

Bandwidth Management Tools

for version devel, 6 April 2005

Written by: Nigel Kukard <nkukard@lbsd.net>

This manual is for BWM Tools (version devel, 6 April 2005)

Copyright © 2005 Linux Based Systems Design.

Permission is granted to copy, distribute and/or modify this document under the terms of the GNU Free Documentation License, Version 1.2 or any later version published by the Free Software Foundation; with no Invariant Sections, no Front-Cover Texts, and no Back-Cover Texts. A copy of the license is included in the section entitled “GNU Free Documentation License”.

Table of Contents

1	Introduction to BWM Tools	1
1.1	BWM Tools Features	1
2	Installing BWM Tools	2
3	Configuring BWM Tools	3
3.1	The <global> section	3
3.2	The <acl> section	4
3.3	The <nat> section	6
3.4	The <traffic> section	10
4	Integrating BWM Tools with your system ..	15
5	Graphing	18
6	Examples	21
6.1	Basic configuration examples	21
6.2	Advanced configuration examples	21
Appendix A	Copying This Manual	30
A.1	GNU Free Documentation License	30
A.1.1	ADDENDUM: How to use this License for your documents	36
Index	37

1 Introduction to BWM Tools

Bandwidth Management Tools was designed to provide a full suite of bandwidth management applications, able to shape, log and graph traffic.

Seeing as BWM Tools uses iptables for matching traffic, the complexity of traffic control is limitless.

BWM Tools is a set of userspace utilities, no kernel patches are required. As long as your iptables supports the ‘-j QUEUE’ target, traffic shaping will work.

1.1 BWM Tools Features

This section lists a few features which make BWM Tools a good solution for small to large enterprises. . .

- Traffic Shaping
 - Hierarchical flows Allows you to embed flows within flows to form complex traffic shaping rules.
 - Parent burst thresholds Parent burst thresholds allow child flows to burst until their parent flow has reached a specific utilization threshold.
- Graphing
 - RRD Tool file support Generation of rrdtool files which can be used to create custom graphs.
 - Builtin RRD Tool graphing support BWM Tools can generate pretty looking graphs all by itself. Parameters for graphing are discussed in the Graphing section.
- Logging
 - Logging of traffic BWM Tools logs can log traffic stats to file at pre-defined intervals for use in reporting or graphing.

2 Installing BWM Tools

Before you can use BWM Tools, you must make sure you have all the dependencies installed. . .

- glib2 >= 2.2.0
- libxml2 >= 2.5.0
- rrdtool >= 1.0.49 (required for graphing)

Next you need to download BWM Tools, compile it and install it.

Here is step-by-step instructions on how to do this. . .

1. Download the latest version of BWM Tools, the latest version can be found on the project homepage: <http://bwm-tools.pr.linuxrulz.org>
2. Uncompress the archive using either `tar jxvf <archive name>.tar.bz2` or `tar zxvf <archive name>.tar.gz` depending weather its a .tar.bz2 or .tar.gz respectively.
3. Run `./configure` in the source directory. Optionally a `'--prefix=...'` parameter can be passed which will determine where BWM Tools will be installed.
4. Once the configure process is complete, issue a `make` command, this will compile BWM Tools.
5. When BWM Tools has finished compiling, type `make install`. This will by default install BWM Tools into `/usr/local`, unless of course if you specified a `'--prefix=...'` above.

3 Configuring BWM Tools

Configuration of BWM Tools is done via an XML configuration file, this file is normally located in `/etc/bwm_tools/firewall.xml`

The layout of the file is pretty simple and is split up into various sections, these are detailed in the following sections...

3.1 The `<global>` section

This section contains global tags pertaining to either the operation of BWM Tools or definitions used in other sections. These tags are detailed below...

- Module management in the `<modules>` section

This section is used to load modules when `bwmd` starts. The syntax to load a module is as follows...

```
<load name="kernel_module_name" />
```

The `<load />` tag takes the following parameters...

- `name="..."` - This is the name of the module to load
- `params="..."` - Parameters to load module with

Here is how it can be used to load the `ip_queue` kernel module required by `bwmd` for shaping. Including `ftp` connection tracking to allow users to `ftp` through a tightly secured firewall.

```
<firewall>
  <global>
    <modules>
      <load name="ip_queue"/>
      <load name="ip_nat_ftp"/>
      <load name="ip_conntrack_ftp"/>
    </modules>
  </global>
  .
  .
  .
</firewall>
```

- Class definition in the `<class>` section

This section is used to define classes used in both firewalling and network address translation. The basic syntax is as follows...

```
<class name="traffic_from_support">
  <address name="pete_in" src="192.168.0.100" />
</class>
```

The `<class>` tag has got no other options apart from name.

The `<address />` tag on the other hand has the following options. . .

- `name="..."` - This is a descriptive name for the address, isn't really used anywhere
- `cmd-line="..."` - Optional command line arguments for iptables, for example `cmd-line="-m helper --helper <string>"`
- `dst="..."` - Optional destination IP address
- `dst-iface="..."` - Optional destination interface
- `dst-port="..."` - Optional destination port
- `proto="..."` - Optional protocol specification, any valid protocol in `'/etc/protocols'`
- `src="..."` - Optional source IP address
- `src-iface="..."` - Optional source interface
- `src-port="..."` - Optional source port

Here is an example how it can be used to match connections over a specific number. . .

```

<firewall>
  <global>
  .
  .
  .
    <class name="excess_connections_to_webserver">
      <address name="excess_to_server1" dst="192.168.0.100" proto="tcp" dst-
port="80" cmd-line="-m connlimit --connlimit-above 10"/>
    </class>
  </global>
  .
  .
  .
</firewall>

```

3.2 The `<acl>` section

This is basically the firewall section, you can add all your firewall rules here or just leave it blank to use your current firewall.

The syntax for this section is a little more complex and is as follows. . .

```

<firewall>
.
.
.
  <acl>
    <table name="filter">
      <chain name="INPUT" default="ACCEPT">
        <rule name="excess_connections" target="DROP">
          excess_connections_to_webserver
        </rule>
      </chain>
    </table>
  </acl>
.
.
.
</firewall>

```

Explaining the above example, this will add 1 rule to the INPUT chain under the `filter` table which will drop all new packets that arrive if the concurrent connections on port 80 is higher than 10.

It is the equivalent to...

```
iptables -t filter -A INPUT -d 192.168.0.10 -p tcp -dport 80 -m connlimit
--connlimit-above 10 -j DROP
```

The following tags and parameters are available...

- Specify the table with `<table> ... </table>`

The `<table>` tag is used to enclose the directives you plan to use with a specific table. Examples of tables are... `filter`, `nat`, `mangle`

The `<table>` tag takes the following parameters...

- `name="..."` - This is the name of the table we will be working with

- Specify a chain with `<chain> ... </chain>`

The `<chain>` tag is used to specify what chain the rules defined between the starting and ending tags apply to. Examples of already defined chains are `INPUT`, `OUTPUT` and `FORWARD`.

The `<chain>` tag takes the following parameters...

- `name="..."` - This is the name of the chain we will be working with
- `default="..."` - This specifies the default target for the chain

- Specify a rule with `<rule> ... </rule>`

The `<rule>` tag is used to specify what classes apply to what rule, and are in order inserted into the actual iptables chains as iptables rules.

The `<rule>` tag takes the following parameters...

- `name="..."` - Optional name of rule
- `cmd-line="..."` - Optional extra command line parameters to pass to iptables
- `target="..."` - This is the target for the rule, used as the `'-j <target>'` parameter when generating iptables rules.

Between the opening and closing tags, classes defined in the `<global>` section are listed, these classify which traffic applies to which rule.

Multiple classes can be listed, one per line.

Using the above, here is an example of a simple firewall which allows http and ssh traffic, assuming your IP address is 10.0.0.2 of course...

```
<firewall>
  # Global configuration and access classes
  <global>
    <class name="http_traffic">
      <address dst="10.0.0.2" proto="tcp" dst-port="80"/>
    </class>
    <class name="ssh_traffic">
      <address dst="10.0.0.2" proto="tcp" dst-port="22"/>
    </class>
  </global>

  # Access control lists
  <acl>
    <table name="filter">
      <chain name="INPUT" default="DROP">
        <rule name="allowed_traffic" target="ACCEPT">
          http_traffic
          ssh_traffic
        </rule>
      </chain>
      <chain name="FORWARD" default="DROP">
      </chain>
      <chain name="OUTPUT" default="ACCEPT">
      </chain>
    </table>
  </acl>
</firewall>
```

3.3 The <nat> section

The NAT section is used to define network address translation rules, these rules allow one to translate the source or destination IP address within packets. A common use for this is when a webserver is behind a firewall, requests are made to a globally routable IP address and translated to the internal IP address of the webserver and visa versa.

This section has the following syntax...

```

<firewall>
.
.
.
  <nat>
    <snat>
      <rule name="traf_from_webserver"
        to-src="<globally routable IP here>"
        traffic_from_webserver
      </rule>
    </snat>
    <dnat>
      <rule name="traf_to_webserver" to-dst="192.168.1.100">
        traffic_to_webserver
      </rule>
    </dnat>
    <masq>
      <rule name="traf_to_from_inside">
        internal_dsl_ips
      </rule>
    </masq>
  </nat>
.
.
.
</firewall>

```

There are 3 tags available, <snat>, <dnat> and <masq>, these three tags are used for source network address translation, destination address translation and masquerading respectively.

Valid options for these tags are as follows...

- Source network address translation using <snat>

SNAT is used for source network address translation, an example of which is again a webserver behind a firewall. Where SNAT comes in handy is when the webserver makes a query through the firewall, instead of the traffic on the internet coming from the webserver's internal IP 192.168.1.100 which is not going to work, the firewall translates 192.168.1.100 to a globally routable IP address.

There are no parameters for this tag, although the following sub-tags and parameters are available...

- Specify a rule with `<rule> ... </rule>`

The `<rule>` tag is used to specify what classes apply to what rule, and are in order inserted into the actual iptables chains as iptables rules.

The `<rule>` tag takes the following parameters...

- `name="..."` - Optional name of rule
- `to-src"..."` - Translate all traffic matched in the class specification to this source IP address.

Between the opening and closing tags, classes defined in the `<global>` section are listed, these classify which traffic applies to which rule.

Multiple classes can be listed, one per line.

- Destination network address translation using `<dnat>`

DNAT is used for destination network address translation, an example of which is yet again a webserver behind a firewall. Where DNAT comes in handy is when requests are made to the webserver's globally routable IP, this IP address is routed through the firewall and translated to the webserver's internal IP address. Optional traffic filtering can be carried out on the traffic, this is in most instances the case and prevents a lot of harmful traffic from interfering with the webserver's operation.

There are no parameters for this tag, although the following sub-tags and parameters are available...

- Specify a rule with `<rule> ... </rule>`

The `<rule>` tag is used to specify what classes apply to what rule, and are in order inserted into the actual iptables chains as iptables rules.

The `<rule>` tag takes the following parameters...

- `name="..."` - Optional name of rule
- `to-dst"..."` - Translate all traffic matched in the class specification to this destination IP address.

Between the opening and closing tags, classes defined in the `<global>` section are

listed, these classify which traffic applies to which rule.

Multiple classes can be listed, one per line.

- Masquerading using `<masq>`

Masquerading is normally used for source address translation in the scenario where you have a dynamic IP and never know what address to do the translation to. An example of which is a home PC acting as a DSL router.

There are no parameters for this tag, although the following sub-tags and parameters are available...

- Specify a rule with `<rule> ... </rule>`

The `<rule>` tag is used to specify what classes apply to what rule, and are in order inserted into the actual iptables chains as iptables rules.

The `<rule>` tag takes the following parameters...

- `name="..."` - Optional name of rule
- `to-ports"..."` - This specifies a range of source ports to use, overriding the default SNAT source port-selection heuristics. For this parameter to work you MUST have defined a protocol in all the classes specified. For example `proto="tcp"`.

Between the opening and closing tags, classes defined in the `<global>` section are listed, these classify which traffic applies to which rule.

Multiple classes can be listed, one per line.

An example using the above definitions would look something like this...

```
<firewall>
# Global configuration and access classes
<global>
  <class name="traf_from_webserver">
    <address src="192.168.0.100"/>
  </class>
  <class name="traf_to_webserver">
    <address dst="globally routable IP here"/>
  </class>
</global>

# Network address translation
<nat>
  <snat>
```

```

        <rule to-src="<globally routable IP here>">
            traf_from_webserver
        </rule>
    </snat>
    <dnat>
        <rule to-dst="192.168.0.100">
            traf_to_webserver
        </rule>
    </dnat>
</nat>
</firewall>

```

Here is an example if you pc is acting as a DSL router...

```

<firewall>
    # Global configuration and access classes
    <global>
        <class name="traf_going_to_dsl">
            <address src="192.168.0.0/24"/>
        </class>
    </global>

    # Network address translation
    <nat>
        <masq>
            <rule name="masq_traffic_going_out">
                traf_going_to_dsl
            </rule>
        </masq>
    </nat>
</firewall>

```

3.4 The <traffic> section

This section is used to define traffic shaping rules. These traffic shaping rules are called *flows*, the concept of flows is a single-parent child relationship. For instance you can define 1 major flow, within this flow you can define separate priorities and limits for different traffic such as mail, browsing and p2p traffic. This example setup might be used for a DSL internet connection where one would like to prioritize internet browsing.

The syntax of this section follows...

```

<firewall>
.
.
.
    # Traffic flows
    <traffic>
        <flow name="dsl_line_in" max-rate="64000" report-timeout="60">
            <flow name="http_in" max-rate="32000" burst-rate="64000" nfmark="100">
                http_traffic_in
            </flow>
            <flow name="smtp_in" max-rate="8000" burst-rate="32000" nfmark="101">
                smtp_traffic_in
            </flow>
            <flow name="p2p_in" max-rate="24000" burst-rate="32000"
                nfmark="102">
                p2p_traffic_in
        </flow>
    </traffic>

```

```

        </flow>
    </flow>
    <flow name="dsl_line_out" max-rate="64000" report-timeout="60">
        <flow name="http_out" max-rate="32000" burst-rate="64000" nfmark="200">
            http_traffic_out
        </flow>
        <flow name="smtp_out" max-rate="8000" burst-rate="32000" nfmark="201">
            smtp_traffic_out
        </flow>
        <flow name="p2p_out" max-rate="24000" burst-rate="32000" nfmark="202">
            p2p_traffic_out
        </flow>
    </flow>
</traffic>
.
.
.
</firewall>

```

The `<traffic> ... </traffic>` tags have no parameters.

Valid sub-tags and their parameters are detailed below...

- Specify a flow with `<flow> ... </flow>`

The `<flow>` tag is used to specify a traffic flow and takes the following parameters...

- **name="..."** - Mandatory flow name, this is used to identify the flow when reporting and monitoring
- **nfmark="..."** - Mandatory/Optional parameter to specify the NFMARK of the traffic that applies to this flow. This must be used at the deepest level of flow embedding to match traffic. Each **nfmark** value MUST be unique!
- **stats-len="..."** - Optional parameter to specify the period in seconds that the average bandwidth rate and packet rate is based on. If 0 is specified here there will be no average
- **queue-size="..."** - Optional parameter to specify the size of the entire packet queue. If 0 is specified, queue size is unlimited. If -1 is specified, the queue will not be used.
- **queue-len="..."** - Optional parameter to specify the maximum number of packets that can be in the entire queue at any one time. If -1 is specified the queue will not be used.
- **max-rate="..."** - Optional parameter to specify the maximum rate in bytes/s before packets are queued, packets are not queued if they can be bursted. If 0 is specified, no traffic limiting will occur. If however the **report-timeout="..."**

parameter is also specified then only logging will occur.

- **burst-rate**="..." - Optional parameter to specify the maximum rate in bytes/s which packets can be bursted. Bursting can only occur until the parent has maxed out its **max-rate**. Unlimited bursting will occur when **burst-rate** = 0, remember unlimited meaning until the parent has maxed its **max-rate**. This value must be greater than **max-rate**.
- **burst-threshold**="..." - Optional parameter to specify at what percentage we will stop bursting to our parent flow with regards to the parents current rate of usage. If this is set to 75, bursting to our parent will only be allowed until parent has maxed out 75% of its allowed maximum bandwidth utilization. If other flows max 70% of the parents bandwidth, we will be allowed to max our **max-rate** and burst until our parent reaches 75% of its **max-rate**. Remember **burst-threshold** pertains to the parents **max-rate** parameter, not the parents **burst-rate**.
- **report-timeout**="..." - Optional parameter to specify if and in what time increments the traffic statistics are logged to file. For example, if this parameter is set to 60, **bwm**d will log traffic stats to file every 60 seconds. Minimum value for this parameter is 30.
- **prio-classifier**="..." - Optional parameter to specify an automatic traffic prioritization classifier. This parameter defaults to the *none* classifier, where no prioritization takes place. Available classifiers are discussed below...
 - The "*port*" classifier With this classification prioritization happens automatically with the following ports mapped to their corresponding priorities. (1 = highest, 100 = lowest)...

TCP Traffic

```
'port 113 (AUTH)'
    'Priority 20'

'port 22, 23 (SSH, TELNET)'
    'Priority 25'

'port 80, 443, 8080, 3128, 3130 (HTTP, HTTPS, PROXY PORTS)'
    'Priority 65'

'port 2401 (CVS)'
    'Priority 70'

'port 110, 143 (POP3, IMAP4)'
    'Priority 75'

'port 20, 21 (FTP)'
    'Priority 80'
```

```

UDP Traffic
'port 53 (DNS)'
    'Priority 10'
'port 123 (NTP)'
    'Priority 15'
'port 1645/6, 1812/3 (RADIUS)'
    'Priority 30'
'port 33434-33465 (Normally traceroute)'
    'Priority 5'

```

The default priority for traffic not matching any of the above is 50.

- The *"none"* classifier This is the default classifier, no prioritization will occur and all traffic will be dumped in the default priority 50 queue.

Between the opening and closing tags, classes defined in the `<global>` section can be listed, if you want to list multiple classes use one per line, these classes classify which traffic applies to which rule.

Please note listing classes is required only if you are using BWM Tools to generate your firewall for you, otherwise just make sure you *MARK* your traffic correctly and the *MARK* value matches the `nfmark="..."` parameter value used above.

Alternatively `<flow> ... </flow>` tags can be embedded to form a more complex hierarchy.

On a last note, if you are infact not using BWM Tools to generate your firewall and don't want to embed flows in multiple hierarchical levels you can specify the flow tag quickly in the following way `<flow ... />`.

To continue on the line of complexity, one can specify the following sub-tags, within the `<flow> ... </flow>` tags...

- The `<queue> ... </queue>` tag is used to finer tune queuing

This tag can be specified to finer tune into which queue the traffic is put and has the following parameters...

- `prio="..."` - Mandatory parameter to specify the priority of the matched traffic. (1 = highest, 100 = lowest).
- `nfmark="..."` - Mandatory parameter to specify the mark value of the traffic.

Below is an example of using the `<queue> ... </queue>` tags to give VNC traffic highest priority...

```
<flow name="line_in" max-rate="32000">
  <flow name="p2p_traffic_in" max-rate="8000" burst-rate="24000" nfmark="100">
    class_p2p_traffic_in
  </flow>
  <flow name="vnc_in" max-rate="24000" burst-rate="32000">
    <queue prio="1" nfmark="101">
      class_vnc_in
    </queue>
  </flow>
</flow>
```

Between the opening and closing tags, classes defined in the `<global>` section can be listed, if you want to list multiple classes use one per line, these classes classify which traffic applies to which rule.

Please note listing classes is required only if you are using BWM Tools to generate your firewall for you, otherwise just make sure you *MARK* your traffic correctly and the *MARK* value matches the `nfmark="..."` parameter value used above.

On a last note, if you are infact not using BWM Tools to generate your firewall and want to specify a queue quickly, you can do so in the following way `<queue ... />`.

- Specify a group of flows with `<group> ... </group>`

The `<group>` tag is used for reporting only. It is for grouping flows together into 1 reporting name. This tag takes the following parameters...

- `name="..."` - Mandatory flow name, this is used to identify the flow when reporting and monitoring
- `report-timeout="..."` - Optional parameter to specify if and in what time increments the traffic statistics are logged to file. For example, if this parameter is set to 60, `bwmd` will log traffic stats to file every 60 seconds. Minimum value for this parameter is 30.
- `stats-len="..."` - Optional parameter to specify the period in seconds that the average bandwidth rate and packet rate is based on. If 0 is specified here there will be no average

4 Integrating BWM Tools with your system

This section will describe how to integrate BWM Tools into your system, be it you use BWM Tools to entirely manage your firewall, NAT and traffic shaping or just to do the traffic shaping.

There are two possible scenarios here detailed below. . .

- You want to use BWM Tools for both your firewall and traffic shaping.

This is the easiest scenario to deal with, only having 4 steps below to get your firewall, NAT and traffic shaping up and running. . .

1. Configure your classes, ACL's, NAT and traffic shaping rules as described in the previous sections. The end target for all accepted traffic must be *bwmd* in the *INPUT* chain or *OUTPUT* chain if you doing single box or a router configuration respectively.
2. Run BWM Firewall with the below possible arguments to generate an `iptables-restore` compatible configuration file. . .

```
Usage: bwm_firewall <options>

Options:
  -c, --config=<config_file>  Specify non-default BWM Tools config file
  -f, --file[=<output_file>]  Generate iptables-restore file from
                              BWM Tools firewall
  -l, --load                   Load BWM Tools firewall directly into
                              kernel
  -h, --help                   Display this page
  -r, --reset-counters        Reset iptables counters, usable with
                              "iptables-restore -c"
```

BWM Firewall takes the BWM Tools XML configuration file and translates the various sections and tags into a firewall which can be loaded directly with `iptables-restore`.

BWM Firewall defaults to writing the `iptables-restore` configuration file to `/etc/sysconfig/iptables`.

3. Once you've generated the `iptables` restore file you must load it atomically into the kernel with the following command. . .


```
iptables-restore < /etc/sysconfig/iptables
```
4. The last step is to fire up `bwmd` with your choice of the available options below. . .

```
Usage: bwmd <options>

Options:
```

```

    -c, --config=<config_file>    Specify non-default BWM Tools config file
    -f, --foreground              Run in foreground and print debug infoma-
tion to the screen
    -h, --help                    Display this page

```

BWMD defaults to using the configuration file in
 ‘/etc/bwm_tools/firewall.xml’.

- You want to use another firewalling application and have BWM Tools do only the traffic shaping.

Here there are a few things to remember. . .

- BWM Tools works with the *NFMARK* parameter attached to packets. Marking packets can only be done in the *mangle* table in *iptables*.
- BWM Tools uses the userpace queueing mechanism, all packets to be shaped must be targetted at *QUEUE* in the *filter* table. This is done by either adding a rule to the *INPUT* and *OUTPUT* chain in the case of a single box which you need to shape traffic to and from respectively. While in the case of a firewall where traffic passes through you would add a rule to the *FORWARD* chain.
- Therefore in order for BWM Tools to shape traffic, packets must be MARK'ed with a number corosponding to the number specified in the *nfmark="..."* parameter defined in the <flow> tag and targetted in *iptables* to *QUEUE* instead of *ACCEPT* as per above.

Imagine you would like your linux router to rate limit all traffic from and to IP 192.168.1.100, an example of this can be found below. . .

- Configuring *iptables*

```

iptables -t filter -A FORWARD -m mark ! --mark 0x0 -j QUEUE
iptables -t mangle -A FORWARD -s 192.168.1.100 -j MARK --set-mark 100
iptables -t mangle -A FORWARD -d 192.168.1.100 -j MARK --set-mark 101

```

- Configuring *bwmd*

```

<firewall>
  <global>
    <modules>
      <load name="ip_queue"/>
    </modules>
  </global>

  # Traffic flows

```

```
<traffic>
  <flow name="pc_in" max-rate="64000" report-timeout="60"
    nfmark="100" />
  <flow name="pc_out" max-rate="64000" report-timeout="60"
    nfmark="101" />
</traffic>
</firewall>
```

5 Graphing

BWM Tools supports graphing of traffic flows which have been specified with the `report-timeout=""`.

Generating a graph can be achieved using `bwm_graph` or by using the RRD files generated by `bwm_graph`.

These two methods are discussed below...

- Generating RRD files

The following section will explain how to have `bwm_graph` generate only RRD files and not graphs. This can be done quickly and simply using the following 3 commandline options...

1. The `-f` and `--flows` mandatory option

This option is used to specify the flows to include when generating the RRD files. An example of this option can be found below...

```
bwm_firewall --flows="flow_name_1,flow_name2,flow_name3" ...
```

There is an optional parameter to specify which counter will be used when outputting the RRD file. For this there are 3 possibilities, all 3 are the totals per `report-timeout=""` seconds specified in the relevant flow tag.

```
'pkt'      'Number of packets processed'

'size_bit'  'Bits transferred in above period'

'size_byte' 'Bytes transferred'

'dropped'  'Packets dropped'

'burst'    'Packets bursted'
```

The counter to use is specified in the following manner...

```
bwm_firewall --flows="flow_name_1(size_bit),flow_name_2(size_byte)" ...
```

2. The '-s' and '--start' mandatory option

This option is used to specify the date and/or time which to start our report from. The format for date and/or time specification is `yyyy/mm/dd hh:mm:ss`. An example of this option is as follows...

```
bwm_firewall ... --start="2003/01/20 01:20" ...
```

3. The '-e' and '--end' mandatory option

This option is used to specify the date and/or time which our report will end. The format for this option is the same as the '-s' and '--start' options.

An example of how to use all 3 above options to specify both the flows to work on and the reporting period can be done something like this...

```
bwm_firewall --flows="flow_name_1(size_bit),flow_name_2(size_bit)" --start="2003/01/20" -  
-end="2003/01/21" ■
```

- Creating a pretty graph using `bwm_graph`

`bwm_graph` has a builtin interface to `rrdtool`. Using this interface one can easily have `bwm_graph` generate pretty looking graphs itself.

The graphing capability of `bwm_graph` is in addition to the generation of RRD files, meaning that you are required to use all 3 mandatory options discussed in "Generating RRD files" above.

The following graphing options can be used...

- '--graph-filename=<filename>'

This parameter is used to specify an output filename for the generated .png image.
- '--graph-avg'

Write counter averages on the graph
- '--graph-date'

Write the start datetime and end datetime of the reporting period on the graph
- '--graph-title=<graph_title>'

Specify a title for your graph

- `--graph-total`
Write out counter totals on the graph

- `--graph-vert-title=<graph_title>`
Specify a vertical title for the graph

6 Examples

6.1 Basic configuration examples

6.2 Advanced configuration examples

1. This example demonstrates a firewall configuration which is used for an organization connected to a Cisco router, which in turn is used as the gateway to the internet. The server is configured to accept SMTP traffic from outside including incoming POP3 connections. This firewall will block all smtp traffic sourcing from inside going outside, this blocks most mass mailing worms.

```
<firewall>

#
#   Global configuration and access classes
#

<global>
  # Modules we need to load
  <modules>
    <load name="ip_queue"/>
    <load name="ip_contrack_ftp"/>
    <load name="ip_nat_ftp"/>
  </modules>

#
# BEGIN - STANDARD CLASSES
#
<class name="local_iface">
  <address src-iface="lo"/>
</class>

<class name="valid_connections">
  <address cmd-line="-m state --state ESTABLISHED,RELATED"/>
</class>

<class name="syn_packets">
  <address proto="tcp" cmd-line="--syn -m state --state NEW"/>
</class>

<class name="udp_packets">
  <address proto="udp"/>
</class>

<class name="icmp_packets">
  <address proto="icmp"/>
</class>

<class name="rsvp_packets">
  <address proto="2"/>
</class>
```

```
<class name="invalid_tcp_packets">
  <address proto="tcp" cmd-line="--tcp-flags ALL FIN,URG,PSH"/>
  <address proto="tcp" cmd-line="--tcp-flags ALL ALL"/>
  <address proto="tcp" cmd-line="--tcp-flags ALL SYN,RST,ACK,FIN,URG"/>
  <address proto="tcp" cmd-line="--tcp-flags ALL NONE"/>
  <address proto="tcp" cmd-line="--tcp-flags SYN,RST SYN,RST"/>
  <address proto="tcp" cmd-line="--tcp-flags SYN,FIN SYN,FIN"/>
</class>

<class name="valid_icmp_packets">
  <address proto="icmp" cmd-line="--icmp-type 0"/>
  <address proto="icmp" cmd-line="--icmp-type 3"/>
  <address proto="icmp" cmd-line="--icmp-type 8"/>
  <address proto="icmp" cmd-line="--icmp-type 11"/>
</class>

<class name="traceroute_packets">
  <address proto="udp" dst-port="33434:33465"/>
</class>

<class name="service_ftp">
  <address proto="tcp" dst-port="21"/>
</class>

<class name="service_ssh">
  <address proto="tcp" dst-port="22"/>
</class>

<class name="service_smtp">
  <address proto="tcp" dst-port="25"/>
</class>

<class name="service_dns">
  <address proto="tcp" dst-port="53"/>
  <address proto="udp" dst-port="53"/>
</class>

<class name="service_http">
  <address proto="tcp" dst-port="80"/>
</class>

<class name="service_https">
  <address proto="tcp" dst-port="443"/>
</class>

<class name="service_pop3">
  <address proto="tcp" dst-port="110"/>
</class>

<class name="service_tinc">
  <address proto="udp" dst-port="655"/>
  <address proto="tcp" dst-port="655"/>
</class>

<class name="service_ident">
  <address proto="tcp" dst-port="113"/>
</class>
```

```
<class name="service_imap">
  <address proto="tcp" dst-port="143"/>
</class>

<class name="service_pserver">
  <address proto="tcp" dst-port="2401"/>
</class>

<class name="service_httpproxy">
  <address proto="tcp" dst-port="3128"/>
  <address proto="tcp" dst-port="8080"/>
</class>

<class name="service_postgresql">
  <address proto="tcp" dst-port="5432"/>
</class>

<class name="service_time">
  <address proto="udp" dst-port="123" src-port="123"/>
</class>

<class name="service_rip">
  <address proto="udp" dst-port="520" src-port="520"/>
</class>

<class name="service_datametrics">
  <address proto="udp" dst-port="1645"/>
  <address proto="udp" dst-port="1646"/>
</class>

<class name="service_radius">
  <address proto="udp" dst-port="1812"/>
  <address proto="udp" dst-port="1813"/>
</class>

<class name="service_dhcp">
  <address proto="udp" dst-port="67:68"/>
</class>

<class name="30_per_min">
  <address cmd-line="-m limit --limit 30/min --limit-burst 10"/>
</class>

<class name="blank">
  <address />
</class>
#
# END - STANDARD CLASSES
#

<class name="valid_internal_traffic">
  <address src-iface="eth1" src="192.168.101.0/26" dst-iface="eth0"/>
</class>
```

```

<class name="nat_internal_traffic">
  <address src="192.168.101.0/26" dst="! 192.168.101.0/24"/>
</class>

<class name="internal_traffic">
  <address src-iface="eth1" dst-iface="eth0"/>
</class>

<class name="proxy_redirect">
  <address src="192.168.101.0/24" proto="tcp" dst="! 192.168.101.0/24"
    dst-port="80"/>
</class>

<class name="internal_local">
  <address src="192.168.101.0/24" />
</class>

# eth0 loop is normally used when doing strange NAT stuff
<class name="eth0_loop">
  <address src-iface="eth0" dst-iface="eth0"/>
</class>

</global>

#
# Access control lists
#
<acl>
  <table name="filter">

    #
    # CUSTOM RULES
    #

    <chain name="accept_input_all">
</chain>

    <chain name="accept_input_tcp">
      <rule target="accept_traffic">
        service_smtp;
        service_pop3;
      </rule>
</chain>

    <chain name="accept_input_udp">
</chain>

    <chain name="accept_input_icmp">
</chain>

    <chain name="invalid_forwarding">

```

```
        <rule target="REJECT">
            service_smtp;
        </rule>
    </chain>

    <chain name="accept_forward_all">
        <rule target="invalid_forwarding">
            internal_traffic;
        </rule>
    </chain>

    <chain name="accept_forward_tcp">
        <rule target="accept_traffic">
            valid_internal_traffic;
        </rule>
    </chain>

    <chain name="accept_forward_udp">
        <rule target="accept_traffic">
            valid_internal_traffic;
        </rule>
    </chain>

    <chain name="accept_forward_icmp">
        <rule target="accept_traffic">
            valid_internal_traffic;
        </rule>
    </chain>

    <chain name="accept_output_all">
        <rule target="accept_traffic">
            blank;
        </rule>
    </chain>

    <chain name="accept_output_tcp">
    </chain>

    <chain name="accept_output_udp">
    </chain>

    <chain name="accept_output_icmp">
    </chain>

#
# SYSTEM INPUT RULES - CUSTOMIZE ABOVE
#
    <chain name="accept_input_all">
        <rule target="accept_traffic">
            local_iface;
    </chain>
```

```
        </rule>
</chain>
<chain name="accept_input_tcp">
    <rule target="accept_traffic">
        service_ssh;
    </rule>
</chain>

<chain name="accept_input_udp">
</chain>

<chain name="accept_input_icmp">
    <rule target="accept_traffic">
        valid_icmp_packets;
        traceroute_packets;
    </rule>
</chain>

#
# SYSTEM FORWARD RULES - CUSTOMIZE ABOVE
#
<chain name="accept_forward_all">
</chain>

<chain name="accept_forward_tcp">
</chain>

<chain name="accept_forward_udp">
</chain>

<chain name="accept_forward_icmp">
</chain>

#
# SYSTEM LOGGING RULES
#
<chain name="log_input">
    <rule target='LOG --log-prefix "FW:filter:INPUT "'>
        30_per_min;
    </rule>
</chain>

<chain name="log_forward">
    <rule target='LOG --log-prefix "FW:filter:FORWARD "'>
        30_per_min;
    </rule>
</chain>

<chain name="log_output">
    <rule target='LOG --log-prefix "FW:filter:OUTPUT "'>
        30_per_min;
    </rule>
</chain>
```

```
<chain name="log_drop_packets">
  <rule target='LOG --log-prefix "FW:filter:check_packets "'>
    30_per_min;
  </rule>
  <rule target="DROP">
    blank;
  </rule>
</chain>

#
# MAIN SYSTEM RULES
#

# Remove bwmd rule if you not using it
<chain name="accept_traffic">
  <rule target="ACCEPT">
    blank;
  </rule>
</chain>

<chain name="accept_state">
  <rule target="accept_traffic">
    valid_connections;
  </rule>
</chain>

<chain name="check_packets">
  <rule target="log_drop_packets">
    invalid_tcp_packets;
  </rule>
</chain>

#
# MAIN SYSTEM CHAINS
#
<chain name="INPUT" default="DROP">
  <rule target="check_packets">
    blank;
  </rule>
  <rule target="accept_state">
    blank;
  </rule>
  <rule target="accept_input_all">
    blank;
  </rule>
  <rule target="accept_input_tcp">
    syn_packets;
  </rule>
  <rule target="accept_input_udp">
    udp_packets;
  </rule>
  <rule target="accept_input_icmp">
    icmp_packets;
  </rule>
  <rule target="log_input">
```

```
        blank;
    </rule>
</chain>

<chain name="FORWARD" default="DROP">
    <rule target="check_packets">
        blank;
    </rule>
    <rule target="accept_state">
        blank;
    </rule>
    <rule target="accept_forward_all">
        blank;
    </rule>
    <rule target="accept_forward_tcp">
        syn_packets;
    </rule>
    <rule target="accept_forward_udp">
        udp_packets;
    </rule>
    <rule target="accept_forward_icmp">
        icmp_packets;
    </rule>
    <rule target="log_forward">
        blank;
    </rule>
</chain>

<chain name="OUTPUT" default="DROP">
    <rule target="check_packets">
        blank;
    </rule>
    <rule target="accept_state">
        blank;
    </rule>
    <rule target="accept_output_all">
        blank;
    </rule>
    <rule target="accept_output_tcp">
        syn_packets;
    </rule>
    <rule target="accept_output_udp">
        udp_packets;
    </rule>
    <rule target="accept_output_icmp">
        icmp_packets;
    </rule>
    <rule target="log_output">
        blank;
    </rule>
</chain>
</table>
</acl>

<nat>
    <snat>
        <rule to-src="your.external.ip.here">
            nat_internal_traffic;
```

```
        </rule>
    </snat>
</nat>

</firewall>
```

Appendix A Copying This Manual

A.1 GNU Free Documentation License

Version 1.2, November 2002

Copyright © 2000,2001,2002 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.

0. PREAMBLE

The purpose of this License is to make a manual, textbook, or other functional and useful 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 non-commercially. 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.

1. APPLICABILITY AND DEFINITIONS

This License applies to any manual or other work, in any medium, that contains a notice placed by the copyright holder saying it can be distributed under the terms of this License. Such a notice grants a world-wide, royalty-free license, unlimited in duration, to use that work under the conditions stated herein. The “Document”, below, refers to any such manual or work. Any member of the public is a licensee, and is addressed as “you”. You accept the license if you copy, modify or distribute the work in a way requiring permission under copyright law.

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. (Thus, 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. If a section does not fit the above definition of Secondary then it is not allowed to be designated as Invariant. The Document may contain zero Invariant Sections. If the Document does not identify any Invariant Sections then there are none.

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 Front-Cover Text may be at most 5 words, and a Back-Cover Text may be at most 25 words.

A “Transparent” copy of the Document means a machine-readable copy, represented in a format whose specification is available to the general public, that is suitable for revising the document 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, or absence of markup, has been arranged to thwart or discourage subsequent modification by readers is not Transparent. An image format is not Transparent if used for any substantial amount of text. 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, PostScript or PDF designed for human modification. Examples of transparent image formats include PNG, XCF and JPG. Opaque formats include 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, PostScript or PDF 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.

A section “Entitled XYZ” means a named subunit of the Document whose title either is precisely XYZ or contains XYZ in parentheses following text that translates XYZ in another language. (Here XYZ stands for a specific section name mentioned below, such as “Acknowledgements”, “Dedications”, “Endorsements”, or “History”.) To “Preserve the Title” of such a section when you modify the Document means that it remains a section “Entitled XYZ” according to this definition.

The Document may include Warranty Disclaimers next to the notice which states that this License applies to the Document. These Warranty Disclaimers are considered to be included by reference in this License, but only as regards disclaiming warranties: any other implication that these Warranty Disclaimers may have is void and has no effect on the meaning of this License.

2. 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.

3. COPYING IN QUANTITY

If you publish printed copies (or copies in media that commonly have printed covers) 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 computer-network location from which the general network-using public has access to download using public-standard network protocols a complete Transparent copy of the Document, free of added material. 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.

4. 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 fewer than five), unless they release you from this requirement.
 - 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", Preserve 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. For any section Entitled "Acknowledgements" or "Dedications", Preserve the Title of the section, 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 to be Entitled "Endorsements" or to conflict in title with any Invariant Section.
 - O. Preserve any Warranty Disclaimers.

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.

5. 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, and that you preserve all their Warranty Disclaimers.

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.”

6. 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.

7. 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, is called an “aggregate” if the copyright resulting from the compilation is not used to limit the legal rights of the compilation’s users beyond what the individual works permit. When the Document is included in an aggregate, this License does not apply to the other works in the aggregate which 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 half of the entire aggregate, the Document’s Cover Texts may be placed on covers that bracket the Document within the aggregate, or the electronic equivalent of covers if the Document is in electronic form. Otherwise they must appear on printed covers that bracket the whole aggregate.

8. 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, and all the license notices in the Document, and any Warranty Disclaimers, provided that you also include the original English version of this License and the original versions of those notices and disclaimers. In case of a disagreement between the translation and the original version of this License or a notice or disclaimer, the original version will prevail.

If a section in the Document is Entitled “Acknowledgements”, “Dedications”, or “History”, the requirement (section 4) to Preserve its Title (section 1) will typically require changing the actual title.

9. 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.

10. 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.

A.1.1 ADDENDUM: 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:

```
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 1.2  
or any later version published by the Free Software Foundation;  
with no Invariant Sections, no Front-Cover Texts, and no Back-Cover  
Texts. A copy of the license is included in the section entitled ‘GNU  
Free Documentation License’.
```

If you have Invariant Sections, Front-Cover Texts and Back-Cover Texts, replace the “with...Texts.” line with this:

```
with the Invariant Sections being list their titles, with  
the Front-Cover Texts being list, and with the Back-Cover Texts  
being list.
```

If you have Invariant Sections without Cover Texts, or some other combination of the three, merge those two alternatives to suit the situation.

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.

Index

A

acl 4
address 4

B

burst-rate 12
burst-threshold 12
bwm_firewall 15
bwm_graph 18
bwmd 15

C

chain 5
class 3
cmd-line 6

D

default 5
dnat 8
dst 4
dst-iface 4
dst-port 4

F

FDL, GNU Free Documentation License 30
features 1
firewall 4, 15
flow 11

G

global 3
graphing 18
group 14

I

iptables 16
iptables-restore 15

L

load 3

M

mangle 16
masq 9
max-rate 11
modules 3

N

nat 7
nfmark 10, 11, 13, 16

P

prio 13
prio-classifier 12
proto 4

Q

queue 13, 16
queue-len 11
queue-size 11

R

report-timeout 12, 14
rrd 18
rrdtool 18
rule 6, 8, 9

S

shaping 10
snat 7
src 4
src-iface 4
src-port 4
stats-len 11, 14

T

table 5
target 6
to-dst 8
to-ports 9
to-src 8
traffic priority 12