



*Full credit is given to the above companies including the OS that this PDF file was generated!*

## ***Red Hat Enterprise Linux Release 9.2 Manual Pages on 'ethtool.8' command***

**\$ man ethtool.8**

ETHTOOL(8)                    System Manager's Manual                    ETHTOOL(8)

NAME

ethtool - query or control network driver and hardware settings

SYNOPSIS

ethtool devname

ethtool -h|--help

ethtool --version

ethtool [--debug N] args

ethtool [--json] args

ethtool [-I | --include-statistics] args

ethtool --monitor [ command ] [ devname ]

ethtool -a|--show-pause devname

ethtool -A|--pause devname [autoneg on|off] [rx on|off] [tx on|off]

ethtool -c|--show-coalesce devname

ethtool -C|--coalesce devname [adaptive-rx on|off] [adaptive-tx on|off]

[rx-usecs N] [rx-frames N] [rx-usecs-irq N] [rx-frames-irq N]

[tx-usecs N] [tx-frames N] [tx-usecs-irq N] [tx-frames-irq N]

[stats-block-usecs N] [pkt-rate-low N] [rx-usecs-low N]

[rx-frames-low N] [tx-usecs-low N] [tx-frames-low N]

[pkt-rate-high N] [rx-usecs-high N] [rx-frames-high N]

[tx-usecs-high N] [tx-frames-high N] [sample-interval N]

[cqe-mode-rx on|off] [cqe-mode-tx on|off]

ethtool -g|--show-ring devname

```

ethtool -G|--set-ring devname [rx N] [rx-mini N] [rx-jumbo N] [tx N]
ethtool -i|--driver devname
ethtool -d|--register-dump devname [raw on|off] [hex on|off] [file
    name]
ethtool -e|--eeprom-dump devname [raw on|off] [offset N] [length N]
ethtool -E|--change-eeprom devname [magic N] [offset N] [length N]
    [value N]
ethtool -k|--show-features|--show-offload devname
ethtool -K|--features|--offload devname feature on|off ...
ethtool -p|--identify devname [N]
ethtool -P|--show-permaddr devname
ethtool -r|--negotiate devname
ethtool -S|--statistics devname [--all-groups|--groups [eth-phy]
    [eth-mac] [eth-ctrl] ]
ethtool --phy-statistics devname
ethtool -t|--test devname [offline|online|external_lb]
ethtool -s devname [speed N] [lanes N] [duplex half|full]
    [port tp|aui|bnc|mii] [mdix auto|on|off] [autoneg on|off]
    [advertise N/M] | advertise mode on|off ...] [phyad N]
    [xcvr internal|external] [wol N/M] | wol p|u|m|b|a|g|s|f|d...]
    [sopass xx:yy:zz:aa:bb:cc] [master-slave preferred-
    master|preferred-slave|forced-master|forced-slave] [msglvl
    N/M] | msglvl type on|off ...]
ethtool -n|-u|--show-nfc|--show-ntuple devname
    [ rx-flow-hash tcp4|udp4|ah4|esp4|sctp4|tcp6|udp6|ah6|esp6|sctp6 |
    rule N ]
ethtool -N|-U|--config-nfc|--config-ntuple devname
    rx-flow-hash tcp4|udp4|ah4|esp4|sctp4|tcp6|udp6|ah6|esp6|sctp6
    m|v|t|s|d|f|n|r... |
    flow-type
    ether|ip4|tcp4|udp4|sctp4|ah4|esp4|ip6|tcp6|udp6|ah6|esp6|sctp6
    [src xx:yy:zz:aa:bb:cc [m xx:yy:zz:aa:bb:cc]]
    [dst xx:yy:zz:aa:bb:cc [m xx:yy:zz:aa:bb:cc]] [proto N [m N]]

```

```

[src-ip ip-address [m ip-address]] [dst-ip ip-address [m ip-
address]] [tos N [m N]] [tclass N [m N]] [l4proto N [m N]]
[src-port N [m N]] [dst-port N [m N]] [spi N [m N]]
[l4data N [m N]] [vlan-etype N [m N]] [vlan N [m N]]
[user-def N [m N]] [dst-
mac xx:yy:zz:aa:bb:cc [m xx:yy:zz:aa:bb:cc]] [action N]
[context N] [loc N] |
delete N

ethtool -w|--get-dump devname [data filename]

ethtool -W|--set-dump devname N

ethtool -T|--show-time-stamping devname

ethtool -x|--show-rxfh-indir|--show-rxfh devname

ethtool -X|--set-rxfh-indir|--rxfh devname [hkey xx:yy:zz:aa:bb:cc:...]
[start N] [ equal N | weight W0 W1 ... | default ] [hfunc FUNC]
[context CTX | new] [delete]

ethtool -f|--flash devname file [N]

ethtool -l|--show-channels devname

ethtool -L|--set-channels devname [rx N] [tx N] [other N] [combined N]

ethtool -m|--dump-module-eprom|--module-info devname [raw on|off]
[hex on|off] [offset N] [length N] [page N] [bank N] [i2c N]

ethtool --show-priv-flags devname

ethtool --set-priv-flags devname flag on|off ...

ethtool --show-eee devname

ethtool --set-eee devname [eee on|off] [tx-lpi on|off] [tx-timer N]
[advertise N]

ethtool --set-phy-tunable devname [ downshift on|off [count N] ] [
fast-link-down on|off [msecs N] ] [ energy-detect-power-down
on|off [msecs N] ]

ethtool --get-phy-tunable devname [downshift] [fast-link-down] [energy-
detect-power-down]

ethtool --get-tunable devname [rx-copybreak] [tx-copybreak] [pfc-
prevention-tout]

ethtool --set-tunable devname [rx-copybreak N] [tx-copybreak N]

```

[pfc-prevention-tout N]

ethtool --reset devname [flags N] [mgmt] [mgmt-shared] [irq] [irq-shared] [dma] [dma-shared] [filter] [filter-shared] [offload] [offload-shared] [mac] [mac-shared] [phy] [phy-shared] [ram] [ram-shared] [ap] [ap-shared] [dedicated] [all]

ethtool --show-fec devname

ethtool --set-fec devname encoding auto|off|rs|baser|llrs [...]

ethtool -Q|--per-queue devname [queue\_mask %x] sub\_command ...

.

ethtool --cable-test devname

ethtool --cable-test-tdr devname [first N] [last N] [step N] [pair N]

ethtool --show-tunnels devname

## DESCRIPTION

ethtool is used to query and control network device driver and hardware settings, particularly for wired Ethernet devices.

devname is the name of the network device on which ethtool should operate.

ate.

## OPTIONS

ethtool with a single argument specifying the device name prints current settings of the specified device.

-h --help

Shows a short help message.

--version

Shows the ethtool version number.

--debug N

Turns on debugging messages. Argument is interpreted as a mask:

0x01 Parser information

--json Output results in JavaScript Object Notation (JSON). Only a subset of options support this. Those which do not will continue to

output plain text in the presence of this option.

-l --include-statistics

Include command-related statistics in the output. This option allows displaying relevant device statistics for selected get

commands.

`-a --show-pause`

Queries the specified Ethernet device for pause parameter information.

`-A --pause`

Changes the pause parameters of the specified Ethernet device.

`autoneg on|off`

Specifies whether pause autonegotiation should be enabled.

`rx on|off`

Specifies whether RX pause should be enabled.

`tx on|off`

Specifies whether TX pause should be enabled.

`-c --show-coalesce`

Queries the specified network device for coalescing information.

`-C --coalesce`

Changes the coalescing settings of the specified network device.

`-g --show-ring`

Queries the specified network device for rx/tx ring parameter information.

`-G --set-ring`

Changes the rx/tx ring parameters of the specified network device.

`rx N` Changes the number of ring entries for the Rx ring.

`rx-mini N`

Changes the number of ring entries for the Rx Mini ring.

`rx-jumbo N`

Changes the number of ring entries for the Rx Jumbo ring.

`tx N` Changes the number of ring entries for the Tx ring.

`-i --driver`

Queries the specified network device for associated driver information.

`-d --register-dump`

Retrieves and prints a register dump for the specified network

device. The register format for some devices is known and de-  
coded others are printed in hex. When raw is enabled, then eth-  
tool dumps the raw register data to stdout. If file is speci-  
fied, then use contents of previous raw register dump, rather  
than reading from the device.

`-e --eeprom-dump`

Retrieves and prints an EEPROM dump for the specified network  
device. When raw is enabled, then it dumps the raw EEPROM data  
to stdout. The length and offset parameters allow dumping cer-  
tain portions of the EEPROM. Default is to dump the entire EEPROM.

`raw on|off`

`offset N`

`length N`

`-E --change-eeprom`

If value is specified, changes EEPROM byte for the specified  
network device. offset and value specify which byte and its  
new value. If value is not specified, stdin is read and written  
to the EEPROM. The length and offset parameters allow writing to  
certain portions of the EEPROM. Because of the persistent na-  
ture of writing to the EEPROM, a device-specific magic key must  
be specified to prevent the accidental writing to the EEPROM.

`-k --show-features --show-offload`

Queries the specified network device for the state of protocol  
offload and other features.

`-K --features --offload`

Changes the offload parameters and other features of the speci-  
fied network device. The following feature names are built-in  
and others may be defined by the kernel.

`rx on|off`

Specifies whether RX checksumming should be enabled.

`tx on|off`

Specifies whether TX checksumming should be enabled.

sg on|off

Specifies whether scatter-gather should be enabled.

tso on|off

Specifies whether TCP segmentation offload should be enabled.

ufo on|off

Specifies whether UDP fragmentation offload should be enabled.

gso on|off

Specifies whether generic segmentation offload should be enabled.

gro on|off

Specifies whether generic receive offload should be enabled.

lro on|off

Specifies whether large receive offload should be enabled.

rxvlan on|off

Specifies whether RX VLAN acceleration should be enabled.

txvlan on|off

Specifies whether TX VLAN acceleration should be enabled.

ntuple on|off

Specifies whether Rx ntuple filters and actions should be enabled.

rxhash on|off

Specifies whether receive hashing offload should be enabled.

-p --identify

Initiates adapter-specific action intended to enable an operator to easily identify the adapter by sight. Typically this involves blinking one or more LEDs on the specific network port.

[ N ] Length of time to perform phys-id, in seconds.

-P --show-permaddr

Queries the specified network device for permanent hardware address.

-r --negotiate

Restarts auto-negotiation on the specified Ethernet device, if auto-negotiation is enabled.

**-S --statistics**

Queries the specified network device for standard (IEEE, IETF, etc.), or NIC- and driver-specific statistics. NIC- and driver-specific statistics are requested when no group of statistics is specified.

NIC- and driver-specific statistics and standard statistics are independent, devices may implement either, both or none. There is little commonality between naming of NIC- and driver-specific statistics across vendors.

**--all-groups**

**--groups [eth-phy] [eth-mac] [eth-ctrl] [rmon]**

Request groups of standard device statistics.

**--phy-statistics**

Queries the specified network device for PHY specific statistics.

**-t --test**

Executes adapter selftest on the specified network device. Possible test modes are:

**offline**

Perform full set of tests, possibly interrupting normal operation during the tests,

**online** Perform limited set of tests, not interrupting normal operation,

**external\_lb**

Perform full set of tests, as for offline, and additionally an external-loopback test.

**-s --change**

Allows changing some or all settings of the specified network device. All following options only apply if -s was specified.

**speed N**

Set speed in Mb/s. ethtool with just the device name as an



argument will show you the supported device speeds.

lanes N

Set number of lanes.

duplex half|full

Sets full or half duplex mode.

port tp|aur|bnc|mii

Selects device port.

master-slave preferred-master|preferred-slave|forced-master|forced-slave

Configure MASTER/SLAVE role of the PHY. When the PHY is configured as MASTER, the PMA Transmit function shall source TX\_TCLK from a local clock source. When configured as SLAVE, the PMA Transmit function shall source TX\_TCLK from the clock recovered from data stream provided by MASTER. Not all devices support this.

preferred-master Prefer MASTER role on autonegotiation

preferred-slave Prefer SLAVE role on autonegotiation

forced-master Force the PHY in MASTER role. Can be used without autonegotiation

forced-slave Force the PHY in SLAVE role. Can be used without autonegotiation

mdix auto|on|off

Selects MDI-X mode for port. May be used to override the automatic detection feature of most adapters. An argument of auto means automatic detection of MDI status, on forces MDI-X (crossover) mode, while off means MDI (straight through) mode. The driver should guarantee that this command takes effect immediately, and if necessary may reset the link to cause the change to take effect.

autoneg on|off

Specifies whether autonegotiation should be enabled. Autonegotiation is enabled by default, but in some network devices may have trouble with it, so you can disable it if really necessary.

advertise N

Sets the speed and duplex advertised by autonegotiation.

The argument is a hexadecimal value using one or a combination of the following values:

tion of the following values:

|                    |                    |                                   |
|--------------------|--------------------|-----------------------------------|
| 0x001              | 10baseT Half       |                                   |
| 0x002              | 10baseT Full       |                                   |
| 0x004              | 100baseT Half      |                                   |
| 0x008              | 100baseT Full      |                                   |
| 0x8000000000000000 | 100baseT1 Full     |                                   |
| 0x4000000000000000 | 100baseFX Half     |                                   |
| 0x8000000000000000 | 100baseFX Full     |                                   |
| 0x010              | 1000baseT Half     | (not supported by IEEE standards) |
| 0x020              | 1000baseT Full     |                                   |
| 0x20000            | 1000baseKX Full    |                                   |
| 0x2000000000       | 1000baseX Full     |                                   |
| 0x1000000000000000 | 1000baseT1 Full    |                                   |
| 0x8000             | 2500baseX Full     | (not supported by IEEE standards) |
| 0x800000000000     | 2500baseT Full     |                                   |
| 0x10000000000000   | 5000baseT Full     |                                   |
| 0x1000             | 10000baseT Full    |                                   |
| 0x40000            | 10000baseKX4 Full  |                                   |
| 0x80000            | 10000baseKR Full   |                                   |
| 0x100000           | 10000baseR_FEC     |                                   |
| 0x400000000000     | 10000baseCR Full   |                                   |
| 0x800000000000     | 10000baseSR Full   |                                   |
| 0x10000000000000   | 10000baseLR Full   |                                   |
| 0x20000000000000   | 10000baseLRM Full  |                                   |
| 0x40000000000000   | 10000baseER Full   |                                   |
| 0x200000           | 20000baseMLD2 Full | (not supported by IEEE standards) |
| 0x400000           | 20000baseKR2 Full  | (not supported by IEEE standards) |
| 0x80000000         | 25000baseCR Full   |                                   |
| 0x1000000000       | 25000baseKR Full   |                                   |
| 0x2000000000       | 25000baseSR Full   |                                   |
| 0x8000000          | 40000baseKR4 Full  |                                   |

|                        |                            |
|------------------------|----------------------------|
| 0x1000000              | 40000baseCR4 Full          |
| 0x2000000              | 40000baseSR4 Full          |
| 0x4000000              | 40000baseLR4 Full          |
| 0x400000000            | 50000baseCR2 Full          |
| 0x800000000            | 50000baseKR2 Full          |
| 0x10000000000          | 50000baseSR2 Full          |
| 0x100000000000000      | 50000baseKR Full           |
| 0x200000000000000      | 50000baseSR Full           |
| 0x400000000000000      | 50000baseCR Full           |
| 0x800000000000000      | 50000baseLR_ER_FR Full     |
| 0x1000000000000000     | 50000baseDR Full           |
| 0x8000000              | 56000baseKR4 Full          |
| 0x10000000             | 56000baseCR4 Full          |
| 0x20000000             | 56000baseSR4 Full          |
| 0x40000000             | 56000baseLR4 Full          |
| 0x1000000000           | 100000baseKR4 Full         |
| 0x2000000000           | 100000baseSR4 Full         |
| 0x4000000000           | 100000baseCR4 Full         |
| 0x8000000000           | 100000baseLR4_ER4 Full     |
| 0x2000000000000000     | 100000baseKR2 Full         |
| 0x4000000000000000     | 100000baseSR2 Full         |
| 0x8000000000000000     | 100000baseCR2 Full         |
| 0x10000000000000000    | 100000baseLR2_ER2_FR2 Full |
| 0x20000000000000000    | 100000baseDR2 Full         |
| 0x800000000000000000   | 100000baseKR Full          |
| 0x1000000000000000000  | 100000baseSR Full          |
| 0x2000000000000000000  | 100000baseLR_ER_FR Full    |
| 0x4000000000000000000  | 100000baseCR Full          |
| 0x8000000000000000000  | 100000baseDR Full          |
| 0x4000000000000000000  | 200000baseKR4 Full         |
| 0x8000000000000000000  | 200000baseSR4 Full         |
| 0x10000000000000000000 | 200000baseLR4_ER4_FR4 Full |
| 0x20000000000000000000 | 200000baseDR4 Full         |

|                    |                           |
|--------------------|---------------------------|
| 0x4000000000000000 | 20000baseCR4 Full         |
| 0x1000000000000000 | 20000baseKR2 Full         |
| 0x2000000000000000 | 20000baseSR2 Full         |
| 0x4000000000000000 | 20000baseLR2_ER2_FR2 Full |
| 0x8000000000000000 | 20000baseDR2 Full         |
| 0x1000000000000000 | 20000baseCR2 Full         |
| 0x2000000000000000 | 40000baseKR8 Full         |
| 0x4000000000000000 | 40000baseSR8 Full         |
| 0x8000000000000000 | 40000baseLR8_ER8_FR8 Full |
| 0x1000000000000000 | 40000baseDR8 Full         |
| 0x2000000000000000 | 40000baseCR8 Full         |
| 0x2000000000000000 | 40000baseKR4 Full         |
| 0x4000000000000000 | 40000baseSR4 Full         |
| 0x8000000000000000 | 40000baseLR4_ER4_FR4 Full |
| 0x1000000000000000 | 40000baseDR4 Full         |
| 0x2000000000000000 | 40000baseCR4 Full         |

phyad N

PHY address.

xcvr internal|external

Selects transceiver type. Currently only internal and external can be specified, in the future further types might be added.

wol p|u|m|b|a|g|s|f|d...

Sets Wake-on-LAN options. Not all devices support this. The argument to this option is a string of characters specifying which options to enable.

p Wake on PHY activity

u Wake on unicast messages

m Wake on multicast messages

b Wake on broadcast messages

a Wake on ARP

g Wake on MagicPacket?

s Enable SecureOn? password for MagicPacket?

f Wake on filter(s)

d Disable (wake on nothing). This option clears all previous options.

sopass xx:yy:zz:aa:bb:cc

Sets the SecureOn? password. The argument to this option must be 6 bytes in Ethernet MAC hex format (xx:yy:zz:aa:bb:cc).

msglvl N

msglvl type on|off ...

Sets the driver message type flags by name or number. type names the type of message to enable or disable; N specifies the new flags numerically. The defined type names and numbers are:

|           |        |                              |
|-----------|--------|------------------------------|
| drv       | 0x0001 | General driver status        |
| probe     | 0x0002 | Hardware probing             |
| link      | 0x0004 | Link state                   |
| timer     | 0x0008 | Periodic status check        |
| ifdown    | 0x0010 | Interface being brought down |
| ifup      | 0x0020 | Interface being brought up   |
| rx_err    | 0x0040 | Receive error                |
| tx_err    | 0x0080 | Transmit error               |
| tx_queued | 0x0100 | Transmit queueing            |
| intr      | 0x0200 | Interrupt handling           |
| tx_done   | 0x0400 | Transmit completion          |
| rx_status | 0x0800 | Receive completion           |
| pktdata   | 0x1000 | Packet contents              |
| hw        | 0x2000 | Hardware status              |
| wol       | 0x4000 | Wake-on-LAN status           |

The precise meanings of these type flags differ between drivers.

-n -u --show-nfc --show-ntuple

Retrieves receive network flow classification options or rules.

rx-flow-hash tcp4|udp4|ah4|esp4|sctp4|tcp6|udp6|ah6|esp6|sctp6

Retrieves the hash options for the specified flow type.

tcp4 TCP over IPv4

udp4 UDP over IPv4

ah4 IPSEC AH over IPv4

esp4 IPSEC ESP over IPv4

sctp4 SCTP over IPv4

tcp6 TCP over IPv6

udp6 UDP over IPv6

ah6 IPSEC AH over IPv6

esp6 IPSEC ESP over IPv6

sctp6 SCTP over IPv6

rule N Retrieves the RX classification rule with the given ID.

-N -U --config-nfc --config-ntuple

Configures receive network flow classification options or rules.

rx-flow-hash tcp4|udp4|ah4|esp4|sctp4|tcp6|udp6|ah6|esp6|sctp6  
m|v|t|s|d|f|n|r...

Configures the hash options for the specified flow type.

m Hash on the Layer 2 destination address of the rx packet.

v Hash on the VLAN tag of the rx packet.

t Hash on the Layer 3 protocol field of the rx packet.

s Hash on the IP source address of the rx packet.

d Hash on the IP destination address of the rx packet.

f Hash on bytes 0 and 1 of the Layer 4 header of the rx packet.

n Hash on bytes 2 and 3 of the Layer 4 header of the rx packet.

r Discard all packets of this flow type. When this option is  
set, all other options are ignored.

flow-type

ether|ip4|tcp4|udp4|sctp4|ah4|esp4|ip6|tcp6|udp6|ah6|esp6|sctp6

Inserts or updates a classification rule for the specified  
flow type.

ether Ethernet

ip4 Raw IPv4

tcp4 TCP over IPv4

udp4 UDP over IPv4  
sctp4 SCTP over IPv4  
ah4 IPSEC AH over IPv4  
esp4 IPSEC ESP over IPv4  
ip6 Raw IPv6  
tcp6 TCP over IPv6  
udp6 UDP over IPv6  
sctp6 SCTP over IPv6  
ah6 IPSEC AH over IPv6  
esp6 IPSEC ESP over IPv6

For all fields that allow both a value and a mask to be specified, the mask may be specified immediately after the value using the `m` keyword, or separately using the field name keyword with `-mask` appended, e.g. `src-mask`.

`src xx:yy:zz:aa:bb:cc [m xx:yy:zz:aa:bb:cc]`

Includes the source MAC address, specified as 6 bytes in hexadecimal separated by colons, along with an optional mask. Valid only for flow-type ether.

`dst xx:yy:zz:aa:bb:cc [m xx:yy:zz:aa:bb:cc]`

Includes the destination MAC address, specified as 6 bytes in hexadecimal separated by colons, along with an optional mask. Valid only for flow-type ether.

`proto N [m N]`

Includes the Ethernet protocol number (ethertype) and an optional mask. Valid only for flow-type ether.

`src-ip ip-address [m ip-address]`

Specify the source IP address of the incoming packet to match along with an optional mask. Valid for all IP based flow-types.

`dst-ip ip-address [m ip-address]`

Specify the destination IP address of the incoming packet to match along with an optional mask. Valid for all IP based flow-types.

tos N [m N]

Specify the value of the Type of Service field in the incoming packet to match along with an optional mask. Applies to all IPv4 based flow-types.

tclass N [m N]

Specify the value of the Traffic Class field in the incoming packet to match along with an optional mask. Applies to all IPv6 based flow-types.

I4proto N [m N]

Includes the layer 4 protocol number and optional mask. Valid only for flow-types ip4 and ip6.

src-port N [m N]

Specify the value of the source port field (applicable to TCP/UDP packets) in the incoming packet to match along with an optional mask. Valid for flow-types ip4, tcp4, udp4, and sctp4 and their IPv6 equivalents.

dst-port N [m N]

Specify the value of the destination port field (applicable to TCP/UDP packets) in the incoming packet to match along with an optional mask. Valid for flow-types ip4, tcp4, udp4, and sctp4 and their IPv6 equivalents.

spi N [m N]

Specify the value of the security parameter index field (applicable to AH/ESP packets) in the incoming packet to match along with an optional mask. Valid for flow-types ip4, ah4, and esp4 and their IPv6 equivalents.

I4data N [m N]

Specify the value of the first 4 Bytes of Layer 4 in the incoming packet to match along with an optional mask. Valid for ip4 and ip6 flow-types.

vlan-etype N [m N]

Includes the VLAN tag Ethertype and an optional mask.

vlan N [m N]



Includes the VLAN tag and an optional mask.

user-def N [m N]

Includes 64-bits of user-specific data and an optional mask.

dst-mac xx:yy:zz:aa:bb:cc [m xx:yy:zz:aa:bb:cc]

Includes the destination MAC address, specified as 6 bytes in hexadecimal separated by colons, along with an optional mask. Valid for all IP based flow-types.

action N

Specifies the Rx queue to send packets to, or some other action.

-1 Drop the matched flow

-2 Use the matched flow as a Wake-on-LAN filter

0 or higher Rx queue to route the flow

context N

Specifies the RSS context to spread packets over multiple queues; either 0 for the default RSS context, or a value returned by `ethtool -X ... context new`.

vf N Specifies the Virtual Function the filter applies to. Not compatible with action.

queue N

Specifies the Rx queue to send packets to. Not compatible with action.

loc N Specify the location/ID to insert the rule. This will overwrite any rule present in that location and will not go through any of the rule ordering process.

delete N

Deletes the RX classification rule with the given ID.

-w --get-dump

Retrieves and prints firmware dump for the specified network device. By default, it prints out the dump flag, version and length of the dump data. When data is indicated, then `ethtool` fetches the dump data and directs it to a file.

-W --set-dump

Sets the dump flag for the device.

`-T --show-time-stamping`

Show the device's time stamping capabilities and associated PTP hardware clock.

`-x --show-rxfh-indir --show-rxfh`

Retrieves the receive flow hash indirection table and/or RSS hash key.

`-X --set-rxfh-indir --rxfh`

Configures the receive flow hash indirection table and/or RSS hash key.

`hkey` Sets RSS hash key of the specified network device. RSS hash key should be of device supported length. Hash key format must be in `xx:yy:zz:aa:bb:cc` format meaning both the nibbles of a byte should be mentioned even if a nibble is zero.

`hfunc` Sets RSS hash function of the specified network device. List of RSS hash functions which kernel supports is shown as a part of the `--show-rxfh` command output.

`start N`

For the `equal` and `weight` options, sets the starting receive queue for spreading flows to `N`.

`equal N`

Sets the receive flow hash indirection table to spread flows evenly between the first `N` receive queues.

`weight W0 W1 ...`

Sets the receive flow hash indirection table to spread flows between receive queues according to the given weights. The sum of the weights must be non-zero and must not exceed the size of the indirection table.

`default`

Sets the receive flow hash indirection table to its default value.

`context CTX | new`

Specifies an RSS context to act on; either `new` to allocate a

new RSS context, or CTX, a value returned by a previous  
... context new.

delete Delete the specified RSS context. May only be used in con-  
junction with context and a non-zero CTX value.

-f --flash

Write a firmware image to flash or other non-volatile memory on  
the device.

file Specifies the filename of the firmware image. The firmware  
must first be installed in one of the directories where the  
kernel firmware loader or firmware agent will look, such as  
/lib/firmware.

N If the device stores multiple firmware images in separate  
regions of non-volatile memory, this parameter may be used  
to specify which region is to be written. The default is 0,  
requesting that all regions are written. All other values  
are driver-dependent.

-I --show-channels

Queries the specified network device for the numbers of channels  
it has. A channel is an IRQ and the set of queues that can  
trigger that IRQ.

-L --set-channels

Changes the numbers of channels of the specified network device.

rx N Changes the number of channels with only receive queues.

tx N Changes the number of channels with only transmit queues.

other N

Changes the number of channels used only for other purposes  
e.g. link interrupts or SR-IOV co-ordination.

combined N

Changes the number of multi-purpose channels.

-m --dump-module-EEPROM --module-info

Retrieves and if possible decodes the EEPROM from plugin mod-  
ules, e.g SFP+, QSFP. If the driver and module support it, the  
optical diagnostic information is also read and decoded. When

either one of page, bank or i2c parameters is specified, dumps only of a single page or its portion is allowed. In such a case offset and length parameters are treated relatively to EEPROM page boundaries.

#### --show-priv-flags

Queries the specified network device for its private flags. The names and meanings of private flags (if any) are defined by each network device driver.

#### --set-priv-flags

Sets the device's private flags as specified.

flag on|off Sets the state of the named private flag.

#### --show-eee

Queries the specified network device for its support of Energy-Efficient Ethernet (according to the IEEE 802.3az specifications)

#### --set-eee

Sets the device EEE behaviour.

eee on|off

Enables/disables the device support of EEE.

tx-lpi on|off

Determines whether the device should assert its Tx LPI.

advertise N

Sets the speeds for which the device should advertise EEE capabilities. Values are as for --change advertise

tx-timer N

Sets the amount of time the device should stay in idle mode prior to asserting its Tx LPI (in microseconds). This has meaning only when Tx LPI is enabled.

#### --set-phy-tunable

Sets the PHY tunable parameters.

downshift on|off

Specifies whether downshift should be enabled.

count N

Sets the PHY downshift re-tries count.

fast-link-down on|off

Specifies whether Fast Link Down should be enabled and time until link down (if supported).

msecs N

Sets the period after which the link is reported as down. Note that the PHY may choose the closest supported value. Only on reading back the tunable do you get the actual value.

energy-detect-power-down on|off

Specifies whether Energy Detect Power Down (EDPD) should be enabled (if supported). This will put the RX and TX circuit blocks into a low power mode, and the PHY will wake up periodically to send link pulses to avoid any lock-up situation with a peer PHY that may also have EDPD enabled. By default, this setting will also enable the periodic transmission of TX pulses.

msecs N

Some PHYs support configuration of the wake-up interval to send TX pulses.

This setting allows the control of this interval, and 0 disables TX pulses if the PHY supports this. Disabling TX pulses can create a lock-up situation where neither of the PHYs wakes the other one. If unspecified the default value (in milliseconds) will be used by the PHY.

--get-phy-tunable

Gets the PHY tunable parameters.

downshift

For operation in cabling environments that are incompatible with 1000BASE-T, PHY device provides an automatic link speed downshift operation. Link speed downshift after N failed 1000BASE-T auto-negotiation attempts. Downshift is useful where cable does not have the 4 pairs instance.

Gets the PHY downshift count/status.

fast-link-down

Depending on the mode it may take 0.5s - 1s until a broken link is reported as down. In certain use cases a link-down

event needs to be reported as soon as possible. Some PHYs support a Fast Link Down Feature and may allow configuration of the delay before a broken link is reported as being down.

Gets the PHY Fast Link Down status / period.

energy-detect-power-down

Gets the current configured setting for Energy Detect Power Down (if supported).

--get-tunable

Get the tunable parameters.

rx-copybreak

Get the current rx copybreak value in bytes.

tx-copybreak

Get the current tx copybreak value in bytes.

pfc-prevention-tout

Get the current pfc prevention timeout value in msecs.

--set-tunable

Set driver's tunable parameters.

rx-copybreak N

Set the rx copybreak value in bytes.

tx-copybreak N

Set the tx copybreak value in bytes.

pfc-prevention-tout N

Set pfc prevention timeout in msecs. Value of 0 means disable and 65535 means auto.

--reset

Reset hardware components specified by flags and components listed below

flags N

Resets the components based on direct flags mask

mgmt Management processor

irq Interrupt requester

dma DMA engine

filter Filtering/flow direction

offload

Protocol offload

mac Media access controller

phy Transceiver/PHY

ram RAM shared between multiple components ap Application Processor

dedicated

All components dedicated to this interface

all All components used by this interface, even if shared

--show-fec

Queries the specified network device for its support of Forward Error Correction.

--set-fec

Configures Forward Error Correction for the specified network device.

Forward Error Correction modes selected by a user are expected to be persisted after any hotplug events. If a module is swapped that does not support the current FEC mode, the driver or firmware must take the link down administratively and report the problem in the system logs for users to correct.

encoding auto|off|rs|baser|llrs [...]

Sets the FEC encoding for the device. Combinations of options are specified as e.g. encoding auto rs ; the semantics of such combinations vary between drivers.

auto Use the driver's default encoding

off Turn off FEC

RS Force RS-FEC encoding

BaseR Force BaseR encoding

LLRS Force LLRS-FEC encoding

-Q|--per-queue

Applies provided sub command to specific queues.

queue\_mask %x

Sets the specific queues which the sub command is applied

to. If `queue_mask` is not set, the sub command will be applied to all queues.

#### sub\_command

Sub command to apply. The supported sub commands include `--show-coalesce` and `--coalesce`.

#### q.B --cable-test

Perform a cable test and report the results. What results are returned depends on the capabilities of the network interface. Typically open pairs and shorted pairs can be reported, along with pairs being O.K. When a fault is detected the approximate distance to the fault may be reported.

#### --cable-test-tdr

Perform a cable test and report the raw Time Domain Reflectometer data. A pulse is sent down a cable pair and the amplitude of the reflection, for a given distance, is reported. A break in the cable returns a big reflection. Minor damage to the cable returns a small reflection. If the cable is shorted, the amplitude of the reflection can be negative. By default, data is returned for lengths between 0 and 150m at 1m steps, for all pairs. However parameters can be passed to restrict the collection of data. It should be noted, that the interface will round the distances to whatever granularity is actually implemented. This is often 0.8 of a meter. The results should include the actual rounded first and last distance and step size.

#### first N

Distance along the cable, in meters, where the first measurement should be made.

#### last N

Distance along the cable, in meters, where the last measurement should be made.

#### step N

Distance, in meters, between each measurement.

#### pair N



Which pair should be measured. Typically a cable has 4 pairs. 0 = Pair A, 1 = Pair B, ...

--monitor

Listens to netlink notification and displays them.

command

If argument matching a command is used, ethtool only shows notifications of this type. Without such argument or with --all, all notification types are shown.

devname

If a device name is used as argument, only notification for this device are shown. Default is to show notifications for all devices.

--show-tunnels

Show tunnel-related device capabilities and state. List UDP ports kernel has programmed the device to parse as VxLAN, or GENEVE tunnels.

## BUGS

Not supported (in part or whole) on all network drivers.

## AUTHOR

ethtool was written by David Miller.

Modifications by Jeff Garzik, Tim Hockin, Jakub Jelinek, Andre Majorel, Eli Kupermann, Scott Feldman, Andi Kleen, Alexander Duyck, Sucheta Chakraborty, Jesse Brandeburg, Ben Hutchings, Scott Branden.

## AVAILABILITY

ethtool is available from  
[?http://www.kernel.org/pub/software/network/ethtool/?](http://www.kernel.org/pub/software/network/ethtool/)

Ethtool version 5.16      January 2022      ETHTOOL(8)