

⚙️ Address management

In this section `${address}` value should be a host address in dotted decimal format, and `${mask}` can be either a dotted decimal subnet mask or a prefix length. That is, both 192.0.2.10/24 and 192.0.2.10/255.255.255.0 are equally acceptable.

Show all addresses `ip address show` All "show" commands can be used with "-4" or "-6" options to show only IPv4 or IPv6 addresses.

Show addresses for a single interface `ip address show ${interface name}` `ip address show eth0`

Show addresses only for running interfaces `ip address show up`

Show only statically configured addresses `ip address show [dev ${interface}] permanent`

Show only addresses learnt via autoconfiguration `ip address show [dev ${interface}] dynamic`

Add an address to an interface `ip address add ${address}/${mask} dev ${interface name}` `ip address add 192.0.2.10/27 dev eth0` `ip address add 2001:db8:1::48 dev tun10`

You can add as many addresses as you want. The first address will be primary and will be used as source address by default.

Add an address with human-readable description `ip address add ${address}/${mask} dev ${interface name} label ${interface name}:${description}` `ip address add 192.0.2.1/24 dev eth0 label eth0:my_wan_address` Interface name with a colon before label is required, some backwards compatibility issue.

Delete an address `ip address delete ${address}/${prefix} dev ${interface name}` `ip address delete 192.0.2.1/24 dev eth0` Interface name argument is required. Linux does allow to use the same address on multiple interfaces and it has valid use cases.

Remove all addresses from an interface `ip address flush dev ${interface name}` `ip address flush dev eth1`

Metasyntactic variables are written in shell-style syntax, `${something}`. Optional command parts are in square brackets. Note that there is no way to rearrange addresses and replace the primary address. Make sure you set the primary address first.

↔️ Route management

View all routes `ip route` `ip route show`

View IPv6 routes `ip -6 route`

View routes to a network and all its subnets `ip route show to root ${address}/${mask}` `ip route show to root 192.168.0.0/24`

View routes to a network and all supernets `ip route show to match ${address}/${mask}` `ip route show to match 192.168.0.0/24`

View routes to exact subnet `ip route show to exact ${address}/${mask}` `ip route show to exact 192.168.0.0/24`



By **TME520** (TME520)
cheatography.com/tme520/
tme520.com

Published 10th May, 2015.
Last updated 7th May, 2016.
Page 1 of 11.

Sponsored by **Readable.com**
Measure your website readability!
<https://readable.com>

↩ Route management (cont)

View only the route actually used by the kernel	<code>ip route get \${address}/\${mask}</code>	<code>ip route get 192.168.0.0/24</code>	Note that in complex routing scenarios like multipath routing, the result may be "correct but not complete", as it always shows one route that will be used first.
View route cache (pre 3.6 kernels only)	<code>ip route show cached</code>	Until the version 3.6, Linux used route caching. In older kernels, this command displays the contents of the route cache. It can be used with modifiers described above. In newer kernels it does nothing.	
Add a route via gateway	<code>ip route add \${address}/\${mask} via \${next hop}</code>	<code>ip route add 192.0.2.128/25 via 192.0.2.1</code>	<code>ip route add 2001:db8:1::/48 via 2001:db8:1::1</code>
Add a route via interface	<code>ip route add \${address}/\${mask} dev \${interface name}</code>	<code>ip route add 192.0.2.0/25 dev ppp0</code>	Interface routes are commonly used with point-to-point interfaces like PPP tunnels where next hop address is not required.
Change or replace a route	<code>ip route change 192.168.2.0/24 via 10.0.0.1</code>	<code>ip route replace 192.0.2.1/27 dev tun0</code>	
Delete a route	<code>ip route delete \${rest of the route statement}</code>	<code>ip route delete 10.0.1.0/25 via 10.0.0.1</code>	<code>ip route delete default dev ppp0</code>
Default route	<code>ip route add default via \${address}/\${mask}</code>	<code>ip route add default dev \${interface name}</code>	<code>ip -6 route add default via 2001:db8::1</code>
Blackhole routes	<code>ip route add blackhole \${address}/\${mask}</code>	<code>ip route add blackhole 192.0.2.1/32</code>	Traffic to destinations that match a blackhole route is silently discarded.
Other special routes : unreachable	<code>ip route add unreachable \${address}/\${mask}</code>		Sends ICMP "host unreachable". These routes make the system discard packets and reply with an ICMP error message to the sender.
Other special routes : prohibit	<code>ip route add prohibit \${address}/\${mask}</code>		Sends ICMP "administratively prohibited".



By **TME520** (TME520)
cheatography.com/tme520/tme520.com

Published 10th May, 2015.
 Last updated 7th May, 2016.
 Page 2 of 11.

Sponsored by **Readable.com**
 Measure your website readability!
<https://readable.com>

🔗 Route management (cont)

Other special routes : throw	<code>ip route add throw \${address}/\${mask}</code>	Sends "net unreachable".	
Routes with different metric	<code>ip route add \${address}/\${mask} via \${gateway} metric \${number}</code>	<code>ip route add 192.168.2.0/24 via 10.0.1.1 metric 5</code>	<code>ip route add 192.168.2.0 dev ppp0 metric 10</code>
Multipath routing	<code>ip route add \${address}/\${mask} nexthop via \${gateway 1} weight \${number} nexthop via \${gateway 2} weight \${number}</code>	<code>ip route add default nexthop via 192.168.1.1 weight 1 nexthop dev ppp0 weight 10</code>	

As per the section below, if you set up a static route, and it becomes useless because the interface goes down, it will be removed and never get back on its own. You may not have noticed this behaviour because in many cases additional software (e.g. NetworkManager or rp-pppoe) takes care of restoring routes associated with interfaces.

🔗 Link management

Show information about all links	<code>ip link show</code>	<code>ip link list</code>	
Show information about specific link	<code>ip link show dev \${interface name}</code>	<code>ip link show dev eth0</code>	<code>ip link show dev tun10</code>
Bring a link up or down	<code>ip link set dev \${interface name} [up down]</code>	<code>ip link set dev eth0 down</code>	<code>ip link set dev br0 up</code>
Set human-readable link description	<code>ip link set dev \${interface name} alias "\${description}"</code>	<code>ip link set dev eth0 alias "LAN interface"</code>	
Rename an interface	<code>ip link set dev \${old interface name} name \${new interface name}</code>	<code>ip link set dev eth0 name lan</code>	Note that you can't rename an active interface. You need to bring it down before doing it.
Change link layer address (usually MAC address)	<code>ip link set dev \${interface name} address \${address}</code>	<code>ip link set dev eth0 address 22:ce:e0:99:63:6f</code>	
Change link MTU	<code>ip link set dev \${interface name} mtu \${MTU value}</code>	<code>ip link set dev tun0 mtu 1480</code>	
Delete a link	<code>ip link delete dev \${interface name}</code>		
Enable or disable multicast on an interface	<code>ip link set \${interface name} multicast on</code>	<code>ip link set \${interface name} multicast off</code>	
Enable or disable ARP on an interface	<code>ip link set \${interface name} arp on</code>	<code>ip link set \${interface name} arp off</code>	



By **TME520** (TME520)
cheatography.com/tme520/tme520.com

Published 10th May, 2015.
 Last updated 7th May, 2016.
 Page 3 of 11.

Sponsored by **Readable.com**
 Measure your website readability!
<https://readable.com>

🔗 Link management (cont)

Create a VLAN interface	<i>ip link add name \${VLAN interface name} link \${parent interface name} type vlan id \${tag}</i>	<i>ip link add name eth0.110 link eth0 type vlan id 110</i>	The only type of VLAN supported in Linux is IEEE 802.1q VLAN, legacy implementations like ISL are not supported.
Create a QinQ interface (VLAN stacking)	<i>ip link add name \${service interface} link \${physical interface} type vlan proto 802.1ad id \${service tag}</i>		
	<i>ip link add name \${client interface} link \${service interface} type vlan proto 802.1q id \${client tag}</i>		
	<i>ip link add name eth0.100 link eth0 type vlan proto 802.1ad id 100</i>	Create service tag interface	
	<i>ip link add name eth0.100.200 link eth0.100 type vlan proto 802.1q id 200</i>	Create client tag interface	
Create pseudo-ethernet (aka macvlan) interface	<i>ip link add name \${macvlan interface name} link \${parent interface} type macvlan</i>	<i>ip link add name peth0 link eth0 type macvlan</i>	
Create a dummy interface	<i>ip link add name \${dummy interface name} type dummy</i>	<i>ip link add name dummy0 type dummy</i>	
Create a bridge interface	<i>ip link add name \${bridge name} type bridge</i>	<i>ip link add name br0 type bridge</i>	
Add an interface to bridge	<i>ip link set dev \${interface name} master \${bridge name}</i>	<i>ip link set dev eth0 master br0</i>	
Remove interface from bridge	<i>ip link set dev \${interface name} nomaster</i>	<i>ip link set dev eth0 nomaster</i>	
Create a bonding interface	<i>ip link add name \${name} type bond</i>	<i>ip link add name bond1 type bond</i>	This is not enough to configure bonding (link aggregation) in any meaningful way. You need to set up bonding parameters according to your situation.



By **TME520** (TME520)
cheatography.com/tme520/tme520.com

Published 10th May, 2015.
 Last updated 7th May, 2016.
 Page 4 of 11.

Sponsored by **Readable.com**
 Measure your website readability!
<https://readable.com>

🔗 Link management (cont)

Create an intermediate functional block interface	<code>ip link add \${interface name} type ifb</code>	<code>ip link add ifb10 type ifb</code>	Intermediate functional block devices are used for traffic redirection and mirroring in conjunction with tc.
Create a pair of virtual ethernet devices	<code>ip link add name \${first device name} type veth peer name \${second device name}</code>	<code>ip link add name veth-host type veth peer name veth-guest</code>	Virtual ethernet devices are created in UP state, no need to bring them up manually after creation.

Note that interface name you set with "name \${name}" parameter of "ip link add" and "ip link set" commands may be arbitrary, and even contain unicode characters. It's better however to stick with ASCII because other programs may not handle unicode correctly. Also it's better to use a consistent convention for link names, and use link aliases to provide human descriptions.

🔗 Link group management

Add an interface to a group	<code>ip link set dev \${interface name} group \${group number}</code>	<code>ip link set dev eth0 group 42</code>	<code>ip link set dev eth1 group 42</code>
Remove an interface from a group	<code>ip link set dev \${interface name} group 0</code>	<code>ip link set dev \${interface} group default</code>	<code>ip link set dev tun10 group 0</code>
Assign a symbolic name to a group	<code>echo "10 customer-vlans" >> /etc/iproute2/group</code>	Once you configured a group name, number and name can be used interchangeably in ip commands.	<code>ip link set dev eth0.100 group customer-vlans</code>
Perform an operation on a group	<code>ip link set group \${group number} \${operation and arguments}</code>	<code>ip link set group 42 down</code>	<code>ip link set group uplinks mtu 1200</code>
View information about links from specific group	<code>ip link list group 42</code>	<code>ip address show group customers</code>	

Link groups are similar to port ranges found in managed switches. You can add network interfaces to a numbered group and perform operations on all the interfaces from that group at once.

Links not assigned to any group belong to group 0 aka "default".

Tun and Tap devices

Add an tun/tap device useable by root	<code>ip tuntap add dev \${interface name} mode \${mode}</code>	<code>ip tuntap add dev tun0 mode tun</code>	<code>ip tuntap add dev tap9 mode tap</code>
Tap sends and receives raw Ethernet frames.		Tun sends and receives raw IP packets.	
Add an tun/tap device usable by an ordinary user	<code>ip tuntap add dev \${interface name} mode \${mode} user \${user} group \${group}</code>	<code>ip tuntap add dev tun1 mode tun user me group mygroup</code>	<code>ip tuntap add dev tun2 mode tun user 1000 group 1001</code>



By **TME520** (TME520)
cheatography.com/tme520/
tme520.com

Published 10th May, 2015.
 Last updated 7th May, 2016.
 Page 5 of 11.

Sponsored by **Readable.com**
 Measure your website readability!
<https://readable.com>

Tun and Tap devices (cont)

Add an tun/tap device using an alternate packet format	<code>ip tuntap add dev \${interface name} mode \${mode} pi</code>	<code>ip tuntap add dev tun1 mode tun pi</code>
Add an tun/tap ignoring flow control	<code>ip tuntap add dev \${interface name} mode \${mode} one_queue</code>	<code>ip tuntap add dev tun1 mode tun one_queue</code>
Delete tun/tap device	<code>ip tuntap del dev \${interface name}</code>	<code>ip tuntap del dev tun0 name}</code>

Tun and tap devices allow userspace programs to emulate a network device. When the userspace program opens them they get a file descriptor. Packets routed by the kernel networking stack to the device are read from the file descriptor, data the userspace program writes to the file descriptor are injected as local outgoing packets into the networking stack.

Neighbor (ARP and NDP) tables management

View neighbor tables	<code>ip neighbor show</code>	
View neighbors for single interface	<code>ip neighbor show dev \${interface name}</code>	<code>ip neighbor show dev eth0</code>
Flush table for an interface	<code>ip neighbor flush dev \${interface name}</code>	<code>ip neighbor flush dev eth1</code>
Add a neighbor table entry	<code>ip neighbor add \${network address} lladdr \${link layer address} dev \${interface name}</code>	<code>ip neighbor add 192.0.2.1 lladdr 22:ce:e0:99:63:6f dev eth0</code>
Delete a neighbor table entry	<code>ip neighbor delete \${network address} lladdr \${link layer address} dev \${interface name}</code>	<code>ip neighbor delete 192.0.2.1 lladdr 22:ce:e0:99:63:6f dev eth0</code>

For ladies and gentlemen who prefer UK spelling, this command family supports "neighbour" spelling too.

↔ Tunnel management

Create an IPIP tunnel	<code>ip tunnel add \${interface name} mode ipip local \${local endpoint address} remote \${remote endpoint address}</code>
Create a SIT tunnel	<code>sudo ip tunnel add \${interface name} mode sit local \${local endpoint address} remote \${remote endpoint address}</code>
Create an IPIP6 tunnel	<code>ip -6 tunnel add \${interface name} mode ipip6 local \${local endpoint address} remote \${remote endpoint address}</code>



By **TME520** (TME520)
cheatography.com/tme520/
tme520.com

Published 10th May, 2015.
 Last updated 7th May, 2016.
 Page 6 of 11.

Sponsored by **Readable.com**
 Measure your website readability!
<https://readable.com>

↔ Tunnel management (cont)

Create an IP6IP6 tunnel	<i>ip -6 tunnel add \${interface name} mode ip6ip6 local \${local endpoint address} remote \${remote endpoint address}</i>		
Create a gretap (ethernet over GRE) device	<i>ip link add \${interface name} type gretap local \${local endpoint address} remote \${remote endpoint address}</i>		
Create a GRE tunnel	<i>ip tunnel add \${interface name} mode gre local \${local endpoint address} remote \${remote endpoint address}</i>		
Create multiple GRE tunnels to the same endpoint	<i>ip tunnel add \${interface name} mode gre local \${local endpoint address} remote \${remote endpoint address} key \${key value}</i>		
Create a point-to-multipoint GRE tunnel	<i>ip tunnel add \${interface name} mode gre local \${local endpoint address} key \${key value}</i>		
Create a GRE tunnel over IPv6	<i>ip -6 tunnel add name \${interface name} mode ip6gre local \${local endpoint} remote \${remote endpoint}</i>		
Delete a tunnel	<i>ip tunnel del \${interface name}</i>	<i>ip tunnel del gre1</i>	
Modify a tunnel	<i>ip tunnel change \${interface name} \${options}</i>	<i>ip tunnel change tun0 remote 203.0.113.89</i>	<i>ip tunnel change tun10 key 23456</i>
View tunnel information	<i>ip tunnel show</i>	<i>ip tunnel show \${interface name}</i>	<i>ip tun show tun99</i>

Linux currently supports IP4IP (IPv4 in IPv4), SIT (IPv6 in IPv4), IP6IP6 (IPv6 in IPv6), IP4IP6 (IPv4 in IPv6), GRE (virtually anything in anything), and, in very recent versions, VTI (IPv4 in IPsec).

Note that tunnels are created in DOWN state, you need to bring them up.

In this section `${local endpoint address}` and `${remote endpoint address}` refer to addresses assigned to physical interfaces of endpoint. `${address}` refers to the address assigned to tunnel interface.



By **TME520** (TME520)
cheatography.com/tme520/tme520.com

Published 10th May, 2015.
 Last updated 7th May, 2016.
 Page 7 of 11.

Sponsored by **Readable.com**
 Measure your website readability!
<https://readable.com>

L2TPv3 pseudowire management

Create an L2TPv3 tunnel over UDP	<i>ip l2tp add tunnel tunnel_id \${local tunnel numeric identifier} peer_tunnel_id \${remote tunnel numeric identifier} udp_sport \${source port} udp_dport \${destination port} encap udp local \${local endpoint address} remote \${remote endpoint address}</i>		
	<i>ip l2tp add tunnel tunnel_id 1 peer_tunnel_id 1 udp_sport 5000 udp_dport 5000 encap udp local 192.0.2.1 remote 203.0.113.2</i>		
Create an L2TPv3 tunnel over IP	<i>ip l2tp add tunnel tunnel_id \${local tunnel numeric identifier} peer_tunnel_id {remote tunnel numeric identifier} encap ip local 192.0.2.1 remote 203.0.113.2</i>		
Create an L2TPv3 session	<i>ip l2tp add session tunnel_id \${local tunnel identifier} session_id \${local session numeric identifier} peer_session_id \${remote session numeric identifier}</i>	<i>ip l2tp add session tunnel_id 1 session_id 10 peer_session_id 10</i>	
Delete an L2TPv3 session	<i>ip l2tp del session tunnel_id \${tunnel identifier} session_id \${session identifier}</i>	<i>ip l2tp del session tunnel_id 1 session_id 1</i>	
Delete an L2TPv3 tunnel	<i>ip l2tp del tunnel tunnel_id \${tunnel identifier}</i>	<i>ip l2tp del tunnel tunnel_id 1</i>	
View L2TPv3 tunnel information	<i>ip l2tp show tunnel</i>	<i>ip l2tp show tunnel tunnel_id \${tunnel identifier}</i>	<i>ip l2tp show tunnel tunnel_id 12</i>



By **TME520** (TME520)
cheatography.com/tme520/
tme520.com

Published 10th May, 2015.
 Last updated 7th May, 2016.
 Page 8 of 11.

Sponsored by **Readable.com**
 Measure your website readability!
<https://readable.com>

L2TPv3 pseudowire management (cont)

View L2TPv3 session information	<i>ip l2tp show session</i>	<i>ip l2tp show session session_id \${session identifier} tunnel_id \${tunnel identifier}</i>	<i>ip l2tp show session session_id 1 tunnel_id 12</i>
--	-----------------------------	---	---

Compared to other tunneling protocol implementations in Linux, L2TPv3 terminology is somewhat reversed. You create a tunnel, and then bind sessions to it. You can bind multiple sessions with different identifiers to the same tunnel. Virtual network interfaces (by default named l2tpethX) are associated with sessions.

</> Policy-based routing

Create a policy route	<i>ip route add \${route options} table \${table id or name}</i>	<i>ip route add 192.0.2.0/27 via 203.0.113.1 table 10</i>	<i>ip route add 2001:db8::/48 dev eth1 table 100</i>
View policy routes	<i>ip route show table \${table id or name}</i>	<i>ip route show table 100</i>	<i>ip route show table test</i>
General rule syntax	<i>ip rule add \${options} <lookup \${table id or name} blackhole prohibit unreachable></i>		
Create a rule to match a source network	<i>ip rule add from \${source network} \${action}</i>	<i>ip rule add from 192.0.2.0/24 lookup 10</i>	<i>ip -6 rule add from 2001:db8::/32 prohibit</i>
Create a rule to match a destination network	<i>ip rule add to \${destination network} \${action}</i>	<i>ip rule add to 192.0.2.0/24 blackhole</i>	<i>ip -6 rule add to 2001:db8::/32 lookup 100</i>
Create a rule to match a ToS field value	<i>ip rule add tos \${ToS value} \${action}</i>	<i>ip rule add tos 0x10 lookup 110</i>	
Create a rule to match a firewall mark value	<i>ip rule add fwmark \${mark} \${action}</i>	<i>ip rule add fwmark 0x11 lookup 100</i>	
Create a rule to match inbound interface	<i>ip rule add iif \${interface name} \${action}</i>	<i>ip rule add iif eth0 lookup 10</i>	<i>ip rule add iif lo lookup 20</i>
Create a rule to match outbound interface	<i>ip rule add oif \${interface name} \${action}</i>	<i>ip rule add oif eth0 lookup 10</i>	
Set rule priority	<i>ip rule add \${options} \${action} priority \${value}</i>	<i>ip rule add from 192.0.2.0/25 lookup 10 priority 10</i>	<i>ip rule add from 192.0.2.0/24 lookup 20 priority 20</i>
Show all rules	<i>ip rule show</i>	<i>ip -6 rule show</i>	
Delete a rule	<i>ip rule del \${options} \${action}</i>	<i>ip rule del 192.0.2.0/24 lookup 10</i>	
Delete all rules	<i>ip rule flush</i>	<i>ip -6 rule flush</i>	

Policy-based routing (PBR) in Linux is designed the following way: first you create custom routing tables, then you create rules to tell the kernel it should use those tables instead of the default table for specific traffic.

Some tables are predefined: local (table 255), main (table 254), default (table 253).

netconf (sysctl configuration viewing)

View sysctl configuration for all interfaces	<i>ip netconf show</i>		
View sysctl configuration for specific interface	<i>ip netconf show dev \${interface}</i>	<i>ip netconf show dev eth0</i>	



By **TME520** (TME520)
cheatography.com/tme520/
tme520.com

Published 10th May, 2015.
 Last updated 7th May, 2016.
 Page 9 of 11.

Sponsored by **Readable.com**
 Measure your website readability!
<https://readable.com>

Network namespace management

Create a namespace	<code>ip netns add \${namespace name}</code>	<code>ip netns add foo</code>	
List existing namespaces	<code>ip netns list</code>		
Delete a namespace	<code>ip netns delete \${namespace name}</code>	<code>ip netns delete foo</code>	
Run a process inside a namespace	<code>ip netns exec \${namespace name} \${command}</code>	<code>ip netns exec foo /bin/sh</code>	
List all processes assigned to a namespace	<code>ip netns pids \${namespace name}</code>		The output will be a list of PIDs.
Identify process' primary namespace	<code>ip netns identify \${pid}</code>	<code>ip netns identify 9000</code>	
Assign network interface to a namespace	<code>ip link set dev \${interface name} netns \${namespace name}</code>	<code>ip link set dev \${interface name} netns \${pid}</code>	<code>ip link set dev eth0.100 netns foo</code>
Connect one namespace to another	Create a pair of veth devices:	<code>ip link add name veth1 type veth peer name veth2</code>	
	Move veth2 to namespace foo:	<code>ip link set dev veth2 netns foo</code>	
	Bring veth2 and add an address in "foo" namespace:	<code>ip netns exec foo ip link set dev veth2 up</code>	
		<code>ip netns exec foo ip address add 10.1.1.1/24 dev veth2</code>	
	Add an address to veth1, which stays in the default namespace:	<code>ip address add 10.1.1.2/24 dev veth1</code>	
Monitor network namespace subsystem events	<code>ip netns monitor</code>		

Network namespaces are isolated network stack instances within a single machine. They can be used for security domain separation, managing traffic flows between virtual machines and so on.

Every namespace is a complete copy of the networking stack with its own interfaces, addresses, routes etc. You can run processes inside a namespace and bridge namespaces to physical interfaces.

VXLAN management

Create a VXLAN link	<code>ip link add name \${interface name} type vxlan id <0-16777215> dev \${source interface} group \${multicast address}</code>	<code>ip link add name vxlan0 type vxlan id 42 dev eth0 group 239.0.0.1</code>
---------------------	--	--

VXLAN is a layer 2 tunneling protocol that is commonly used in conjunction with virtualization systems such as KVM to connect virtual machines running on different hypervisor nodes to each other and to outside world. The underlying encapsulation protocol for VXLAN is UDP.



By **TME520** (TME520)
cheatography.com/tme520/
tme520.com

Published 10th May, 2015.
 Last updated 7th May, 2016.
 Page 10 of 11.

Sponsored by **Readable.com**
 Measure your website readability!
<https://readable.com>

⚙ Multicast management

View multicast groups *ip maddress show* *ip maddress show \${interface name}* *ip maddress show dev lo*

Add a link-layer multicast address *ip maddress add \${MAC address} dev \${interface name}* *ip maddress add 01:00:5e:00:00:ab dev eth0*

View multicast routes *ip mroute show* Multicast routes cannot be added manually, so this command can only show multicast routes installed by a routing daemon. It supports the same modifiers to unicast route viewing commands (iif, table, from etc.).

Multicast is mostly handled by applications and routing daemons, so there is not much you can and should do manually here. Multicast-related ip commands are mostly useful for debug.

👁 Network event monitoring

Monitor all events *ip monitor*

Monitor specific events *ip monitor \${event type}* Event type can be: link, address, route, mroute, neigh.

Read a log file produced by rtmon *ip monitor \${event type} file \${path to the log file}* *rtmon [-family <inet|inet6>] [<route|link|address|all>] file \${log file path}* iproute2 includes a program called "rtmon" that serves essentially the same purpose, but writes events to a binary log file instead of displaying them. "ip monitor" command allows you to read files created by the program". rtmon syntax is similar to that of "ip monitor", except event type is limited to link, address, route, and all; and address family is specified in "-family" option.

You can monitor certain network events with iproute2, such as changes in network configuration, routing tables, and ARP/NDP tables.



By **TME520** (TME520)
cheatography.com/tme520/
tme520.com

Published 10th May, 2015.
Last updated 7th May, 2016.
Page 11 of 11.

Sponsored by **Readable.com**
Measure your website readability!
<https://readable.com>