NAME

ip-nexthop − nexthop object management

SYNOPSIS

ip [*ip-OPTIONS*] **nexthop** { *COMMAND* | **help** }

ip nexthop { **show** | **flush** } *SELECTOR*

ip nexthop { **add** | **replace** } id *ID NH*

ip nexthop { **get** | **del** } id *ID*

SELECTOR := [**id** *ID*] [**dev** *DEV*] [**vrf** *NAME*] [**master** *DEV*] [**groups**]

NH := { **blackhole** | [**via** *ADDRESS*] [**dev** *DEV*] [**onlink**] [**encap** *ENCAP*] | **group** *GROUP* }

ENCAP := [*ENCAP_MPLS*]

ENCAP_MPLS := **mpls** [*LABEL*] [**ttl** *TTL*]

GROUP := **id**[,**weight**[/...]

DESCRIPTION

ip nexthop is used to manipulate entries in the kernel's nexthop tables.

ip nexthop add id ID
: add new nexthop entry

ip nexthop replace id ID
: change the configuration of a nexthop or add new one

via *[FAMILY] ADDRESS*
: the address of the nexthop router, in the address family FAMILY. Address family must match address family of nexthop instance.

dev *NAME*
: is the output device.

onlink pretend that the nexthop is directly attached to this link, even if it does not match any interface prefix.

encap *ENCAPTYPE ENCAPHDR*
: attach tunnel encapsulation attributes to this route.

ENCAPTYPE is a string specifying the supported encapsulation type. Namely:

mpls - encapsulation type MPLS

ENCAPHDR is a set of encapsulation attributes specific to the *ENCAPTYPE.*

mpls
: *MPLSLABEL* - mpls label stack with labels separated by **/**

ttl *TTL* - TTL to use for MPLS header or 0 to inherit from IP header

group *GROUP*
create a nexthop group. Group specification is id with an optional weight (id,weight) and a '/' as a separator between entries.

blackhole
create a blackhole nexthop

ip nexthop delete id ID
delete nexthop with given id.

ip nexthop show
show the contents of the nexthop table or the nexthops selected by some criteria.

dev *DEV*
show the nexthops using the given device.

vrf *NAME*
show the nexthops using devices associated with the vrf name

master *DEV*
show the nexthops using devices enslaved to given master device

groups show only nexthop groups

ip nexthop flush
flushes nexthops selected by some criteria. Criteria options are the same as show.

ip nexthop get id ID
get a single nexthop by id

EXAMPLES

ip nexthop ls
Show all nexthop entries in the kernel.

ip nexthop add id 1 via 192.168.1.1 dev eth0
Adds an IPv4 nexthop with id 1 using the gateway 192.168.1.1 out device eth0.

ip nexthop add id 2 encap mpls 200/300 via 10.1.1.1 dev eth0
Adds an IPv4 nexthop with mpls encapsulation attributes attached to it.

ip nexthop add id 3 group 1/2
Adds a nexthop with id 3. The nexthop is a group using nexthops with ids 1 and 2 at equal weight.

ip nexthop add id 4 group 1,5/2,11
Adds a nexthop with id 4. The nexthop is a group using nexthops with ids 1 and 2 with nexthop 1 at weight 5 and nexthop 2 at weight 11.

SEE ALSO

ip(8)

AUTHOR

Original Manpage by David Ahern <dsahern@kernel.org>