

NAME

skbprio – SKB Priority Queue

SYNOPSIS

tc qdisc ... add skbprio [limit packets]

DESCRIPTION

SKB Priority Queue is a queueing discipline intended to prioritize the most important packets during a denial-of-service (**DoS**) attack. The priority of a packet is given by **skb->priority** , where a higher value places the packet closer to the exit of the queue. When the queue is full, the lowest priority packet in the queue is dropped to make room for the packet to be added if it has higher priority. If the packet to be added has lower priority than all packets in the queue, it is dropped.

Without SKB priority queue, queue length limits must be imposed on individual sub-queues, and there is no straightforward way to enforce a global queue length limit across all priorities. SKBprio queue enforces a global queue length limit while not restricting the lengths of individual sub-queues.

While SKB Priority Queue is agnostic to how **skb->priority** is assigned. A typical use case is to copy the 6-bit DS field of IPv4 and IPv6 packets using **tc-skbedit(8)**. If **skb->priority** is greater or equal to 64, the priority is assumed to be 63. Priorities less than 64 are taken at face value.

SKB Priority Queue enables routers to locally decide which packets to drop under a DoS attack. Priorities should be assigned to packets such that the higher the priority, the more expected behavior a source shows. So sources have an incentive to play by the rules.

ALGORITHM

Skbprio maintains 64 lists (priorities go from 0 to 63). When a packet is enqueued, it gets inserted at the **tail** of its priority list. When a packet needs to be sent out to the network, it is taken from the head of the highest priority list. When the queue is full, the packet at the tail of the lowest priority list is dropped to serve the ingress packet - if it is of higher priority, otherwise the ingress packet is dropped. This algorithm allocates as much bandwidth as possible to high priority packets, while only servicing low priority packets when there is enough bandwidth.

PARAMETERS

limit Maximum queue size specified in packets. It defaults to 64. The range for this parameter is [0, UINT32_MAX].

SEE ALSO

tc-prio(8), **tc-skbedit(8)**

AUTHORS

Nishanth Devarajan <devarajn@uci.edu>, Michel Machado <michel@digirati.com.br>

This manpage maintained by Bert Hubert <ahu@ds9a.nl>