You can press X or Q to return to the main menu.

Detailed Interface Statistics

The third menu option displays packet statistics for any selected interface. It provides basically the same information as the *General interface statistics* option, with additional details. This facility provides the following information:

- Total packet and byte counts
- IP packet and byte counts
- TCP packet and byte counts
- UDP packet and byte count
- ICMP packet and byte counts
- Other IP-type packet and byte counts
- Non-IP packet and byte counts
- Checksum error count
- Interface activity
- Broadcast packet and byte counts

All IP byte counts (IP, TCP, UDP, ICMP, other IP) include IP header data and payload. The data link header is not included. The full frame length (including data-link header) is included in the non-IP and Total byte count. All data-link headers are also included in the Total byte counts.




Figure 4-2. The detailed interface statistics screen

The upper portion of the screen contains the packet and byte counts for all IP and non-IP packets intercepted on the interface. The lower portion contains the total, incoming, and outgoing interface data rates.

This facility also displays incoming and outgoing counts and data rates.

An outgoing packet is one that exits your interface, regardless of whether it originated from your machine or came from another machine and was routed through yours. An incoming packet is one that enters your interface, either addressed to you directly, broadcast, multicast, or captured promiscuously.

The rate indicators can be set to display kbits/s or kbytes/s with the *Activity mode* configuration option.

Note: Buffering and some other factors may affect the data rates, notably the outgoing rate, causing it to reflect a higher figure than the actual rate at which the interface is sending.

The figures are logged at regular intervals if logging is enabled. The default log file name at the prompt is `iface_stats_detailed-iface.log` where `iface` is the selected interface for this session (for example, `iface_stats_detailed-eth0.log`).

If you wish to start this facility directly from the command line, you can specify the `-d` parameter and an interface to monitor. For example,

```
iptraf-ng -d eth0
```

starts the statistics for `eth0`. The interface must be specified, or IPtraf-ng will not start the facility.

When started from the command line, the log filename and log interval can be specified with the `-L` and `-I` parameters respectively. See the Command-line Parameters section above for more information.

Note: In both the general and detailed statistics screens, as well as in the IP traffic monitor, the packet counts are for actual network packets (layer 2), not the logical IP packets (layer 3) that may be reconstructed after fragmentation. That means, if a packet was fragmented into four pieces, and these four fragments pass over your interface, the packet counts will indicate four separate packets.

The figure for the IP checksum errors is a packet count only, because the corrupted IP header cannot be relied upon to give a correct IP packet length value.

This facility's output is also affected by IPtraf-ng's filters. See Chapter 7 for more information on filters.

Pressing X or Q takes you back to the main menu (if this facility was started with the command-line option, X or Q drops you back to the shell).

Chapter 5. Statistical Breakdowns

Statistical breakdowns contain two facilities that break down traffic counts by either packet size or TCP/UDP port.

Packet Sizes

The packet size breakdown takes the interface's Maximum Transmission Unit (MTU) size and divides it into 20 brackets, each bracket containing a range of sizes. As a packet is captured, its size is determined and the appropriate bracket is incremented.

This facility provides an idea as to the packet sizes passing over your network, and can aid in network (re)design decisions.

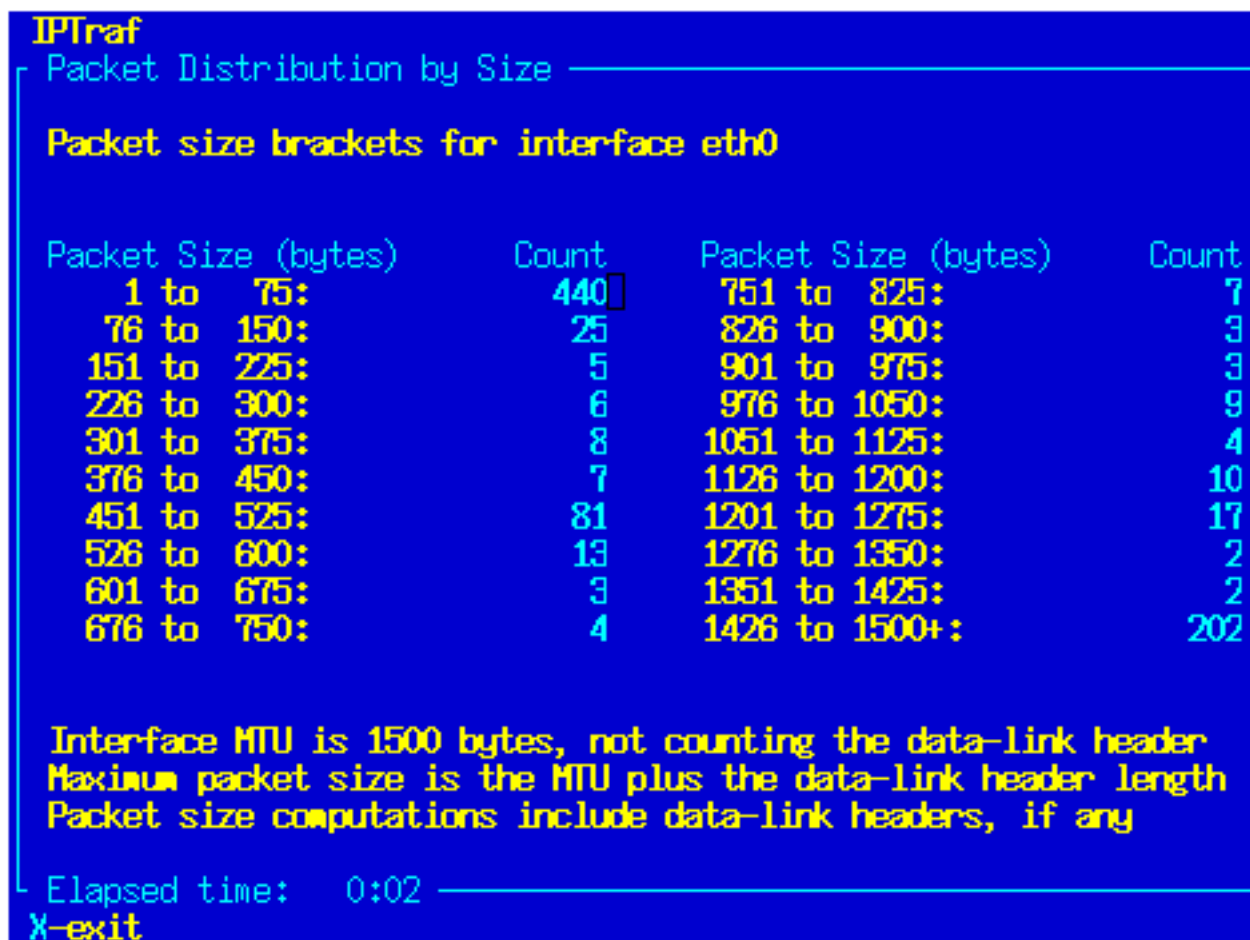


Figure 5-1. The packet size statistical breakdown

If logging is enabled, copies of the statistics are written at regular intervals to a log file. The default log file name is `packet_size-iface.log` where `iface` is the selected interface for this session (for example, `packet_size-eth0.log`).

IPTraf-ng's filters do not affect this facility.

The packet size breakdown can also be invoked straight from the command line by specifying the `-z iface` parameter. The interface parameter is required. For example, this command runs the facility on interface `eth0`.

```
iptraf-ng -z eth0
```

When started from the command line, the log filename and log interval can be specified with the `-L` and `-I` parameters respectively. See the Command-line Parameters section above for more information.

To exit, press `X` or `Ctrl+X`.

TCP and UDP Traffic Statistics

IPTraF-ng also includes a facility that generates statistics on TCP and UDP traffic. This facility displays counts of all TCP and UDP packets with source or destination ports numbered less than 1024. Ports 1 to 1023 are reserved for the TCP/IP application protocols (well-known ports).

| IPTraF | | | | | |
|---|------|---------|--------|---------|----------|
| Proto/Port | Pkts | Bytes | PktsTo | BytesTo | PktsFrom |
| TCP/www | 6064 | 1960227 | 3490 | 387688 | 257 |
| TCP/8088 | 1328 | 411655 | 647 | 71848 | 68 |
| TCP/webcache | 545 | 209710 | 269 | 21707 | 270 |
| TCP/pop3 | 508 | 169510 | 220 | 8952 | 28 |
| TCP/smtp | 177 | 86150 | 88 | 79197 | 8 |
| UDP/domain | 352 | 40643 | 192 | 13357 | 16 |
| TCP/netbios-ss | 160 | 22112 | 86 | 9408 | 7 |
| UDP/netbios-ns | 164 | 15530 | 130 | 10337 | 3 |
| TCP/https | 22 | 7533 | 12 | 1553 | 1 |
| TCP/telnet | 45 | 4649 | 25 | 2052 | 2 |
| TCP/ftp | 25 | 1269 | 13 | 746 | 1 |
| UDP/netbios-dg | 5 | 1177 | 3 | 703 | 1 |
| TCP/nntp | 7 | 578 | 4 | 213 | 1 |
| TCP/74 | 6 | 564 | 6 | 564 | 1 |
| TCP/40 | 9 | 540 | 9 | 540 | 1 |
| UDP/bootps | 1 | 328 | 1 | 328 | 1 |
| UDP/bootpc | 1 | 328 | 0 | 0 | 1 |
| UDP/ntp | 8 | 608 | 4 | 304 | 4 |
| TCP/81 | 7 | 332 | 5 | 252 | 2 |
| TCP/tproxy | 9 | 508 | 9 | 508 | 0 |
| 26 entries | | | | | |
| Elapsed time: 0:00 | | | | | |
| Protocol data rates (kbits/s): 165.25 in 537.00 out 7 | | | | | |
| Up/Down/PgUp/PgDn-scroll window S-sort X-exit | | | | | |

Figure 5-2. The TCP/UDP service monitor

The statistics window indicates the protocol (TCP or UDP), the port number, the total packets and bytes counted for this particular protocol/port combination, the packets and bytes destined for that protocol and port, and the packets and bytes coming from that protocol and port.

Byte counts include the IP header and payload only. The data link header is not included.

The protocol/port indicators are color-coded for easier identification on color terminals. TCP indicators are in yellow, UDP in bright green.

Some network applications or protocols may use port numbers higher than 1023. Examples of these include application proxy servers (HTTP proxy servers typically use values like 8000, 8080, 8888, and the like), and IRC (IRC servers commonly accept connections on ports 6660 to 6669). These ports are by default not included in the counts. If you do want to include a higher-numbered port in the statistics, you can add them yourself from the *Configure.../Additional ports...* menu item. See the section below.

If logging is enabled, The statistics are also written to a log file (the default name is `tcp_udp_services-iface.log`, where `iface` is the selected interface (for example, `tcp_udp_services-eth0.log`).

IPTraf-ng computes the total, incoming, outgoing, and data rates of the protocol currently indicated by the facility's highlight bar. The data rates are indicated at the bottom of the screen. If logging is enabled, the average data rates since the start of the facility are placed in the log file.

The Up and Down cursor keys move the highlight bar. Pressing X or Ctrl+X exits and returns to the main menu (or the shell if it was started from the command line).

Sorting TCP/UDP Entries

Pressing the S key brings up a window which allows you to select the field by which the entries will be sorted. You can press R to sort by port, P to sort by total packets, B to sort by total bytes, T to sort by incoming packets (packets to), O to sort by incoming bytes (bytes to), F to sort by outgoing packets (packets from) and M to sort by outgoing bytes (bytes from). Pressing any other key cancels the sort.

Port numbers are sorted in ascending order (least first) but statistics are sorted in descending order (largest counts first).

As with the IP traffic monitor, sorting is performed only with this sequence. Automatic sorting is not performed so as not to affect performance.

IPtraf

| Proto/Port | Pkts | Bytes | PktsTo | BytesTo | PktsFrom |
|----------------|-------|---------|--------|---------|----------|
| TCP/www | 19978 | 6340546 | 11261 | 1149682 | 871 |
| TCP/8088 | 3420 | 1012240 | 1730 | 173619 | 169 |
| TCP/webcache | 1542 | 636085 | 751 | 55118 | 79 |
| TCP/pop3 | 1137 | 380508 | 518 | 21113 | 61 |
| TCP/smtp | 443 | 140717 | 220 | 121163 | 22 |
| UDP/domain | 1168 | 132623 | | | |
| TCP/netbios-ss | 409 | 52954 | | | |
| TCP/https | 171 | 39711 | | | |
| UDP/netbios-ns | 367 | 34736 | | | |
| TCP/telnet | 113 | 12548 | | | |
| TCP/nntp | 57 | 5157 | | | |
| TCP/auth | 76 | 3870 | | | |
| TCP/ftp | 105 | 5355 | | | |
| UDP/netbios-dg | 10 | 2304 | | | |
| UDP/nntp | 30 | 2280 | | | |
| TCP/40 | 37 | 2220 | | | |
| TCP/81 | 37 | 1844 | | | |
| TCP/74 | 15 | 1410 | | | |
| TCP/tproxy | 26 | 1340 | | | |
| UDP/bootps | 2 | 656 | | | |
| 30 entries | | | 2 | 656 | |

Elapsed time: 0:01

Protocol data rates (kbits/s): 164.00 in 704.50 out 8

Up/Down/PgUp/PgDn-scroll window S-sort X-exit

Select sort criterion

R - port number

P - total packets

B - total bytes

T - packets to

O - bytes to

F - packets from

M - bytes from

Any other key - cancel s

Figure 5-3. The TCP/UDP monitor's sort criteria

Additional Information

IPtraf-ng's filters affect the output of this facility. See Chapter 7, Filters for more information about filters.

If you wish to start this facility from the command line, you can use the `-s` option followed by an interface to monitor. For example,

```
iptraf-ng -s eth0
```

brings up this module for traffic on `eth0`. The interface must be specified, or IPtraf-ng will drop back to the shell.

When started from the command line, the log filename and log interval can be specified with the `-L` and `-I` parameters respectively. See the Command-line Parameters section above for more information.

Chapter 6. LAN Station Statistics

The LAN station monitor (Ethernet station monitor on versions prior to 1.3.0) discovers MAC addresses and displays statistics on the number of incoming, and outgoing packets. It also includes figures for incoming and outgoing kilobits per second for each discovered station.

The entry above each line of statistics is the station's LAN type (Ethernet, PLIP, or FDDI) and the hardware MAC address. Each statistics line consists of the following information:

- Total packets incoming
- IP packets incoming
- Total bytes incoming
- Incoming rate
- Total packets outgoing
- IP packets outgoing
- Total bytes outgoing
- Outgoing rate

The byte counts include the data link header. The activity indicators can be set to display kbits/s or kbytes/s with the *Activity mode* configuration option.

This facility works only for Ethernet, PLIP, and FDDI frames. Loopback. SLIP/PPP networks are not monitored here.

IPTraf

| | PktsIn | IP In | BytesIn | InRate | PktsOut | IP Out | Bytes |
|---|--------|-------|---------|--------|---------|--------|-------|
| Ethernet HW addr: 0050dac010e9 on eth0 | | | | | | | |
| L | 90 | 90 | 6401 | 0.0 | 128 | 127 | 8 |
| Ethernet HW addr: 0030f212f000 on eth0 | | | | | | | |
| L | 133 | 133 | 8958 | 0.0 | 163 | 141 | 15 |
| Ethernet HW addr: 01005e000005 on eth0 | | | | | | | |
| L | 75 | 75 | 11482 | 2.4 | 0 | 0 | |
| Ethernet HW addr: 00d0baccceb47 on eth0 | | | | | | | |
| L | 0 | 0 | 0 | 0.0 | 18 | 0 | |
| Ethernet HW addr: 0180c2000000 on eth0 | | | | | | | |
| L | 18 | 0 | 1152 | 0.2 | 0 | 0 | |
| Ethernet HW addr: 00000c4340a0 on eth0 | | | | | | | |
| L | 0 | 0 | 0 | 0.0 | 26 | 26 | |
| Ethernet HW addr: 006097b77e2e on eth0 | | | | | | | |
| L | 0 | 0 | 0 | 0.0 | 2 | 0 | |
| Ethernet HW addr: ffffffff on eth0 | | | | | | | |
| L | 47 | 4 | 3672 | 0.4 | 0 | 0 | |
| Ethernet HW addr: 0050733f6b21 on eth0 | | | | | | | |
| L | 0 | 0 | 0 | 0.0 | 4 | 0 | |
| Ethernet HW addr: 003094152f01 on eth0 | | | | | | | |
| L | 0 | 0 | 0 | 0.0 | 15 | 15 | |

16 entries — Elapsed time: 0:00 — InRate and OutRate are in
 Up/Down/PgUp/PgDn—scroll window S-sort X-exit

Figure 6-1. The LAN station monitor

Copies of the statistics are written to a log file at regular intervals if logging is enabled. The default log file name is `lan_statistics-n.log`, where `n` is the instance number of this facility (for example, if this is the first instance, the generated default log file name is `lan_statistics-1.log`).

Sorting the LAN Station Monitor Entries

Press `S` to sort the entries. A box will pop up and display the keys you can press to select the field by which the entries will be sorted. Press `P` to sort by total incoming packets, `I` to sort by incoming IP packets, `B` to sort by total incoming bytes, `K` to sort by total outgoing packets, `O` to sort by outgoing IP packets, and `Y` to sort by total outgoing bytes. Pressing any other key cancels the sort.

IPtraf

| PktsIn | IP In | BytesIn | InRate | PktsOut | IP Out | BytesOut |
|---|-------|---------|--------|---------|--------|----------|
| Ethernet HW addr: 0000c4340a0 on eth0 | | | | | | |
| 7113 | 7111 | 3617373 | 1901.4 | 8413 | 8411 | 157 |
| Ethernet HW addr: 0030f212f000 (cebu-7206) on eth0 | | | | | | |
| 8146 | 8146 | 1541936 | 853.4 | 7125 | 7125 | 361 |
| Ethernet HW addr: 00104b0e4bad (proxy.ceu.mozcom.com) on eth0 | | | | | | |
| 218 | 218 | 29774 | 15. | | | |
| Ethernet HW addr: 00d0b7b7ea8d on eth0 | | | | | | |
| 30 | 30 | 2580 | 1. | | | |
| Ethernet HW addr: 00d058890b0f on eth0 | | | | | | |
| 2 | 2 | 444 | 0. | | | |
| Ethernet HW addr: 01005e000005 on eth0 | | | | | | |
| 45 | 45 | 6386 | 3. | | | |
| Ethernet HW addr: 01005e000006 on eth0 | | | | | | |
| 47 | 47 | 9510 | 6. | | | |
| Ethernet HW addr: 00d0bacceb43 on eth0 | | | | | | |
| 0 | 0 | 0 | 0. | | | |
| Ethernet HW addr: 0180c2000000 on eth0 | | | | | | |
| 16 | 0 | 1024 | 0. | | | |
| Ethernet HW addr: 00d0bacceb44 on eth0 | | | | | | |
| 0 | 0 | 0 | 0.0 | 8 | 0 | |
| 17 entries — Elapsed time: 0:00 — InRate and OutRate are in | | | | | | |
| Up/Down/PgUp/PgDn-scroll window S-sort X-exit | | | | | | |

Select sort criterion

P - total packets in

I - IP packets in

B - total bytes in

K - total packets out

O - IP packets out

Y - total bytes out

Any other key - cancel sort

Figure 6-2. The LAN station monitor's sort criteria

When started from the command line, the log filename and log interval can be specified with the `-L` and `-I` parameters respectively. See the Command-line Parameters section above for more information.

Additional Information

The window can be scrolled with the Up and Down cursor keys. Press X or Q to return to the main menu (or the shell if this facility was started with the `-l` command-line option).

The output of this facility is affected by any applied IPtraf-ng filter.

Chapter 7. Filters

Filters are used to control the information displayed by all facilities. You may want to view statistics only on particular traffic so you must restrict the information displayed. The filters also apply to logging activity.

The IPTraf-ng filter management system is accessible through the *Filters...* submenu.




Figure 7-1. The Filters submenu

IP Filters

The *Filters/IP...* menu option allows you to define a set of rules that determine what IP traffic to pass to the monitors. Selecting this option pops up another menu with the tasks used to define and apply custom IP filters.




Figure 7-2. The IP filter menu

Defining a New Filter

A freshly installed program will have no filters defined, so before anything else, you will have to define a filter. You can do this by selecting the *Define new filter...* option.

Selecting this option displays a box asking you to enter a short description of the filter you are going to define. Just enter any text that clearly identifies the nature of the filter.




Figure 7-3. The IP filter name dialog

Press Enter when you're done with that box. As an alternative, you can also press Ctrl+X to cancel the operation.

The Filter Rule Selection Screen

After you enter the filter's description, you will be taken to a blank rule selection box. At this screen you manage the various rules you define for this filter. You can opt to insert, append, edit, or delete rules.




Figure 7-4. The filter rule selection screen. Selecting an entry displays that set for editing

Any rules defined will appear here. You will see the source and destination addresses, masks and ports (long addresses and masks may be truncated) and whether this rule includes or excludes matching packets.

Between the source and destination parameters is an arrow that indicates whether the rule matches packets (single-headed) only exactly or whether it matches packets flowing in the opposite direction (double-headed).

At this screen, press I to insert at the current position of the selection bar, A to append a rule to the end of the list, Enter to edit the highlighted rule and D to delete the selected rule. With an empty list, A or I can be used to add the first rule.

To add the first rule, press A or I. You will then be presented with a dialog box that allows you to enter the rule's parameters.

Entering Filter Rules

You can enter addresses of individual hosts, networks, or a catch-all address. The nature of the address will be determined by the wildcard mask.

You'll notice two sets of fields, marked `Source` and `Destination`. You fill these out with the information about your source and targets.

Fill out the host name or IP address of the hosts or networks in the first field marked `Host name/IP Address`. Enter it in standard dotted-decimal notation. When done, press `Tab` to move to the `Wildcard mask` field. The wildcard mask is similar but not exactly identical to the standard IP subnet mask. The wildcard mask is used to determine which bits to ignore when processing the filter. In most cases, it will work very closely like a subnet mask. Place ones (1) under the bits you want the filter to recognize, and keep zeros (0) under the bits you want the filter to ignore. For example:

To recognize the host 207.0.115.44

| | |
|---------------|-----------------|
| IP address | 207.0.115.44 |
| Wildcard mask | 255.255.255.255 |

To recognize all hosts belonging to network 202.47.132.x

| | |
|---------------|---------------|
| IP address | 202.47.132.0 |
| Wildcard mask | 255.255.255.0 |

To recognize all hosts with any address:

| | |
|---------------|---------|
| IP address | 0.0.0.0 |
| Wildcard mask | 0.0.0.0 |

The IP address/wildcard mask mechanism of the display filter doesn't recognize IP address class. It uses a simple bit- pattern matching algorithm.

The wildcard mask also does not have to end on a byte boundary; you may mask right into a byte itself. For example, 255.255.255.224 masks 27 bits (255 is 11111111, 224 is 11100000 in binary).

IPTraf-ng also accepts host names in place of the IP addresses. IPTraf-ng will resolve the host name when the filter is loaded. When the filter is interpreted, the wildcard mask will also be applied. This can be useful in cases where a single host name may resolve to several IP addresses.

Tip: See the *Linux Network Administrator's Guide* if you need more information on IP addresses and subnet masking.

Tip: IPTraf-ng allows you to specify the wildcard mask in Classless Interdomain Routing (CIDR) format. This format allows you to specify the number of 1-bits that mask the address. CIDR notation is the form *address/bits* where the *address* is the IP address or host name and *bits* is the number of 1-bits in the mask. For example, if you want to mask 10.1.1.0 with 255.255.255.0, note that 255.255.255.0 has 24 1-bits, so instead of specifying 255.255.255.0 in the wildcard mask field, you can just enter 10.1.1.0/24 in the address field. IPTraf-ng will translate the mask bits into an appropriate wildcard mask and fill in the mask field the next time you edit the filter rule.

If you specify the mask in CIDR notation, leave the wildcard mask fields blank. If you fill them up, the wildcard mask fields will take precedence.

The `Port` fields should contain a port number or range of any TCP or UDP service you may be interested in. If you want to match only a single port number, fill in the first field, while leaving the second blank or set to zero. Fill in the second field if you want to match a range of ports (e.g. 80 to 90). Leave the first field blank or set to zero to let the filter ignore the ports altogether. You will most likely be interested in target ports rather than source ports (which are usually unpredictable anyway, perhaps with the exception of FTP data).

Non-TCP and non-UDP packets are not affected by these fields, and these are used only when filtering TCP or UDP packets.

Fill out the second set of fields with the parameters of the opposite end of the connection.

Tip: Any address or mask fields left blank default to 0.0.0.0 while blank `Port` fields default to 0. This makes it easy to define filter rules if you're interested only in either the source or destination, but not the other. For example, you may be interested in traffic originating from network 61.9.88.0, in which case you just enter the source address, mask and port in the `Source` fields, while leaving the `Destination` fields blank.

The next fields let you specify which IP-type protocols you want matched by this filter rule. Any packet whose protocol's corresponding field is marked with a `Y` is matched against the filter's defined IP addresses and ports, otherwise they don't pass through this filter rule.

If you want to evaluate all IP packets just mark with `Y` the `All IP` field.

For example, if you want to see only all TCP traffic, mark the `TCP` field with `Y`.

The long field marked `Additional protocols` allows you to specify other protocols by their IANA number. (You can view the common IP protocol number in the `/etc/protocols` file). You can specify a list of protocol numbers or ranges separated by commas, Ranges have the beginning and ending protocol numbers separated with a hyphen.

For example, to see the RSVP (46), IP mobile (55), and protocols (101 to 104), you use an entry that looks like this:

```
46, 55, 101-104
```

It's certainly possible to specify any of the protocols listed above in this field. Entering `1-255` is functionally identical to marking `All IP` with a `Y`.

The next field is marked `Include/Exclude`. This field allows you to decide whether to include or filter out matching packets. Setting this field to `I` causes the filter to pass matching packets, while setting it to `E` causes the filter to drop them. This field is set to `I` by default.

The last field in the dialog is labeled `Match opposite`. When set to `Y`, the filter will match packets flowing in the opposite direction. Previous versions of IPTraf-ng used to match TCP packets flowing in either direction, so the source and destination address/mask/port combinations were actually interchangeable. Starting with IPTraf 3.0, when filters extended to more than just the IP traffic monitor, this behavior is no longer the default throughout IPTraf-ng except in the IP traffic monitor's TCP window.

Note: For TCP packets, this field is used in all facilities except the IP traffic monitor. Because the IP traffic monitor must capture TCP packets in both directions to properly determine a closed connection, the filter automatically matches packets in the opposite direction, regardless of this field's setting. However in all other facilities, automatic matching of the reverse packets is not performed unless you set this field to `Y`.

Filters for UDP and other IP protocols do not automatically match packets in the opposite direction unless you set the field to `Y`, even in the IP traffic monitor.

Press Enter to accept all parameters when done. The parameters will be accepted and you'll be taken back to the rule selection box. You can then add more rules by pressing A or you can insert new rules at any point by pressing I. Should you make a mistake, you can press Enter to edit the selected filter. You may enter as many sets of parameters as you wish. Press Ctrl+X when done.

Note: Because of the major changes in the filtering system since IPTraf 2.7, old filters will no longer work and will have to be redefined.

IPTraf

Source ————— Destination

IP address 0.0.0.0 61.9.

Wildcard mask 0.0.0.0 255.2

Port to 80

Port fields apply only to TCP and U

Protocols to match (Enter Y beside each protocol to match.)

All IP ☐ TCP ☒ UDP ☐

OSPF ☐ IGP ☐ IGRP ☐

IPSec AH ☐ IPSec ESP ☐

Additional protocols or ranges (e.

Include/Exclude (I/E) I Match opposite (Y/N) N

Tab—next field Enter—accept Ctrl+X—cancel

Up/Down—move ptr I—insert A—add to list D—delete Enter—ed

Figure 7-5. The IP filter parameters dialog

Examples

To see all traffic to/from host 202.47.132.1 from/to 207.0.115.44, regardless of TCP port

| | | |
|----------------------|-----------------|-----------------|
| Host name/IP Address | 202.47.132.2 | 207.0.115.44 |
| Wildcard mask | 255.255.255.255 | 255.255.255.255 |
| Port | 0 | 0 |
| Protocols | TCP: Y | |
| Include/Exclude | I | |
| Match opposite | Y | |

To see all traffic from host 207.0.115.44 to all hosts on network 202.47.132.x

| | | |
|----------------------|-----------------|---------------|
| Host name/IP Address | 207.0.115.44 | 202.47.132.0 |
| Wildcard mask | 255.255.255.255 | 255.255.255.0 |
| Port | 0 | 0 |
| Protocols | All IP: Y | |
| Include/Exclude | I | |
| Match opposite | N | |

To see all Web traffic (to and from port 80) regardless of source or destination

| | | |
|----------------------|---------|---------|
| Host name/IP Address | 0.0.0.0 | 0.0.0.0 |
| Wildcard mask | 0.0.0.0 | 0.0.0.0 |
| Port | 80 | 0 |
| Protocols | TCP: Y | |
| Include/Exclude | I | |
| Match opposite | Y | |

To see all IRC traffic from port 6666 to 6669

| | | |
|----------------------|---------|--------------|
| Host name/IP Address | 0.0.0.0 | 0.0.0.0 |
| Wildcard mask | 0.0.0.0 | 0.0.0.0 |
| Port | 0 | 6666 to 6669 |
| Protocols | TCP: Y | |
| Include/Exclude | I | |
| Match opposite | Y | |

To see all DNS traffic, (TCP and UDP, destination port 53) regardless of source or destination

| | | |
|----------------------|---------------|---------|
| Host name/IP Address | 0.0.0.0 | 0.0.0.0 |
| Wildcard mask | 0.0.0.0 | 0.0.0.0 |
| Port | 0 | 53 |
| Protocols | TCP: Y UDP: Y | |
| Include/Exclude | I | |

Match opposite Y

To see all mail (SMTP) traffic to a single host (202.47.132.2) from anywhere

| | | |
|----------------------|---------|-----------------|
| Host name/IP Address | 0.0.0.0 | 202.47.132.2 |
| Wildcard mask | 0.0.0.0 | 255.255.255.255 |
| Port | 0 | 25 |
| Protocols | TCP: Y | |
| Include/Exclude | I | |
| Match opposite | N | |

To see traffic from from/to host sunsite.unc.edu to/from cebu.mozcom.com

| | | |
|----------------------|-----------------|-----------------|
| Host name/IP Address | sunsite.unc.edu | cebu.mozcom.com |
| Wildcard mask | 255.255.255.255 | 255.255.255.255 |
| Port | 0 | 0 |
| Protocols | All IP: Y | |
| Include/Exclude | I | |
| Match opposite | Y | |

To omit display of traffic to/from 140.66.5.x from/to anywhere

| | | |
|----------------------|---------------|---------|
| Host name/IP Address | 140.66.5.0 | 0.0.0.0 |
| Wildcard mask | 255.255.255.0 | 0.0.0.0 |
| Port | 0 | 0 |
| Protocols | All IP: Y | |
| Include/Exclude | E | |
| Match opposite | Y | |

You can enter as many parameters as you wish. All of them will be interpreted until the first match is found.

Excluding Certain Sites

Filters follow an implicit "no-match" policy, that is, only packets matching defined rules will be matched, others will be filtered out. This is similar to the access-list policy "whatever is not explicitly permitted is denied". If you want to show all traffic to/from everywhere, except certain places, you can specify the sites you wish to exclude, mark them with E in the Include/Exclude field, and define a general catch-all entry with source address 0.0.0.0, mask 0.0.0.0, port 0, and destination 0.0.0.0, mask 0.0.0.0, port 0, tagged with an I in the Include/Exclude field as the last entry.

For example:

To see all traffic except all SMTP (both directions), Web (both directions), and traffic (only) from 207.0.115.44

| | | |
|----------------------|---------|---------|
| Host name/IP address | 0.0.0.0 | 0.0.0.0 |
|----------------------|---------|---------|

| | | |
|----------------------|-----------------|---------|
| Wildcard mask | 0.0.0.0 | 0.0.0.0 |
| Port | 25 | 0 |
| Protocols | TCP: Y | |
| Include/Exclude | E | |
| Match opposite | Y | |
| Host name/IP address | 0.0.0.0 | 0.0.0.0 |
| Wildcard mask | 0.0.0.0 | 0.0.0.0 |
| Port | 80 | 0 |
| Protocols | TCP: Y | |
| Include/Exclude | E | |
| Match opposite | Y | |
| Host name/IP address | 207.0.115.44 | 0.0.0.0 |
| Wildcard mask | 255.255.255.255 | 0.0.0.0 |
| Port | 0 | 0 |
| Protocols | All IP: Y | |
| Include/Exclude | E | |
| Match opposite | N | |
| Host name/IP address | 0.0.0.0 | 0.0.0.0 |
| Wildcard mask | 0.0.0.0 | 0.0.0.0 |
| Port | 0 | 0 |
| Protocols | All IP: Y | |
| Include/Exclude | I | |
| Match opposite | N | |

Tip: To filter out all TCP, define a filter with a single entry, with a source of 0.0.0.0 mask 0.0.0.0 port 0, and a destination of 0.0.0.0 mask 0.0.0.0 port 0, with the Include/Exclude field marked E (exclude). Then apply this filter.

Applying a Filter

The above steps only add the filter to a defined list. To actually apply the filter, you must select *Apply filter...* from the menu. You will be presented with a list of filters you already defined. Select the one you want to apply, and press Enter.

The applied filter stays in effect over exits and restarts of the IPTraf-ng program until it is detached.

Editing a Defined Filter

Select *Edit filter...* to modify an existing filter. Once you select this option, you will be

presented with the list of defined filters. Select the filter you want to edit by moving the selection bar and press Enter.

Edit the description if you wish. Pressing Ctrl+X at this point will abort the operation, and the filter will remain unmodified. Press Enter to accept any changes to the filter description.

After pressing Enter, you will see the filter's rules. To edit an existing filter rule, move the selection bar to the desired entry and press Enter. A prefilled dialog box will appear. Edit its contents as desired. Press Enter to accept the changes or Ctrl+X to discard.

You can add a new filter rule by pressing I to insert at the selection bar's current position. When you press I, you will be presented with a dialog box asking you to enter the new rule data. Pressing A results in a similar operation, except the rule will be appended as the last entry in the rule list.

Pressing D deletes the currently pointed entry.

Press X or Ctrl+X to end the edit and save the changes.

Note: If you're editing the currently applied filter, you will need to re-apply the filter for the changes to take effect.

Note: Be aware that the filter processes the rules in order. In other words, if a packet matches more than one rule, only the first matching rule is followed.

Deleting a Defined Filter

Select *Delete filter...* from the menu to remove a filter from the list. Just move the selection bar to the filter you want to delete, and press Enter.

Detaching a Filter

The *Detach filter* option deactivates the filter currently in use. Selecting this option causes all TCP traffic to be passed to the monitors.

When you're done with the menu, just select the Exit menu option.

ARP, RARP, and other Non-IP Packet Filters

The *Non-IP* filter option toggles the display and logging of all non-IP packets, except ARP and RARP, which are toggled separately.

Chapter 8. Configuring IPTraf-ng

IPTraf-ng can be easily configured with the *Configure...* item in the main menu. The configuration is stored in the `/var/local/iptraf-ng/iptraf.cfg` file. If the file is not found, IPTraf-ng uses the default settings. Any changes to the configuration immediately get stored in the configuration file.

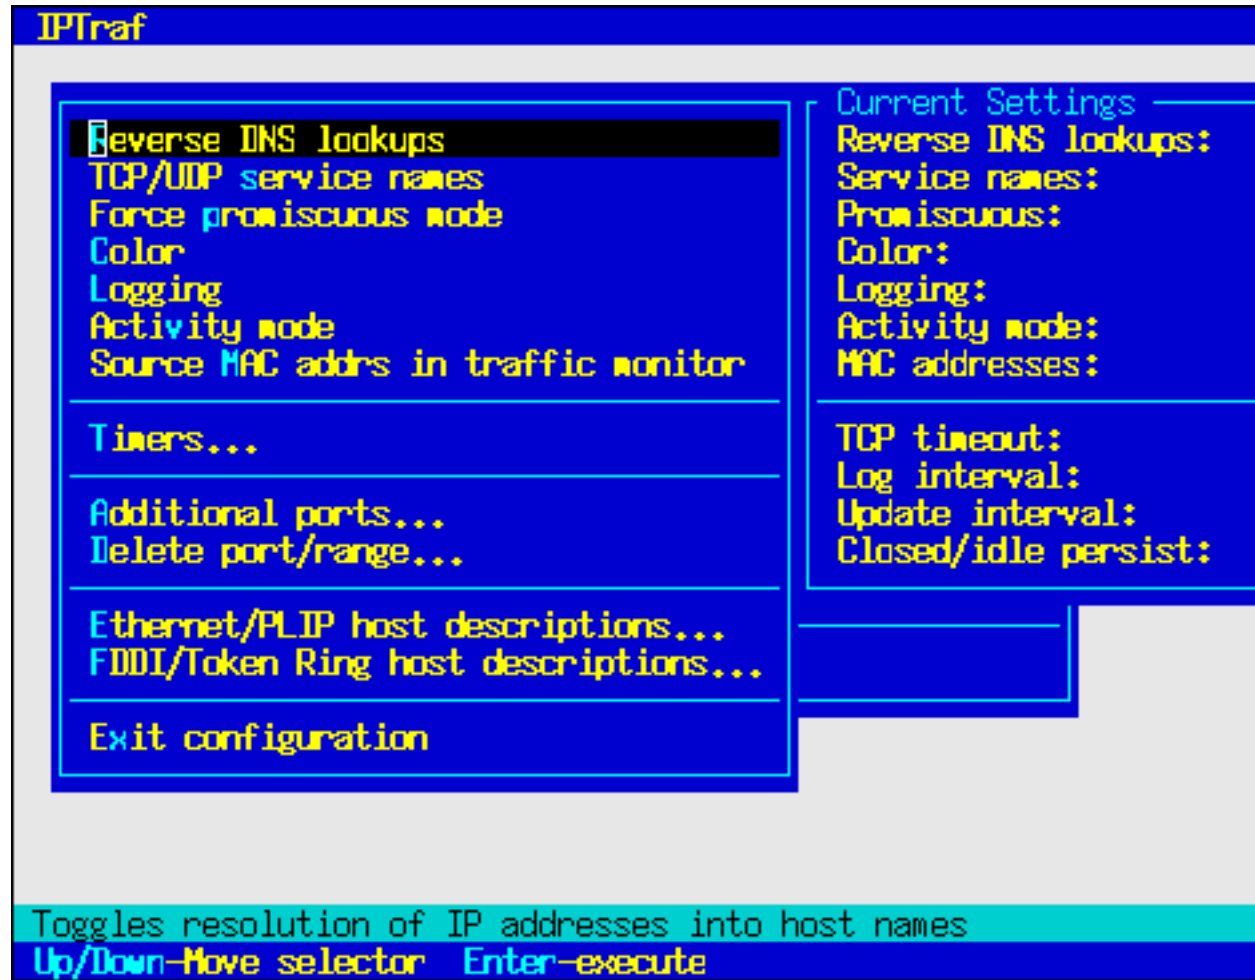


Figure 8-1. The IPTraf-ng configuration menu

Toggles

Reverse DNS Lookups

Activating reverse lookup causes IPTraf-ng to find out the name of the hosts with the addresses in the IP packets. When this option is enabled, IPTraf-ng's IP traffic monitor starts the DNS lookup server to help resolve IP addresses in the background while allowing IPTraf-ng to continue capturing packets.

This option is off by default.

TCP/UDP Service Names

This option, when on, causes IPTraf-ng to display the TCP/UDP service names (smtp, www, pop3, etc.) instead of their numeric ports (25, 80, 110, etc). The number-to-name mappings will depend on the systems services database file (usually `/etc/services`). Should there be no corresponding service name for the port number, the numeric form will still be displayed.

This setting is off by default.

Note: Reverse lookup and service name lookup take some time and may impact performance and increase the chances of dropped packets. Performance and results are best (albeit more cryptic) with both these settings off.

Force promiscuous

If this option is enabled, your LAN interfaces will capture all packets on your LAN. Using this option enables you to see all TCP connections and packets passing your LAN segment, even if they're not from or for your machine. When this option is active in the statistics windows, the Activity indicators will show a good estimate of the load on your LAN segment.

When this option is disabled, you'll only receive information about packets coming from and entering your machine.

The setting of this option affects all LAN (Ethernet, FDDI) interfaces on your machine, if you have more than one.

The interface's promiscuous flag is set only when a facility is started, and turned off when it exits. However, if promiscuous mode was already set when a facility was started, it remains set on exit.

If multiple instances of IPTraf-ng are started, the promiscuous setting is restored only upon exit of the last facility.

Note: Do not use other programs that change the interface's promiscuous flag at the same time you're using IPTraf-ng. The programs can interfere with each other's expected operations. While IPTraf-ng tries to obtain the initial setting of any promiscuous flags for restoration upon exit, other programs may not be as well-behaved, and they may turn off the promiscuous flags while IPTraf-ng is still monitoring.

Color

Turn this on with color monitors. Turn it off with black-and- white monitors or non-color terminals (like xterms). Changes to this setting will take effect the next time the program is started.

Color is on by default on consoles and color xterms, off on non-color terminals like xterms and VT100s.

Logging

When this option is active, IPTraf-ng will log information to a disk file, which can be examined or analyzed later. Since IPTraf-ng 2.4.0, IPTraf-ng prompts you for the name of the file to which to write the logs. It will provide a default name, which you are free to accept or change. The IP traffic monitor and LAN station monitor will

generate a log file name that is based on what instance they are (first, second, and so on). The general interface statistics' default log file name is constant, because it listens to all interfaces at once, and only one instance can run at one time.

The other facilities generate a log file name based on the interface they're listening on.

See the descriptions on the facilities above for the default log file names.

Press Enter to accept the log file name, or Ctrl+X to cancel. Canceling will turn logging off for that session.

The IP traffic monitor will write the following pieces of information to its log file:

- Start of the traffic monitor
- Receipt of the first TCP packet for a connection. If that packet is a SYN, (SYN) will be indicated in the log entry. (Of course, the traffic monitor may start in the middle of established connections. It will still count those packets. This also explains why some connection entries may become idle if the traffic monitor is started in the middle of a half-closed connection, and miss the first FIN. Such entries time out in a while.)
- Receipt of a FIN (with average flow rate)
- ACK of a FIN
- Timeouts of TCP entries (with average flow rate)
- Reset connections (with average flow rate)
- Everything that appears in the bottom window of the traffic monitor
- Stopping of the traffic monitor

Each log entry includes the date and time the entry was written. Logging is also affected by the defined filters.

Log files can grow very fast, so be prepared with plenty of free space and delete unneeded logs. Log write errors are not indicated.

Copies of the interface statistics, TCP/UDP statistics, packet size statistics, and LAN host statistics are also written to the log files at regular intervals. See *Log Interval...* in this chapter.

IPTraf-ng closes and reopens the active log file when it receives a `USR1` signal. This is useful in cases where a facility is run for long periods of time but the log files have to be cleared or moved.

To clear or move an active log file, rename it first. IPTraf-ng will continue to write to the file despite the new name. Then use the UNIX kill command to send the running IPTraf-ng process a `USR1` signal. IPTraf-ng will then close the log file and open another with the original name. You can then safely remove or delete the renamed file.

Do not delete an open log file. Doing so will only result in a file just as large but filled with null characters (ASCII code 0).

Logging comes disabled by default. The `USR1` signal is caught only if logging is enabled, it is ignored otherwise.

A valid specification of `-L` on the command line will automatically enable logging for that particular session. The saved configuration setting is not affected.

Activity mode

Toggles activity indicators in the interface and LAN statistics facilities between kilobits per second (kbits/s) or kilobytes per second (kbytes/s).

The default setting is kilobits per second.

Source MAC addrs in traffic monitor

When enabled, the IP traffic monitor retrieves the packets' source MAC addresses if they came in on an Ethernet, FDDI, or PLIP interface. The addresses appear in the lower window for non-TCP packets, while for TCP connections, they can be viewed by pressing M.

No such information is displayed if the network interface doesn't use MAC addresses (such as PPP interfaces).

This can be used to determine the actual source of the packets on your local LAN.

The traffic monitor also logs the MAC addresses with this option enabled. The default setting is off.

Timers

The *Timers...* submenu allows you to IPTraf-ng's interval and timeout functions.




Figure 8-2. The Timers configuration submenu

TCP Timeout

This figure determines the amount of time (in minutes) a connection entry may remain idle before it becomes eligible for replacement by a new connection. The default is 15 minutes. You may want to reduce this on an isolated (not connected to the Internet) LAN or a LAN connected to the Internet with high-speed links. Just enter the new value and press Enter. You can press Ctrl+X to leave the current value unchanged.

Log Interval

This figure determines the number of minutes between logging of interface statistics, TCP/UDP figures, and LAN host statistics. The default is 60 minutes. This figure is meaningless if logging is disabled.

This configuration item can be overridden with the `-I` when a facility is directly invoked from the command line (not accessed via the main menu), and remains effective for that particular session. The configured value is not affected.

Screen Update Interval

This value determines the rate in seconds at which the screen is updated. The default is 0, which means the screen is updated as fast as possible, giving close-to-realtime reflection of network activity. However, this high-speed update can cause incredible amounts of traffic if IPTraf-ng is run on a remote terminal (e.g. a Telnet or Secure Shell session). You can set this to a higher value, such as 1 or 2 seconds to slow down the updates.

This figure does not affect the rate of data capture. Only the screen refresh is affected. The figures are still updated as fast as possible, although the figure display will no longer be as close to realtime.

The default setting is 0, which shouldn't be a problem on the console. Set it to a slightly higher value on remote terminals or slow links. The setting affects all monitoring facilities.

Note: Updating the screen is one of the slowest operations in a program. Older versions of IPTraf-ng had a problem once network activity became very high. Because each packet caused a screen update, IPTraf-ng began spending more time with the screen updates, causing a loss of packets once network activity reached a certain point.

However, since many users like rapid counts on their screen, a compromise was incorporated. Even when the screen update interval is set to 0, there is still a 50ms delay between screen updates (except the LAN station monitor, which has a 100 ms delay). This is still visually fast, but provides more time to the packet capture routine. Higher delays may result in better accuracy of counts and activity.

In any case, this setting only affects screen updates. Capture still proceeds as fast as possible.

TCP closed/idle persistence

This parameter determines the interval (in minutes) at which the IP Traffic Monitor clears from the TCP display window all closed, idle, and timed out entries. Enter 0 to keep such entries on the screen indefinitely, disappearing only when replaced by new connections.

Note: The *TCP timeout...* option only tells IPTraf-ng how long it should take before a connection should be considered idle and open to replacement by new connections. This does not determine how long it remains onscreen. The *TCP closed/idle persistence...* parameter flushes entries that have been closed or reset, or idle for the number of minutes defined by the *TCP timeout...* option.

Custom Information

The remaining configuration items allow you to enter information which IPTraf-ng uses for its displays and logs.

Additional ports

Select this item to enter a port number to be included in the TCP/UDP counts in the TCP/UDP service statistics main menu item described above. By default, port numbers above 1023 are not monitored. If you do have a higher-numbered port to monitor, enter it here.

You will see two fields. If you have only one port to enter, just fill up the first field. To specify a range, fill both fields, the first port in the first field, the last port in the second field.

You can select this option multiple times to add more values or ranges.

Delete port/range

Select this item to remove a higher-numbered port number or port range you entered earlier with the *Additional ports...* option. A window will come up containing the entered ports and ranges. Select the entry you want delete and press Enter.

LAN Station Identifiers

The LAN station statistics facility monitors stations based on their respective MAC addresses. The hexadecimal notation of these addresses make them even more difficult to remember than the dotted-decimal IP addresses, so these facilities were added to help you better determine which station is which.

Selecting the *Ethernet/PLIP host descriptions...* or *FDDI host descriptions...* options brings up a submenu asking you to add, edit, or delete descriptions.

To add a new description, select the *Add description...* option. A dialog box will appear, asking you for the MAC address and an appropriate description. Type in the address in hexadecimal notation with no punctuation of any kind. The dialog box is case-insensitive for the address; the alphabetical digits A to F will be stored in lowercase.

Use the Tab key to move between fields and Enter to accept. Press Ctrl+X to discard this dialog and return to the main menu.

The description may be anything: the IP address, a fully-qualified domain name, or a description of your liking as long as the field can hold.

Enter as many descriptions as you need. Press Ctrl+X at a blank dialog after you have entered the last entry

These descriptions will be displayed alongside the MAC addresses in the LAN station monitor, together with the type of frame (Ethernet, PLIP, or FDDI).

An existing address or description may be edited by selecting the *Edit description...* option from the submenu. A panel will appear with a list of existing address descriptions. Select the one you wish to edit and press Enter. A dialog box identical to that when you add a description will appear with prefilled fields. Just backspace over and edit the fields. Press Enter to accept or Ctrl+X to cancel.

Selecting the *Delete description...* submenu item brings up the selection panel. Select the description you want to delete and press Enter. You can also press Ctrl+X to cancel the operation.

IPTraf 2.4 and later also recognizes the `/etc/ethers` file. Should a hardware address be present in the IPTraf-ng definition files and in `/etc/ethers`, the IPTraf-ng definition will be used.

Note: The description file for Ethernet and PLIP is `ethernet.desc`, while the FDDI mappings are stored in `fddi.desc` in the IPTraf-ng working directory. These files are in colon-delimited text format. Database engines or custom scripts can be told to append data lines to those files. Each line follows this simple format:

Chapter 8. Configuring IPTracing

address:description

For example

00201e457e:Cisco 3640 gateway

Do not put colons, periods, or any invalid characters in the MAC address.

Chapter 9. Background Operation

IPtraf-ng's facilities can be placed in the background solely for logging. When running in the background, it doesn't display any output on the screen, and doesn't receive input from the keyboard, and drops you back to the shell.

Before starting a statistical facility in the background, configure IPtraf-ng in the usual way (set filters, add TCP/UDP ports, etc).

Once that's done, exit all instances of IPtraf-ng on the system, then invoke IPtraf-ng from the command line with the parameter to start the facility you want, the timeout (-t) parameter if you wish, and the -B parameter to actually daemonize the program. For example, to run the IP traffic monitor in the background for all interfaces, issue the command

```
iptraf-ng -i all -B
```

To run the detailed interface statistics on interface `eth0` for 5 minutes in the background:

```
iptraf-ng -d eth0 -t 5 -B
```

If the timeout parameter is not specified, the facility will run until the process receives a USR2 signal. To stop a facility in the background, do a

```
ps x
```

at the command line, and find the process id (pid) of the iptraf-ng process you're looking for. Then send that process a USR2 signal with the kill command:

```
kill -USR2 pid
```

Since IPtraf-ng cannot send error messages to the terminal, all messages are written to the file `daemon.log` in the IPtraf-ng logging directory.

The -B parameter automatically enables logging regardless of its configured setting. The parameter is ignored if not used with one of the parameters to start a facility from the command line.

The log file can be specified with the -L command-line parameter. If this parameter is not specified, the default log file name for the facility will be used (see the descriptions of the facilities above for the default log name patterns). If you don't specify a path, the log file will be placed in `/var/log/iptraf-ng`.

The logging interval for all facilities (except the IP traffic monitor) can also be overridden with the -I command-line parameter.

Appendix A. Messages

IPTraf-ng's messages are presented in two ways. In interactive mode, messages are displayed in a distinctive message box. In daemon (background) mode, appropriate messages are written to the `iptraf-ng.log` file in the IPTraf-ng log directory (normally `/var/log/iptraf-ng`).

IPTraf-ng Messages

Unable to create config file

IPTraf-ng cannot create the configuration file. The most likely cause of this is that you didn't properly install the program, and the necessary directory `/var/local/iptraf-ng` does not exist. Can also be generated if you have a disk problem or if you have too many files open.

Unable to read config file

The configuration record cannot be read. You most likely have a disk problem.

Unable to write config file

The configuration file cannot be written. You either have a disk problem, or (more likely), your disk is full.

Enter an appropriate description for this filter

Enter something to clearly describe the filter you are defining.

Error loading filter list file

IPTraf-ng cannot access the list of defined TCP or UDP filters. Can also be an indicator of a bad disk.

Error writing filter list file

The filter list file cannot be written to. You may have trouble accessing your filters.

Unable to read TCP/UDP/misc IP filter file

IPTraf-ng cannot read the filter data off the file. Could be caused by a bad disk.

Error opening filter data file

IPTraf-ng cannot open the filter file. Could be caused by a shortage of file descriptors or a bad disk.

Unable to write filter data

IPTraf-ng cannot add the newly defined filter to the filter list. This may be due to a bad disk.

Cannot create filter data file

IPTraf-ng cannot create the filter record file. The defined filter is lost.

Unable to save filter changes

IPTraf-ng cannot save the changes you made to the filter. You probably have a disk error.

Unable to write filter state information

The current state of the filters cannot be saved. IPTraf-ng will be unable to correctly reload the filters the next time it's started. This can be caused by a bad disk or improper installation.

Unable to save interface flags

IPTraf-ng was unable to save the flags of the network interfaces. This is probably due to a bad installation or full filesystem.

Unable to retrieve saved interface flags

IPTraf-ng was unable to retrieve the save interface flags. Probably again due to a bad installation or full filesystem.

protocol filter data file in use; try again later

Filter state file in use; try again later

Another IPTraf-ng process is modifying the TCP, UDP or miscellaneous IP filter data or the filter state file and has locked the files or file. Try again once the other IPTraf-ng process has terminated or completed its modifications and unlocked the files.

Unable to resolve hostname

The indicated host name in the filter cannot be resolved into an IP address. Check the local hosts database `/etc/hosts` or your machine's DNS configuration or DNS server. The filter parameters will not be used.

Unable to open host description file

IPTraf-ng cannot open the file containing the descriptions for Ethernet or FDDI addresses. Could be due to a bad disk or a hit on the file descriptor limit.

Unable to write host description

IPTraf-ng was unable to write the description record for this Ethernet or FDDI address. Could be due to a bad disk or corrupted filesystem.

No descriptions

You tried to edit or delete a description with no previous descriptions defined.

Cannot open log file

There is a problem opening the log file. There is most likely a problem with the disk, or there are too many open files.

Unable to obtain interface list

IPTraf-ng was unable to retrieve the list of network interfaces from the `/proc` filesystem. This may be due to a badly configured kernel. IPTraf-ng needs `/proc` filesystem support.

No active interfaces. Check their status or the `/proc` filesystem.

IPTraf-ng found no active interfaces. Either all interfaces are down or the `/proc/net/dev` file was empty or unavailable. Activate at least one interface or check the `/proc/net/dev` file.

Unable to obtain interface parameters for interface

The system call to retrieve the interface's flags failed. Check your interface or kernel driver.

Promisc change failed for interface

The system call to change the promiscuous flag failed. Check your interface or its kernel driver.

Unable to open raw socket for flag change

IPTraf-ng was unable to open the necessary socket for the promiscuous change operation. May be due to a shortage of file descriptors.

Unable to open socket for MTU determination

Returned by the facility for detailed interface statistics if the raw socket's opening sequence failed. The facility will abort.

Unable to open raw socket

IPTraf-ng was unable to open the raw socket for packet capture. May be due to a shortage of file descriptors.

Reminder: IPTraf 2.x.x requires Linux kernel 2.2.x, with the Packet Socket option compiled in or installed as a module. IPTraf 2.x will return this error on a pre-2.2 kernel or on a 2.2 kernel without Packet Socket.

Unable to obtain interface MTU

The detailed statistics facility was unable to obtain the maximum transmission unit (MTU) for the selected interface. The facility will abort.

Specified interface not supported

The interface specified with the `-i`, `-d`, `-s`, `-l`, or `-z` command-line parameters is not supported by IPTraf-ng.

Specified interface not active

The interface specified with the `-i`, `-d`, `-s`, `-l`, or `-z` command-line parameters is supported, but not currently activated.

Fatal: memory allocation error

May occur if you have too little memory to allocate for windows, the menu system, or dialog boxes. IPTraf-ng tries to prevent further allocations if memory runs out during a monitor. However, this could also mean a bug if you're reasonably sure you're not out of memory. An instructional message on bug reporting follows this message.

Technical note: This is actually a response to the segmentation fault error (SIGSEGV).

This program can be run only by the system administrator

IPtraf-ng normally does not allow anybody but uid 0 (root) to run it. This measure is included for safety reasons. See the section on recompiling the program below if you want to override this. This feature is built in, and not part of the configuration

Your TERM variable is not set

The TERM (terminal type) environment variable must be set to a valid terminal type so that the screen management routines can function properly. Set it to the appropriate terminal type. Linux consoles typically have their TERM variables set to `linux`.

Received TERM signal

Not related to the previous message. The TERM (terminate) signal is normally used to gracefully shut down a program. This message simply indicates that the TERM signal was caught and IPtraf-ng is attempting to shut down as gracefully as possible.

Invalid option or missing parameter, use `iptraf-ng -h` for help

The `-i`, `-d`, `-s`, `-l`, or `-z` options were specified but no interface was specified on the command line. These parameters require a valid interface name (or `all` for `-i` or `-l`). This message also appears if an unknown option is passed to the **iptraf-ng** command.

Warning: unable to tag this process

IPtraf-ng normally tags itself when it runs to prevent multiple instances of the statistical facilities from running. This message means the program was unable to create the necessary tag file. This may be due to a bad or improper installation. Try running the **make install** procedure or the **Setup** in the distribution's top-level directory.

Warning: unable to tag facility

IPtraf-ng was unable to create the tag file for the facility you started. The facility will still run, but other instances of IPtraf-ng that may be running simultaneously will allow the same facility to run. This may cause both instances of the facility to malfunction. This could be due to a bad disk or bad installation.

facility already running/listening on interface

The facility you tried to start is currently running on the indicated interface in another IPTraf-ng process on the machine. This restriction is placed to prevent conflicts involving internal sockets or the log files.

General interface statistics already active in another process

Only one instance of the general interface statistics can run at a time.

Duplicate port/range entry

You entered a port number or range that was already added to the list of additional ports to be monitored by the TCP/UDP service monitor

No custom ports

There are no ports or port ranges earlier added. There's nothing to delete.

Can't get communication sockets; lookups will block

IPTraf-ng cannot communicate with the resolving process. IPTraf-ng will fall back to blocking lookups.

Can't spawn new process; lookups will block

IPTraf-ng cannot start a new process. This may be due to memory shortage. IPTraf-ng will fall back to blocking lookups.

Fork error, IPTraf-ng cannot run in background

IPTraf-ng cannot start a new process, and can go into the background. This may be due to memory shortage. IPTraf-ng aborts.

No memory for new filter entry

IPTraf-ng was unable to allocate memory for a new filter entry. Most likely due to memory shortage.

Memory Low

This indicator appears if memory runs low due to a lot of entries in a facility. Should critical functions fail (window creation, internal allocation), the program could terminate with a segmentation violation.

Note: Any message or indicator about low memory means that your system does not have enough memory to handle the entries. It is almost certain that sooner or

later, IPTraf-ng or other applications will abort due to the failure of important system calls or library functions. Memory must be added right away.

IPC Error

This indicator appears if an error occurs receiving data from the resolving process (IPC stands for Interprocess Communication). This indication should not occur under normal circumstances. Report instances of this condition and the circumstances under which it happens. You may also include data from the `rvnamed-ng.log` file.

Error opening terminal: `terminal`

The screen management routines cannot find the `terminfo` entry for your terminal. IPTraf-ng expects the `terminfo` database located in `/usr/share/terminfo`. This error could occur when your `terminfo` database is located somewhere else. See the section on controlling the `terminfo` search path.

This will end your IPTraf-ng session

In interactive mode IPTraf-ng asks you to confirm your exit command. Press Enter to return to the shell or any other key to cancel your command and return to the main menu.

Resolving Process Messages

Resolving process does not send messages to the screen. It writes its messages to the file `rvnamed-ng.log` in the IPTraf-ng log directory.

Unable to open child communication socket

Resolving process was unable to open the communication endpoint for data reception from the children it creates. This is highly unusual and should it occur, report the circumstances.

Unable to open client communication socket

Resolving process was unable to open the communication endpoint for data exchange with the IPTraf-ng program. This is highly unusual and should it occur, report the circumstances.

Error binding client communication socket Error binding child communication socket

Appendix A. Messages

Resolving process was unable to assign a name to the indicated communication socket. This may be due to a bad, full or corrupted filesystem.

Fatal error: no memory for descriptor monitoring

Resolving process ran out of memory. IPTraf-ng will resort to blocking and may freeze.

Error on fork, returning IP address

Resolving process had a problem spawning a copy of itself to resolve the IP address; it will simply return the IP address in its literal, dotted-decimal notation. IPTraf-ng will still function normally. This may be due to lack of memory or a process limit hit.

Maximum child process limit reached

Resolving process has reached its maximum number of child processes. This is intended as a "brake" to prevent too many children from hogging your computer's resources and possibly crashing it. Unless IPTraf-ng is monitoring an extremely busy network without filters, this shouldn't happen, at least, not that often. If you notice this message, try applying filters or check your DNS server. Many times, this can happen when the DNS server goes down for whatever reason and you have resolving process children taking too long to resolve.

Appendix B. GNU Free Documentation License

Version 1.1, March 2000

Copyright (C) 2000 Free Software Foundation, Inc. 59 Temple Place, Suite 330, Boston, MA 02111-1307 USA Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed.

PREAMBLE

The purpose of this License is to make a manual, textbook, or other written document "free" in the sense of freedom: to assure everyone the effective freedom to copy and redistribute it, with or without modifying it, either commercially or noncommercially. Secondly, this License preserves for the author and publisher a way to get credit for their work, while not being considered responsible for modifications made by others.

This License is a kind of "copyleft", which means that derivative works of the document must themselves be free in the same sense. It complements the GNU General Public License, which is a copyleft license designed for free software.

We have designed this License in order to use it for manuals for free software, because free software needs free documentation: a free program should come with manuals providing the same freedoms that the software does. But this License is not limited to software manuals; it can be used for any textual work, regardless of subject matter or whether it is published as a printed book. We recommend this License principally for works whose purpose is instruction or reference.

APPLICABILITY AND DEFINITIONS

This License applies to any manual or other work that contains a notice placed by the copyright holder saying it can be distributed under the terms of this License. The "Document", below, refers to any such manual or work. Any member of the public is a licensee, and is addressed as "you".

A "Modified Version" of the Document means any work containing the Document or a portion of it, either copied verbatim, or with modifications and/or translated into another language.

A "Secondary Section" is a named appendix or a front-matter section of the Document that deals exclusively with the relationship of the publishers or authors of the Document to the Document's overall subject (or to related matters) and contains nothing that could fall directly within that overall subject. (For example, if the Document is in part a textbook of mathematics, a Secondary Section may not explain any mathematics.) The relationship could be a matter of historical connection with the subject or with related matters, or of legal, commercial, philosophical, ethical or political position regarding them.

The "Invariant Sections" are certain Secondary Sections whose titles are designated, as being those of Invariant Sections, in the notice that says that the Document is released under this License.

The "Cover Texts" are certain short passages of text that are listed, as Front-Cover Texts or Back-Cover Texts, in the notice that says that the Document is released under this License.

A "Transparent" copy of the Document means a machine-readable copy, represented in a format whose specification is available to the general public, whose contents can be viewed and edited directly and straightforwardly with generic text editors or (for images composed of pixels) generic paint programs or (for drawings) some widely available drawing editor, and that is suitable for input to text formatters or for automatic translation to a variety of formats suitable for input to text formatters. A

copy made in an otherwise Transparent file format whose markup has been designed to thwart or discourage subsequent modification by readers is not Transparent. A copy that is not "Transparent" is called "Opaque".

Examples of suitable formats for Transparent copies include plain ASCII without markup, Texinfo input format, LaTeX input format, SGML or XML using a publicly available DTD, and standard-conforming simple HTML designed for human modification. Opaque formats include PostScript, PDF, proprietary formats that can be read and edited only by proprietary word processors, SGML or XML for which the DTD and/or processing tools are not generally available, and the machine-generated HTML produced by some word processors for output purposes only.

The "Title Page" means, for a printed book, the title page itself, plus such following pages as are needed to hold, legibly, the material this License requires to appear in the title page. For works in formats which do not have any title page as such, "Title Page" means the text near the most prominent appearance of the work's title, preceding the beginning of the body of the text.

VERBATIM COPYING

You may copy and distribute the Document in any medium, either commercially or noncommercially, provided that this License, the copyright notices, and the license notice saying this License applies to the Document are reproduced in all copies, and that you add no other conditions whatsoever to those of this License. You may not use technical measures to obstruct or control the reading or further copying of the copies you make or distribute. However, you may accept compensation in exchange for copies. If you distribute a large enough number of copies you must also follow the conditions in section 3.

You may also lend copies, under the same conditions stated above, and you may publicly display copies.

COPYING IN QUANTITY

If you publish printed copies of the Document numbering more than 100, and the Document's license notice requires Cover Texts, you must enclose the copies in covers that carry, clearly and legibly, all these Cover Texts: Front-Cover Texts on the front cover, and Back-Cover Texts on the back cover. Both covers must also clearly and legibly identify you as the publisher of these copies. The front cover must present the full title with all words of the title equally prominent and visible. You may add other material on the covers in addition. Copying with changes limited to the covers, as long as they preserve the title of the Document and satisfy these conditions, can be treated as verbatim copying in other respects.

If the required texts for either cover are too voluminous to fit legibly, you should put the first ones listed (as many as fit reasonably) on the actual cover, and continue the rest onto adjacent pages.

If you publish or distribute Opaque copies of the Document numbering more than 100, you must either include a machine-readable Transparent copy along with each Opaque copy, or state in or with each Opaque copy a publicly-accessible computer-network location containing a complete Transparent copy of the Document, free of added material, which the general network-using public has access to download anonymously at no charge using public-standard network protocols. If you use the latter option, you must take reasonably prudent steps, when you begin distribution of Opaque copies in quantity, to ensure that this Transparent copy will remain thus accessible at the stated location until at least one year after the last time you distribute an Opaque copy (directly or through your agents or retailers) of that edition to the public.

It is requested, but not required, that you contact the authors of the Document well before redistributing any large number of copies, to give them a chance to provide you with an updated version of the Document.

MODIFICATIONS

You may copy and distribute a Modified Version of the Document under the conditions of sections 2 and 3 above, provided that you release the Modified Version under precisely this License, with the Modified Version filling the role of the Document, thus licensing distribution and modification of the Modified Version to whoever possesses a copy of it. In addition, you must do these things in the Modified Version:

- A. Use in the Title Page (and on the covers, if any) a title distinct from that of the Document, and from those of previous versions (which should, if there were any, be listed in the History section of the Document). You may use the same title as a previous version if the original publisher of that version gives permission.
- B. List on the Title Page, as authors, one or more persons or entities responsible for authorship of the modifications in the Modified Version, together with at least five of the principal authors of the Document (all of its principal authors, if it has less than five).
- C. State on the Title page the name of the publisher of the Modified Version, as the publisher.
- D. Preserve all the copyright notices of the Document.
- E. Add an appropriate copyright notice for your modifications adjacent to the other copyright notices.
- F. Include, immediately after the copyright notices, a license notice giving the public permission to use the Modified Version under the terms of this License, in the form shown in the Addendum below.
- G. Preserve in that license notice the full lists of Invariant Sections and required Cover Texts given in the Document's license notice.
- H. Include an unaltered copy of this License.
- I. Preserve the section entitled "History", and its title, and add to it an item stating at least the title, year, new authors, and publisher of the Modified Version as given on the Title Page. If there is no section entitled "History" in the Document, create one stating the title, year, authors, and publisher of the Document as given on its Title Page, then add an item describing the Modified Version as stated in the previous sentence.
- J. Preserve the network location, if any, given in the Document for public access to a Transparent copy of the Document, and likewise the network locations given in the Document for previous versions it was based on. These may be placed in the "History" section. You may omit a network location for a work that was published at least four years before the Document itself, or if the original publisher of the version it refers to gives permission.
- K. In any section entitled "Acknowledgements" or "Dedications", preserve the section's title, and preserve in the section all the substance and tone of each of the contributor acknowledgements and/or dedications given therein.
- L. Preserve all the Invariant Sections of the Document, unaltered in their text and in their titles. Section numbers or the equivalent are not considered part of the section titles.
- M. Delete any section entitled "Endorsements". Such a section may not be included in the Modified Version.

N. Do not retitle any existing section as "Endorsements" or to conflict in title with any Invariant Section.

If the Modified Version includes new front-matter sections or appendices that qualify as Secondary Sections and contain no material copied from the Document, you may at your option designate some or all of these sections as invariant. To do this, add their titles to the list of Invariant Sections in the Modified Version's license notice. These titles must be distinct from any other section titles.

You may add a section entitled "Endorsements", provided it contains nothing but endorsements of your Modified Version by various parties—for example, statements of peer review or that the text has been approved by an organization as the authoritative definition of a standard.

You may add a passage of up to five words as a Front-Cover Text, and a passage of up to 25 words as a Back-Cover Text, to the end of the list of Cover Texts in the Modified Version. Only one passage of Front-Cover Text and one of Back-Cover Text may be added by (or through arrangements made by) any one entity. If the Document already includes a cover text for the same cover, previously added by you or by arrangement made by the same entity you are acting on behalf of, you may not add another; but you may replace the old one, on explicit permission from the previous publisher that added the old one.

The author(s) and publisher(s) of the Document do not by this License give permission to use their names for publicity for or to assert or imply endorsement of any Modified Version.

COMBINING DOCUMENTS

You may combine the Document with other documents released under this License, under the terms defined in section 4 above for modified versions, provided that you include in the combination all of the Invariant Sections of all of the original documents, unmodified, and list them all as Invariant Sections of your combined work in its license notice.

The combined work need only contain one copy of this License, and multiple identical Invariant Sections may be replaced with a single copy. If there are multiple Invariant Sections with the same name but different contents, make the title of each such section unique by adding at the end of it, in parentheses, the name of the original author or publisher of that section if known, or else a unique number. Make the same adjustment to the section titles in the list of Invariant Sections in the license notice of the combined work.

In the combination, you must combine any sections entitled "History" in the various original documents, forming one section entitled "History"; likewise combine any sections entitled "Acknowledgements", and any sections entitled "Dedications". You must delete all sections entitled "Endorsements."

COLLECTIONS OF DOCUMENTS

You may make a collection consisting of the Document and other documents released under this License, and replace the individual copies of this License in the various documents with a single copy that is included in the collection, provided that you follow the rules of this License for verbatim copying of each of the documents in all other respects.

You may extract a single document from such a collection, and distribute it individually under this License, provided you insert a copy of this License into the extracted document, and follow this License in all other respects regarding verbatim copying of that document.

AGGREGATION WITH INDEPENDENT WORKS

A compilation of the Document or its derivatives with other separate and independent documents or works, in or on a volume of a storage or distribution medium, does not as a whole count as a Modified Version of the Document, provided no compilation copyright is claimed for the compilation. Such a compilation is called an "aggregate", and this License does not apply to the other self-contained works thus compiled with the Document, on account of their being thus compiled, if they are not themselves derivative works of the Document.

If the Cover Text requirement of section 3 is applicable to these copies of the Document, then if the Document is less than one quarter of the entire aggregate, the Document's Cover Texts may be placed on covers that surround only the Document within the aggregate. Otherwise they must appear on covers around the whole aggregate.

TRANSLATION

Translation is considered a kind of modification, so you may distribute translations of the Document under the terms of section 4. Replacing Invariant Sections with translations requires special permission from their copyright holders, but you may include translations of some or all Invariant Sections in addition to the original versions of these Invariant Sections. You may include a translation of this License provided that you also include the original English version of this License. In case of a disagreement between the translation and the original English version of this License, the original English version will prevail.

TERMINATION

You may not copy, modify, sublicense, or distribute the Document except as expressly provided for under this License. Any other attempt to copy, modify, sublicense or distribute the Document is void, and will automatically terminate your rights under this License. However, parties who have received copies, or rights, from you under this License will not have their licenses terminated so long as such parties remain in full compliance.

FUTURE REVISIONS OF THIS LICENSE

The Free Software Foundation may publish new, revised versions of the GNU Free Documentation License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns. See <http://www.gnu.org/copyleft/>.

Each version of the License is given a distinguishing version number. If the Document specifies that a particular numbered version of this License "or any later version" applies to it, you have the option of following the terms and conditions either of that specified version or of any later version that has been published (not as a draft) by the Free Software Foundation. If the Document does not specify a version number of this License, you may choose any version ever published (not as a draft) by the Free Software Foundation.

How to use this License for your documents

To use this License in a document you have written, include a copy of the License in the document and put the following copyright and license notices just after the title page:

Appendix B. GNU Free Documentation License

Copyright (c) YEAR YOUR NAME. Permission is granted to copy, distribute and/or modify this document under the terms of the GNU Free Documentation License, Version 1.1 or any later version published by the Free Software Foundation; with the Invariant Sections being LIST THEIR TITLES, with the Front-Cover Texts being LIST, and with the Back-Cover Texts being LIST. A copy of the license is included in the section entitled "GNU Free Documentation License".

If you have no Invariant Sections, write "with no Invariant Sections" instead of saying which ones are invariant. If you have no Front-Cover Texts, write "no Front-Cover Texts" instead of "Front-Cover Texts being LIST"; likewise for Back-Cover Texts.

If your document contains nontrivial examples of program code, we recommend releasing these examples in parallel under your choice of free software license, such as the GNU General Public License, to permit their use in free software.

Notes

1. <http://www.gnu.org/copyleft/>