

NAME

ip-link - network device configuration

SYNOPSIS

ip [*OPTIONS*] **link** { *COMMAND* | **help** }

OPTIONS := { **-V**[*ersion*] | **-s**[*tatistics*] | **-r**[*esolve*] | **-f**[*amily*] { **inet** | **inet6** | **ipx** | **dnet** | **link** } | **-o**[*neline*] }

ip link add [**link** *DEVICE*] [**name**] *NAME*
 [**txqueuelen** *PACKETS*]
 [**address** *LLADDR*] [**broadcast** *LLADDR*]
 [**mtu** *MTU*] [**index** *IDX*]
 [**numtxqueues** *QUEUE_COUNT*] [**numrxqueues** *QUEUE_COUNT*]
type *TYPE* [*ARGS*]

TYPE := [**bridge** | **bond**] **can** | **dummy** | **hsr** | **ifb** | **ipoib** | **macvlan** | **macvtap** | **vcan** | **veth** | **vlan** | **vxlan** | **ip6tnl** | **ipip** | **sit** | **gre** | **gretap** | **ip6gre** | **ip6gretap**]

ip link delete *DEVICE* **type** *TYPE* [*ARGS*]

ip link set { *DEVICE* | **group** *GROUP* } { **up** | **down** | **arp** { **on** | **off** } | **promisc** { **on** | **off** } | **allmulticast** { **on** | **off** } | **dynamic** { **on** | **off** } | **multicast** { **on** | **off** } | **txqueuelen** *PACKETS* | **name** *NEWNAME* | **address** *LLADDR* | **broadcast** *LLADDR* | **mtu** *MTU* | **netns** *PID* | **netns** *NETNSNAME* | **alias** *NAME* | **vf** *NUM* [**mac** *LLADDR*] [**vlan** *VLANID* [**qos** *VLAN-QOS*]] [**rate** *TXRATE*] [**max_tx_rate** *TXRATE*] [**min_tx_rate** *TXRATE*] [**spoofchk** { **on** | **off** }] [**state** { **auto** | **enable** | **disable** }] | **master** *DEVICE* | **nomaster** }

ip link show [*DEVICE* | **group** *GROUP*]

DESCRIPTION

ip link add - add virtual link

link *DEVICE*

specifies the physical device to act operate on.

NAME specifies the name of the new virtual device.

TYPE specifies the type of the new device.

Link types:

bridge - Ethernet Bridge device

bond - Bonding device

dummy - Dummy network interface

ifb - Intermediate Functional Block device

ipoib - IP over Infiniband device

macvlan - Virtual interface base on link layer address (MAC)

macvtap - Virtual interface based on link layer address (MAC) and TAP.

vcan - Virtual Controller Area Network interface

veth - Virtual ethernet interface

vlan - 802.1q tagged virtual LAN interface

vxlan - Virtual eXtended LAN

ip6tnl - Virtual tunnel interface IPv4|IPv6 over IPv6

ipip - Virtual tunnel interface IPv4 over IPv4

sit - Virtual tunnel interface IPv6 over IPv4

gre - Virtual tunnel interface GRE over IPv4

gretap - Virtual L2 tunnel interface GRE over IPv4

ip6gre - Virtual tunnel interface GRE over IPv6

ip6gretap - Virtual L2 tunnel interface GRE over IPv6

numtxqueues *QUEUE_COUNT*

specifies the number of transmit queues for new device.

numrxqueues *QUEUE_COUNT*

specifies the number of receive queues for new device.

index *IDX*

specifies the desired index of the new virtual device. The link creation fails, if the index is busy.

VXLAN Type Support

For a link of type *VXLAN* the following additional arguments are supported:

```
ip link add DEVICE type vxlan id ID [ dev PHYS_DEV ] [ { group | remote }
IPADDR ] [ local IPADDR ] [ ttl TTL ] [ tos TOS ] [ port MIN MAX ] [ [no]learning ]
[ [no]proxy ] [ [no]rsc ] [ [no]l2miss ] [ [no]l3miss ]
```

id *VNI* - specifies the VXLAN Network Identifier (or VXLAN Segment Identifier) to use.

dev *PHYS_DEV* - specifies the physical device to use for tunnel end-point communication.

group *IPADDR* - specifies the multicast IP address to join. This parameter cannot be specified with the **remote** parameter.

remote *IPADDR* - specifies the unicast destination IP address to use in outgoing packets when the destination link layer address is not known in the VXLAN device forwarding database. This parameter cannot be specified with the **group** parameter.

local *IPADDR* - specifies the source IP address to use in outgoing packets.

ttl *TTL* - specifies the TTL value to use in outgoing packets.

tos *TOS* - specifies the TOS value to use in outgoing packets.

port *MIN MAX* - specifies the range of port numbers to use as UDP

source ports to communicate to the remote VXLAN tunnel endpoint.

[no]learning - specifies if unknown source link layer addresses and IP addresses are entered into the VXLAN device forwarding database.

[no]rsc - specifies if route short circuit is turned on.

[no]proxy - specifies ARP proxy is turned on.

[no]l2miss - specifies if netlink LLADDR miss notifications are generated.

[no]l3miss - specifies if netlink IP ADDR miss notifications are generated.

IP6GRE/IP6GRETAP Type Support

For a link of type *IP6GRE/IP6GRETAP* the following additional arguments are supported:

```
ip link add DEVICE type { ip6gre | ip6gretap } remote ADDR local ADDR [ [i]o/seq ]
[ [i]o/key KEY ] [ [i]o/csum ] [ hoplimit TTL ] [ encaplimit ELIM ] [ tclass TCLASS ]
[ flowlabel FLOWLABEL ] [ dscp inherit ] [ dev PHYS_DEV ]
```

remote *ADDR* - specifies the remote IPv6 address of the tunnel.

local *ADDR* - specifies the fixed local IPv6 address for tunneled packets. It must be an address on another interface on this host.

[i]oseq - serialize packets. The **oseq** flag enables sequencing of outgoing packets. The **iseq** flag requires that all input packets are serialized.

[i]okey *KEY* - use keyed GRE with key *KEY*. *KEY* is either a number or an IPv4 address-like dotted quad. The **key** parameter specifies the same key to use in both directions. The **ikey** and **okey** parameters specify different keys for input and output.

[i]ocsum - generate/require checksums for tunneled packets. The **ocsum** flag calculates checksums for outgoing packets. The **icsum** flag requires that all input packets have the correct checksum. The **csum** flag is equivalent to the combination **icsum ocsum**.

hoplimit *TTL* - specifies Hop Limit value to use in outgoing packets.

encaplimit *ELIM* - specifies a fixed encapsulation limit. Default is 4.

flowlabel *FLOWLABEL* - specifies a fixed flowlabel.

tclass *TCLASS* - specifies the traffic class field on tunneled packets, which can be specified as either a two-digit hex value (e.g. c0) or a predefined string (e.g. internet). The value **inherit** causes the field to be copied from the original IP header. The values **inherit/STRING** or **inherit/00..ff** will set the field to *STRING* or *00..ff* when tunneling non-IP packets. The default value is 00.

ip link delete - delete virtual link

DEVICE specifies the virtual device to act operate on. *TYPE* specifies the type of the device.

dev *DEVICE*

specifies the physical device to act operate on.

ip link set - change device attributes**dev** *DEVICE*

DEVICE specifies network device to operate on. When configuring SR-IOV Virtual Function (VF) devices, this keyword should specify the associated Physical Function (PF) device.

group *GROUP*

GROUP has a dual role: If both group and dev are present, then move the device to the specified group. If only a group is specified, then the command operates on all devices in that group.

up and **down**

change the state of the device to **UP** or **DOWN**.

arp on or **arp off**

change the **NOARP** flag on the device.

multicast on or **multicast off**

change the **MULTICAST** flag on the device.

dynamic on or **dynamic off**

change the **DYNAMIC** flag on the device.

name *NAME*

change the name of the device. This operation is not recommended if the device is running or has some addresses already configured.

txqueuelen *NUMBER***txqlen** *NUMBER*

change the transmit queue length of the device.

mtu *NUMBER*

change the *MTU* of the device.

address *LLADDRESS*

change the station address of the interface.

broadcast *LLADDRESS***brd** *LLADDRESS***peer** *LLADDRESS*

change the link layer broadcast address or the peer address when the interface is *POINTOPOINT*.

netns *PID*

move the device to the network namespace associated with the process *PID*.

netns *NETNSNAME*

move the device to the network namespace associated with name *NETNSNAME*.

alias *NAME*

give the device a symbolic name for easy reference.

group *GROUP*

specify the group the device belongs to. The available groups are listed in file */etc/iproute2/group*.

vf *NUM*

specify a Virtual Function device to be configured. The associated PF device must be specified using the **dev** parameter.

mac *LLADDRESS* - change the station address for the specified VF. The **vf** parameter must be specified.

vlan *VLANID* - change the assigned VLAN for the specified VF. When specified, all traffic sent from the VF will be tagged with the specified VLAN ID. Incoming traffic will be filtered for the specified VLAN ID, and will have all VLAN tags stripped before being passed to the VF. Setting this parameter to 0 disables VLAN tagging and filtering. The **vf** parameter must be specified.

qos *VLAN-QOS* - assign VLAN QOS (priority) bits for the VLAN tag. When specified, all VLAN tags transmitted by the VF will include the specified priority bits in the VLAN tag. If not specified, the value is assumed to be 0. Both the **vf** and **vlan** parameters must be specified. Setting both **vlan** and **qos** as 0 disables VLAN tagging and filtering for the VF.

rate *TXRATE* -- change the allowed transmit bandwidth, in Mbps, for the specified VF. Setting this parameter to 0 disables rate limiting. **vf** parameter must be specified. Please use new API **max_tx_rate** option instead.

max_tx_rate *TXRATE* - change the allowed maximum transmit bandwidth, in Mbps, for the specified VF. **vf** parameter must be specified.

min_tx_rate *TXRATE* - change the allowed minimum transmit bandwidth, in Mbps, for the specified VF. Minimum TXRATE should be always <= Maximum TXRATE. **vf** parameter must be specified.

spoofchk *on|off* - turn packet spoof checking on or off for the specified VF.

state *auto|enable|disable* - set the virtual link state as seen by the specified VF. Setting to auto means a reflection of the PF link state, enable lets the VF to communicate with other VFs on this host even if the PF link state is down, disable causes the HW to drop any packets sent by the VF.

master *DEVICE*

set master device of the device (enslave device).

nomaster

unset master device of the device (release device).

Warning: If multiple parameter changes are requested, **ip** aborts immediately after any of the changes have failed. This is the only case when **ip** can move the system to an unpredictable state. The solution is to avoid changing several parameters with one **ip link set** call.

ip link show - display device attributes**dev** *NAME* (default)

NAME specifies the network device to show. If this argument is omitted all devices in the default group are listed.

group *GROUP*
GROUP specifies what group of devices to show.

up only display running interfaces.

EXAMPLES

`ip link show`
Shows the state of all network interfaces on the system.

`ip link set dev ppp0 mtu 1400`
Change the MTU the ppp0 device.

`ip link add link eth0 name eth0.10 type vlan id 10`
Creates a new vlan device eth0.10 on device eth0.

`ip link delete dev eth0.10`
Removes vlan device.

SEE ALSO

[ip\(8\)](#)

AUTHOR

Original Manpage by Michail Litvak <mci@owl.openwall.com>