

ntop User's Guide

Network Usage Monitor for Unix and Windows Systems

Version 2.3

© 1998-2003

Luca Deri <deri@ntop.org>

1. Introduction

Everyday I have to monitor the traffic flowing across the network backbone. In order to control the network activities I run several network monitor tools. Due to this monitoring activity, many people accused me be the cause of the frequent network slowdowns. Since I had no way to proof that this wasn't the case, I decided to create an application portable on (virtually) every Unix and Windows system, that allowed me to control the network activity hence to find out who was the real net assassin. This is why ntop had birth.

ntop is an application for Unix and Windows systems that allows people to monitor the network activity. Similar to the popular top program, it shows the network activity. This is implemented capturing and analyzing the network traffic that flows on the specified network interface. ntop relies on libpcap for packet capture, a public-domain portable capture library.

The following sections describe how to compile and take advantage ntop.

Happy reading.

Luca Deri, November 2003.

2. Compiling and Installing ntop

ntop is can be downloaded from <http://www.ntop.org/> and several other mirrors (e.g. <http://www.sourceforge.net/projects/ntop/>) in both source and binary (either application binary or binary package) format. However in order to:

- take advantage of the latest ntop features.
- report problems we can fix

It is strongly recommended that you fetch the ntop code using CVS as described here <http://www.ntop.org/download.html>.

In order to compile ntop you need to install some prerequisite libraries (see appendix).

Supposing that you have already installed the prerequisite packages and downloaded the ntop source code, in order to compile ntop do:

```
# cd <your path to ntop>/ntop
# ./configure
# make
# su          <you need root password>
# make install
# exit
```

ntop requires superuser (root) capability. In order to allow non-root users to use ntop please do:

```
# su <you need root password>
# cd <directory where you have installed ntop (e.g.
  /usr/sbin/)>
# chown root.root ntop
# chmod 6111 ntop
# exit
```

Under windows systems, in order to compile ntop you must first get a compiler (e.g. MS Visual C++ or .NET) then use the project you can find in the ntop/packages/Win32 directory part of the ntop source distribution. It is possible to get a binary ntop package from <http://shop.ntop.org/> at little cost: the money we gather is reinvested in the project for paying expenses and purchasing network hardware and software.

At this point ntop should be installed properly and ready to use. If you have experienced problems while compiling ntop, please report the problem.

3. Starting ntop

ntop shows the current network usage. It displays a list of hosts that are currently using the network and reports information concerning the (IP and non-IP) traffic generated by each host and much more. ntop can be started either in a terminal window or as a service (Window NT/2K/XP only). ntop may operate as a front-end collector (sFlow and/or NetFlow plugins) or as a stand-alone collector/display program. ntop is a hybrid layer 2 / layer 3 network monitor, that is by default it uses the layer 2 Media Access Control (MAC) addresses AND the layer 3 tcp/ip addresses. ntop is capable of associating the two, so that ip and non-ip traffic (e.g. arp, rarp) are combined for a complete picture of network activity.

A browser is used for connecting to ntop and browsing the traffic reports. The traffic is sorted to various criteria including to the host and network protocol. As ntop is accessed using a web browser, multiple remote users can access it simultaneously.

3.1 ntop Command Line Options

In order to start ntop, open a terminal window and type `ntop -h` in order to see an online help. The available options are:

```
ntop [@filename] [-a|--access-log-path <path>] [-b|--disable-decoders]
[-c|--sticky-hosts] [-f|--traffic-dump-file file>] [-g|--track-local-
hosts] [-h|--help] [-k|--filter-expression-in-extra-frame] [-l|--pcap-
log <path>] [-m|--local-subnets <addresses>] [-n|--numeric-ip-
addresses] [-o|--no-mac] [-p|--protocols <list>] [-q|--create-suspi-
cious-packets] [-r|--refresh-time <number>] [-s|--no-promiscuous]
[-t|--trace-level <number>] [-x <max_num_hash_entries>] [-w|--http-
server <port>] [-z|--disable-sessions] [-A|--set-admin-password pass-
word] [-B|--filter-expression expression] [-D|--domain <name>]
[-F|--flow-spec <specs>] [-M|--no-interface-merge] [-O|----output-
packet-path] [-P|--db-file-path <path>] [-Q|--spool-file-path <path>]
[-R|--filter-rule <file>] <number>] [-U|--mapper <URL>] [-V|--version]
[-X <max_num_TCP_sessions>] [--disable-stopcap] [--log-extra <number>]
[--disable-instantsessionpurge] [--disable-schedyield] [--disable-
mutextrainfo]
```

Unix options:

```
[-d|--daemon] [-i|--interface <name>] [-u|--user <user>] [-K|--enable-
debug] [-L] [-use-syslog= <facility>]
```

Windows option:

```
[-i|--interface <number|name>]
```

SSL options:

[-W|--https-server <port>]

These are the command line options (specified on the command line) accepted by ntop:

@filename

The text of filename is copied ignoring line breaks and comment lines (anything following a #) into the command line. ntop behaves as if all of the text had simply been typed directly on the command line. For example, if the command line is "-t 3 @d -u ntop" and file d contains just the line '-d', then the effective command line is "-t 3 -d -u ntop. Multiple @s are permitted. Nested @s (an @ inside the file) are not permitted.

Remember, most ntop options are "sticky", that is they just set an internal flag. Invoking them multiple times doesn't change ntop's behavior. However, options that set a value, such as --trace-level, will use the LAST value given: --trace-level 2 --trace-level 3 will run as --trace-level 3.

-a | --access-log-path

By default ntop does not maintain a log of HTTP requests to the internal web server. Use this parameter to request logging and to specify the location of the file where these HTTP requests are logged.

Each log entry is in Apache-like style. The only difference between Apache and ntop logs is that an additional column has been added which has the time (in milliseconds) that ntop needed to serve the request. Log entries look like this:

```
192.168.1.1 [04/Sep/2003:20:38:55 -0500] "GET / HTTP/1.1" 200 1489 4
192.168.1.1 [04/Sep/2003:20:38:55 -0500] "GET /index_top.html HTTP/1.1" 200 1854 4
192.168.1.1 [04/Sep/2003:20:38:55 -0500] "GET /index_inner.html HTTP/1.1" 200 1441 7
192.168.1.1 [04/Sep/2003:20:38:56 -0500] "GET /index_left.html HTTP/1.1" 200 1356 4
192.168.1.1 [04/Sep/2003:20:38:56 -0500] "GET /home_.html HTTP/1.1" 200 154/617 9
192.168.1.1 [04/Sep/2003:20:38:56 -0500] "GET /home.html HTTP/1.1" 200 1100/3195 10
192.168.1.1 [04/Sep/2003:20:38:56 -0500] "GET /About.html HTTP/1.1" 200 2010 10
```

Although this parameter is called a 'path', it is actually the complete file name of the access log.

-b | --disable-decoders

This parameter disables protocol decoders.

Protocol decoders examine and collect information about layer 2 protocols such as NetBIOS or Network SAP, as well as about specific tcp/ip (layer 3) protocols, such as DNS, http and ftp.

This support is specifically coded for each protocol and is different from the capability to count raw information (packets and bytes) by protocol specified by the -p | --protocols parameter, below.

Decoding protocols is a significant consumer of resources. If the ntop host is underpowered or monitoring a very busy network, you may wish to disable protocol decoding via this parameter. It may also be appropriate to use this parameter if you believe that ntop has problems handling some protocols that occur on your network.

Even if decoding is disabled, ftp-data traffic is still decoded to look for passive ftp port commands.

-c | --sticky-hosts

Use this parameter to prevent idle hosts from being purged from memory.

By default idle hosts are periodically purged from memory. An idle host is identified when no packets from or to that host have been monitored for the period of time defined by the value of PARM_HOST_PURGE_MINIMUM_IDLE in globals-defines.h.

If you use this option, all hosts active and idle are retained in memory for the duration of the ntop run.

P2P users, port scans, popular web servers and other activity will cause ntop to record data about a large number of hosts. On an active network, this will consume a significant and always growing amount of memory. It is strongly recommended that you use a filtering expression to limit the hosts which are stored if you use --sticky hosts.

The idle purge is a statistical one a random selection of the eligible hosts will be purged during each cycle. Thus it is possible on a busy system for an idle host to remain in the ntop tables and appear 'active' for some considerable time after it is truly idle.

-d | --daemon

This parameter causes ntop to become a daemon, i.e. a task which runs in the background without connection to a specific terminal. To use ntop other than as a casual monitoring tool, you probably will want to use this option.

WARNING: If you are running as a daemon, the messages from ntop will be 'printed' on to stdout and thus dropped. You probably don't want to do this. So remember to also use the -L or --use-syslog options to save the messages into the system log.

-e | --max-table-rows

This defines the maximum number of lines that ntop will display on each generated HTML page. If there are more lines to be displayed than this setting permits, only part of the data will be displayed. There will be page forward/back arrows placed at the bottom of the page for navigation between pages.

-f | --traffic-dump-file

By default, ntop captures traffic from network interface cards (NICs) or from NetFlow/sFlow probes. However, ntop can also read data from a file typically a tcpdump capture or the output from one of the ntop packet captures options.

if you specify -f, ntop will not capture any traffic from NICs during or after the file has been read. NetFlow/sFlow capture if enabled would still be active.

This option is mostly used for debug purposes.

-g | --track-local-hosts

By default, ntop tracks all hosts that it sees from packets captured on the various NICs. Use this parameter to tell ntop to capture data only about local hosts. Local hosts are defined based on the addresses of the NICs and those networks identified as local via the -m | --local-subnets parameter.

This parameter is useful on large networks or those that see many hosts, (e.g. a border router or gateway), where information about remote hosts is not desired/required to be tracked.

-h | --help

Print help information for ntop, including usage and parameters.

-i | --interface

Specifies the network interface or interfaces to be used by ntop for network monitoring.

If multiple interfaces are used (this feature is available only if ntop is compiled with thread support) their names must be separated with a comma. For instance `-i "eth0,lo"`.

If not specified, the default is the first Ethernet device, e.g. `eth0`. The specific device that is 'first' is highly system dependent. Especially on systems where the device name reflects the driver name instead of the type of interface.

By default, traffic information obtained by all the interfaces is merged together as if the traffic was seen by only one interface. Use the `-M` parameter to keep traffic separate by interface.

Under Windows, the parameter value is either the number of the interface or its name, e.g. `{6252C14C-44C9-49D9-BF59-B2DC18C7B811}`. Run `ntop -h` to see a list of interface name-number mappings (at the end of the help information).

-k | --filter-expression-in-extra-frame

When this parameter is used, the current filter expression is displayed in an extra frame and thus is always visible. This extra frame contains other information, including the report creation date, ntop version information and the active interfaces.

-l | --pcap-log

This parameter causes a dump file to be created of the network traffic captured by ntop in tcpdump (pcap) format. This file is useful for debug, and may be read back into ntop by the `-f | --traffic-dump-file` parameter. The dump is made after processing any filter expression (never even sees filtered packets).

The output file will be named `<path>/<log>.<device>.pcap` (Windows: `<path>/<log>.pcap`), where `<path>` is defined by the `-O | --outputpacket-path` parameter and `<log>` is defined by this `-l | --pcap-log` parameter.

-m | --local-subnets

ntop determines the ip addresses and netmasks for each active interface. Any traffic on those networks is considered local. This parameter allows the user to define additional networks and subnetworks whose traffic is also considered local in ntop reports. All other hosts are considered remote.

Commas separate multiple network values. Both netmask and CIDR notation may be used even mixed together, for instance `"131.114.21.0/24,10.0.0.0/255.0.0.0"`.

The local subnet as defined by the interface address(es) is/are always local and do not need to be specified. If you do give the same value as a NIC's local address, a harmless warning message is issued.

-n | --numeric-ip-addresses

By default, ntop resolves IP addresses using a combination of active (explicit) DNS queries and passive sniffing. Sniffing of DNS responses occurs when ntop receives a network packet containing the response to some other user's DNS query. ntop captures this information and enters it into ntop's DNS cache, in expectation of shortly seeing traffic addressed to that host. This way ntop significantly reduces the number of DNS queries it makes.

This parameter causes ntop to skip DNS resolution, showing only numeric IP addresses instead of the symbolic names. This option can be useful when the DNS is not present or quite slow.

-o | --no-mac

ntop is a hybrid layer 2/3 network monitor. That is, it uses both the lower level, physical device address the MAC (Media Access Control) address and the higher level, logical, tcp/ip address (the familiar `www.ntop.org` or `131.114.21.9` address). This allows ntop to link the logical addresses to a physical machine with multiple addresses (This occurs with virtual hosts or additional addresses assigned to the interface, etc.) to present consolidated reporting.

This parameter specifies that ntop should not trust the MAC addresses but just use the IP addresses.

Normally, since the MAC address must be globally unique, the dual nature of ntop is a benefit and provides far better information about the network than is available via a pure layer 2 or pure layer 3 monitor.

Under certain circumstances whenever ntop is started on an interface where MAC addresses cannot be really trusted you may require this option.

Situations which may require this option include port/VLAN mirror, some cases with switches and spanning tree protocol, and (reportedly) some specific models of Ethernet switches which re-write MAC addresses of the packets they process. Normally, you discover that this option is necessary when you observe that hosts seem to change their addresses or information about different machines get lumped together.

Note that with this option, information which is dependent upon the MAC addresses (non tcp/ip protocols like IPX) will not be collected nor displayed.

-p | --protocols

This parameter is used to specify the TCP/UDP protocols that ntop will monitor. The format is <label>=<protocol list> [, <label>=<protocol list>], where label is used to symbolically identify the <protocol list>. The format of <protocol list> is <protocol>[<protocol>], where <protocol> is either a valid protocol specified inside the /etc/services file or a numeric port range (e.g. 80, or 6000-6500).

A simple example is --protocols="HTTP=http|www|https|3128,FTP=ftplftpdata", which reduces the protocols displayed on the "IP" pages to three:

```
Host Domain Data HTTP FTP Other IP
ns2.attbi.com <flag> 954 63.9 % 0 0 954
64.124.83.112.akamai.com <flag> 240 16.1 % 240 0 0
64.124.83.99.akamai.com <flag> 240 16.1 % 240 0 0
toolbarqueries.google.com <flag> 60 4.0 % 60 0 0
```

If the <protocol list> is very long you may store it in a file (for instance protocol.list). To do so, specify the file name instead of the <protocol list> on the command line. e.g. ntop -p protocol.list

If the -p parameter is omitted the following default value is used:

FTP=ftplftp-data

HTTP=http|www|https|3128 3128 is Squid, the HTTP cache

DNS=namelddomain

Telnet=telnet|login

NBios-IP=netbios-ns|netbios-dgmlnetbios-ssn

Mail=pop-2|pop-3|pop3|kpop|smtp|imap|imap2

DHCP-BOOTP=67-68

SNMP=snmp|snmp-trap

NNTP=nnntp

NFS=mount|pcnfs|bwnfs|nfsd|nfsd-status

X11=6000-6010

SSH=22

Peer-to-Peer Protocols

Gnutella=63461634716348

Kazaa=1214

WinMX=669917730

DirectConnect=0 Dummy port as this is a pure P2P protocol

eDonkey=4661-4665

Instant Messenger

Messenger=1863150001500115190-5193

NOTE: to resolve protocol names to port numbers, they must be specified in the system file used to list tcp/udp protocols and ports, which is typically /etc/services file. You will have to match the names in that file, exactly. Missing or unspecified (non-standard) ports must be specified by number, such as 3128 in our examples above.

If you have a file named /etc/protocols, don't get confused by it, as that's the Ethernet protocol numbers, which are not what you're looking for.

-q | --create-suspicious-packets

This parameter tells ntop to create a dump file of suspicious packets. There are many, many, things that cause a packet to be labeled as 'suspicious', including:

Detected ICMP fragment

Detected Land Attack against host

Detected overlapping/tiny packet fragment

Detected traffic on a diagnostic port

Host performed ACK/FIN/NULL scan

Host rejected TCP session

HTTP/FTP/SMTP/SSH detected at wrong port

Malformed TCP/UDP/ICMP packet (packet too short)

Packet # %u too long

Received a ICMP protocol Unreachable from host

Sent ICMP Administratively Prohibited packet to host

Smurf packet detected for host

TCP connection with no data exchanged

TCP session reset without completing 3-way handshake

Two MAC addresses found for the same IP address

UDP data to a closed port

Unknown protocol (no HTTP/FTP/SMTP/SSH) detected (on port 80/21/25/22)

Unusual ICMP options

When this parameter is used, one file is created for each network interface where suspicious packets are found. The file is in tcpdump (pcap) format and is named `<path>/ntop-suspicious-pkts.<device>.pcap`, where `<path>` is defined by the `-O | --output-packet-path` parameter.

-r | --refresh-time

Specifies the delay (in seconds) between automatic screen updates for those generated HTML pages, which support them. This parameter allows you to leave your browser window open and have it always displaying nearly real-time data from ntop.

The default is 3 seconds. Please note that if the delay is very short (1 second for instance), ntop might not be able to process all of the network traffic.

-s | --no-promiscuous

Use this parameter to prevent from setting the interface(s) into promiscuous mode.

An interface in promiscuous mode will accept all Ethernet frames, regardless of whether they directed (addressed) to the specific network interface (NIC) or not. This is an essential part of enabling ntop to monitor an entire network. Without promiscuous mode, ntop will only see traffic directed to the specific host it is running on, plus broadcast traffic such as the arp and dhcp protocols.

Even if you use this parameter, the interface could well be in promiscuous mode if another application enabled it.

ntop passes this setting on to libpcap, the packet capture library. On many systems, a non-promiscuous open of the network interface will fail, since the libpcap function on most systems require it to capture raw packets (ntop captures raw packets so that we may view and analyze the layer 2 MAC information).

Thus on most systems, ntop must probably still be started as root, and this option is largely ornamental. If it fails, you will see a *****FATALERROR***** message referring to `pcap_open_live()` and then an information message, "Sorry, but on this system, even with -s, it appears that ntop must be started as root".

-t | --trace-level

This parameter specifies the 'information' level of messages that you wish ntop to display (on stdout or to the log). The higher the trace levels number the more information that is displayed. The trace level ranges between 0 (no trace) and 5 (full debug tracings).

The default trace value is 3.

Trace level 0 is not quite zero messages. Fatal errors and certain startup/shutdown messages are always displayed. Trace level 1 is used to display errors only, level 2 for both errors and warnings, and level 3 displays error, warning and informational messages.

Trace level 4 is called 'noisy' and it is generating many messages about the internal functioning of ntop. Trace level 5 is 'noisy' plus `--log-extra 1`, i.e. all possible messages, with a file:line tag prepended to every message.

-u | --user

Specifies the user ntop should run as after it initializes.

ntop must normally be started as root so that it has sufficient privileges to open the network interfaces in promiscuous mode and to receive raw frames. See the discussion of `-s l --no-promiscuous` above, if you wish to try starting ntop as a non-root user.

Shortly after starting up, ntop becomes the user you specify here, which normally has substantially reduced privileges, such as no login shell. This is the userid which owns ntop's database and output files.

The value specified may be either a username or a numeric user id. The group id used will be the primary group of the user specified.

If this parameter is not specified, ntop will try to switch first to 'nobody' and then to 'anonymous' before giving up.

NOTE: This should not be root unless you really understand the security risks. In order to prevent this by accident, the only way to run ntop as root is to explicitly specify `-u root`. Don't do it.

-x

-X

ntop creates a new hash/list entry for each new host/TCP session seen. In case of DOS (Denial Of Service) an attacker can easily exhaust all the host available memory because ntop is creating entries for dummy hosts. In order to avoid this you can set an upper limit in order to limit the memory ntop can use.

-w l --http-server

-W l --https-server

ntop offers an embedded web server to present the information that has been so painstakingly gathered. An external HTTP server is NOT required NOR supported. The ntop web server is embedded into the application. These parameters specify the port (and optionally the address (i.e. interface)) of the ntop web server.

For example, if started with `-w 3000` (the default port), the URL to access ntop is `http://hostname:3000/`. If started with a full specification, e.g. `-w 192.168.1.1:3000`, ntop listens on only that address/port combination.

If `-w` is set to 0 the web server will not listen for `http://` connections.

`-W` operates similarly, but controls the port for the `https://` connections.

Some examples:

`ntop -w 3000 -W 0` (this is the default setting) HTTP requests on port 3000 and no HTTPS.

`ntop -w 80 -W 443` Both HTTP and HTTPS have been enabled on their most common ports.

`ntop -w 0 -W 443` HTTP disabled, HTTPS enabled on the common port.

Certain sensitive, configuration pages of the ntop web server are protected by a userid/password. By default, these are the user/URL administration, filter, shutdown and reset stats are password protected and are accessible initially only to user admin with a password set during the first run of ntop.

Users can modify/add/delete users/URLs using ntop itself see the Admin tab.

The passwords, userids and URLs to protect with passwords are stored in a database file. Passwords are stored in an encrypted form in the database for further security. Best practices call for securing that database so that only the ntop user can read it.

There is a discussion in docs/FAQ about further securing the ntop environment.

-z | --disable-sessions

This parameter disables TCP session tracking. Use it for better performance or when you don't really need/care to track sessions.

-A | --set-admin-password

This parameter is used to start ntop , set the admin password and quit. It is quite useful for installers that need to automatically set the password for the admin user.

-A and --set-admin-password (without a value) will prompt the user for the password.

You may also use this parameter to set a specific value using --setadmin-password=value. The = is REQUIRED and no spaces are permitted!

If you attempt to run ntop as a daemon without setting a password, a FATAL ERROR message is generated and ntop stops.

-B | --filter-expression

Filters allows the user to restrict the traffic seen by ntop on just about any imaginable item.

The filter expression is set at run time by this parameter, but it may be changed during the ntop run on the Admin | Change Filter web page.

The basic format is -B filter , where the quotes are REQUIRED

The syntax of the filter expression uses the same BPF (Berkeley Packet Filter) expressions used by other packages such as tcpdump

For instance, suppose you are interested only in the traffic generated/received by the host jake.unipi.it. ntop can then be started with the following filter:

```
ntop -B src host jake.unipi.it or dst host jake.unipi.it
```

or in shorthand:

```
ntop -B host jake.unipi.it or host jake.unipi.it
```

See the 'expression' section of the tcpdump man page usually available at http://www.tcpdump.org/tcpdump_man.html for further information and the best quick guide to BPF filters currently available.

-D | --domain

This identifies the local domain suffix, e.g. ntop.org. It may be necessary, if ntop is having difficulty determining it from the interface.

-F | --flow-spec

t is used to specify network flows similar to more powerful applications such as NeTraMet. A flow is a stream of captured packets that match a specified rule. The format is

```
<flow-label>='<matching expression>',[<flow-label>='<matching expression>']
```

, where the label is used to symbolically identify the flow specified by the expression. The expression format is specified in the appendix. If an expression is specified, then the information concerning flows can be accessed following the HTML link named 'List NetFlows'.

For instance define two flows with the following expression LucaHosts='host jake.unipi.it or host pisanino.unipi.it',GatewayRoutedPkts='gateway gateway.unipi.it' .

All the traffic sent/received by hosts `jake.unipi.it` or `pisanino.unipi.it` is collected by ntop and added to the `LucaHosts` flow, whereas all the packet routed by the gateway `gateway.unipi.it` are added to the `GatewayRoutedPkts` flow. If the flows list is very long you may store in a file (for instance `flows.list`) and specify the file name instead of the actual flows list (in the above example, this would be `'ntop -F flows.list'`).

Note that the double quotations around the entire flow expression are required.

-K | --enable-debug

Use this parameter to simplify application debug. It does three things: 1. Does not fork() on the "read only" html pages. 2. Displays mutex values on the configuration (`info.html`) page. 3. (If available `glibc/gcc`) Activates an automated backtrace on application errors.

-L | --use-syslog=facility

Use this parameter to send log messages to the system log instead of stdout.

`-L` and the simple form `--use-syslog` use the default log facility, defined as `LOG_DAEMON` in the `#define` symbol `DEFAULT_SYSLOG_FACILITY` in `globals-defines.h`.

The complex form, `--use-syslog=facility` will set the log facility to whatever value (e.g. `local3`, `security`) you specify. The `=` is REQUIRED and no spaces are allowed!

This setting applies both to ntop and to any child fork()ed for reporting. If this parameter is not specified, any fork()ed child will use the default value and will log it's messages to the system log (this occurs because the fork()ed child must give up it's access to the parents stdout).

Because various systems do not make the permissible names available, we have a table at the end of `globals-core.c`. Look for `myFacilityNames`.

-M | --no-interface-merge

By default, ntop merges the data collected from all of the interfaces (NICs) it is monitoring into a single set of counters.

If you have a simple network, say a small LAN with a connection to the internet, merging data is good as it gives you a better picture of the whole network. For larger, more complex networks, this may not be desirable. You may also have other reasons for wishing to monitor each interface separately, for example DMZ vs. LAN traffic.

This option instructs ntop not to merge network interfaces together. This means that ntop will collect statistics for each interface and report them separately.

Only ONE interface may be reported on at a time use the `Admin | Switch NIC` option on the web server to select which interface to report upon.

Note that activating either the NetFlow and/or sFlow plugins will force the setting of `-M`. Once enabled, you cannot go back.

-O | --output-packet-path

This parameter defines the base path for the `ntop-suspiciouspkts.XXX.pcap` and normal packet dump files.

If this parameter is not specified, the default value is the `config.h` parameter `CFG_DBFILE_DIR`, which is set during `./configure` from the `--localstatedir=` parameter. If `--localstatedir` is not specified, it defaults to the `--prefix` value plus `/var` (e.g. `/usr/local/var`).

Be aware that this may not be what you expect when running ntop as a daemon or Windows service. Setting an explicit and absolute path value is STRONGLY recommended if you use this facility.

-P | --db-file-path

-Q | --pool-file-path

These parameters specify where ntop stores database files.

There are two types, 'temporary' that is ones which need not be retained from ntop run to ntop run, and 'permanent', which must be retained (or recreated).

The 'permanent' databases are the preferences, "prefsCache.db" and the password file, "ntop_pw.db". These are stored in the -P | --db-filepath specified location.

Certain plugins use the -P | --db-file-path specified location for their database ("LsWatch.db") or (as a default value) for files (.../rrd/...).

The 'temporary' databases are the address queue, "addressQueue.db", the cached DNS resolutions, "dnsCache.db" and the MAC prefix (vendor table), "macPrefix.db".

If only -P | --db-file-path is specified, it is used for both types of databases.

The directories named must allow read/write and file creation by the ntop user. For security, nobody else should have even read access to these files.

Note that the default value is the config.h parameter CFG_DBFILE_DIR. This is set during ./configure from the --localstatedir= parameter. If --localstatedir is not specified, it defaults to the --prefix value plus /var (e.g. /usr/local/var).

This may not be what you expect when running ntop as a daemon or Windows service.

Note that on versions of ntop prior to 2.3, these parameters defaulted to "." (the current working directory, e.g. the value returned by the pwd command) and caused havoc as it was different when ntop was run from the command line, vs. run via cron, vs. run from an initialization script.

Setting an explicit and absolute path value is STRONGLY recommended.

-U | --mapper

Specifies the URL of the mapper.pl utility.

If provided, ntop creates a clickable hyperlink on the 'Info about host xxxxxx' page to this URL by appending ?host=xxxxx. Any type of host lookup could be performed, but this is intended to lookup the geographical location of the host.

A cgi-based mapper interface to <http://www.multimap.com> is part of the ntop distribution [see [www/Perl/mapper.pl](#)].

-V | --version

Prints ntop version information and then exits.

-W | --https-server

(See the joint documentation with the -w parameter, above)

--disable-stopcap

Return ntop to the old (v2.1) behavior on a memory error. The default of stopcap enabled makes the web interface available albeit with static content until ntop is shutdown.

--log-extra

This optional parameter controls the addition of more information to each log message. Both choices are useful in different ways, for debugging and for using log watching and filtering packages.

Setting 1 adds a [file:line] to the beginning of the log message. Setting 2 adds a [MSGIDnnnnnnn] tag at the end of the log message. The nnnnnnn value should be unique number for every message and should be stable across ntop releases.

--disable-instantsessionpurge

ntop sets completed sessions as 'timed out' and then purge them almost instantly, which is not the behavior you might expect from the discussions about purge timeouts. This switch makes ntop respect the timeouts for completed sessions. It is NOT the default because a busy web server may have 100s or 1000s of completed sessions and this would significantly increase the amount of memory ntop uses.

--disable-schedyield

ntop uses sched_yield() calls for better interactive performance. Under some situations, primarily under RedHat Linux 8.0, this can deadlock, causing the ntop web server to stop responding, although ntop appears to still be operational according to the ps command. Use this switch to disable these calls, IF you are seeing deadlocks.

--disable-mutexextrainfo

ntop stores extra information about the locks and unlocks of the protective mutexes it uses. Since ntop uses fine-grained locking, this information is updated frequently. On some OSes, the system calls used to collect this informatio (getpid() and gettimeofday()) are expensive. This option disables the extra information. It should have no processing impact on ntop however should ntop actually deadlock, we would lose the information that sometimes tells us why.

3.2 Web Views

While ntop is running, multiple users can access the traffic information using their web browsers. ntop does not generate 'fancy' or 'complex' html, although it does use frame, shallowly nested tables and makes minimal use of Cascading Style Sheets.

We do not expect problems with any current web browser, but our ability to test less common ones is very limited.

The main HTML page is divided into three frames. Beginning with release 2.3, the menus have been compacted to small text selections stacked on top of each other.

The top frame is a 'tabbed' navigation bar, containing broad items such as 'Total', 'Sent' and 'IP Protos'. The middle frame is the detailed navigation or menu bar, containing the items relevant to the top selection, for example "IP" traffic statistics from a "Totals" menu. The resulting data is displayed in the bottom frame. In documentation and this man page, when we refer to a page such as Admin | Switch NIC, we mean the Broad category "Admin" and the detailed item "Switch NIC" on that Admin menu.

3.3 Starting ntop on Windows

As explained before, the differences between ntop on Unix and Windows are very little. The main difference is the way ntop is started. In fact under Win NT/2K/XP there's the concept of service that's similar to the concept of daemon on Unix. A service is started/stopped from the Services control panel. In order to tell Windows about the presence of a new service (ntop in our case), it is necessary to add ntop to the Windows registry, that's basically a database used by Windows for keeping track of configurations. The Windows ntop package automatically registers ntop at installation, and removes it when the package is removed.

If ntop is started on the console (Win 95 and above), the /c flag needs to be used (e.g. ntop /c -P . -u ntop). If used as service (Win NT and above), the command line options need to be specified at service registration and can be modified only removing and adding the service. The package uses the default options that should be fine for most users:

The protected area of web interface can be accessed using admin as both user and password.

The default interface from which packets are captured is the first one in the interface list.

The default port where ntop listens for http connections is 3000 (hence you should point your browser to `http://<your host>:3000/`).

If you need to change the ntop setup, you need to do as follows:

```
> ntop /r Remove the service
```

```
> ntop /i <your options> Install the service with the specified options.
```

Services are started and stopped using the Services application part of the Windows administrative tools.

As network interfaces on Windows can have long names, a numeric index is associated to the interface in order to ease the ntop configuration. The association interface name and index is shows typing the `'ntop /c -h'`

For instance:

```
> ntop /c -h
```

```
lt-ntop v.2.2.96 MT (SSL) [Win32] (10/29/03 01:10:39 PM build)
```

```
Copyright 1998-2003 by Luca Deri <deri@ntop.org>.
```

```
...
```

```
Available interfaces:
```

```
[index=0] Realtek 8139-series PCI NIC
```

```
[index=1] B2C2 Broadband Receiver PCI Adapter
```

```
[index=2] NdisWan Adapter
```

```
[index=3] VMware Virtual Ethernet Adapter
```

```
[index=4] VMware Virtual Ethernet Adapter
```

```
[index=5] NET IP/1394 Miniport
```

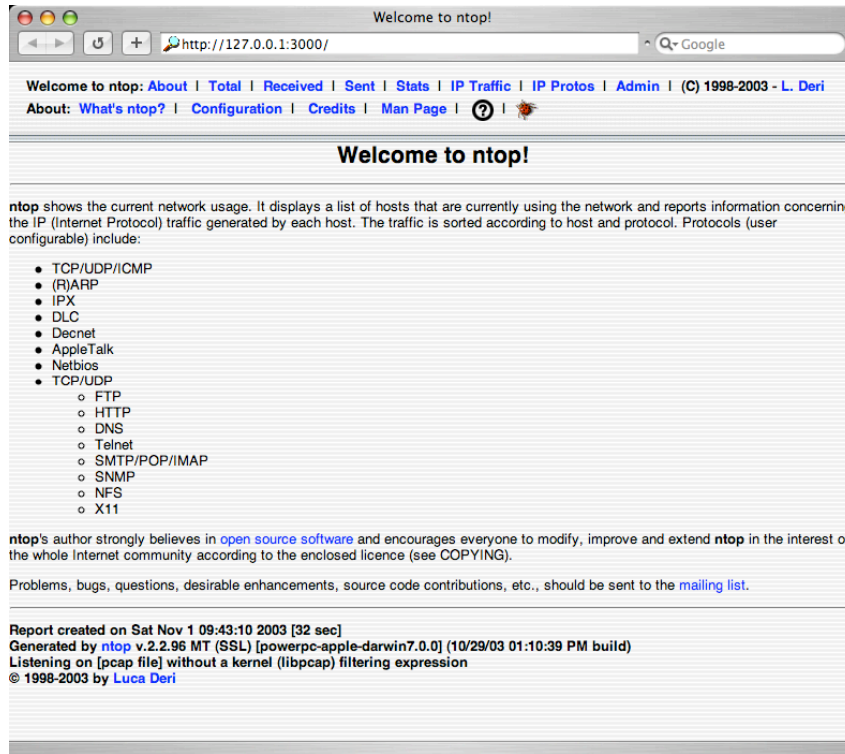
```
...
```

For instance, in the above example the index 0 is associated to the interface Realtek 8139-series PCI NIC, hence in order to select this interface ntop needs to be started with `-i 0` option. Therefore you should type the following commands from the command line for storing this configuration:

- Stop the ntop service from the Services control panel (if started)
- `ntop /r`
- `ntop /I -I 0`
- Start the ntop service

4. Using ntop

As stated before, you need a web browser for connecting to a running ntop. Supposing that ntop runs on host jake you should type `http://jake:3000/` in order to connect to it.



As shown above, the web view is divided in two frames: the top frame contains the menus divided in two rows, the bottom frame contains the output of the selected menu. When the user selects an entry from the top menu row, the bottom menu row is updated with submenus entries. Every hyperlink can be selected in order to let ntop show you detailed view of the object.

Before glancing through the menus, it is worth to explain some conventions used by ntop:

- **Host Lifetime**
ntop stores information in a hash. As this information increases with the number of host, it is necessarily to periodically purge the hash in order to remove information about hosts that do not perform any network activity from some time. Due to this, it is normal to see hosts entries disappear (when no traffic is detected for some time) and reappear (when the host makes some network activity). Ntop automatically purges entries from hash so you don't need to specify special command line options to tune it.
- **Local vs. Remote Host**
ntop considers a host local if it belongs to the local network of the adapter from which the packets are captured. In case you want to add further networks to the list of local networks, use the `-m` flag. Make sure you know what you're doing when doing this because ntop reports you wrong figures if the list of networks are not properly configured.
- **MAC vs. IP Address**
Due to the way the IP/networks work, the MAC address is used at L2 and the IP address at L3. For this reason ntop uses the MAC address for identifying local hosts and the IP address for remote ones. In case ntop captures traffic that's not really flowing in the network trunk being analyzed, make sure you use the `-o` option to tell ntop not to trust MAC addresses, as they are not reliable at all.

- **Host Icons**
Whenever ntop detects some special host activity (passive host fingerprinting) worth to highlight to the user, it places some special icons next to the host name that can indicate a service being provided (e.g. DNS or DHCP server), a role (e.g. router or L2 bridge) or a security level (e.g. a host has a duplicated address). Note that if you see too many red/yellow flags on your network you're probably capturing traffic from a network trunk that's not "real" so you better add `-o` to the list of options you pass to ntop at startup.

4.1 Ntop Menus

The ntop menu structure is the following:

About

This menu reports general information about the ntop community and the currently active ntop instance.

What's ntop

It displays a quick summary of the ntop features

Configuration

Summary of the ntop configuration. This is very useful for reporting problems or keeping track of the resources (memory, disk) being used by ntop.

Credits

List of people who contribute to ntop with code, testing, ideas.

Man page

The ntop man page you can access on Unix systems.

Help

Pointers to mailing lists and other resources available to ntop users.

Bug Report

Form for reporting a bug (in case you find one!).

Total

This menu reports information about total traffic (i.e. sent plus received) captured by ntop.

All Protocols

Total (sent+received, IP+non IP) traffic sorted according to the active hosts. It is possible to restrict the view to all hosts, local hosts only, or remote hosts only, by selecting the links placed at the top right of the page.

IP

Same as above with reports only about IP traffic (non IP traffic is not represented). You can customize the list of protocols ntop tracks by using the `-p` flag.

Throughput

ntop keeps track of the total traffic sent/received by a host. The drawback is that a host that performed a significant amount of traffic in past hours is still on top of the list although now is maybe doing very little traffic. This menu allows you to see what's the current bandwidth used by a host in order to have a real-time view of the top senders/receivers

Host Activity

This is a colored map of the traffic being sent/received by known hosts exploded according to the hour of the day. It is very useful for detecting strange host activities (e.g. traffic during the night), or for finding out when your secretary turns on her PC © (please remember that ntop has not been written for being a spy, so please do not misuse the data ntop reports).

NetFlows

A network flow is some traffic that satisfies some rules otherwise difficult to measure. For instance if you want to know the traffic between sent by host A to

host B and C of protocol Y. If you want to define your own flows please have a look at the `-F` flag.

Received

This menu is the same as the Total menu with the exception that all the views are restricted only to the traffic received by hosts (i.e. sent traffic is not shown).

All Protocols

IP

Throughput

Host Activity

Sent

This menu is the same as the Total menu with the exception that all the views are restricted only to the traffic sent by hosts (i.e. received traffic is not shown).

All Protocols

IP

Throughput

Host Activity

Stats

This menu reports general traffic statistics.

Multicast

List of all the hosts that generated or received multicast traffic.

Traffic

Detailed view of the the traffic received so far completed with charts, RMON-like stats, historical view (RRD), interface statistics.

Hosts

Reports similar to Total->Traffic with more information about hosts (e.g. AS information) and no protocol detail.

Local Info

ntop performs passive OS fingerprinting¹ and detects/guesses the operating system of local hosts. The OS detection requires that the host performs TCP traffic while ntop is active. When the OS is detected, ntop reports in this page all the hosts whose OS has been detected, sorted according to the OS type. You can easily use this page for host inventory, asset tracking and users who installed an OS that's prohibited by the current network policy. Note that the algorithm guesses the OS, so it's likely that in some cases (e.g. when multiple OSs have the same fingerprint) the detection fails or it's incorrect. In this case please update the `etter.passive.os.fp.gz` file that's part of the ntop distribution.

Network load

Charts about the network load.

Domain

List of active hosts sorted according to the Internet domain they belong to,

IP Traffic

Remote -> Local

- _____

¹ Code and algorithm courtesy of the Ettercap project (<http://ettercap.sourceforge.net/>).

Traffic statistics about traffic generated by remote hosts and sent to local hosts.

Local -> Remote

Traffic statistics about traffic generated by local hosts and sent to remote hosts.

Local <->Local

Traffic statistics about traffic generated by local hosts and sent to local hosts.

Remote -> Remote

Traffic statistics about traffic generated by remote hosts and sent to remote hosts. Note that you'll not see much traffic here unless you analyze a network trunk where flows non-local traffic.

Matrix

Matrix representing traffic sent by local hosts to local hosts. Each cell contains the total traffic sent and received, and it's colored according to the amount of traffic.

IP Protos

Distribution

Distribution of IP protocols with respect to various criteria (e.g. TCP vs. UDP) and chart representing the proportions of local vs. remote traffic.

Usage

List of IP ports and hosts that use such ports. This report can be used for obtaining the list of all hosts that use port Y.

Sessions

List of active TCP sessions.

Routers

List of network routers detected by ntop completed with the name of hosts that use them

ASs

List of hosts sorted to the origin ASs (Autonomous System) they belong to. The mapping between IP address and AS is stored in the file AS-list.txt.gz that comes with ntop.

VLANs

List of hosts sorted to the VLAN they belong to. Note that this report is empty unless ntop captures traffic from a network adapter where flows 802.11Q traffic.

Admin

Plugins

Manage the state and configuration of plugins. See below for exhaustive descriptions of the options available for each plugin.

Switch NIC

ntop can simultaneously capture traffic from various adapter. However the web interface reports information of only one interface at time. Use this menu entry for switching among the various network interfaces. Have a look at the `-M` option if you want to learn more about how to keep sperated traffic belonging to various interfaces.

Dump Data

External applications can take advantage of ntop by periodically fetching data from it. ntop has been designed to export traffic data in various formats including text, XML and also code ready to use coded in programming languages such as Perl, PHP, Python.

Log

Display the last 50 entries that have been added to the ntop log.

Change Filter [Requires Authorization]

ntop allows interface filters to be specified in order to restrict the amount of traffic that can be analyzed by a network interface. You can specify the filter via the command line with the `-B` option or at runtime using this menu entry. Note that the syntax of the filter accepted by ntop is BPF (see appendix).

Reset Stats [Requires Authorization]

ntop traffic statistics are never reset (i.e. counters are always incremented) unless users explicitly require it by using this menu entry. Note that if you want to reset the ntop stats every day, you can write a simple script that daily calls this URL.

Users [Requires Authorization]

Manage web users: add/remove user, change password. Note that the web users have nothing to do (actually keep them separate) with the local host users.

URLs [Requires Authorization]

For each configured user, it is possible to restrict access only to selected URLs. This prevents users from having access to all the reports produced by ntop but just to the information that refers to them.

Shutdown [Requires Authorization]

Shutdown ntop.

In order to allow users to sort table data according to the various columns, users can click on the column name in order to sort table data according to the selected column. The web view represents much more information in a more natural way than the one accessible in interactive mode. In particular the web interfaces enables users to visualize:

- multicast information
- network flows
- local IP subnet traffic matrix
- active TCP sessions
- traffic distribution (local vs. remote)

The web interface supports multiple concurrent connections. Each time an HTTP request is received by ntop, a new ntop instance is forked (Unix only, unless the `-K` option is used). While the main ntop application processed packets, the child instance serves the HTTP request and then quits.

4.2 Ntop Plugins

Currently ntop ships with seven plugins. Plugins can be activated/deactivated/configured at runtime using the web interface. ntop saves the plugin preferences persistently so they don't have to be reconfigured if ntop is restarted. Due to the open architecture of ntop, users can code new plugins (e.g. for extending report capabilities or for handling new protocols) and add them to ntop without changing anything into the main ntop code.

The available plugins are:

sFlow

This plugin is used to add ntop probe/collector capabilities for the sFlow protocol². In particular the following options can be set:

Incoming Flows: the local port where a remote sFlow collector sends packets in sFlow format. Set this value to zero to disable sFlow collection. This option allows turning ntop into an sFlow collector. All the collected data is associated with a

- _____

² See <http://www.sflow.org> for more information about this protocol.

virtual sFlow interface.

Outgoing Flows: specify the IP address of the remote host where the sFlow collector is listening at port 6343. This option allows to turn ntop into an sFlow probe similar to the probes that run on network switches.

RrdPlugin

This plugin is used to persistently save traffic data on disk exploiting the RRD library³. This plugin allows to tune some parameters including:

Dump Interval: how often ntop stores traffic data into RRD. Please do not set this parameter to a value too little as RRD savings is CPU and disk intensive. The default is 300 seconds.

Dump Hours/Days/Months: parameters that are used to tune the RRD data aggregation capabilities. Do not change them unless you're familiar with them.

Data to Dump: select what you want to save into RRD. Note that saving data to RRD is costly in term of CPU time and disk space. Make sure you save only the information you really care about.

Hosts Filter: list of networks to which the hosts (see the previous option) that will be saved to RRD must belong. An empty list means that every host will be saved on RRD.

RRD Detail: specify the detail of the information you save into RRD. Changing detail level increases the number of RRD counters saved to disk.

RRD Files Path: specify the path of the directory that contains the RRD database files.

File/Directory Permissions: specify the file permissions of the RRD database files.

PDAplugin

This plugin shows a subset of the information reported by ntop so that it can be easily rendered on a PDA or mobile phone.

NfsWatch

Display detailed NFS traffic information (if any).

NetFlow

This plugin turns ntop into a NetFlow collector/probe. Due to the ntop design, ntop performs better as collector other than probe: if you need a reliable, software-only, GPL probe give nProbe⁴ a try. Currently can emit flows into NetFlow V5 format and collect flows in V5/V7/V9 formats. The plugin supports the following options:

Incoming Flows (Collector Mode)

Flow Collection: local IP port where ntop receives NetFlow packets. All the collected data is associated with a virtual NetFlow interface.

Virtual Interface Address: specify the network to which the virtual NetFlow interface belongs. This information is used by ntop to compute whether hosts are local/remote with respect to the virtual interface.

Flow Aggregation: enable/disable flow aggregation.

White List: if specified, instructs ntop to accept flows from hosts that belong to the specified network list.

• _____

³ See <http://www.rrdtool.org/>.

⁴ nProbe is the ntop NetFlow companion, available from <http://www.ntop.org/nProbe.html>

Black List: if specified, instruments ntop to reject flows received from hosts that belong to the specified network list.

Outgoing Flows (Probe Mode)

Interfaces: enable/disable the interfaces from which flows are emitted. You can tweak the interface status by clicking on the Yes/No hyperlink.

Remote IP Address: address of the remote collector that will collect flows generated by ntop.

Last Seen

Display information about when a certain host was seen for the last time.

IcmpWatch

Display detailed ICMP traffic information (if any). Note that the ICMP traffic is very important in a network, so keep your eyes on this report.

5. Hints and Tips

Large Networks: Memory and CPU Usage

As explained before, ntop stores information on a hash whose size increases with the number of monitored hosts. When ntop monitors a large network, it is likely that it spends a lot of time adding/purging hosts that are not really interesting from the monitoring point of view. Additionally this activity requires a lot of CPU time that can lead to packet loss as the time allocated to packet processing decreases dramatically. In order to optimize your resources it is a good idea to start ntop using the following command line options: `-g` for focusing ntop only on local hosts and `-m`, if necessary, for specifying the list of local networks. If you want you can also disable protocol decoders (`-b`) and disable TCP session tracking (`-z`).

Packet Loss

ntop sometimes is not able to process all the packets (see menu Stats -> Traffic) simply because the time it spends processing a packet is longer than the intra-packet arrival time. Additionally the libpcap (on which ntop is built) loses packets whenever the OS is not able to capture all the packets. Some solutions are:

- Get a better CPU (this can help partially, see below).
- Reduce work ntop has to do (see previous tip).
- Linux-only: use libpcap with mmap support (<http://public.lanl.gov/cpw/>).
- Use kernel device polling to prevent interrupt livelock. On Linux you need to use kernel 2.6 and use a NIC whose driver supports NAPI. On FreeBSD you need to enable polling into your kernel and do `sysctl kern.polling.enable=1`.

Strange Host Activities

ntop tracks host activities and adds a red/yellow flag next to an host that is suspicious. However ntop cannot put too many yellow flags because the interface would become horrible on some networks. For this reason, if you want to know whether there are hosts that perform strange activities on your network, you should look at the following parameters:

- Ratio between the number of SYN sent and SYN/ACK received, and ICMP echo request sent and ICMP echo responses received.
- Host activity indexes too high (note that hosts that have P2P programs enable increment this index significantly but this does not necessarily mean that this host is suspicious).
- Too much traffic generated by different hosts towards the same AS (Autonomous System).

6. Staying in Touch with ntop Users

If you have a question concerning ntop, please subscribe to the ntop mailing list: ntop and ntop-dev. Please send bug reports to the ntop-dev <ntop-dev@ntop.org> mailing list. The ntop <ntop@ntop.org> mailing list is used for discussing ntop usage issues. In order to post messages on the lists a (free) subscription is required in order to limit/avoid spam.

To subscribe to the ntop list, please have a look at <http://listgateway.unipi.it/mailman/listinfo/ntop>. Before sending a mail, please check the mailing list archive where you might find the answer to your question.

Please do NOT contact the author directly unless this is a personal question. Commercial support is available under request. Please see the ntop (<http://www.ntop.org/>) site for further info. Please send code patches to <patch@ntop.org>.

7. Acknowledgments

The author acknowledges the Centro Serra of the University of Pisa, Italy (<http://www-serra.unipi.it/>) for hosting the ntop sites (both web and mailing lists), and Burton Strauss <burton@ntopsupport.com> for his help and user assistance. Many thanks to Stefano Suin <stefano@ntop.org> and Rocco Carbone <rocco@ntop.org> for contributing to the project.

8. References

Lawrence Berkeley National Laboratory , *Libpcap v.0.4*, <ftp://ftp.ee.lbl.gov/libpcap.tar.Z>

Lawrence Berkeley National Laboratory , *tcpdump*, <ftp://ftp.ee.lbl.gov/tcpdump.tar.Z>

Tobias Oetiker, *RRDtools*, <http://www.rrdtools.org>

Appendix

A. ntop Prerequisite Packages

ntop requires a number of external tools and libraries to operate. Certain other tools are optional, but add to the program's capabilities

Required libraries include:

- libpcap from <http://www.tcpdump.org/>
The Windows version makes use of WinPcap (libpcap for Windows) which may be downloaded from <http://winpcap.polito.it/install/default.htm>.
WARNING: The 2.x releases of WinPcap will NOT support SMP machines.
- gdbm from <http://www.gnu.org/software/gdbm/gdbm.html>
- ntop requires a POSIX threads library. Although a single-threaded version of ntop can be built from the source if requested during `./configure`, it is not recommended for more than trivial usage.
- The gd library, for the creation of png files, available at <http://www.boutell.com/gd/>. ntop supports both gd 1.X and 2.X. The libpng library, for the creation of png files, available at <http://www.libpng.org/pub/png/libpng.html>. ntop supports both the 1.0.x series and the 1.2.x series of libpng, but cautions users that there are incompatibilities if you compile with one and run with the other. Please read the discussion in docs/FAQ before reporting ANY problem with libpng.
- The OpenSSL library provides HTTPS support to ntop. The OpenSSL project is available at <http://www.openssl.org>.
- The rrdtool library is required by the rrd plugin. rrdtool creates 'Round-Robin databases' which are used to store and graph historical data in a format that permits long duration retention without growing larger over time. The rrdtool home page is <http://people.ee.ethz.ch/~oetiker/webtools/rrdtool/>. ntop includes a patched and frozen version of rrdtool in the `myrrd/` directory. Users of ntop v2.3 should not need to specifically install rrdtool.
- The sflow Plugin is courtesy of and supported by InMon Corporation, <http://www.inmon.com/sflowTools.htm>.

There are other optional libraries. See the output of `./configure` for a fuller listing. Tool locations are current as of July 2003 please send email to report new locations or dead links.

B. BPF Packet Filtering Expressions

This section has been extracted from the `tcpdump` man page. The *expression* consists of one or more *primitives*. Primitives usually consist of an *id* (name or number) preceded by one or more qualifiers. There are three different kinds of qualifier:

type

qualifiers say what kind of thing the *id* name or number refers to. Possible types are **host**, **net** and **port**. E.g., `host foo`, `net 128.3`, `port 20`. If there is no type qualifier, **host** is assumed.

dir

qualifiers specify a particular transfer direction to and/or from *id*. Possible directions are **src**, **dst**, **src or dst** and **src and dst**. E.g., `src foo`, `dst net 128.3`, `src or dst port ftp-data`. If there is no *dir* qualifier, **src or dst** is assumed.

proto

qualifiers restrict the match to a particular protocol. Possible protos are: **ether**, **fdi**, **ip**, **arp**, **rarp**, **decnet**, **lat**, **moprc**, **mopdl**, **tcp** and **udp**. E.g., `ether src foo`, `arp net 128.3`, `tcp port 21`. If there is no *proto* qualifier, all protocols consistent with the type are assumed. E.g., `src foo` means `(ip or arp or rarp) src foo` (except the latter is not legal syntax), `net bar` means `(ip or arp or rarp) net bar` and `port 53` means `(tcp or udp) port 53`.

[`fdi` is actually an alias for `ether`; the parser treats them identically as meaning "the data link level used on the specified network interface." FDDI headers contain Ethernet-like source and destination addresses, and often contain Ethernet-like packet types, so you can filter on these FDDI fields just as with the analogous Ethernet fields. FDDI headers also contain other fields, but you cannot name them explicitly in a filter expression.]

In addition to the above, there are some special 'primitive' keywords that don't follow the pattern: **gateway**, **broadcast**, **less**, **greater** and arithmetic expressions. All of these are described below.

More complex filter expressions are built up by using the words **and**, **or** and **not** to combine primitives. E.g., `host foo and not port ftp and not port ftp-data`. To save typing, identical qualifier lists can be omitted. E.g., `tcp dst port ftp or ftp-data or domain` is exactly the same as `tcp dst port ftp or tcp dst port ftp-data or tcp dst port domain`.

Allowable primitives are:

dst host *host*

True if the IP destination field of the packet is *host*, which may be either an address or a name.

src host *host*

True if the IP source field of the packet is *host*.

host *host*

True if either the IP source or destination of the packet is *host*. Any of the above *host* expressions can be prepended with the keywords, **ip**, **arp**, or **rarp** as in: **ip host *host*** which is equivalent to: **ether proto *ip* and *host* *host*** If *host* is a name with multiple IP addresses, each address will be checked for a match.

ether dst *ehost*

True if the ethernet destination address is *ehost*. *Ehost* may be either a name from `/etc/ethers` or a number (see [ethers\(3N\)](#) for numeric format).

ether src *ehost*

True if the ethernet source address is *ehost*.

ether host *ehost*

True if either the ethernet source or destination address is *ehost*.

gateway *host*

True if the packet used *host* as a gateway. I.e., the ethernet source or destination address was *host* but neither the IP source nor the IP destination was *host*. *Host* must be a name and must be found in both `/etc/hosts` and `/etc/ethers`. (An equivalent expression is **ether host *ehost* and not host host** which can be used with either names or numbers for *host* / *ehost*.)

dst net *net*

True if the IP destination address of the packet has a network number of *net*, which may be either an address or a name.

src net *net*

True if the IP source address of the packet has a network number of *net*.

net *net*

True if either the IP source or destination address of the packet has a network number of *net*.

dst port *port*

True if the packet is ip/tcp or ip/udp and has a destination port value of *port*. The *port* can be a number or a name used in `/etc/services` (see [tcp\(4P\)](#) and [udp\(4P\)](#)). If a name is used, both the port number and protocol are checked. If a number or ambiguous name is used, only the port number is checked (e.g., **dst port 513** will print both tcp/login traffic and udp/who traffic, and **port domain** will print both tcp/domain and udp/domain traffic).

src port *port*

True if the packet has a source port value of *port*.

port *port*

True if either the source or destination port of the packet is *port*. Any of the above port expressions can be prepended with the keywords, **tcp** or **udp**, as in: **tcp src port *port*** which matches only tcp packets.

less *length*

True if the packet has a length less than or equal to *length*. This is equivalent to: **len <= *length***.

greater *length*

True if the packet has a length greater than or equal to *length*. This is equivalent to: **len >= *length***.

ip proto *protocol*

True if the packet is an ip packet (see *ip4P*) of protocol type *protocol*. *Protocol* can be a number or one of the names *icmp*, *udp*, *nd*, or *tcp*. Note that the identifiers *tcp*, *udp*, and *icmp* are also keywords and must be escaped via backslash (\), which is \\ in the C-shell.

ether broadcast

True if the packet is an ethernet broadcast packet. The *ether* keyword is optional.

ip broadcast

True if the packet is an IP broadcast packet. It checks for both the all-zeroes and all-ones broadcast conventions, and looks up the local subnet mask.

ether multicast

True if the packet is an ethernet multicast packet. The *ether* keyword is optional. This is shorthand for **ether[0] & 1 != 0**.

ip multicast

True if the packet is an IP multicast packet.

ether proto *protocol*

True if the packet is of ether type *protocol*. *Protocol* can be a number or a name like *ip*, *arp*, or *rarp*. Note these identifiers are also keywords and must be escaped via backslash (\). [In the case of FDDI (e.g., **fddi protocol arp**), the protocol identification comes from the 802.2 Logical Link Control (LLC) header, which is usually layered on top of the FDDI header. *ntop* assumes, when filtering on the protocol identifier, that all FDDI packets include an LLC header, and that the LLC header is in so-called SNAP format.]

decnet src *host*

True if the DECNET source address is *host*, which may be an address of the form "0.123", or a DECNET host name. [DECNET host name support is only available on Ultrix systems that are configured to run DECNET.]

decnet dst *host*

True if the DECNET destination address is *host*.

decnet host *host*

True if either the DECNET source or destination address is *host*.

ip, arp, rarp, decnet

Abbreviations for: **ether proto *p***
where *p* is one of the above protocols.

lat, moprc, mopdl

Abbreviations for: **ether proto *p***
where *p* is one of the above protocols. Note that *ntop* does not currently know how to parse these protocols.

tcp, udp, icmp

Abbreviations for: **ip proto *p***
where *p* is one of the above protocols.

expr relop expr

True if the relation holds, where *relop* is one of >, <, >=, <=, =, !=, and *expr* is an

arithmetic expression composed of integer constants (expressed in standard C syntax), the normal binary operators [+ , - , * , / , & , |], a length operator, and special packet data accessors. To access data inside the packet, use the following syntax: *proto [expr : size]*

Proto is one of **ether**, **fddi**, **ip**, **arp**, **rarp**, **tcp**, **udp**, or **icmp**, and indicates the protocol layer for the index operation. The byte offset, relative to the indicated protocol layer, is given by *expr*. *Size* is optional and indicates the number of bytes in the field of interest; it can be either one, two, or four, and defaults to one. The length operator, indicated by the keyword **len**, gives the length of the packet.

For example, `ether[0] & 1 != 0` catches all multicast traffic. The expression `ip[0] & 0xf != 5` catches all IP packets with options. The expression `ip[6:2] & 0x1fff = 0` catches only unfragmented datagrams and frag zero of fragmented datagrams. This check is implicitly applied to the **tcp** and **udp** index operations. For instance, `tcp[0]` always means the first byte of the TCP *header*, and never means the first byte of an intervening fragment.

Primitives may be combined using:

A parenthesized group of primitives and operators

(parentheses are special to the Shell and must be escaped).

Negation (`!` or `not`).

Concatenation (`&&` or `and`).

Alternation (`|` or `or`).

Negation has highest precedence. Alternation and concatenation have equal precedence and associate left to right. Note that explicit **and** tokens, not juxtaposition, are now required for concatenation. If an identifier is given without a keyword, the most recent keyword is assumed. For example, **not host vs and ace** is short for **not host vs and host ace** which should not be confused with **not (host vs or ace)**. Expression arguments can be passed to ntop as either a single argument or as multiple arguments, whichever is more convenient. Generally, if the expression contains Shell metacharacters, it is easier to pass it as a single, quoted argument. Multiple arguments are concatenated with spaces before being parsed.

Examples

To select all packets arriving at or departing from *sundown*:

```
ntop host sundown
```

To select traffic between *helios* and either *hot* or *ace*:

```
ntop host helios and \( hot or ace \)
```

To select all IP packets between *ace* and any host except *helios*:

```
ntop ip host ace and not helios
```

To select all traffic between local hosts and hosts at Berkeley:

```
ntop net ucb-ether
```

To select all ftp traffic through internet gateway *snu*: (note that the expression is quoted to prevent the shell from (mis-)interpreting the parentheses):

```
ntop 'gateway snu and (port ftp or ftp-data)'
```

To select traffic neither sourced from nor destined for local hosts (if you gateway to one other net, this stuff should never make it onto your local net).

ntop ip and not net *localnet*

To select the start and end packets (the SYN and FIN packets) of each TCP conversation that involves a non-local host.

ntop 'tcp[13] & 3 != 0 and not src and dst net *localnet*

To select IP packets longer than 576 bytes sent through gateway *snoop*:

ntop 'gateway snoop and ip[2:2] > 576'

To select IP broadcast or multicast packets that were *not* sent via ethernet broadcast or multicast:

ntop 'ether[0] & 1 = 0 and ip[16] >= 224'

To select all ICMP packets that are not echo requests/replies (i.e., not ping packets):

ntop 'icmp[0] != 8 and icmp[0] != 0'

C. ntop Licence

ntop is open-source software (see <http://www.opensource.org/>) and is distributed under the GNU GPL2 licence.

GNU GENERAL PUBLIC LICENSE

Version 2, June 1991

Copyright (C) 1989, 1991 Free Software Foundation, Inc.

59 Temple Place Suite 330, Boston, MA, 02111-1307, USA.

GNU GENERAL PUBLIC LICENSE

TERMS AND CONDITIONS FOR COPYING, DISTRIBUTION AND MODIFICATION

This License applies to any program or other work which contains a notice placed by the copyright holder saying it may be distributed under the terms of this General Public License. The "Program", below, refers to any such program or work, and a "work based on the Program" means either the Program or any derivative work under copyright law: that is to say, a work containing the Program or a portion of it, either verbatim or with modifications and/or translated into another language. (Hereinafter, translation is included without limitation in the term "modification".) Each licensee is addressed as "you".

Activities other than copying, distribution and modification are not covered by this License; they are outside its scope. The act of running the Program is not restricted, and the output from the Program is covered only if its contents constitute a work based on the Program (independent of having been made by running the Program).

Whether that is true depends on what the Program does.

You may copy and distribute verbatim copies of the Program's source code as you receive it, in any medium, provided that you conspicuously and appropriately publish on each copy an appropriate copyright notice and disclaimer of warranty; keep intact all the notices that refer to this License and to the absence of any warranty; and give any other recipients of the Program a copy of this License along with the Program.

You may charge a fee for the physical act of transferring a copy, and you may at your option offer warranty protection in exchange for a fee.

You may modify your copy or copies of the Program or any portion of it, thus forming a work based on the Program, and copy and distribute such modifications or work under the terms of Section 1 above, provided that you also meet all of these conditions:

You must cause the modified files to carry prominent

notices stating that you changed the files and the date of any change.

You must cause any work that you distribute or publish, that in whole or in part contains or is derived from the Program or any part thereof, to be licensed as a whole at no charge to all third parties under the terms of this License.

If the modified program normally reads commands interactively when run, you must cause it, when started running for such interactive use in the most ordinary way, to print or display an announcement including an appropriate copyright notice and a notice that there is no warranty (or else, saying that you provide a warranty) and that users may redistribute the program under these conditions, and telling the user how to view a copy of this License. (Exception: if the Program itself is interactive but does not normally print such an announcement, your work based on the Program is not required to print an announcement.)

These requirements apply to the modified work as a whole. If identifiable sections of that work are not derived from the Program, and can be reasonably considered independent and separate works in themselves, then this License, and its terms, do not apply to those sections when you distribute them as separate works. But when you distribute the same sections as part of a whole which is a work based on the Program, the distribution of the whole must be on the terms of this License, whose permissions for other licensees extend to the entire whole, and thus to each and every part regardless of who wrote it.

Thus, it is not the intent of this section to claim rights or contest your rights to work written entirely by you; rather, the intent is to exercise the right to control the distribution of derivative or collective works based on the Program.

In addition, mere aggregation of another work not based on the Program with the Program (or with a work based on the Program) on a volume of a storage or distribution medium does not bring the other work under the scope of this License.

You may copy and distribute the Program (or a work based on it, under Section 2) in object code or executable form under the terms of Sections 1 and 2 above provided that you also do one of the following:

Accompany it with the complete corresponding machine-readable source code, which must be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange; or,

Accompany it with a written offer, valid for at least three years, to give any third party, for a charge no more than your cost of physically performing source distribution, a complete machine-readable copy of the corresponding source code, to be distributed under the

terms of Sections 1 and 2 above on a medium customarily used for software interchange; or,

Accompany it with the information you received as to the offer to distribute corresponding source code. (This alternative is allowed only for non commercial distribution and only if you received the program in object code or executable form with such an offer, in accord with Subsection b above.)

The source code for a work means the preferred form of the work for making modifications to it. For an executable work, complete source code means all the source code for all modules it contains, plus any associated interface definition files, plus the scripts used to control compilation and installation of the executable. However, as a special exception, the source code distributed need not include anything that is normally distributed (in either source or binary form) with the major components (compiler, kernel, and so on) of the operating system on which the executable runs, unless that component itself accompanies the executable.

If distribution of executable or object code is made by offering access to copy from a designated place, then offering equivalent access to copy the source code from the same place counts as distribution of the source code, even though third parties are not compelled to copy the source along with the object code.

You may not copy, modify, sublicense, or distribute the Program except as expressly provided under this License. Any attempt otherwise to copy, modify, sublicense or distribute the Program 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.

You are not required to accept this License, since you have not signed it. However, nothing else grants you permission to modify or distribute the Program or its derivative works. These actions are prohibited by law if you do not accept this License. Therefore, by modifying or distributing the Program (or any work based on the Program), you indicate your acceptance of this License to do so, and all its terms and conditions for copying, distributing or modifying the Program or works based on it.

Each time you redistribute the Program (or any work based on the Program), the recipient automatically receives a license from the original licensor to copy, distribute or modify the Program subject to these terms and conditions. You may not impose any further restrictions on the recipients' exercise of the rights granted herein. You are not responsible for enforcing compliance by third parties to this License.

If, as a consequence of a court judgment or allegation of patent infringement or for any other reason (not limited to patent issues), conditions are imposed on you (whether by court

order, agreement or otherwise) that contradict the conditions of this License, they do not excuse you from the conditions of this License. If you cannot distribute so as to satisfy simultaneously your obligations under this License and any other pertinent obligations, then as a consequence you may not distribute the Program at all. For example, if a patent license would not permit royalty-free redistribution of the Program by all those who receive copies directly or indirectly through you, then the only way you could satisfy both it and this License would be to refrain entirely from distribution of the Program. If any portion of this section is held invalid or unenforceable under any particular circumstance, the balance of the section is intended to apply and the section as a whole is intended to apply in other circumstances.

It is not the purpose of this section to induce you to infringe any patents or other property right claims or to contest validity of any such claims; this section has the sole purpose of protecting the integrity of the free software distribution system, which is implemented by public license practices. Many people have made generous contributions to the wide range of software distributed through that system in reliance on consistent application of that system; it is up to the author/donor to decide if he or she is willing to distribute software through any other system and a licensee cannot impose that choice. This section is intended to make thoroughly clear what is believed to be a consequence of the rest of this License.

If the distribution and/or use of the Program is restricted in certain countries either by patents or by copyrighted interfaces, the original copyright holder who places the Program under this License may add an explicit geographical distribution limitation excluding those countries, so that distribution is permitted only in or among countries not thus excluded. In such case, this License incorporates the limitation as if written in the body of this License.

The Free Software Foundation may publish revised and/or new versions of the General Public 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.

Each version is given a distinguishing version number. If the Program specifies a version number of this License which applies to it and "any later version", you have the option of following the terms and conditions either of that version or of any later version published by the Free Software Foundation. If the Program does not specify a version number of this License, you may choose any version ever published by the Free Software Foundation.

If you wish to incorporate parts of the Program into other free programs whose distribution conditions are different, write to the author to ask for permission. For software which is copyrighted by the Free Software Foundation, write to the Free Software Foundation; we sometimes make exceptions for this. Our decision will be guided by the two goals of preserving the free status of all derivatives of our free software and of promoting the sharing and reuse of software generally.

NO WARRANTY

BECAUSE THE PROGRAM IS LICENSED FREE OF CHARGE, THERE IS NO WARRANTY FOR THE PROGRAM, TO THE EXTENT PERMITTED BY APPLICABLE LAW. EXCEPT WHEN OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR OTHER PARTIES PROVIDE THE PROGRAM "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE PROGRAM IS WITH YOU. SHOULD THE PROGRAM PROVE DEFECTIVE, YOU ASSUME THE COST OF ALL NECESSARY SERVICING, REPAIR OR CORRECTION.

IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN WRITING WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MAY MODIFY AND/OR REDISTRIBUTE THE PROGRAM AS PERMITTED ABOVE, BE LIABLE TO YOU FOR DAMAGES, INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING OUT OF THE USE OR INABILITY TO USE THE PROGRAM (INCLUDING BUT NOT LIMITED TO LOSS OF DATA OR DATA BEING RENDERED INACCURATE OR LOSSES SUSTAINED BY YOU OR THIRD PARTIES OR A FAILURE OF THE PROGRAM TO OPERATE WITH ANY OTHER PROGRAMS), EVEN IF SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.