



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

### ***Rocky Enterprise Linux 9.2 Manual Pages on command 'rtnetlink.7'***

***\$ man rtnetlink.7***

RTNETLINK(7)          Linux Programmer's Manual          RTNETLINK(7)

NAME

rtnetlink - Linux IPv4 routing socket

SYNOPSIS

```
#include <asm/types.h>
```

```
#include <linux/netlink.h>
```

```
#include <linux/rtnetlink.h>
```

```
#include <sys/socket.h>
```

```
rtnetlink_socket = socket(AF_NETLINK, int socket_type, NETLINK_ROUTE);
```

DESCRIPTION

Rtnetlink allows the kernel's routing tables to be read and altered.

It is used within the kernel to communicate between various subsystems, though this usage is not documented here, and for communication with user-space programs. Network routes, IP addresses, link parameters, neighbor setups, queueing disciplines, traffic classes and packet classifiers may all be controlled through NETLINK\_ROUTE sockets. It is based on netlink messages; see netlink(7) for more information.

Some rtnetlink messages have optional attributes after the initial header:

```
struct rtattr {
    unsigned short rta_len; /* Length of option */
    unsigned short rta_type; /* Type of option */
    /* Data follows */
};
```

These attributes should be manipulated using only the RTA\_\* macros or libnetlink, see rtnetlink(3).

### Messages

Rtnetlink consists of these message types (in addition to standard netlink messages):

RTM\_NEWLINK, RTM\_DELLINK, RTM\_GETLINK

Create, remove, or get information about a specific network interface. These messages contain an ifinfomsg structure followed by a series of rtattr structures.

```
struct ifinfomsg {
    unsigned char ifi_family; /* AF_UNSPEC */
    unsigned short ifi_type; /* Device type */
    int ifi_index; /* Interface index */
    unsigned int ifi_flags; /* Device flags */
    unsigned int ifi_change; /* change mask */
};
```

ifi\_flags contains the device flags, see netdevice(7); ifi\_index is the unique interface index (since Linux 3.7, it is possible to feed a nonzero value with the RTM\_NEWLINK message, thus creating a link with the given ifindex); ifi\_change is reserved for future use and should be always set to 0xFFFFFFFF.

#### Routing attributes

rta_type	Value type	Description
IFLA_UNSPEC	-	unspecified
IFLA_ADDRESS	hardware address	interface L2 address

IFLA\_BROADCAST hardware address L2 broadcast address

IFLA\_IFNAME asciiz string Device name

IFLA\_MTU unsigned int MTU of the device

IFLA\_LINK int Link type

IFLA\_QDISC asciiz string Queueing discipline

IFLA\_STATS see below Interface Statistics

The value type for IFLA\_STATS is struct rtnl\_link\_stats (struct net\_device\_stats in Linux 2.4 and earlier).

### RTM\_NEWADDR, RTM\_DELADDR, RTM\_GETADDR

Add, remove, or receive information about an IP address associated with an interface. In Linux 2.2, an interface can carry multiple IP addresses, this replaces the alias device concept in 2.0. In Linux 2.2, these messages support IPv4 and IPv6 addresses. They contain an ifaddrmsg structure, optionally followed by rtattr routing attributes.

```
struct ifaddrmsg {
    unsigned char ifa_family; /* Address type */
    unsigned char ifa_prefixlen; /* Prefixlength of address */
    unsigned char ifa_flags; /* Address flags */
    unsigned char ifa_scope; /* Address scope */
    unsigned int ifa_index; /* Interface index */
};
```

ifa\_family is the address family type (currently AF\_INET or AF\_INET6), ifa\_prefixlen is the length of the address mask of the address if defined for the family (like for IPv4), ifa\_scope is the address scope, ifa\_index is the interface index of the interface the address is associated with. ifa\_flags is a flag word of IFA\_F\_SECONDARY for secondary address (old alias interface), IFA\_F\_PERMANENT for a permanent address set by the user and other undocumented flags.

#### Attributes

rta\_type Value type Description

??

IFA\_UNSPEC - unspecified  
 IFA\_ADDRESS raw protocol address interface address  
 IFA\_LOCAL raw protocol address local address  
 IFA\_LABEL asciiz string name of the interface  
 IFA\_BROADCAST raw protocol address broadcast address  
 IFA\_ANYCAST raw protocol address anycast address  
 IFA\_CACHEINFO struct ifa\_cacheinfo Address information

RTM\_NEWROUTE, RTM\_DELROUTE, RTM\_GETROUTE

Create, remove, or receive information about a network route.

These messages contain an rtmmsg structure with an optional sequence of rtattr structures following. For RTM\_GETROUTE, setting rtm\_dst\_len and rtm\_src\_len to 0 means you get all entries for the specified routing table. For the other fields, except rtm\_table and rtm\_protocol, 0 is the wildcard.

```
struct rtmmsg {
    unsigned char rtm_family; /* Address family of route */
    unsigned char rtm_dst_len; /* Length of destination */
    unsigned char rtm_src_len; /* Length of source */
    unsigned char rtm_tos; /* TOS filter */
    unsigned char rtm_table; /* Routing table ID;
                             see RTA_TABLE below */
    unsigned char rtm_protocol; /* Routing protocol; see below */
    unsigned char rtm_scope; /* See below */
    unsigned char rtm_type; /* See below */
    unsigned int rtm_flags;
};
```

rtm\_type Route type

??

RTN\_UNSPEC unknown route  
 RTN\_UNICAST a gateway or direct route  
 RTN\_LOCAL a local interface route  
 RTN\_BROADCAST a local broadcast route (sent as a broadcast)

RTN\_ANYCAST a local broadcast route (sent as a uni-  
cast)

RTN\_MULTICAST a multicast route

RTN\_BLACKHOLE a packet dropping route

RTN\_UNREACHABLE an unreachable destination

RTN\_PROHIBIT a packet rejection route

RTN\_THROW continue routing lookup in another table

RTN\_NAT a network address translation rule

RTN\_XRESOLVE refer to an external resolver (not im-  
plemented)

rtm\_protocol Route origin

??

RTPROT\_UNSPEC unknown

RTPROT\_REDIRECT by an ICMP redirect (cur-  
rently unused)

RTPROT\_KERNEL by the kernel

RTPROT\_BOOT during boot

RTPROT\_STATIC by the administrator

Values larger than RTPROT\_STATIC are not interpreted by the ker-  
nel, they are just for user information. They may be used to  
tag the source of a routing information or to distinguish be-  
tween multiple routing daemons. See <linux/rtnetlink.h> for the  
routing daemon identifiers which are already assigned.

rtm\_scope is the distance to the destination:

RT\_SCOPE\_UNIVERSE global route

RT\_SCOPE\_SITE interior route in the lo-  
cal autonomous system

RT\_SCOPE\_LINK route on this link

RT\_SCOPE\_HOST route on the local host

RT\_SCOPE\_NOWHERE destination doesn't exist

The values between RT\_SCOPE\_UNIVERSE and RT\_SCOPE\_SITE are  
available to the user.

The rtm\_flags have the following meanings:

RTM\_F\_NOTIFY if the route changes, notify the user via  
rtnetlink

RTM\_F\_CLONED route is cloned from another route

RTM\_F\_EQUALIZE a multipath equalizer (not yet implemented)

rtm\_table specifies the routing table

RT\_TABLE\_UNSPEC an unspecified routing table

RT\_TABLE\_DEFAULT the default table

RT\_TABLE\_MAIN the main table

RT\_TABLE\_LOCAL the local table

The user may assign arbitrary values between RT\_TABLE\_UNSPEC and  
RT\_TABLE\_DEFAULT.

#### Attributes

rtm_type	Value type	Description
RTA_UNSPEC	-	ignored
RTA_DST	protocol address	Route destination address
RTA_SRC	protocol address	Route source address
RTA_IIF	int	Input interface index
RTA_OIF	int	Output interface index
RTA_GATEWAY	protocol address	The gateway of the route
RTA_PRIORITY	int	Priority of route
RTA_PREFSRC	protocol address	Preferred source address
RTA_METRICS	int	Route metric
RTA_MULTIPATH		Multipath nexthop data br (see below).
RTA_PROTOINFO		No longer used
RTA_FLOW	int	Route realm
RTA_CACHEINFO	struct rta_cacheinfo	(see linux/rtnetlink.h)
RTA_SESSION		No longer used
RTA_MP_ALGO		No longer used
RTA_TABLE	int	Routing table ID; if set, rtm_table is ignored
RTA_MARK	int	

RTA\_MFC\_STATS struct rta\_mfc\_stats (see linux/rtnetlink.h)

RTA\_VIA struct rtvia Gateway in different AF  
(see below)

RTA\_NEWDST protocol address Change packet destination  
address

RTA\_PREF char RFC4191 IPv6 router pref?  
erence (see below)

RTA\_ENCAP\_TYPE short Encapsulation type for  
lwtunnels (see below)

RTA\_ENCAP Defined by RTA\_ENCAP\_TYPE

RTA\_EXPIRES int Expire time for IPv6  
routes (in seconds)

RTA\_MULTIPATH contains several packed instances of struct rtnex?

thop together with nested RTAs (RTA\_GATEWAY):

```
struct rtnexthop {
    unsigned short rtnh_len; /* Length of struct + length
                             of RTAs */
    unsigned char rtnh_flags; /* Flags (see
                              linux/rtnetlink.h) */
    unsigned char rtnh_hops; /* Nexthop priority */
    int rtnh_ifindex; /* Interface index for this
                     nexthop */
}
```

There exist a bunch of RTNH\_\* macros similar to RTA\_\* and NL?

HDR\_\* macros useful to handle these structures.

```
struct rtvia {
    unsigned short rtvia_family;
    unsigned char rtvia_addr[0];
};
```

rtvia\_addr is the address, rtvia\_family is its family type.

RTA\_PREF may contain values ICMPV6\_ROUTER\_PREF\_LOW,  
ICMPV6\_ROUTER\_PREF\_MEDIUM, and ICMPV6\_ROUTER\_PREF\_HIGH defined  
incw <linux/icmpv6.h>.

RTA\_ENCAP\_TYPE may contain values LWTUNNEL\_ENCAP\_MPLS, LWTUNNEL\_ENCAP\_IP, LWTUNNEL\_ENCAP\_ILA, or LWTUNNEL\_ENCAP\_IP6 defined in <linux/lwtunnel.h>.

Fill these values in!

RTM\_NEWNEIGH, RTM\_DELNEIGH, RTM\_GETNEIGH

Add, remove, or receive information about a neighbor table entry (e.g., an ARP entry). The message contains an ndmsg structure.

```
struct ndmsg {
    unsigned char ndm_family;

    int          ndm_ifindex; /* Interface index */

    __u16        ndm_state; /* State */

    __u8         ndm_flags; /* Flags */

    __u8         ndm_type;
};

struct nda_cacheinfo {
    __u32        ndm_confirmed;

    __u32        ndm_used;

    __u32        ndm_updated;

    __u32        ndm_refcnt;
};
```

ndm\_state is a bit mask of the following states:

NUD\_INCOMPLETE a currently resolving cache entry  
NUD\_REACHABLE a confirmed working cache entry  
NUD\_STALE an expired cache entry  
NUD\_DELAY an entry waiting for a timer  
NUD\_PROBE a cache entry that is currently reprobed  
NUD\_FAILED an invalid cache entry  
NUD\_NOARP a device with no destination cache  
NUD\_PERMANENT a static entry

Valid ndm\_flags are:

NTF\_PROXY a proxy arp entry  
NTF\_ROUTER an IPv6 router

The rtaattr struct has the following meanings for the rta\_type



field:

NDA\_UNSPEC unknown type

NDA\_DST a neighbor cache n/w layer destination address

NDA\_LLADDR a neighbor cache link layer address

NDA\_CACHEINFO cache statistics

If the rta\_type field is NDA\_CACHEINFO, then a struct nda\_cache?

info header follows.

RTM\_NEWRULE, RTM\_DELRULE, RTM\_GETRULE

Add, delete, or retrieve a routing rule. Carries a struct rtmmsg

RTM\_NEWQDISC, RTM\_DELQDISC, RTM\_GETQDISC

Add, remove, or get a queueing discipline. The message contains a struct tcmsg and may be followed by a series of attributes.

```

struct tcmsg {
    unsigned char    tcm_family;

    int              tcm_ifindex; /* interface index */

    __u32           tcm_handle; /* Qdisc handle */

    __u32           tcm_parent; /* Parent qdisc */

    __u32           tcm_info;

};

```

Attributes

rta_type	Value type	Description
??		
TCA_UNSPEC	-	unspecified
TCA_KIND	asciiz string	Name of queueing discipline
TCA_OPTIONS	byte sequence	Qdisc-specific options follow
TCA_STATS	struct tc_stats	Qdisc statistics
TCA_XSTATS	qdisc-specific	Module-specific statistics
TCA_RATE	struct tc_estimator	Rate limit

In addition, various other qdisc-module-specific attributes are allowed. For more information see the appropriate include files.

RTM\_NEWTCLASS, RTM\_DELTCLASS, RTM\_GETTCLASS

Add, remove, or get a traffic class. These messages contain a

struct tcmsg as described above.

RTM\_NEWTFILTER, RTM\_DELTFILTER, RTM\_GETTFILTER

Add, remove, or receive information about a traffic filter.

These messages contain a struct tcmsg as described above.

## VERSIONS

rtnetlink is a new feature of Linux 2.2.

## BUGS

This manual page is incomplete.

## SEE ALSO

cmsg(3), rtnetlink(3), ip(7), netlink(7)

## COLOPHON

This page is part of release 5.10 of the Linux man-pages project. A description of the project, information about reporting bugs, and the latest version of this page, can be found at <https://www.kernel.org/doc/man-pages/>.

Linux

2020-06-09

RTNETLINK(7)