

NAME

ixl - Intel Ethernet 700 Series Driver

SYNOPSIS

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

```
device iflib  
device ixl
```

To load the driver as a module at boot time, place the following lines in loader.conf(5):

```
if_ixl_load="YES"
```

DESCRIPTION**Features**

The **ixl** driver provides support for any PCI Express adapter or LOM (LAN On Motherboard) in the Intel Ethernet 700 Series. As of this writing, the series includes devices with these model numbers:

- ◆ XL710 (40G)
- ◆ X710 (10G)
- ◆ XXV710 (25G)
- ◆ X722 (10G)

The driver supports Jumbo Frames, TX/RX checksum offload, TCP segmentation offload (TSO), Large Receive Offload (LRO), VLAN tag insertion/extraction, VLAN checksum offload, VLAN TSO, and Receive Side Steering (RSS), all for both IPv4 and IPv6. For further hardware information and questions related to hardware requirements, see <http://support.intel.com/>.

Support for Jumbo Frames is provided via the interface MTU setting. Selecting an MTU larger than 1500 bytes with the `ifconfig(8)` utility configures the adapter to receive and transmit Jumbo Frames. The maximum MTU size for Jumbo Frames is 9706.

Offloads are also controlled via the interface, for instance, checksumming for both IPv4 and IPv6 can be set and unset, TSO4 and/or TSO6, and finally LRO can be set and unset.

For more information on configuring this device, see `ifconfig(8)`.

Additional Utilities

There are additional tools available from Intel to help configure and update the adapters covered by this driver. These tools can be downloaded directly from Intel at <https://downloadcenter.intel.com>, by

searching for their names, or by installing certain packages:

- To change the behavior of the QSFP+ ports on XL710 adapters, use the Intel QCU (QSFP+ configuration utility); installed by the *sysutils/intel-qcu* package.
- To update the firmware on an adapter, use the Intel Non-Volatile Memory (NVM) Update Utility; installed by the *sysutils/intel-nvmupdate-10g*, *sysutils/intel-nvmupdate-40g*, or *sysutils/intel-nvmupdate-100g*, package.
- Drivers are provided by Intel outside of the FreeBSD kernel; install the *net/intel-ixl-kmod* package for the latest driver.

HARDWARE

Most adapters in the Intel Ethernet 700 Series with SFP+/SFP28/QSFP+ cages have firmware that requires that Intel qualified modules are used; these qualified modules are listed below. This qualification check cannot be disabled by the driver.

The **ixl** driver supports 40Gb Ethernet adapters with these QSFP+ modules:

- Intel 4x10G/40G QSFP+ 40GBASE-SR4 E40GQSFP4SR
- Intel 4x10G/40G QSFP+ 40GBASE-LR4 E40GQSFP4LR

The **ixl** driver supports 25Gb Ethernet adapters with these SFP28 modules:

- Intel 10G/25G SFP28 25GBASE-SR E25GSFP28SR
- Intel 10G/25G SFP28 25GBASE-SR E25GSFP28SRX (Extended Temp)

The **ixl** driver supports 25Gb and 10Gb Ethernet adapters with these SFP+ modules:

- Intel 1G/10G SFP+ SR FTLX8571D3BCV-IT
- Intel 1G/10G SFP+ SR AFBR-703SDZ-IN2
- Intel 1G/10G SFP+ LR FTLX1471D3BCV-IT
- Intel 1G/10G SFP+ LR AFCT-701SDZ-IN2
- Intel 1G/10G SFP+ 10GBASE-SR E10GSFP4SR
- Intel 10G SFP+ 10GBASE-SR E10GSFP4SRX (Extended Temp)
- Intel 1G/10G SFP+ 10GBASE-LR E10GSFP4LR

Note that adapters also support all passive and active limiting direct attach cables that comply with SFF-8431 v4.1 and SFF-8472 v10.4 specifications.

This is not an exhaustive list; please consult product documentation for an up-to-date list of supported media.

LOADER TUNABLES

Tunables can be set at the loader(8) prompt before booting the kernel or stored in loader.conf(5).

hw.ixl.rx_itr

The RX interrupt rate value, set to 62 (124 usec) by default.

hw.ixl.tx_itr

The TX interrupt rate value, set to 122 (244 usec) by default.

hw.ixl.i2c_access_method

Access method that driver will use for I2C read and writes via sysctl(8) or verbose ifconfig(8) information display:

- 0 - best available method
- 1 - bit bang via I2CPARAMS register
- 2 - register read/write via I2CCMD register
- 3 - Use Admin Queue command (default best)

Using the Admin Queue is only supported on 710 devices with FW version 1.7 or newer. Set to 0 by default.

hw.ixl.enable_tx_fc_filter

Filter out packets with Ethertype 0x8808 from being sent out by non-adapter sources. This prevents (potentially untrusted) software or iavf(4) devices from sending out flow control packets and creating a DoS (Denial of Service) event. Enabled by default.

hw.ixl.enable_head_writeback

When the driver is finding the last TX descriptor processed by the hardware, use a value written to memory by the hardware instead of scanning the descriptor ring for completed descriptors. Enabled by default; disable to mimic the TX behavior found in ixgbe(4).

SYSCTL PROCEDURES

dev.ixl.#.fc

Sets the 802.3x flow control mode that the adapter will advertise on the link. A value of 0 disables flow control, 3 enables full, 1 is RX, and 2 is TX pause.

The negotiated flow control setting can be viewed in ifconfig(8), in the interface's media field.

dev.ixl#.advertise_speed

Set the speeds that the interface will advertise on the link. *dev.ixl#.supported_speeds* contains the speeds that are allowed to be set.

dev.ixl#.current_speed

Displays the current speed.

dev.ixl#.fw_version

Displays the current firmware and NVM versions of the adapter.

dev.ixl#.debug.switch_vlans

Set the Ethertype used by the hardware itself to handle internal services. Frames with this Ethertype will be dropped without notice. Defaults to 0x88a8, which is a well known number for IEEE 802.1ad VLAN stacking. If you need 802.1ad support, set this number to any another Ethertype i.e. 0xffff.

INTERRUPT STORMS

It is important to note that 40G operation can generate high numbers of interrupts, often incorrectly being interpreted as a storm condition in the kernel. It is suggested that this be resolved by setting:

hw.intr_storm_threshold: 0

IOVCTL OPTIONS

The driver supports additional optional parameters for created VFs (Virtual Functions) when using `iovctl(8)`:

mac-addr (unicast-mac)

Set the Ethernet MAC address that the VF will use. If unspecified, the VF will use a randomly generated MAC address.

mac-anti-spoof (bool)

Prevent the VF from sending Ethernet frames with a source address that does not match its own.

allow-set-mac (bool)

Allow the VF to set its own Ethernet MAC address

allow-promisc (bool)

Allow the VF to inspect all of the traffic sent to the port.

num-queues (uint16_t)

Specify the number of queues the VF will have. By default, this is set to the number of MSI-X vectors supported by the VF minus one.

An up to date list of parameters and their defaults can be found by using `iovctl(8)` with the `-S` option.

SUPPORT

For general information and support, go to the Intel support website at: <http://support.intel.com/>.

If an issue is identified with this driver with a supported adapter, email all the specific information related to the issue to freebsd@intel.com.

SEE ALSO

`arp(4)`, `iavf(4)`, `iflib(4)`, `netintro(4)`, `vlan(4)`, `ifconfig(8)`, `iovctl(8)`

HISTORY

The `ixl` device driver first appeared in FreeBSD 10.1. It was converted to use `iflib(9)` in FreeBSD 12.

AUTHORS

The `ixl` driver was written by Jack Vogel <jfv@freebsd.org> and Eric Joyner <erj@freebsd.org>.