

**NAME**

netlink - communication between kernel and user space (AF\_NETLINK)

**SYNOPSIS**

```
#include <asm/types.h>
#include <sys/socket.h>
#include <linux/netlink.h>
```

```
netlink_socket = socket(AF_NETLINK, socket_type, netlink_family);
```

**DESCRIPTION**

Netlink is used to transfer information between the kernel and user-space processes. It consists of a standard sockets-based interface for user space processes and an internal kernel API for kernel modules. The internal kernel interface is not documented in this manual page. There is also an obsolete netlink interface via netlink character devices; this interface is not documented here and is provided only for backward compatibility.

Netlink is a datagram-oriented service. Both **SOCK\_RAW** and **SOCK\_DGRAM** are valid values for *socket\_type*. However, the netlink protocol does not distinguish between datagram and raw sockets.

*netlink\_family* selects the kernel module or netlink group to communicate with. The currently assigned netlink families are:

**NETLINK\_ROUTE**

Receives routing and link updates and may be used to modify the routing tables (both IPv4 and IPv6), IP addresses, link parameters, neighbor setups, queueing disciplines, traffic classes and packet classifiers (see [rtnetlink\(7\)](#)).

**NETLINK\_W1** (Linux 2.6.13 to 2.16.17)

Messages from 1-wire subsystem.

**NETLINK\_USERSOCK**

Reserved for user-mode socket protocols.

**NETLINK\_FIREWALL** (up to and including Linux 3.4)

Transport IPv4 packets from netfilter to user space. Used by *ip\_queue* kernel module. After a long period of being declared obsolete (in favor of the more advanced *nfnfnetlink\_queue* feature), **NETLINK\_FIREWALL** was removed in Linux 3.5.

**NETLINK\_INET\_DIAG** (since Linux 2.6.14)

Query information about sockets of various protocol families from the kernel (see [sock\\_diag\(7\)](#)).

**NETLINK\_SOCK\_DIAG** (since Linux 3.3)

A synonym for **NETLINK\_INET\_DIAG**.

**NETLINK\_NFLOG** (up to and including Linux 3.16)

Netfilter/iptables ULOG.

**NETLINK\_XFRM**

IPsec.

**NETLINK\_SELINUX** (since Linux 2.6.4)

SELinux event notifications.

**NETLINK\_ISCSI** (since Linux 2.6.15)

Open-iSCSI.

**NETLINK\_AUDIT** (since Linux 2.6.6)

Auditing.

**NETLINK\_FIB\_LOOKUP** (since Linux 2.6.13)

Access to FIB lookup from user space.

**NETLINK\_CONNECTOR** (since Linux 2.6.14)

Kernel connector. See *Documentation/connector/\** in the Linux kernel source tree for further information.

**NETLINK\_NETFILTER** (since Linux 2.6.14)

Netfilter subsystem.

**NETLINK\_SCSITRANSPORT** (since Linux 2.6.19)

SCSI Transports.

**NETLINK\_RDMA** (since Linux 3.0)

Infiniband RDMA.

**NETLINK\_IP6\_FW** (up to and including Linux 3.4)

Transport IPv6 packets from netfilter to user space. Used by *ip6\_queue* kernel module.

**NETLINK\_DNRTMSG**

DECnet routing messages.

**NETLINK\_KOBJECT\_UEVENT** (since Linux 2.6.10)

Kernel messages to user space.

**NETLINK\_GENERIC** (since Linux 2.6.15)

Generic netlink family for simplified netlink usage.

**NETLINK\_CRYPTO** (since Linux 3.2)

Netlink interface to request information about ciphers registered with the kernel crypto API as well as allow configuration of the kernel crypto API.

Netlink messages consist of a byte stream with one or multiple *nlmsg\_hdr* headers and associated payload. The byte stream should be accessed only with the standard **NLMSG\_\*** macros. See [netlink\(3\)](#) for further information.

In multipart messages (multiple *nlmsg\_hdr* headers with associated payload in one byte stream) the first and all following headers have the **NLM\_F\_MULTI** flag set, except for the last header which has the type **NLMSG\_DONE**.

After each *nlmsg\_hdr* the payload follows.

```
struct nlmsg_hdr {
    __u32 nlmsg_len; /* Length of message including header */
    __u16 nlmsg_type; /* Type of message content */
    __u16 nlmsg_flags; /* Additional flags */
    __u32 nlmsg_seq; /* Sequence number */
    __u32 nlmsg_pid; /* Sender port ID */
};
```

*nlmsg\_type* can be one of the standard message types: **NLMSG\_NOOP** message is to be ignored, **NLMSG\_ERROR** message signals an error and the payload contains an *nlmsgerr* structure, **NLMSG\_DONE** message terminates a multipart message.

```
struct nlmsgerr {
    int error; /* Negative errno or 0 for acknowledgements */
    struct nlmsg_hdr msg; /* Message header that caused the error */
};
```

A netlink family usually specifies more message types, see the appropriate manual pages for that, for example, [rtnetlink\(7\)](#) for **NETLINK\_ROUTE**.

Standard flag bits in *nlmsg\_flags*

---

**NLM\_F\_REQUEST** Must be set on all request messages.

<b>NLM_F_MULTI</b>	The message is part of a multipart message terminated by <b>NLMSG_DONE</b> .
<b>NLM_F_ACK</b>	Request for an acknowledgment on success.
<b>NLM_F_ECHO</b>	Echo this request.

Additional flag bits for GET requests

<b>NLM_F_ROOT</b>	Return the complete table instead of a single entry.
<b>NLM_F_MATCH</b>	Return all entries matching criteria passed in message content. Not implemented yet.
<b>NLM_F_ATOMIC</b>	Return an atomic snapshot of the table.
<b>NLM_F_DUMP</b>	Convenience macro; equivalent to (NLM_F_ROOT NLM_F_MATCH).

Note that **NLM\_F\_ATOMIC** requires the **CAP\_NET\_ADMIN** capability or an effective UID of 0.

Additional flag bits for NEW requests

<b>NLM_F_REPLACE</b>	Replace existing matching object.
<b>NLM_F_EXCL</b>	Don't replace if the object already exists.
<b>NLM_F_CREATE</b>	Create object if it doesn't already exist.
<b>NLM_F_APPEND</b>	Add to the end of the object list.

*nlmsg\_seq* and *nlmsg\_pid* are used to track messages. *nlmsg\_pid* shows the origin of the message. Note that there isn't a 1:1 relationship between *nlmsg\_pid* and the PID of the process if the message originated from a netlink socket. See the **ADDRESS FORMATS** section for further information.

Both *nlmsg\_seq* and *nlmsg\_pid* are opaque to netlink core.

Netlink is not a reliable protocol. It tries its best to deliver a message to its destination(s), but may drop messages when an out-of-memory condition or other error occurs. For reliable transfer the sender can request an acknowledgement from the receiver by setting the **NLM\_F\_ACK** flag. An acknowledgment is an **NLMSG\_ERROR** packet with the error field set to 0. The application must generate acknowledgements for received messages itself. The kernel tries to send an **NLMSG\_ERROR** message for every failed packet. A user process should follow this convention too.

However, reliable transmissions from kernel to user are impossible in any case. The kernel can't send a netlink message if the socket buffer is full: the message will be dropped and the kernel and the user-space process will no longer have the same view of kernel state. It is up to the application to detect when this happens (via the **ENOBUFS** error returned by `recvmsg(2)`) and resynchronize.

### Address formats

The *sockaddr\_nl* structure describes a netlink client in user space or in the kernel. *Asoc kaddr\_nl* can be either unicast (only sent to one peer) or sent to netlink multicast groups (*nl\_groups* not equal 0).

```
struct sockaddr_nl {
    sa_family_t nl_family; /* AF_NETLINK */
    unsigned short nl_pad; /* Zero */
    pid_t nl_pid; /* Port ID */
    __u32 nl_groups; /* Multicast groups mask */
};
```

*nl\_pid* is the unicast address of netlink socket. It's always 0 if the destination is in the kernel. For a user-space process, *nl\_pid* is usually the PID of the process owning the destination socket. However, *nl\_pid* identifies a netlink socket, not a process. If a process owns several netlink sockets, then *nl\_pid* can be equal to the process ID only for at most one socket. There are two ways to assign *nl\_pid* to a netlink socket. If the application sets *nl\_pid* before calling `bind(2)`, then it is up to the application to make sure that *nl\_pid* is unique. If the application sets it to 0, the kernel takes care of assigning it. The kernel assigns the process ID to the first netlink socket the process opens and assigns a unique *nl\_pid* to every netlink socket that the process subsequently creates.

*nl\_groups* is a bit mask with every bit representing a netlink group number. Each netlink family has a set of 32 multicast groups. When `bind(2)` is called on the socket, the *nl\_groups* field in the *sockaddr\_nl* should be set to a bit mask of the groups which it wishes to listen to. The default value for this field is zero which means that no multicasts will be received. A socket may multicast messages to any of the multicast groups by setting *nl\_groups* to a bit mask of the groups it wishes to send to when it calls `sendmsg(2)` or does a `connect(2)`. Only processes with an effective UID of 0 or the `CAP_NET_ADMIN` capability may send or listen to a netlink multicast group. Since Linux 2.6.13, messages can't be broadcast to multiple groups. Any replies to a message received for a multicast group should be sent back to the sending PID and the multicast group. Some Linux kernel subsystems may additionally allow other users to send and/or receive messages. As at Linux 3.0, the `NETLINK_KOBJECT_UEVENT`, `NETLINK_GENERIC`, `NETLINK_ROUTE`, and `NETLINK_SELINUX` groups allow other users to receive messages. No groups allow other users to send messages.

### Socket options

To set or get a netlink socket option, call `getsockopt(2)` to read or `setsockopt(2)` to write the option with the option level argument set to `SOL_NETLINK`. Unless otherwise noted, *optval* is a pointer to an *int*.

**NETLINK\_PKTINFO** (since Linux 2.6.14)

Enable `nl_pktinfo` control messages for received packets to get the extended destination group number.

**NETLINK\_ADD\_MEMBERSHIP**, **NETLINK\_DROP\_MEMBERSHIP** (since Linux 2.6.14)

Join/leave a group specified by *optval*.

**NETLINK\_LIST\_MEMBERSHIPS** (since Linux 4.2)

Retrieve all groups a socket is a member of. *optval* is a pointer to `__u32` and *optlen* is the size of the array. The array is filled with the full membership set of the socket, and the required array size is returned in *optlen*.

**NETLINK\_BROADCAST\_ERROR** (since Linux 2.6.30)

When not set, `netlink_broadcast()` only reports `ESRCH` errors and silently ignore `NOBUFS` errors.

**NETLINK\_NO\_ENOBUFS** (since Linux 2.6.30)

This flag can be used by unicast and broadcast listeners to avoid receiving `ENOBUFS` errors.

**NETLINK\_LISTEN\_ALL\_NSID** (since Linux 4.2)

When set, this socket will receive netlink notifications from all network namespaces that have an *nsid* assigned into the network namespace where the socket has been opened. The *nsid* is sent to user space via an ancillary data.

**NETLINK\_CAP\_ACK** (since Linux 4.2)

The kernel may fail to allocate the necessary room for the acknowledgment message back to user space. This option trims off the payload of the original netlink message. The netlink message header is still included, so the user can guess from the sequence number which message triggered the acknowledgment.

### VERSIONS

The socket interface to netlink first appeared Linux 2.2.

Linux 2.0 supported a more primitive device-based netlink interface (which is still available as a compatibility option). This obsolete interface is not described here.

### NOTES

It is often better to use netlink via *libnetlink* or *libnl* than via the low-level kernel interface.

### BUGS

This manual page is not complete.

### EXAMPLE

The following example creates a `NETLINK_ROUTE` netlink socket which will listen to the `RTMGRP_LINK` (network interface create/delete/up/down events) and `RTMGRP_IPV4_IFADDR` (IPv4

addresses add/delete events) multicast groups.

```
struct sockaddr_nl sa;

memset(&sa, 0, sizeof(sa));
sa.nl_family = AF_NETLINK;
sa.nl_groups = RTMGRP_LINK | RTMGRP_IPV4_IFADDR;

fd = socket(AF_NETLINK, SOCK_RAW, NETLINK_ROUTE);
bind(fd, (struct sockaddr *) &sa, sizeof(sa));
```

The next example demonstrates how to send a netlink message to the kernel (pid 0). Note that the application must take care of message sequence numbers in order to reliably track acknowledgements.

```
struct nlmsg_hdr *nh; /* The nlmsg_hdr with payload to send */
struct sockaddr_nl sa;
struct iovec iov = { nh, nh->nlmsg_len };
struct msghdr msg;

msg = { &sa, sizeof(sa), &iov, 1, NULL, 0, 0 };
memset(&sa, 0, sizeof(sa));
sa.nl_family = AF_NETLINK;
nh->nlmsg_pid = 0;
nh->nlmsg_seq = ++sequence_number;
/* Request an ack from kernel by setting NLM_F_ACK */
nh->nlmsg_flags |= NLM_F_ACK;

sendmsg(fd, &msg, 0);
```

And the last example is about reading netlink message.

```
int len;
char buf[4096];
struct iovec iov = { buf, sizeof(buf) };
struct sockaddr_nl sa;
struct msghdr msg;
struct nlmsg_hdr *nh;

msg = { &sa, sizeof(sa), &iov, 1, NULL, 0, 0 };
len = recvmsg(fd, &msg, 0);

for (nh = (struct nlmsg_hdr *) buf; NLMSG_OK (nh, len);
     nh = NLMSG_NEXT (nh, len)) {
    /* The end of multipart message */
    if (nh->nlmsg_type == NLMSG_DONE)
        return;

    if (nh->nlmsg_type == NLMSG_ERROR)
        /* Do some error handling */

    /* Continue with parsing payload */
}
```

## SEE ALSO

[cmsg\(3\)](#), [netlink\(3\)](#), [capabilities\(7\)](#), [rtnetlink\(7\)](#), [sock\\_diag\(7\)](#)

[information about libnetlink](#)

[information about libnl](#)

RFC 3549 "Linux Netlink as an IP Services Protocol"

**COLOPHON**

This page is part of release 4.10 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/>.