



Full credit is given to the above companies including the OS that this PDF file was generated!

Red Hat Enterprise Linux Release 9.2 Manual Pages on 'tc-police.8' command

\$ man tc-police.8

Policing action in tc(8) Linux Policing action in tc(8)

NAME

police - policing action

SYNOPSIS

```
tc ... action police [ rate RATE burst BYTES[/BYTES] ] [ pkts_rate RATE
pkts_burst PACKETS ] [ mtu BYTES[/BYTES] ] [ peakrate RATE ] [
overhead BYTES ] [ linklayer TYPE ] [ CONTROL ]
```

```
tc ... filter ... [ estimator SAMPLE AVERAGE ] action police avrate
RATE [ CONTROL ]
```

CONTROL := conform-exceed EXCEEDACT[/NOTEXCEEDACT

EXCEEDACT/NOTEXCEEDACT := { pipe | ok | reclassify | drop | continue |
goto chain CHAIN_INDEX }

DESCRIPTION

The police action allows limiting of the byte or packet rate of traffic matched by the filter it is attached to.

There are two different algorithms available to measure the byte rate:

The first one uses an internal dual token bucket and is configured using the rate, burst, mtu, peakrate, overhead and linklayer parameters. The second one uses an in-kernel sampling mechanism. It can be fine-tuned using the estimator filter parameter.

There is one algorithm available to measure packet rate and it is similar to the first algorithm described for byte rate. It is configured using the pkt_rate and pkt_burst parameters.

At least one of the rate and pkt_rate parameters must be configured.

OPTIONS

rate RATE

The maximum byte rate of packets passing this action. Those exceeding it will be treated as defined by the conform-exceed option.

burst BYTES[/BYTES]

Set the maximum allowed burst in bytes, optionally followed by a slash (/) sign and cell size which must be a power of 2.

pkt_rate RATE

The maximum packet rate or packets passing this action. Those exceeding it will be treated as defined by the conform-exceed option.

pkt_burst PACKETS

Set the maximum allowed burst in packets.

mtu BYTES[/BYTES]

This is the maximum packet size handled by the policer (larger ones will be handled like they exceeded the configured rate).

Setting this value correctly will improve the scheduler's precision.

sion. Value formatting is identical to burst above. Defaults to unlimited.

peakrate RATE

Set the maximum bucket depletion rate, exceeding rate.

avrate RATE

Make use of an in-kernel bandwidth rate estimator and match the given RATE against it.

overhead BYTES

Account for protocol overhead of encapsulating output devices when computing rate and peakrate.

linklayer TYPE

Specify the link layer type. TYPE may be one of ethernet (the default), atm or adsl (which are synonyms). It is used to align the precomputed rate tables to ATM cell sizes, for ethernet no action is taken.

estimator SAMPLE AVERAGE

Fine-tune the in-kernel packet rate estimator. SAMPLE and AVERAGE are time values and control the frequency in which samples are taken and over what timespan an average is built.

conform-exceed EXCEEDACT[/NOTEXCEEDACT]

Define how to handle packets which exceed or conform the configured bandwidth limit. Possible values are:

continue

Don't do anything, just continue with the next action in line.

drop Drop the packet immediately.

shot This is a synonym to drop.

ok Accept the packet. This is the default for conforming packets.

pass This is a synonym to ok.

reclassify

Treat the packet as non-matching to the filter this action is attached to and continue with the next filter in line (if any). This is the default for exceeding packets.

pipe Pass the packet to the next action in line.

EXAMPLES

A typical application of the police action is to enforce ingress traffic rate by dropping exceeding packets. Although better done on the sender's side, especially in scenarios with lack of peer control (e.g. with dial-up providers) this is often the best one can do in order to keep latencies low under high load. The following establishes input bandwidth policing to 1mbit/s using the ingress qdisc and u32 filter:

```
# tc qdisc add dev eth0 handle ffff: ingress
# tc filter add dev eth0 parent ffff: u32 \
    match u32 0 0 \
    police rate 1mbit burst 100k
```

As an action can not live on it's own, there always has to be a filter involved as link between qdisc and action. The example above uses u32 for that, which is configured to effectively match any packet (passing it to the police action thereby).

SEE ALSO

tc(8)

iproute2

20 Jan 2015

Policing action in tc(8)