# NAME

**ng\_l2tp** - L2TP protocol netgraph node type

# SYNOPSIS

#include <sys/types.h>
#include <netgraph/ng\_l2tp.h>

# DESCRIPTION

The **l2tp** node type implements the encapsulation layer of the L2TP protocol as described in RFC 2661. This includes adding the L2TP packet header for outgoing packets and verifying and removing it for incoming packets. The node maintains the L2TP sequence number state and handles control session packet acknowledgment and retransmission.

# HOOKS

The **l2tp** node type supports the following hooks:

*lower* L2TP frames.

*ctrl* Control packets.

session\_hhhh Session 0xhhhh data packets.

L2TP control and data packets are transmitted to, and received from, the L2TP peer via the lower hook. Typically this hook would be connected to the inet/dgram/udp hook of an ng\_ksocket(4) node for L2TP over UDP.

The ctrl hook connects to the local L2TP management entity. L2TP control messages (without any L2TP headers) are transmitted and received on this hook. Messages written to this hook are guaranteed to be delivered to the peer reliably, in order, and without duplicates.

Packets written to the ctrl hook must contain a two byte session ID prepended to the frame (in network order). This session ID is copied to the outgoing L2TP header. Similarly, packets read from the ctrl hook will have the received session ID prepended.

Once an L2TP session has been created, the corresponding session hook may be used to transmit and receive the session's data frames: for the session with session ID 0xabcd, the hook is named session\_abcd.

# CONTROL MESSAGES

This node type supports the generic control messages, plus the following:

# NGM\_L2TP\_SET\_CONFIG (setconfig)

This command updates the configuration of the node. It takes a *struct ng\_l2tp\_config* as an argument:

```
/* Configuration for a node */
```

struct ng\_l2tp\_config {

u_char	enabled;	/* enables traffic flow */
u_char	match_id;	/* tunnel id must match 'tunnel_id' */
uint16_t	tunnel_id;	/* local tunnel id */
uint16_t	peer_id;	/* peer's tunnel id */
uint16_t	peer_win;	/* peer's max recv window size */
uint16_t	rexmit_max	; /* max retransmits before failure */
uint16_t	rexmit_max	_to; /* max delay between retransmits */
};		

The *enabled* field enables packet processing. Each time this field is changed back to zero the sequence number state is reset. In this way, reuse of a node is possible.

The *tunnel\_id* field configures the local tunnel ID for the control connection. The *match\_id* field determines how incoming L2TP packets with a tunnel ID field different from *tunnel\_id* are handled. If *match\_id* is non-zero, they will be dropped; otherwise, they will be dropped only if the tunnel ID is non-zero. Typically *tunnel\_id* is set to the local tunnel ID as soon as it is known and *match\_id* is set to non-zero after receipt of the SCCRP or SCCCN control message.

The peer's tunnel ID should be set in *peer\_id* as soon as it is learned, typically after receipt of a SCCRQ or SCCRP control message. This value is copied into the L2TP header for outgoing packets.

The *peer\_win* field should be set from the "Receive Window Size" AVP received from the peer. The default value for this field is one; zero is an invalid value. As long as *enabled* is non-zero, this value may not be decreased.

The *rexmit\_max* and *rexmit\_max\_to* fields configure packet retransmission. *rexmit\_max\_to* is the maximum retransmission delay between packets, in seconds. The retransmit delay will start at a small value and increase exponentially up to this limit. The *rexmit\_max* sets the maximum number of times a packet will be retransmitted without being acknowledged before a failure condition is declared. Once a failure condition is declared, each additional retransmission will cause the **l2tp** node to send a NGM\_L2TP\_ACK\_FAILURE (**ackfailure**) control message back to the node that sent the last NGM\_L2TP\_SET\_CONFIG. Appropriate action should then be taken to shutdown the control connection.

# NGM\_L2TP\_GET\_CONFIG (getconfig)

Returns the current configuration as a *struct ng\_l2tp\_config*.

## NGM\_L2TP\_SET\_SESS\_CONFIG (setsessconfig)

This control message configures a single data session. The corresponding hook must already be connected before sending this command. The argument is a *struct ng\_l2tp\_sess\_config*:

```
/* Configuration for a session hook */
struct ng_l2tp_sess_config {
    uint16_t session_id; /* local session id */
    uint16_t peer_id; /* peer's session id */
    u_char control_dseq; /* whether we control data sequencing */
    u_char enable_dseq; /* whether to enable data sequencing */
    u_char include_length; /* whether to include length field */
};
```

The session\_id and peer\_id fields configure the local and remote session IDs, respectively.

The *control\_dseq* and *enable\_dseq* fields determine whether sequence numbers are used with L2TP data packets. If *enable\_dseq* is zero, then no sequence numbers are sent and incoming sequence numbers are ignored. Otherwise, sequence numbers are included on outgoing packets and checked on incoming packets.

If *control\_dseq* is non-zero, then the setting of *enable\_dseq* will never change except by another NGM\_L2TP\_SET\_SESS\_CONFIG control message. If *control\_dseq* is zero, then the peer controls whether sequence numbers are used: if an incoming L2TP data packet contains sequence numbers, *enable\_dseq* is set to one, and conversely if an incoming L2TP data packet does not contain sequence numbers, *enable\_dseq* is set to zero. The current value of *enable\_dseq* is always accessible via the NGM\_L2TP\_GET\_SESS\_CONFIG control message (see below). Typically an LNS would set *control\_dseq* to one while a LAC would set *control\_dseq* to zero (if the Sequencing Required AVP were not sent), thus giving control of data packet sequencing to the LNS.

The *include\_length* field determines whether the L2TP header length field is included in outgoing L2TP data packets. For incoming packets, the L2TP length field is always checked when present.

#### NGM\_L2TP\_GET\_SESS\_CONFIG (getsessconfig)

This command takes a two byte session ID as an argument and returns the current configuration for the corresponding data session as a *struct ng\_l2tp\_sess\_config*. The corresponding session hook must be connected.

# NGM\_L2TP\_GET\_STATS (getstats)

This command returns a *struct ng\_l2tp\_stats* containing statistics of the L2TP tunnel.

#### NGM\_L2TP\_CLR\_STATS (clrstats)

This command clears the statistics for the L2TP tunnel.

# NGM\_L2TP\_GETCLR\_STATS (getclrstats)

Same as NGM\_L2TP\_GET\_STATS, but also atomically clears the statistics as well.

# NGM\_L2TP\_GET\_SESSION\_STATS (getsessstats)

This command takes a two byte session ID as an argument and returns a *struct*  $ng_l2tp_session_stats$  containing statistics for the corresponding data session. The corresponding session hook must be connected.

# NGM\_L2TP\_CLR\_SESSION\_STATS (clrsessstats)

This command takes a two byte session ID as an argument and clears the statistics for that data session. The corresponding session hook must be connected.

# NGM\_L2TP\_GETCLR\_SESSION\_STATS (getclrsessstats)

Same as NGM\_L2TP\_GET\_SESSION\_STATS, but also atomically clears the statistics as well.

#### NGM\_L2TP\_SET\_SEQ (setsequence)

This command sets the sequence numbers of a not yet enabled node. It takes a *struct*  $ng_l2tp\_seq\_config$  as argument, where *xack* and *nr* respectively *ns* and *rack* must be the same. This option is particularly useful if one receives and processes the first packet entirely in userspace and wants to hand over further processing to the node.

#### **SHUTDOWN**

This node shuts down upon receipt of a NGM\_SHUTDOWN control message, or when all hooks have been disconnected.

#### SEE ALSO

netgraph(4), ng\_ksocket(4), ng\_ppp(4), ng\_pptpgre(4), ngctl(8)

W. Townsley, A. Valencia, A. Rubens, G. Pall, G. Zorn, and B. Palter, *Layer Two Tunneling Protocol L2TP*, RFC 2661.

#### HISTORY

The l2tp node type was developed at Packet Design, LLC, http://www.packetdesign.com/.

# AUTHORS

Archie Cobbs <archie@packetdesign.com>