

**NAME**

**ieee802.3** -- IEEE 802.3 Ethernet parameters and statistics

**DESCRIPTION**

The IEEE 802.3 standard specifies the details for Ethernet networking. This page describes the various statistics and tunables that device drivers supporting Ethernet commonly offer. Note that not every device or driver supports every one of these values, and many devices offer additional statistics and tunables that are specific to that hardware. See the device driver's documentation for those specific details.

Values that are statistics are visible `kstat(1M)`, whereas properties are visible using the `dladm(1M)` **show-linkprop** subcommand. Tunables are properties that can be changed using the `dladm(1M)` **set-linkprop** subcommand. A more useful summary of current operational state can be seen with the `dladm(1M)` **show-ether** subcommand.

**Statistics**

The following statistics are accessible with `kstat(1M)`. Note that some statistics are available in both 32- and 64-bit counters, in which case the name of the 64 bit statistic will be the same as the 32-bit, but with "64" appended". For example, **ipackets64** is the 64-bit version of the **ipackets** statistic. These are indicated with the special suffix [64] in the table below.

**adv\_cap\_1000fdx**

Advertises 1000 Mbps full-duplex support.

**adv\_cap\_1000hdx**

Advertises 1000 Mbps half-duplex support.

**adv\_cap\_100fdx** Advertises 100 Mbps full-duplex support.

**adv\_cap\_100hdx** Advertises 100 Mbps half-duplex support.

**adv\_cap\_100T4** Advertises 100BASE-T4 support.

**adv\_cap\_10fdx** Advertises 10 Mbps full-duplex support.

**adv\_cap\_10gfdx** Advertises 10 Gbps support.

**adv\_cap\_10hdx** Advertises 10 Mbps half-duplex support.

**adv\_cap\_autoneg** Advertises auto-negotiation support.

<b>adv_cap_asmpause</b>	Advertises asymmetric flow control support.
<b>adv_cap_pause</b>	Advertises flow control support.
<b>adv_rem_fault</b>	Remote fault status sent to peer.
<b>align_errors</b>	Mis-aligned frames received.
<b>brdcstrev</b>	Broadcast frames received.
<b>brdcstxmt</b>	Broadcast frames transmitted.
<b>cap_1000fdx</b>	Device supports 1000 Mbps full-duplex.
<b>cap_1000hdx</b>	Device supports 1000 Mbps half-duplex.
<b>cap_100fdx</b>	Device supports 100 Mbps full-duplex.
<b>cap_100hdx</b>	Device supports 100 Mbps half-duplex.
<b>cap_100T4</b>	Device supports 100BASE-T4.
<b>cap_10fdx</b>	Device supports 10 Mbps full-duplex.
<b>cap_10gfdx</b>	Device supports 10 Gpbs.
<b>cap_10hdx</b>	Device supports 10 Mbps half-duplex.
<b>cap_asmpause</b>	Device supports asymmetric flow control.
<b>cap_autoneg</b>	Device supports auto-negotiation.
<b>cap_pause</b>	Device supports symmetric flow control.
<b>cap_rem_fault</b>	Device supports remote fault notification.
<b>carrier_errors</b>	Frames dropped due to loss of link.
<b>collisions</b>	Collisions.

<b>defer_xmts</b>	Transmits deferred due to link activity.															
<b>ex_collisions</b>	Frames dropped due to too many collisions.															
<b>fcs_errors</b>	Frames received with bad frame checksum.															
<b>first_collisions</b>	Frames with at least one collision.															
<b>iererrors</b>	Receive errors.															
<b>ifspeed</b>	Link speed in bits per second.															
<b>ipackets[64]</b>	Frames received successfully.															
<b>jabber_errors</b>	Jabber errors.															
<b>link_asmpause</b>	Asymmetric flow control; works together with <b>link_pause</b> . See the the description for it below.															
<b>link_autoneg</b>	Link was auto-negotiated.															
<b>link_duplex</b>	Link duplex status, values as follows:  0   Unknown. 1   Half-duplex. 2   Full-duplex.															
<b>link_pause</b>	Link flow control available; works together with <b>link_asmpause</b> . The meanings of these bits are:  <table><tr><td><b>pause</b></td><td><b>asmpause</b></td><td><b>meaning</b></td></tr><tr><td>0</td><td>0</td><td>No flow control.</td></tr><tr><td>1</td><td>0</td><td>Symmetric flow control.</td></tr><tr><td>0</td><td>1</td><td>Honor received pause frames.</td></tr><tr><td>1</td><td>1</td><td>Send pause frames when congested.</td></tr></table>	<b>pause</b>	<b>asmpause</b>	<b>meaning</b>	0	0	No flow control.	1	0	Symmetric flow control.	0	1	Honor received pause frames.	1	1	Send pause frames when congested.
<b>pause</b>	<b>asmpause</b>	<b>meaning</b>														
0	0	No flow control.														
1	0	Symmetric flow control.														
0	1	Honor received pause frames.														
1	1	Send pause frames when congested.														
<b>link_state</b>	Link state; 0 for down, 1 for up.															
<b>link_up</b>	Link is up if 1.															

<b>lp_cap_1000fdx</b>	Peer supports 1000 Mbps full-duplex.
<b>lp_cap_1000hdx</b>	Peer supports 1000 Mbps half-duplex.
<b>lp_cap_100fdx</b>	Peer supports 100 Mbps full-duplex.
<b>lp_cap_100hdx</b>	Peer supports 100 Mbps half-duplex.
<b>lp_cap_100T4</b>	Peer supports 100BASE-T4.
<b>lp_cap_10fdx</b>	Peer supports 10 Mbps full-duplex.
<b>lp_cap_10gfdx</b>	Peer supports 10 Gbps.
<b>lp_cap_10hdx</b>	Peer supports 10 Mbps half-duplex.
<b>lp_cap_asmpause</b>	Peer supports asymmetric flow control.
<b>lp_cap_autoneg</b>	Peer supports auto-negotiation.
<b>lp_cap_pause</b>	Peer advertises flow control support.
<b>lp_rem_fault</b>	Peer announces a remote fault.
<b>macrv_errors</b>	Generic receive errors.
<b>macxmt_errors</b>	Generic transmit errors.
<b>multi_collisions</b>	Frames with more than one collision.
<b>multircv</b>	Multicast frames received.
<b>multixmt</b>	Multicast frames transmitted.
<b>norecvbuf</b>	Receive frames dropped due to lack of resources.
<b>noxmtbuf</b>	Transmit frames dropped due to lack of resources.
<b>obytes[64]</b>	Bytes (octets) transmitted successfully.

<b>oerrors</b>	Transmit errors.
<b>oflo</b>	Overflow errors.
<b>opackets[64]</b>	Frames successfully transmitted.
<b>promisc</b>	Interface is in promiscuous mode.
<b>rbytes[64]</b>	Bytes (octets) received successfully.
<b>runt_errors</b>	Frames received that were too short.
<b>sqe_errors</b>	Squelch errors.
<b>toolong_errors</b>	Frames received that were too long.
<b>tx_late_collisions</b>	Late collisions on transmit.
<b>uflo</b>	Underflow errors.
<b>unknowns</b>	Frames received with no local recipient.
<b>xcvr_addr</b>	Transceiver address.
<b>xcvr_id</b>	Transceiver vendor and device ID.
<b>xcvr_inuse</b>	Identifies the type of transceiver in use. Values are as follows:
	0 Unknown or undefined.
	1 None.
	2 10 Mbps
	3 100BASE-T4
	4 100BASE-X
	5 100BASE-T2
	6 1000BASE-X
	7 1000BASE-T

### Properties

The following parameters are accessible with `dladm(1M)`. Some of these are normally only read-only. Other properties that are not specific to IEEE 802.3 / Ethernet links are also available via `dladm(1M)`,

and are documented in its man page rather than here.

<b>speed</b>	Link speed, in Mbps per second (dladm only).
<b>duplex</b>	Link duplex, either "full" or "half".
<b>state</b>	Link state, either "up" or "down".
<b>mtu</b>	Maximum link frame size in bytes. See <i>Jumbo Frames</i> .
<b>flowctrl</b>	Flow control setting, one of "no", "tx", "rx", or "bi". See <i>Flow Control</i> .
<b>adv_10gfdx_cap</b>	Advertising 10 Gbps support.
<b>en_10gfdx_cap</b>	Enable 10 Gbps support.
<b>adv_1000fdx_cap</b>	Advertising 1000 Mbps full-duplex support.
<b>en_1000fdx_cap</b>	Enable 1000 Mbps full-duplex.
<b>adv_1000hdx_cap</b>	Advertising 1000 Mbps half-duplex support.
<b>en_1000hdx_cap</b>	Enable 1000 Mbps half-duplex.
<b>adv_100fdx_cap</b>	Advertising 100 Mbps full-duplex support.
<b>en_100fdx_cap</b>	Enable 100 Mbps full-duplex.
<b>adv_100hdx_cap</b>	Advertising 100 Mbps half-duplex support.
<b>en_100hdx_cap</b>	Enable 100 Mbps half-duplex.
<b>adv_10fdx_cap</b>	Advertising 10 Mbps full-duplex support.
<b>en_10fdx_cap</b>	Enable 10 Mbps full-duplex.
<b>adv_10hdx_cap</b>	Advertising 10 Mbps half-duplex support.
<b>en_10hdx_cap</b>	Enable 10 Mbps half-duplex.

### Auto-negotiation

With modern devices, auto-negotiation is normally handled automatically. With 10 Gbps and 1000 Gbps, it is mandatory. (10GBASE-T also requires full-duplex operation.) It is also *strongly* recommended for use whenever possible; without auto-negotiation the link will usually not operate unless both partners are configured to use the same link mode.

Auto-negotiation, when enabled, takes place by comparing the local capabilities that have been advertised (which must also be supported by the local device), with the capabilities that have been advertised by the link partner (peer). The first of the following modes that is supported by both partners is selected as the link negotiation result:

- o 10 Gbps (10gfdx)
- o 1000 Mbps full-duplex (1000fdx)
- o 1000 Mbps half-duplex (1000hdx)
- o 100 Mbps full-duplex (100fdx)
- o 100BASE-T4 (100T4)
- o 100 Mbps half-duplex (100hdx)
- o 10 Mbps full-duplex (10fdx)
- o 10 Mbps half-duplex (10hdx)

Advertisement of these modes can be enabled or disabled by setting the appropriate **en\_** property in `dladm(1M)`.

Auto-negotiation may also be disabled, by setting the **adv\_autoneg\_cap** property to 0. In this case, the highest enabled link mode (using the above list) is “forced” for the link.

### Flow Control

Link layer flow control is available on many modern devices, and is mandatory for operation at 10 Gbps. It requires that the link be auto-negotiated, and that the link be full-duplex, in order to function.

Flow control is applied when a receiver becomes congested. In this case the receiver can send a special frame, called a pause frame, to request its partner cease transmitting for a short period of time.

Flow control can be said to be either symmetric, in which case both partners can send and honor pause frames, or asymmetric, in which case one partner may not transmit pause frames.

The flow control mode used is driven by the **flowctrl** property. It has the following meanings:

- "no" Neither send, nor honor pause frames.
- "tx" Send pause frames, provided that the peer can support them, but do not honor them.

"rx" Receive and honor pause frames.

"bi" Both send and receive (and honor) pause frames.

The statistics for flow control (**adv\_cap\_pause**, **adv\_cap\_asmpause**, **lp\_cap\_pause**, **lp\_cap\_asmpause**, **link\_pause**, and **link\_asmpause**) are based on the properties exchanged in the auto-negotiation and are confusing as a result. Administrators are advised to use the **flowctrl** property instead.

### Jumbo Frames

The IEEE 802.3 standard specifies a standard frame size of 1518 bytes, which includes a 4-byte frame checksum, a 14-byte header, and 1500 bytes of payload. Most devices support larger frame sizes than this, and when all possible parties on the same local network can do so, it may be advantageous to choose a larger frame size; 9000 bytes is the most common option, as it allows a transport layer to convey 8 KB (8192) of data, while leaving room for various link, network, and transport layer headers.

Note that the use of frames carrying more than 1500 bytes of payload is not standardized, even though it is common practice.

The **mtu** property is used to configure the frame size. Note that this is the size of the payload, and excludes the preamble, checksum, and header. It also excludes the tag for devices that support tagging (see *Virtual LANs* below).

Care must be taken to ensure that all communication parties agree on the same size, or communication may cease to function properly.

Note that the **mtu** property refers to the link layer property. It may be necessary to configure upper layer protocols such as IP to use a different size when this changes. See `ifconfig(1M)`.

### Virtual LANs

Most devices support virtual LANs (and also priority control tagging) though the use of a 4-byte tag inserted between the frame header and payload. The details of configuration of this are covered in the `dladm(1M)` manual.

### Data Link Provider Interface (DLPI) Details

The correct method for applications to access Ethernet devices directly is to use the DLPI. See `dlpi(7P)` and `libdlpi(3LIB)` for further information.

The following DLPI parameters are presented to applications.

Maximum SDU	1500 (or larger, as determined by the <b>mtu</b> property.)
Minimum SDU	0



Address length	6
MAC type	DL_ETHER
SAP length	-2
Service mode	DL_CLDLS
Broadcast address	ff:ff:ff:ff:ff:ff (6 bytes with all bits set)

Note that if the application binds to SAP of 0, then standard IEEE 802.3 mode is assumed and the frame length is stored in place of the Ethernet type. Frames that arrive with the type field set to 1500 or less, are delivered to applications that bind to SAP 0.

Ethernet drivers on the support both DLPI style 1 and style 2 operation. Additionally, it is possible to configure provide “vanity” names to interfaces using the `dladm(1M)` **rename-link** subcommand. Such vanity names are only accessible using DLPI style 1.

## NOTES

There may be other mechanisms available to configure link layer properties. Historically the `ndd(1M)` command, and `driver.conf(4)` files could be used to do this. These methods are deprecated in favor of `dladm(1M)` properties.

## INTERFACE STABILITY

When present, the statistics and properties presented here are **Committed**. However, note that not every Ethernet device supports all of these, and some devices may support additional statistics and properties.

The DLPI and IEEE 802.3 itself are **Standard**.

## SEE ALSO

`dladm(1M)`, `ifconfig(1M)`, `kstat(1M)`, `netstat(1M)`, `ndd(1M)`, `libdlpi(3LIB)`, `driver.conf(4)`, `dlpi(7P)`

*IEEE 802.3: Ethernet*, IEEE Standards Association.

*Data Link Provider Interface (DLPI)*, The Open Group, 1997.

*STREAMs Programming Guide*, Sun Microsystems, Inc., January 2005.