

NAME

if_ntb - Virtual Ethernet interface for Non-Transparent Bridges

SYNOPSIS

To compile this driver into your kernel, place the following lines in your kernel configuration file:

```
device ntb
device ntb_transport
device if_ntb
```

Or, to load the driver as a module at boot, place the following line in loader.conf(5):

```
if_ntb_load="YES"
```

The following tunables are settable from the loader(8):

hw.if_ntb.num_queues

Limits maximal number of queues per interface. Default is unlimited.

DESCRIPTION

The **if_ntb** driver attaches on top of the **ntb_transport(4)** driver to utilize one or more of its packet queues to create virtual Ethernet network interface between the systems. Typical MTU for the interface is about 64KB to reduce overhead. Default MAC address for the interface is randomly generated.

The **if_ntb** driver does not implement any real hardware offload, but since PCIe link is protected by CRC32, in some situations it may be possible to save some CPU cycles by enabling fake checksum offload on both link sides via setting **rxchecksum** and **txchecksum** interface options.

SEE ALSO

ntb_transport(4)

AUTHORS

The **if_ntb** driver was developed by Intel and originally written by Carl Delsey <carl@FreeBSD.org>. Later improvements were done by Conrad E. Meyer <cem@FreeBSD.org> and Alexander Motin <mav@FreeBSD.org>.

BUGS

Linux supports only one queue per interface, so manual configuration may be required for compatibility.