



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-codel.8' command

\$ man tc-codel.8

CoDel(8) Linux CoDel(8)

NAME

CoDel - Controlled-Delay Active Queue Management algorithm

SYNOPSIS

```
tc qdisc ... codel [ limit PACKETS ] [ target TIME ] [ interval TIME ]
[ ecn | noecn ] [ ce_threshold TIME ]
```

DESCRIPTION

CoDel (pronounced "coddle") is an adaptive "no-knobs" active queue management algorithm (AQM) scheme that was developed to address the shortcomings of RED and its variants. It was developed with the following goals in mind:

- o It should be parameterless.
- o It should keep delays low while permitting bursts of traffic.
- o It should control delay.
- o It should adapt dynamically to changing link rates with no impact on utilization.
- o It should be simple and efficient and should scale from simple to complex routers.

ALGORITHM

CoDel comes with three major innovations. Instead of using queue size or queue average, it uses the local minimum queue as a measure of the standing/persistent queue. Second, it uses a single state-tracking variable of the minimum delay to see where it is relative to the stand?

ing queue delay. Third, instead of measuring queue size in bytes or packets, it is measured in packet-sojourn time in the queue.

CoDel measures the minimum local queue delay (i.e. standing queue delay) and compares it to the value of the given acceptable queue delay target. As long as the minimum queue delay is less than target or the buffer contains fewer than MTU worth of bytes, packets are not dropped.

CoDel enters a dropping mode when the minimum queue delay has exceeded target for a time greater than interval. In this mode, packets are dropped at different drop times which is set by a control law. The control law ensures that the packet drops cause a linear change in the throughput. Once the minimum delay goes below target, packets are no longer dropped.

Additional details can be found in the paper cited below.

PARAMETERS

limit

hard limit on the real queue size. When this limit is reached, incoming packets are dropped. If the value is lowered, packets are dropped so that the new limit is met. Default is 1000 packets.

target

is the acceptable minimum standing/persistent queue delay. This minimum delay is identified by tracking the local minimum queue delay that packets experience. Default and recommended value is 5ms.

interval

is used to ensure that the measured minimum delay does not become too stale. The minimum delay must be experienced in the last epoch of length interval. It should be set on the order of the worst-case RTT through the bottleneck to give endpoints sufficient time to react. Default value is 100ms.

ecn | noecn

can be used to mark packets instead of dropping them. If ecn has been enabled, noecn can be used to turn it off and vice-a-versa. By default, ecn is turned off.

ce_threshold

sets a threshold above which all packets are marked with ECN Congestion Experienced. This is useful for DCTCP-style congestion control algorithms that require marking at very shallow queueing thresholds.

EXAMPLES

```
# tc qdisc add dev eth0 root codel
# tc -s qdisc show
qdisc codel 801b: dev eth0 root refcnt 2 limit 1000p target 5.0ms
interval 100.0ms
Sent 245801662 bytes 275853 pkt (dropped 0, overlimits 0 requeues
24)
backlog 0b 0p requeues 24
count 0 lastcount 0 ldelay 2us drop_next 0us
maxpacket 7306 ecn_mark 0 drop_overlimit 0
# tc qdisc add dev eth0 root codel limit 100 target 4ms interval 30ms
ecn
# tc -s qdisc show
qdisc codel 801c: dev eth0 root refcnt 2 limit 100p target 4.0ms in?
interval 30.0ms ecn
Sent 237573074 bytes 268561 pkt (dropped 0, overlimits 0 requeues
5)
backlog 0b 0p requeues 5
count 0 lastcount 0 ldelay 76us drop_next 0us
maxpacket 2962 ecn_mark 0 drop_overlimit 0
```

SEE ALSO

tc(8), tc-red(8)

SOURCES

o Kathleen Nichols and Van Jacobson, "Controlling Queue Delay", ACM Queue, <http://queue.acm.org/detail.cfm?id=2209336>

AUTHORS

CoDel was implemented by Eric Dumazet and David Taht. This manpage was written by Vijay Subramanian. Please report corrections to the Linux Networking mailing list <netdev@vger.kernel.org>.