

NAME

CBS – Credit Based Shaper (CBS) Qdisc

SYNOPSIS

```
tc qdisc ... dev dev parent classid [ handle major: ] cbs idleslope idleslope sendslope sendslope hicredit
hicredit locredit locredit [ offload 0|1 ]
```

DESCRIPTION

The CBS (Credit Based Shaper) qdisc implements the shaping algorithm defined by the IEEE 802.1Q-2014 Section 8.6.8.2, which applies a well defined rate limiting method to the traffic.

This queueing discipline is intended to be used by TSN (Time Sensitive Networking) applications, the CBS parameters are derived directly by what is described by the Annex L of the IEEE 802.1Q-2014 Specification. The algorithm and how it affects the latency are detailed there.

CBS is meant to be installed under another qdisc that maps packet flows to traffic classes, one example is **mqprio(8)**.

PARAMETERS

idleslope

Idleslope is the rate of credits that is accumulated (in kilobits per second) when there is at least one packet waiting for transmission. Packets are transmitted when the current value of credits is equal or greater than zero. When there is no packet to be transmitted the amount of credits is set to zero. This is the main tunable of the CBS algorithm and represents the bandwidth that will be consumed. Note that when calculating idleslope, the entire packet size must be considered, including headers from all layers (i.e. MAC framing and any overhead from the physical layer), as described by IEEE 802.1Q-2014 section 34.4.

As an example, for an ethernet frame carrying 284 bytes of payload, and with no VLAN tags, you must add 14 bytes for the Ethernet headers, 4 bytes for the Frame check sequence (CRC), and 20 bytes for the L1 overhead: 12 bytes of interpacket gap, 7 bytes of preamble and 1 byte of start of frame delimiter. That results in 322 bytes for the total packet size, which is then used for calculating the idleslope.

sendslope

Sendslope is the rate of credits that is depleted (it should be a negative number of kilobits per second) when a transmission is occurring. It can be calculated as follows, (IEEE 802.1Q-2014 Section 8.6.8.2 item g):

$$\text{sendslope} = \text{idleslope} - \text{port_transmit_rate}$$

hicredit Hicredit defines the maximum amount of credits (in bytes) that can be accumulated. Hicredit depends on the characteristics of interfering traffic, 'max_interference_size' is the maximum size of any burst of traffic that can delay the transmission of a frame that is available for transmission for this traffic class, (IEEE 802.1Q-2014 Annex L, Equation L-3):

$$\text{hicredit} = \text{max_interference_size} * (\text{idleslope} / \text{port_transmit_rate})$$

locredit Locredit is the minimum amount of credits that can be reached. It is a function of the traffic flowing through this qdisc (IEEE 802.1Q-2014 Annex L, Equation L-2):

$$\text{locredit} = \text{max_frame_size} * (\text{sendslope} / \text{port_transmit_rate})$$

offload When **offload** is 1, **cbs(8)** will try to configure the network interface so the CBS algorithm runs in the controller. The default is 0.

EXAMPLES

CBS is used to enforce a Quality of Service by limiting the data rate of a traffic class, to separate packets into traffic classes the user may choose **mqprio(8)**, and configure it like this:

```
# tc qdisc add dev eth0 handle 100: parent root mqprio num_tc 3 \  
    map 2 2 1 0 2 2 2 2 2 2 2 2 2 2 2 2 \  
    queues 1@0 1@1 2@2 \  
    hw 0
```

To replace the current queuing discipline by CBS in the current queueing discipline connected to traffic class number 0, issue:

```
# tc qdisc replace dev eth0 parent 100:4 cbs \  
    locredit -1470 hicredit 30 sendslope -980000 idleslope 20000
```

These values are obtained from the following parameters, idleslope is 20mbit/s, the transmission rate is 1Gbit/s and the maximum interfering frame size is 1500 bytes.

AUTHORS

Vinicius Costa Gomes <vinicius.gomes@intel.com>