Network Device Initialization Example

Initialization begins in the device driver.

```
JMW:eepro100_init_one initializing net_device
JMW:speedo_found1 (eepro100) about to call init_etherdev
```

The *init_etherdev()* function calls *register_netdevice()*

```
JMW:register_netdevice registered device c1840800
```

The *register_netdevice()* function calls the *netdev_chain* of notifiers. This results in an indirect call to *rtnetlink_event()* with the notifier code of 5 representing NETDEV_REGISTER.

```
JMW:rtnetlink_event was called with 5
```

The *rtnetlink_fill_ifinfo()* function is called to compose and send a netlink notifier message that the interface is now up. The caller of *rtnetlink_fill_ifinfo()* here is *rtmsg_ifinfo()*.

```
JMW:rtnetlink_fill_ifinfo was called from c0254814 with c1840800
```

After the message is sent to the netlink socket, control returns to the device driver where initialization of the hardware continues.

```
JMW:speedo_found1 (eepro100) dev at c1840800
```

The netlink notifier message is received asyncronously by *rtnetlink_rcv()* which then calls *rtnetlink_rcv_skb()*.

```
JMW:rtnetlink_rcv_skb skb was received
```

The message is forwarded to *rtnetlink_rcv_msg()*.

```
JMW:rtnetlink_rcv_msg message was received
JMW:rtnetlink_rcv_msg message family is 11 type is 2
```

Here *rtnetlink_fill_ifinfo* is called from *rtnetlink_dump_ifinfo()*

```
JMW:rtnetlink_fill_ifinfo was called from c0254715 with c0313840
JMW:rtnetlink_fill_ifinfo was called from c0254715 with c1840800
```

This appears to be the end of the chain of events that are triggered be NETDEV_REGISTER

Here the devinet_ioctl() is invoked from the socket layer with a request to change the flags associated with the interface. We assume call is the result of the ifconfig command being issued. The device here is most likely the loopback.

```
JMW:devinet_ioctl called from c0247fce
JMW:devinet_ioctl SIOCSIFFLAGS calling change flags for c0313840
JMW:dev_change_flags called for c0313840 from c027eaaf w/ 9
JMW:dev_open called for c0313840 from c024f65d
JMW:dev_open calling notifier for c0313840
```

The dev_open() function eventually calls the NETDEV notifier chain with event code 1 (NETDEV_UP) and this leads to the activation of rtnetlink_event().

```
JMW:rtnetlink_event was called with 1
```

Oddly, rtnetlink_event() responds identically to UP and DOWN events.

```
497 case NETDEV_UP:
498 case NETDEV_DOWN:
499     rtmsg_ifinfo(RTM_NEWLINK, dev, IFF_UP|IFF_RUNNING);
500 break;
```

The caller here is rtmsg_ifinfo().

```
JMW:rtnetlink_fill_ifinfo was called from c0254814 with c0313840
JMW:rtnetlink_fill_ifinfo was called from c0254814 with c0313840
```

Eventually the message constructed here is broadcast on the netlink socket.

```
JMW:rtnetlink_rcv_skb skb was received
JMW:rtnetlink_rcv_msg message was received
JMW:rtnetlink_rcv_msg message family is 0 type is 2
```

Here rtnetlink_fill_ifinfo is called from rtnetlink_dump_ifinfo(). The values c0313840and c1840800 are pointers to the struct netdevice associated with interfaces lo and eth0 respectively.

```
JMW:rtnetlink_fill_ifinfo was called from c0254715 with c0313840
JMW:rtnetlink_fill_ifinfo was called from c0254715 with c1840800
```

Message family 0 is AF_UNSPEC. Message types 6 and 2 are RTM_GETADDR and RTM_GETLINK respectively.

```
JMW:rtnetlink_rcv_skb skb was received
JMW:rtnetlink_rcv_msg message was received
JMW:rtnetlink_rcv_msg message family is 0 type is 6
JMW:rtnetlink_rcv_skb skb was received
JMW:rtnetlink_rcv_msg message was received
JMW:rtnetlink_rcv_msg message family is 0 type is 2
```

Again rtnetlink_fill_ifinfo is called from rtnetlink_dump_ifinfo(). The values c0313840and c1840800 are pointers to the struct netdevice associated with interfaces lo and eth0 respectively.

```
JMW:rtnetlink_fill_ifinfo was called from c0254715 with c0313840
JMW:rtnetlink_fill_ifinfo was called from c0254715 with c1840800
```

This call appears to produce messages of family 2 (AF_INET) and type 4 RTM_NEWADDR.

```
JMW:rtnetlink_rcv_skb skb was received
JMW:rtnetlink_rcv_msg message was received
JMW:rtnetlink_rcv_msg message family is 2 type is 4
```

Receipt of the RTM_NEWADDR message finally triggers a call to the doit() routine which in this case is inet_rtm_newaddr().

```
JMW:rtnetlink_rcv_msg calling do it at c027e590
The doit function is: c027e590 T inet_rtm_newaddr

JMW:inet_rtm_newaddr dev is c0313840
JMW:inet_rtm_newaddr calling inetdev_init
JMW:inet_rtm_newaddr in_dev is dfed47c0
JMW:inet_rtm_newaddr ifa is c184a4e0
JMW:inet_rtm_newaddr label is lo
```

The next sequence of events appears to have been triggered by the call to dev_open() which called the netdev notifier chain with parameter NETDEV_UP. There are only two known callers of fib_add_ifaddr(). These are fib_inetaddr_event() and fib_netdev_event(). As shown on the next page entry to fib_inetaddr_event() appears to generate a message. Probably the printk was placed after the call but this needs to be checked. The parameter c184a4e0 is a pointer to the associated struct in_ifaddr. The dst fields are IP addresses and here it can be seen that all of these belong to the loopback (lo) interface.

```
JMW:fib_add_ifaddr called from c028327b with c184a4e0
JMW:fib_magic called from c0283005 with c184a4e0
JMW:fib_magic called with cmd 18, type 2, dst 100007f, len 20
JMW:fib_magic called from c0283033 with c184a4e0
JMW:fib_magic called with cmd 18, type 3, dst ffffff7f, len 20
JMW:fib_magic called from c028306e with c184a4e0
JMW:fib_magic called with cmd 18, type 2, dst 7f, len 8
JMW:fib_magic called from c028308c with c184a4e0
JMW:fib_magic called with cmd 18, type 3, dst 7f, len 20
JMW:fib_magic called from c0283099 with c184a4e0
JMW:fib_magic called with cmd 18, type 3, dst ffffff7f, len 20

JMW:fib_inetaddr_event called for ifa c184a4e0 w/ UP
```

```
JMW:rtnetlink_rcv_skb skb was received
JMW:rtnetlink_rcv_msg message was received
JMW:rtnetlink_rcv_msg message family is 0 type is 2
```

Message type 2 is RTM_GETLINK

```
JMW:rtnetlink_fill_ifinfo was called from c0254715 with c0313840
JMW:rtnetlink_fill_ifinfo was called from c0254715 with c1840800
JMW:rtnetlink_rcv_skb skb was received
JMW:rtnetlink_rcv_msg message was received
```

The message family here is AF_INET and the type of 8 is RTM_NEWROUTE

```
JMW:rtnetlink_rcv_msg message family is 2 type is 8
```

Here the "doit" routine at `c0282d50` is actually inet_rtm_newroute()

```
JMW:rtnetlink_rcv_msg calling doit at c0282d50
```

This call to devinet_ioctl() is presumably the result of running the ifconfig command on the eth0 interface.

```
JMW:devinet_ioctl called from c0247fce
```

The source of these messages remains unclear.

```
JMW:rtnetlink_rcv_skb skb was received
JMW:rtnetlink_rcv_msg message was received
JMW:rtnetlink_rcv_msg message family is 11 type is 2
JMW:rtnetlink_fill_ifinfo was called from c0254715 with c0313840
JMW:rtnetlink_fill_ifinfo was called from c0254715 with c1840800
JMW:rtnetlink_rcv_skb skb was received
JMW:rtnetlink_rcv_msg message was received
JMW:rtnetlink_rcv_msg message family is 11 type is 2
JMW:rtnetlink_fill_ifinfo was called from c0254715 with c0313840
JMW:rtnetlink_fill_ifinfo was called from c0254715 with c1840800
```

As seen earlier with the loopback device we see the calls to dev_change_flags() and dev_open().

```
JMW:devinet_ioctl called from c0247fce
JMW:devinet_ioctl SIOCSIFFLAGS calling change flags for c1840800
JMW:dev_change_flags called for c1840800 from c027eaaf w/ 1003
JMW:dev_open called for c1840800 from c024f65d
JMW:dev_open calling notifier for c1840800
```

This parameter value of 1 is NETDEV_UP and the struct netdevice is that of eth0.

```
JMW:rtnetlink_event was called with 1
JMW:rtnetlink_fill_ifinfo was called from c0254814 with c1840800
JMW:rtnetlink_fill_ifinfo was called from c0254814 with c1840800
```

As before receipt of the first message appears to trigger synthesis of additional messages.

```
JMW:rtnetlink_rcv_skb skb was received
JMW:rtnetlink_rcv_msg message was received
JMW:rtnetlink_rcv_msg message family is 0 type is 2

JMW:rtnetlink_fill_ifinfo was called from c0254715 with c0313840
JMW:rtnetlink_fill_ifinfo was called from c0254715 with c1840800

JMW:rtnetlink_rcv_skb skb was received
JMW:rtnetlink_rcv_msg message was received
JMW:rtnetlink_rcv_msg message family is 0 type is 6

JMW:rtnetlink_rcv_skb skb was received
JMW:rtnetlink_rcv_msg message was received
JMW:rtnetlink_rcv_msg message family is 0 type is 2

JMW:rtnetlink_fill_ifinfo was called from c0254715 with c0313840
JMW:rtnetlink_fill_ifinfo was called from c0254715 with c1840800
```

And after the exchange of mystery messages finally the RTM_NEWADDR message appears.

```
JMW:rtnetlink_rcv_skb skb was received
JMW:rtnetlink_rcv_msg message was received
JMW:rtnetlink_rcv_msg message family is 2 type is 4
JMW:rtnetlink_rcv_msg calling do it at c027e590
JMW:inet_rtm_newaddr dev is c1840800
JMW:inet_rtm_newaddr calling inetdev_init
JMW:inet_rtm_newaddr in_dev is df52e6e0
JMW:inet_rtm_newaddr ifa is c184a460
JMW:inet_rtm_newaddr label is eth0
```

The next sequence of events appears to have been triggered by the call to dev_open() which called the netdev notifier chain with parameter NETDEV_UP. There are only two known callers of fib_add_ifaddr(). These are fib_inetaddr_event() and fib_netdev_event(). The caller here appears to be fib_inetaddr_event() as it is located at address c0283260. Here we see fib_magic registering the addresses associated with the eth0 interface. The network address actually being used is 192.168.2.32 and the device addresses used are .36, .63, .32, .32, and .63

```
JMW:fib_add_ifaddr called from c028327b with c184a460
JMW:fib_magic called from c0283005 with c184a460
JMW:fib_magic called with cmd 18, type 2, dst 2402a8c0, len 20
JMW:fib_magic called from c0283033 with c184a460
JMW:fib_magic called with cmd 18, type 3, dst 3f02a8c0, len 20
JMW:fib_magic called from c028306e with c184a460
JMW:fib_magic called with cmd 18, type 1, dst 2002a8c0, len 1b
JMW:fib_magic called from c028308c with c184a460
JMW:fib_magic called with cmd 18, type 3, dst 2002a8c0, len 20
JMW:fib_magic called from c0283099 with c184a460
JMW:fib_magic called with cmd 18, type 3, dst 3f02a8c0, len 20
```

Presumably the printk that generated this was accidentally placed after the call to fib_add_ifaddr().

```
JMW:fib_inetaddr_event called for ifa c184a460 w/ UP

JMW:rtnetlink_rcv_skb skb was received
JMW:rtnetlink_rcv_msg message was received
JMW:rtnetlink_rcv_msg message family is 0 type is 2

JMW:rtnetlink_fill_ifinfo was called from c0254715 with c0313840
JMW:rtnetlink_fill_ifinfo was called from c0254715 with c1840800

JMW:rtnetlink_rcv_skb skb was received
JMW:rtnetlink_rcv_msg message was received
JMW:rtnetlink_rcv_msg message family is 2 type is 8
JMW:rtnetlink_rcv_msg calling do it at c0282d50

JMW:rtnetlink_rcv_skb skb was received
JMW:rtnetlink_rcv_msg message was received
JMW:rtnetlink_rcv_msg message family is 2 type is 8
JMW:rtnetlink_rcv_msg calling do it at c0282d50

JMW:devinet_ioctl called from c0247fce
JMW:devinet_ioctl called from c0247fce
```