

NAME

red – Random Early Detection

SYNOPSIS

tc qdisc ... red limit bytes **min** bytes **max** bytes **avpkt** bytes **burst** packets [**ecn**] [**bandwidth** rate] **probability** chance

DESCRIPTION

Random Early Detection is a classless qdisc which manages its queue size smartly. Regular queues simply drop packets from the tail when they are full, which may not be the optimal behaviour. RED also performs tail drop, but does so in a more gradual way.

Once the queue hits a certain average length, packets enqueued have a configurable chance of being marked (which may mean dropped). This chance increases linearly up to a point called the **max** average queue length, although the queue might get bigger.

This has a host of benefits over simple taildrop, while not being processor intensive. It prevents synchronous retransmits after a burst in traffic, which cause further retransmits, etc.

The goal is to have a small queue size, which is good for interactivity while not disturbing TCP/IP traffic with too many sudden drops after a burst of traffic.

Depending on if ECN is configured, marking either means dropping or purely marking a packet as overlimit.

ALGORITHM

The average queue size is used for determining the marking probability. This is calculated using an Exponential Weighted Moving Average, which can be more or less sensitive to bursts.

When the average queue size is below **min** bytes, no packet will ever be marked. When it exceeds **min**, the probability of doing so climbs linearly up to **probability**, until the average queue size hits **max** bytes. Because **probability** is normally not set to 100%, the queue size might conceivably rise above **max** bytes, so the **limit** parameter is provided to set a hard maximum for the size of the queue.

PARAMETERS

- min** Average queue size at which marking becomes a possibility.
- max** At this average queue size, the marking probability is maximal. Should be at least twice **min** to prevent synchronous retransmits, higher for low **min**.
- probability** Maximum probability for marking, specified as a floating point number from 0.0 to 1.0. Suggested values are 0.01 or 0.02 (1 or 2%, respectively).
- limit** Hard limit on the real (not average) queue size in bytes. Further packets are dropped. Should be set higher than **max+burst**. It is advised to set this a few times higher than **max**.
- burst** Used for determining how fast the average queue size is influenced by the real queue size. Larger values make the calculation more sluggish, allowing longer bursts of traffic before marking starts. Real life experiments support the following guideline: $(\text{min} + \text{min} + \text{max}) / (3 * \text{avpkt})$.
- avpkt** Specified in bytes. Used with **burst** to determine the time constant for average queue size calculations. 1000 is a good value.

bandwidth

This rate is used for calculating the average queue size after some idle time. Should be set to the bandwidth of your interface. Does not mean that RED will shape for you! Optional.

ecn

As mentioned before, RED can either 'mark' or 'drop'. Explicit Congestion Notification allows RED to notify remote hosts that their rate exceeds the amount of bandwidth available. Non-ECN capable hosts can only be notified by dropping a packet. If this parameter is specified, packets which indicate that their hosts honor ECN will only be marked and not dropped, unless the queue size hits **limit** bytes. Needs a tc binary with RED support compiled in. Recommended.

SEE ALSO

tc(8)

SOURCES

- o Floyd, S., and Jacobson, V., Random Early Detection gateways for Congestion Avoidance. <http://www.aciri.org/floyd/papers/red/red.html>
- o Some changes to the algorithm by Alexey N. Kuznetsov.

AUTHORS

Alexey N. Kuznetsov, <kuznet@ms2.inr.ac.ru>, Alexey Makarenko <makar@phoenix.kharkov.ua>, J Hadi Salim <hadi@nortelnetworks.com>. This manpage maintained by bert hubert <ahu@ds9a.nl>