

NAME

ctinfo – tc connmark processing action

SYNOPSIS

```
tc ... action ctinfo [ dscp MASK [STATEMASK] ] [ cpmark [MASK] ] [ zone ZONE ] [ CONTROL ] [ index <INDEX> ]
```

DESCRIPTION

CTINFO (Conntrack Information) is a tc action for retrieving data from conntrack marks into various fields. At present it has two independent processing modes which may be viewed as sub-functions.

DSCP mode copies a DSCP stored in conntrack's connmark into the IPv4/v6 diffserv field. The copying may conditionally occur based on a flag also stored in the connmark. DSCP mode was designed to assist in restoring packet classifications on ingress, classifications which may then be used by qdiscs such as CAKE. It may be used in any circumstance where ingress classification needs to be maintained across links that otherwise bleach or remap according to their own policies.

CPMARK (copymark) mode copies the conntrack connmark into the packet's mark field. Without additional parameters it is functionally completely equivalent to the existing connmark action. An optional mask may be specified to mask which bits of the connmark are restored. This may be useful when DSCP and CPMARK modes are combined.

Simple statistics (tc -s) on DSCP restores and CPMARK copies are maintained where values for set indicate a count of packets altered for that mode. DSCP includes an error count where the destination packet's diffserv field was unwritable.

PARAMETERS**DSCP mode parameters:**

mask A mask of 6 contiguous bits indicating where the DSCP value is located in the 32 bit conntrack mark field. A mask must be provided for this mode. mask is a 32 bit unsigned value.

statemask

A mask of at least 1 bit indicating where a conditional restore flag is located in the 32 bit conntrack mark field. The statemask bit/s must NOT overlap the mask bits. The DSCP will be restored if the conntrack mark logically ANDed with the statemask yields a non-zero result. statemask is an optional unsigned 32 bit value.

CPMARK mode parameters:

mask Store the logically ANDed result of conntrack mark and mask into the packet's mark field. Default is 0xffffffff i.e. the whole mark field. mask is an optional unsigned 32 bit value

Overall action parameters:

zone Specify the conntrack zone when doing conntrack lookups for packets. zone is a 16bit unsigned decimal value. Default is 0.

CONTROL

The following keywords allow to control how the tree of qdisc, classes, filters and actions is further traversed after this action.

reclassify

Restart with the first filter in the current list.

pipe

Continue with the next action attached to the same filter.

drop

Drop the packet.

shot

synonym for **drop**

continue

Continue classification with the next filter in line.

- pass** Finish classification process and return to calling qdisc for further packet processing. This is the default.
- index** Specify an index for this action in order to being able to identify it in later commands. index is a 32bit unsigned decimal value.

EXAMPLES

Example showing conditional restoration of DSCP on ingress via an IFB

```
#Set up the IFB interface
tc qdisc add dev ifb4eth0 handle ffff: ingress

#Put CAKE qdisc on it
tc qdisc add dev ifb4eth0 root cake bandwidth 40mbit

#Set interface UP
ip link set dev ifb4eth0 up

#Add 2 actions, ctinfo to restore dscp & mirred to redirect the packets to IFB
tc filter add dev eth0 parent ffff: protocol all prio 10 u32 \
    match u32 0 0 flowid 1:1 action \
    ctinfo dscp 0xfc000000 0x01000000 \
    mirred egress redirect dev ifb4eth0

tc -s qdisc show dev eth0 ingress

filter parent ffff: protocol all pref 10 u32 chain 0
filter parent ffff: protocol all pref 10 u32 chain 0 fh 800: ht divisor 1
filter parent ffff: protocol all pref 10 u32 chain 0 fh 800::800 order 2048 k
    match 00000000/00000000 at 0
    action order 1: ctinfo zone 0 pipe
    index 2 ref 1 bind 1 dscp 0xfc000000 0x01000000 installed 72 sec used 0 se
    Action statistics:
    Sent 658484 bytes 1833 pkt (dropped 0, overlimits 0 requeues 0)
    backlog 0b 0p requeues 0

    action order 2: mirred (Egress Redirect to device ifb4eth0) stolen
    index 1 ref 1 bind 1 installed 72 sec used 0 sec
    Action statistics:
    Sent 658484 bytes 1833 pkt (dropped 0, overlimits 0 requeues 0)
    backlog 0b 0p requeues 0
```

Example showing conditional restoration of DSCP on egress

This may appear nonsensical since iptables marking of egress packets is easy to achieve, however the iptables flow classification rules may be extensive and so some sort of set once and forget may be useful especially on cpu constrained devices.

```
# Send unmarked connections to a marking chain which needs to store a DSCP
and set statemask bit in the connmark
iptables -t mangle -A POSTROUTING -o eth0 -m connmark \
    --mark 0x00000000/0x01000000 -g CLASS_MARKING_CHAIN

# Apply marked DSCP to the packets
tc filter add dev eth0 protocol all prio 10 u32 \
    match u32 0 0 flowid 1:1 action \
```

```
ctinfo dscp 0xfc000000 0x01000000

tc -s filter show dev eth0
filter parent 800e: protocol all pref 10 u32 chain 0
filter parent 800e: protocol all pref 10 u32 chain 0 fh 800: ht divisor 1
filter parent 800e: protocol all pref 10 u32 chain 0 fh 800::800 order 2048 k
match 00000000/00000000 at 0
  action order 1: ctinfo zone 0 pipe
  index 1 ref 1 bind 1 dscp 0xfc000000 0x01000000 installed 7414 sec used 0
Action statistics:
Sent 32890260 bytes 120441 pkt (dropped 0, overlimits 0 requeues 0)
backlog 0b 0p requeues 0
```

SEE ALSO

tc(8), **tc-cake(8)** **tc-connmark(8)** **tc-mirred(8)**

AUTHORS

ctinfo was written by Kevin Darbyshire-Bryant.