

# LINX Protocols

## Abstract

*This document describes the Enea **LINX protocol version 1.0**. See the Revision History section (and document header line) for the version of this document.*

*LINX is a distributed communication protocol stack for transparent inter node and inter process communication for a heterogeneous mix of systems.*

*Copyright © 2006 Enea Software AB.*

*LINX, OSE, OSEck, OSE Epsilon, Optima, NASP, Element, Polyhedra are trademarks of Enea Software AB. Linux is a registered trademark of Linus Torvalds. All other copyrights, trade names and trademarks used herein are the property of their respective owners and are used for identification purposes only.*

*Disclaimer. The information in this document is subject to change without notice and should not be construed as a commitment by Enea Software AB.*

## Table of Contents

- [1. Introduction](#)
  - [1.1 Purpose](#)
  - [1.2 Revision History](#)
  - [1.3 References](#)
  - [1.4 Definitions and Acronyms](#)
- [2. Overview](#)
- [3. The RLNH Protocol](#)
  - [3.1 RLNH Link Creation and Initialization](#)
  - [3.2 Publication of Names](#)
  - [3.3 Remote Name Lookup](#)
  - [3.4 Link Supervision](#)
  - [3.5 Protocol Messages](#)
    - [3.5.1 RLNH\\_INIT](#)
    - [3.5.2 RLNH\\_INIT\\_REPLY](#)
    - [3.5.3 RLNH\\_PUBLISH](#)
    - [3.5.4 RLNH\\_QUERY\\_NAME](#)
    - [3.5.5 RLNH\\_UNPUBLISH](#)
    - [3.5.6 RLNH\\_UNPUBLISH\\_ACK](#)
- [4. Enea LINX Connection Manager Protocols](#)
  - [4.1 Connection Establishment](#)
  - [4.2 Reliable Message Passing](#)
  - [4.3 Connection Supervision](#)
- [5. Enea LINX Ethernet Connection Manager](#)
  - [5.1 Protocol Descriptions](#)
    - [5.1.1 Enea LINX Ethernet Connection Manager Headers](#)
      - [5.1.1.1 ETHCM\\_MAIN Header](#)
    - [5.1.2 Enea LINX Connect Protocol](#)
      - [5.1.2.1 Connect Protocol](#)
      - [5.1.2.2 Connect Protocol Description](#)
      - [5.1.2.3 ETHCM\\_CONN Header](#)
    - [5.1.3 Enea LINX User Data Protocol](#)
      - [5.1.3.1 User data and Fragmentation Protocol](#)
      - [5.1.3.2 ETHCM\\_UDATA Header](#)
      - [5.1.3.3 ETHCM\\_FRAG Header](#)
    - [5.1.4 Enea LINX Reliability Protocol](#)
      - [5.1.4.1 Reliability Protocol Description](#)
      - [5.1.4.2 ETHCM\\_ACK header](#)
      - [5.1.4.3 ETHCM\\_NACK header](#)
    - [5.1.5 Enea LINX Connection Supervision Protocol](#)
      - [5.1.5.1 Connection Supervision protocol description](#)
  - [5.2 LINX Discovery Daemon](#)
    - [5.2.1 Linxdisc protocol](#)

# 1. Introduction

## 1.1 Purpose

This document describes the Enea LINX protocol.

## 1.2 Revision History

Current version is also seen in document header when printed (HTML title line).

<b>Revision</b>	<b>Author</b>	<b>Date</b>	<b>Status and Description of purpose for new revision</b>
9	lejo	2006-10-25	Added copyright on front page. Cleaned out prerelease entries of document history.
8	lejo	2006-10-13	Converted document from Word to XHTML.
7	jonj	2006-09-13	Added reserved field in UDATA Header, Fixed bug in MAIN header.
6	jonj	2006-09-11	Fixed a few bugs, improved RLNH protocol description.
5	jonj	2006-08-05	<b>Approved.</b> Updated after review (internal reference ida 010209).

## 1.3 References

## 1.4 Definitions and Acronyms

### Definitions

- A**  
The active (initiating) side in a protocol exchange.
- B**  
The passive (responding) side in a protocol exchange.
- CM**  
Connection Manager, the entity implementing the transport layer of the Enea LINX protocol.
- Endpoint**  
A (part of) an application that uses the Enea LINX messaging services.
- Connection**  
An association between Connection Managers.
- Connection ID**  
A key used by Ethernet Connection Managers to quickly lookup the destination of an incoming packet. Connection ID based lookup is much efficient than MAC-address based lookup.
- Connection Manager**  
A Connection Manager provides reliable communication for message passing. The connection layer roughly corresponds to layer four, the transport layer, in the OSI model.
- Connection Supervision**  
A function within the Connection Manager responsible for detection of connection failure.
- Header**  
Protocol Data filled in by the Connection Manager. A PDU consists of several headers.
- Link**  
An association between link handlers using the Enea LINX protocol.
- Link Handler**  
A concept from the OSE world. A Link Handler makes the OSE messaging IPC mechanism available in a distributed context.
- MAC-address**  
Link layer address on Ethernet.
- Message**  
A unit of information transported over LINX. Messages are transformed by Connection Managers in order to fit the media.
- Packet**  
Protocol Data Unit sent over a connection.
- RLNH**  
Rapid Link Handler, the link handler of Enea LINX.

### Abbreviation

- IPC**  
Inter Process Communication.
- PDU**  
Protocol Data Unit.

## 2. Overview

Enea LINX is an open technology for distributed system IPC which is platform and interconnect independent, scales well to large systems with any topology, but that still has the performance needed for high traffic bearing components of the system. It is based on a transparent message passing method.

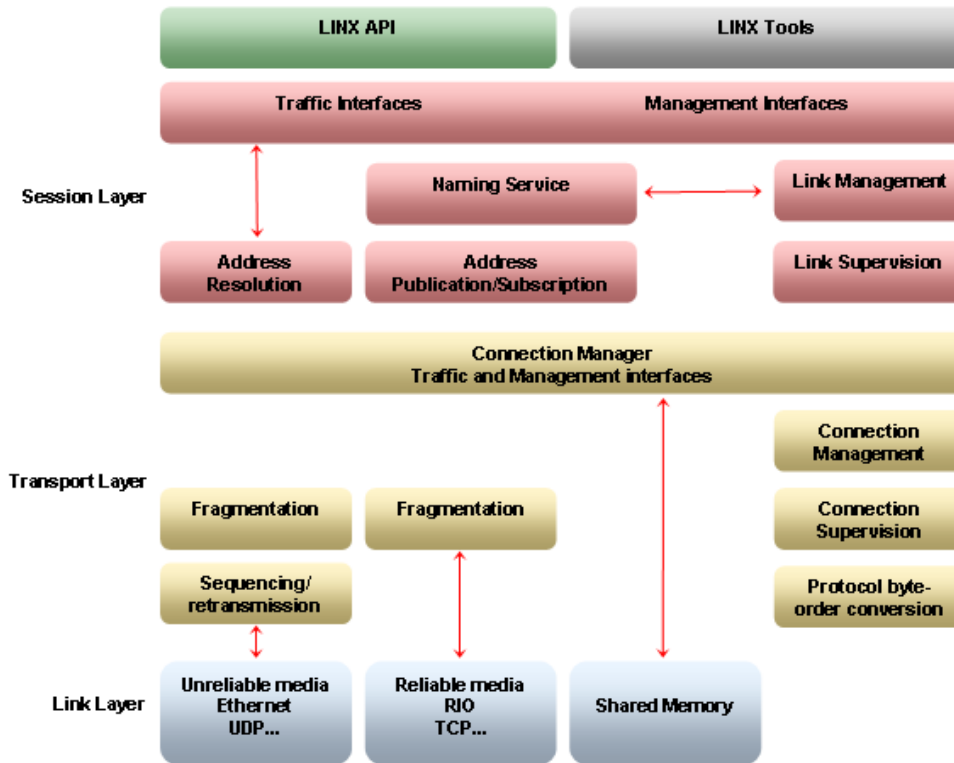


Figure 1 LINX Architecture

Enea LINX provides a solution for inter process communication for the growing class of heterogeneous systems using a mixture of operating systems, CPUs, microcontrollers DSPs and media interconnects such as shared memory, RapidIO, Gigabit Ethernet or network stacks. Architectures like this poses obvious problems, endpoints on one CPU typically uses the IPC mechanism native to that particular platform and they are seldom usable on platform running other OSes. For distributed IPC other methods, such as TCP/IP, must be used but that comes with rather high overhead and TCP/IP stacks may not be available on small systems like DSPs. Enea LINX solves the problem since it can be used as the sole IPC mechanism for local and remote communication in the entire heterogeneous distributed system.

The Enea LINX protocol stack has two layers - the RLNH and the Connection Manager, or CM, layers. RLNH corresponds to the session layer in the OSI model and implements IPC functions including methods to look up endpoints by name and to supervise to get asynchronous notifications if they die. The Connection Manager layer corresponds to the transport layer in the OSI model and implements reliable in order transmission of arbitrarily sized messages over any media.

RLNH is responsible for resolving remote endpoint names, and for setting up and removing local representations of remote endpoints. The RLNH layer provides translation of endpoint OS IDs from the source to the destination system. It also handles the interaction with the local OS and applications that use the Enea LINX messaging services. RLNH consists of a common, OS independent protocol module and an OS adaptation layer that handles OS specific interactions.

The Connection Manager provides a reliable transport mechanism for RLNH. When the media is unreliable, such as Ethernet, or have other quirks the Connection Manager must implement means like flow control, retransmission of dropped packets, peer supervision, and becomes much more complex. On reliable media, such as shared memory or RapidIO, Connection Managers can be quite simple.

The rest of this document contains a detailed description of the protocols used by the RLNH and the CM layers of Enea LINX. Chapter 3 describes the RLNH protocol. Chapter 4 describes features shared by all Connection Managers, i.e. what RLNH requires from a Connection Manager. The description is intentionally kept on a high level since the functionality actually implemented in any instance of a CM depends very much on the properties of the underlying media. It is the intention that when Connection Managers for new media are implemented the protocols will be described in the following chapters. Finally Chapter 5 describes version two of the Ethernet Connection Manager protocol.

The table below shows how protocol headers are described. PDUs are sent in network byte order (e.g. big endian). Bits and bytes are numbered in the order they are transmitted on the media. For example:

0										1											2											3	
0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1		
0								1								2								3									
One								Two								Three								Reserved									
Four																																	

Table 1 Protocol Message legend

The first two rows show bit numbers, the third row byte numbers for those more comfortable with that view, and the last two rows example protocol fields. In the header above, the fields in the header are sent in order: one, two, three, reserved, four.

## 3. The RLNH Protocol

The RLNH protocol is designed to be light-weight and efficient. There are six RLNH-to-RLNH control PDUs described in detail below. The required overhead associated with user signal transmission is not carried in any dedicated RLNH PDU. Instead it is passed as arguments along with the signal data to the Connection Manager layer, where an optimized transmission scheme and message layout can be implemented based on knowledge of the underlying media and/or protocols. In particular the source and destination link addresses are sent this way, the addresses identifies the sending and receiving endpoints respectively. RLNH protocol messages are sent with source and destination link addresses set to zero (0).

### 3.1 RLNH Link Creation and Initialization

The Connection Manager is initialized by RLNH. It is responsible for providing reliable, in-order delivery of messages and for notifying RLNH when the connection to the peer becomes available / unavailable. An RLNH link is created maps a unique name to a Connection Manager object.

As soon as the Connection Manager has indicated that the connection is up, RLNH transmits an RLNH\_INIT message to the peer carrying its protocol version number. Upon receiving this message, the peer responds with an RLNH\_INIT\_REPLY message indicating whether it supports the given protocol version or not. When this simple message exchange has been successfully completed, RLNH is ready to provide messaging services for the link name.

The RLNH protocol is stateless after this point.

### 3.2 Publication of Names

RLNH assigns each endpoint a link address identifier. The association between the name of an endpoint and the link address that shall be used to refer to it is published to the peer in an RLNH\_PUBLISH message. RLNH\_PUBLISH is always the first RLNH message transmitted when a new endpoint starts using the link.

Upon receiving an RLNH\_PUBLISH message, RLNH creates a local representation of the remote name. Local applications can communicate transparently with such a representation, but the signals are in reality forwarded by RLNH to the peer system where the real destination endpoint resides (and the destination will receive the signals from a representation of the true sender).

The link addresses of source and destination for each signal are given to the Connection Manager for transmission along with the signal data and delivered to the peer RLNH. The link addressing scheme is designed to enable O(1) translation between link addresses and their corresponding local OS IDs.

### 3.3 Remote Name Lookup

When a hunt call is performed for the name of a remote endpoint, the request is passed on by RLNH to the peer as a RLNH\_QUERY\_NAME message. If the hunter has not used the link before, an RLNH\_PUBLISH message is sent first.

Upon receiving a RLNH\_QUERY\_NAME message, RLNH resolves the requested endpoint name locally and returns an RLNH\_PUBLISH message when it has been found and assigned a link address. As described above, this triggers the creation of a representation at the remote node. This in turn resolves the hunt call - the hunter is provided with the OS ID of the representation.

### 3.4 Link Supervision

RLNH supervises all endpoints that use the link. When a published name is terminated, a RLNH\_UNPUBLISH message with its link address is sent to the peer.

When a RLNH\_UNPUBLISH message is received, RLNH removes its local representation of the endpoint referred to by the given link address. When RLNH no longer associates the link address with any resources, it returns a RLNH\_UNPUBLISH\_ACK message to the peer.

The link address can be reused after a RLNH\_UNPUBLISH\_ACK message.



### 3.5 Protocol Messages

#### 3.5.1 RLNH\_INIT

The RLNH\_INIT message initiates link establishment between two peers. It is sent when the Connection Manager indicates to RLNH that the connection is up.

0										1										2									3		
0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1
0								1								2								3							
Sig_no																															
Version																															

Table 2 RLNH Init Header

Field	Definition
Sig_no	Message type. The value of RLNH_INIT is 5.
Version	The RLNH protocol version. The current version is 1.

Table 3 RLNH Init Header Description

#### 3.5.2 RLNH\_INIT\_REPLY

During link set up, RLNH responds to the RLNH\_INIT message by sending an RLNH\_INIT\_REPLY message. The status field indicates whether the protocol is supported or not.

0										1											2									3	
0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1
0								1								2								3							
Sig_no																															
Status																															

Table 4 RLNH Init Reply Header

Field	Definition
Sig_no	Message type. The value of RLNH_INIT_REPLY is 6.
Status	Status code indicating whether the protocol version received in the RLNH_INIT message is supported (0) or not (1).

Table 5 RLNH Init Reply Header Description



### 3.5.5 RLNH\_UNPUBLISH

The RLNH\_UNPUBLISH message tells the remote RLNH that the endpoint previously assigned the given link-address has been closed.

0										1										2										3	
0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1
0								1								2								3							
Sig_no																															
Linkaddr																															

Table 10 RLNH Unpublish Header

Field	Definition
Sig_no	Message type. The value of RLNH_UNPUBLISH is 3.
Linkaddr	The address of the closed endpoint.

Table 11 RLNH Unpublish Header Description

### 3.5.6 RLNH\_UNPUBLISH\_ACK

The RLNH\_UNPUBLISH\_ACK message tells the remote RLNH that all associations regarding an unpublished link address have been terminated. This indicates to the peer that it is ok to reuse the link address.

0										1										2										3	
0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1
0								1								2								3							
Sig_no																															
Linkaddr																															

Table 12 RLNH Unpublish Ack Header

Field	Definition
Sig_no	Message type. The value of RLNH_UNPUBLISH_ACK is 4.
Linkaddr	The link address of the unpublished endpoint.

Table 13 RLNH Unpublish Ack Header Description

## 4. Enea LINX Connection Manager Protocols

This chapter describes in general terms the functionality a Connection Manager must provide.

A Connection Manager shall hide details of the underlying media, e.g. addressing, general media properties, how to go about to establish connections, media aggregation, media redundancy and so on. A Connection Manager shall support creation of associations, known as connections that are suitable for reliable message passing of arbitrarily sized messages. A Connection Manager must not rely on implicit connection supervision since this can cause long delay between peer failure and detection if the link is idle. A message accepted by a Connection Manager must be delivered to its destination.

If, for any reason, a Connection Manager is unable to deliver a message RLNH must be notified and the connection restarted since its state at this point is inconsistent.

Since Enea LINX runs on systems with very differing size, the Connection Manager protocol must be scalable. A particular implementation may ignore this requirement either because the underlying media doesn't support scalability or that a non-scalable implementation has been chosen. However, interfaces to the Connection Manager and protocol design when the media supports big configuration shall not make design decisions that prevents the system to grow to big configurations should the need arise.

Enea LINX must be able to coexist with other protocols when sharing media.

### 4.1 Connection Establishment

Creation of connections always requires some sort of hand shake protocol to ensure that the Connection Manager at the endpoints agrees on communication parameters and the properties of the media.

### 4.2 Reliable Message Passing

A channel suitable for reliable message passing must appear to have the following properties:

- Messages are delivered in the order they are sent.
- Messages can be of arbitrary size.
- Messages are never lost.
- The channel has infinite bandwidth.

Although no medium has all these properties some come rather close and others simulate them by implementing functions like fragmentation of messages larger than the media can handle, retransmission of lost messages, and flow control i.e. don't send more than the other side can consume. If a Connection Manager is unable to continue deliver messages according to these rules the connection must be reset and a notification sent to RLNH.

### 4.3 Connection Supervision

A distributed IPC mechanism should support means for applications to be informed when a remote endpoint fails. Connection Managers are responsible for detecting when the media fails to deliver messages and when the host where a remote endpoint runs has crashed. There are of course other reasons an endpoint might fail to respond but in those situations the communication link continues to function and higher layers in the LINX architecture manages supervision.

## 5. Enea LINX Ethernet Connection Manager

This chapter describes version 2 of the Enea LINX Ethernet Connection Manager Protocol.

### 5.1 Protocol Descriptions

Enea LINX PDUs are stacked in front of, possible, user data to form an Enea LINX Ethernet packet. All PDUs contain a field next header which contain the protocol number of following headers or the value -1 (1111b) if this is the last PDU. All headers except Enea LINX main header are optional. Everything from the transmit() down call, including possible control plane signaling, from the RLNH layer is sent reliably as user data. The first field in all headers is the next header field, having this field in the same place simplifies implementation and speeds up processing.

If a malformed packet is received the Ethernet Connection Manager resets the connection and informs RLNH.

When the Ethernet Connection Manager encounters problems which prevents delivery of a message or part of a message it must reset the connection. Notification of RLNH is implicit in the Ethernet CM, when the peer replies with RESET or, if the peer has crashed, the Connection Supervision timer fires.

#### 5.1.1 Enea LINX Ethernet Connection Manager Headers

Version 2 of the Enea LINX Ethernet Connection Manager protocol defines these headers.

Protocol number	Definition
ETHCM_MAIN	Main header sent first in all Enea LINX messages.
ETHCM_CONN	Connect header. Used to establish and tear down connections.
ETHCM_UDATA	User data header. All messages generated outside the Connection Manager are sent as UDATA.
ETHCM_FRAG	Fragment header. Messages bigger than the MTU are sent fragmented and PDUs following the first carries the ETHCM_FRAG header instead of ETHCM_UDATA.
ETHCM_ACK	Reliability header. Carries seqno and ackno. ACK doubles as empty acknowledge PDU and ACK-request PDU, in sliding window management and connection supervision.
ETHCM_NACK	Request retransmission of one or more packets.
ETHCM_NONE	Indicates that the current header is the last in the PDU.

Table 14 Ethernet Connection Manager Protocol Headers




### 5.1.2 Enea LINX Connect Protocol

The Enea LINX Connect Protocol is used to establish a connection on the Connection Manager level between two peers A and B. A Connection Manager will only try to establish a connection or accept connection attempt from a peer if it has been explicitly configured to do. After configuration a CM will maintain connection with the peer until explicitly told to destroy the connection or an un-recoverable error occurs.

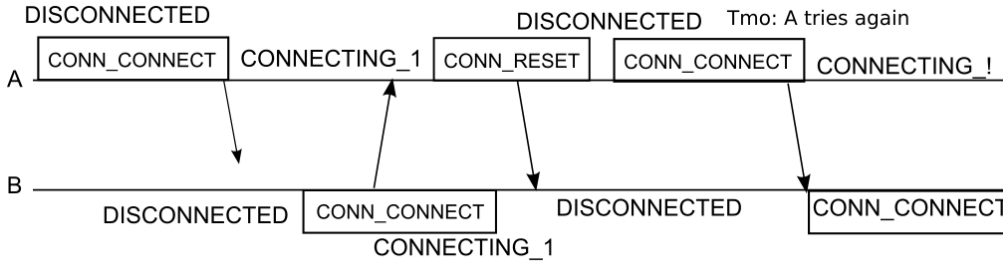
If a Connect-message is received, with a version number different from 2, the Ethernet CM refuses to connect.

The Connect Protocol determines Connection IDs to be used for this connection. Connection IDs are small keys used by receivers to quickly lookup a packets destination. Each side selects the Connection ID to be used by the peer when sending packets over the connection. The peer saves the ID and will use it for all future communication. If a node has a lot of connections it may run out of available Connection IDs. In this case the node sends Connection ID 0, which means no Connection ID and reverts to the slower MAC-addressing mode to determine destination for incoming packets.

A window size options may be sent in the CONNECT-header to indicate how the sender configuration deviates from the default. Deviations from default values are configured per connection in the create\_conn()-call.



Below, A starts first and tries to connect. B is not active, so the PDU is lost. A tries again..



Above, when B tries to connect to A, A is in the wrong state and sends RESET to synchronize  
 Figure 2 Successful Connect.

### 5.1.2.1 Connect Protocol

In the following protocol description the side initiating the connection is called A and the side responding to the request is called B. All protocol transitions are supervised by a timer, if the timer fires before the next step in the protocol have been completed the state machine reverts to state disconnected.

The Connect Protocol is symmetrical, there is no master and both sides try to initiate the connection. Collisions, i.e. when the initial CONNECT PDU is sent simultaneously from both sides, are handled by the protocol and the connection restarted after a randomized back of timeout.

### 5.1.2.2 Connect Protocol Description

Unless a Connection Manager has been configured to create a connection to a peer no messages are sent and the Connection Manager doesn't respond to connection attempts.

#### A-side

1. When a Connection Object is created at A by calling create\_conn() A starts from DISCONNECTED state, sends a CONNECT PDU to B and enters state CONNECTING\_1.
  - a. If tmo, go back to step 1 and try again after a grace period.
  - b. If B replies with a CONNECT\_ACK PDU, move to CONNECTED state, send ACK PDU to B, notify RLNH that a connection have been established, and start the Connection Supervision function.
  - c. If any other PDU is received from B send RESET and go back to step 1 and try again after a short grace period.

#### B-side

1. After configuration, B waits in state DISCONNECTED for a CONNECT PDU from A.
  - a. If a CONNECT PDU arrives, send a CONNECT\_ACK PDU and go to state CONNECTING\_2
  - b. If some other PDU arrives, send RESET to A and go back to step 2.
2. In state CONNECTING\_2, B waits for an ACK PDU from A.
  - a. When CONN\_ACK arrives, the connect protocol is complete, the Connection Manager notifies higher layers that a connection have been established, and start the Connection Supervision function.
  - b. If some other PDU is received or the timer fires, send RESET and go to state DISCONNECTED



Allowed messages and state changes are summarized in this state diagram. The notation [xxx/yyy] means: event xxx causes action yyy.

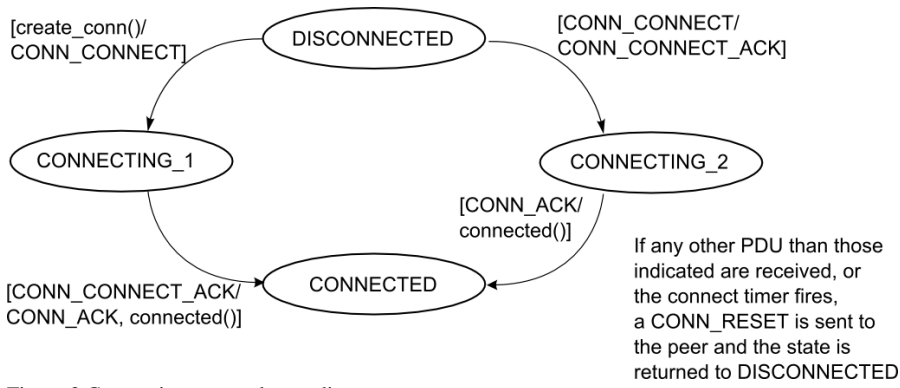


Figure 3 Connection protocol state diagram.

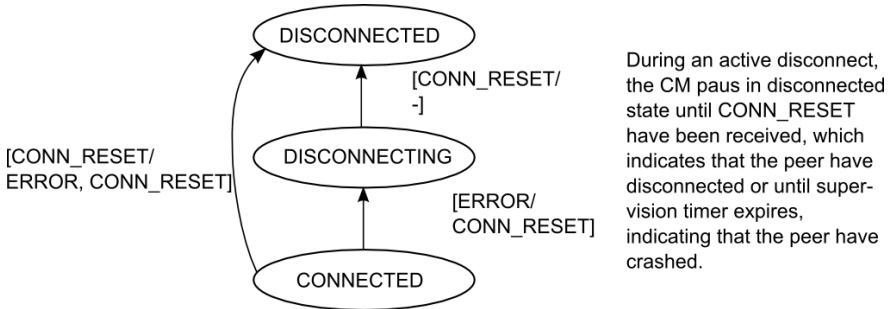


Figure 4 Disconnect state diagram.

### 5.1.2.3 ETHCM\_CONN Header

The Connection header varies in size depending on the size of the address on the media, on Ethernet it is 16 bytes.




Table 17 Connect Header

Field	Definition
Next	Next header, the protocol number of the following Enea LINX header or 1111b if last header.
Type	CONNECT. Start a connect transaction. CONNECT ACK. Reply from passive side. ACK. Confirm that the connection have been created. RESET. Sent if any error occurs. Also sent if the next step in connect protocol fails to complete within allowed time.
Size	Media address size
Window	Window size. Power of 2, thus Window == 5 means a window of $2^5 == 32$ packets.
Reserved	Reserved must be 0.
Conn ID	Use this Connection ID. Informs peer which connection ID to use when sending packets over this connection. Conn ID == 0, means don't use Connection ID.
Dst and src media addresses	Dst media address immediately followed by src media address,

Table 18 Connect Header Description

### 5.1.3 Enea LINX User Data Protocol

All messages originating outside the Connection Manger are sent as USER\_DATA. There are two types of USER\_DATA header. The first type is used for messages not requiring fragmentation and for the first fragment of fragmented messages. The second type is used for all remaining fragments.

#### 5.1.3.1 User data and Fragmentation Protocol

##### A-side

1. Accept a new message from RLNH. Calculate how many fragments are required to send this message.
2. Frag\_cnt = 0.
3. For each fragment in the message:
  - a. If first fragment send as USER\_DATA else send as FRAG.
  - b. If only fragment set fragno to -1 else set fragno to frag\_cnt and increment frag\_cnt.
  - c. If last fragment set MORE to 0 else set MORE to 1.
  - d. Forward the packet to lower layer for transmission.
  - e. If last fragment go back to step 1.

##### B-side

1. When a USER\_DATA or a FRAG packet arrives:
2. If fragno = -1 (0x7fff) deliver() to RLNH since this is a complete message and wait for next packet.
3. If fragno ≠ -1 find the reassembly queue for this message and add the packet to the tail of the queue (lower layers doesn't emit packets out-of-order).  
If the packet is the last fragment deliver the complete message to RLNH else wait for next packet.

#### 5.1.3.2 ETHCM\_UDATA Header

0										1													2													3			
0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1								
0										1										2										3									
Next										Reserved_1										M										Frag no									
Reserved_2																																							
Dst addr																	Src addr																						

Table 19 User Data Header

Field	Definition
Next	Next header, the protocol number of the following Enea LINX header or 1111b if last header.
Reserved_1	Reserved for future use, must be 0.
M	More fragment follows.
Frag no	Number of this fragment. Fragments are numbered 0 to (number of fragments - 1). Un-fragmented messages have fragment number -1 (0x7fff).
Reserved_2	Reserved for future use, must be 0.
Dst addr	Opaque address (to CM) identifying the receiver.
Src addr	Opaque address identifying the sender.

Table 20 User Data Header Description



## 5.1.4 Enea LINX Reliability Protocol

The Connection Manager uses a selective repeat sliding window protocol with modulo  $m$  sequence numbers. The ack header carries sequence and request numbers for all reliable messages sent over the connection. Sequence and Request numbers are 12 bits wide and  $m$  is thus 4096.

In the Selective Retransmit Sliding Window algorithm the B-side explicitly request retransmission of dropped packets and remembers packets received out of order. If the sliding window is full, the A-side queues outgoing packets in a deferred queue. When space becomes available in the window (as sent packets are acknowledged by B) packets are from the front of the deferred queue and sent as usual. A strict ordering is maintained, as long as there are packets waiting in the deferred queue new packets from RLNH are deferred, even if there is space available in the window.

### 5.1.4.1 Reliability Protocol Description

For sliding window operation the sender A have the variables  $SN_{min}$  and  $SN_{max}$ ,  $SN_{min}$  points to the first unacknowledged packet in the sliding window and  $SN_{max}$  points at the next packet to be sent. In addition the sliding window size is denoted  $n$  and size of sequence numbers are modulo  $m$ . The receiver side B maintains a variable  $RN$  which denotes the next expected sequence number. In the protocol description, the sequence number and the request number of a packet is called  $sn$  and  $rn$  respectively. The module number  $m$  must be  $\geq 2n$ . In version two of the Enea LINX Ethernet protocol  $m$  is 4096 and  $n$  is a power of  $2 \leq 128$ . Drawing sequence number from a much larger space than the size of the window allows the reliability protocol to detect random packets with bad sequence numbers.

At A, the algorithm works as follows:

1. Set modulo variables  $SN_{min}$  and  $SN_{max}$  to 0.
2. In A if a message is sent from higher layer or there are packets in the defer queue, and  $(SN_{max} - SN_{min}) \bmod m < n$ , accept a packet into the sliding window set  $sn$  to  $SN_{max}$  and increment  $SN_{max}$  to  $(SN_{max} + 1) \bmod m$ . If  $(SN_{max} - SN_{min}) \bmod m \geq n$  defer sending the packet, i.e. queue the packet until there is room in the sliding window.
3. If an error free frame is received from B containing  $rn$ , and  $(RN - SN_{min}) \bmod m \leq (RN - SN_{max}) \bmod m$ , set  $SN_{min}$  to  $RN$  and remove packets with  $sn < SN_{min} \bmod m$  from send queue.
4. If a NACK frame is received, retransmit NACKed packets in-order with  $rn$  set to  $RN$ .
5. At arbitrary times but within bounded delay after receiving a reliable packet from B and if there are unacknowledged packets in the sliding window, send the first un-acked packet with  $rn$   $RN$  and the request bit set.

The selective repeat algorithm at B:

1. Set the modulo  $m$  variable  $RN$  to 0.
2. When an error free frame is received from A containing  $sn$  equal to  $RN$ , release the packet as well as following queued packets with consecutive  $sn$  to higher layer and increment  $RN$  to  $(last\ released\ sn + 1) \bmod m$ .
3. When an error free frame is received from A containing  $sn$  in the interval  $RN < sn < RN + n$ , put the packet in sequence number order in the receive queue and send NACK requesting retransmission of sequentially dropped packets to A. The seqno field in the NACK frame shall contain  $RN$  and the count field shall contain the number of missing packets.
4. At arbitrary times but within a bounded delay after receiving an error free frame from A transmit a frame containing  $RN$  to A. If there are frames in the receive queue send a NACK indicating missing frames.

Note that only user data is sent reliable, i.e. consume sequence numbers. ACKR, NACK and empty ACK are unreliable messages sequence numbers are not incremented as they are sent.

### 5.1.4.2 ETHCM\_ACK header

0										1										2												3	
0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1		
0								1								2								3									
Next				R	Res				Ackno								Seqno																

Table 23 Ack Header

Field	Definition
Next	Next header, the protocol number of the following Enea LINX header or 1111b if last header.
R	ACK-request, peer shall respond with an ACK of its own as soon as possible. (Used during connection supervision.)
Res	Reserved for future use, must be 0.
Ackno	Shall be set to last consecutive received seqno from peer.
Seqno	Incremented for every sent packet. Note that seqnos are set based on sent packets rather than sent bytes as in TCP.

Table 24 Ack Header Description

### 5.1.4.3 ETHCM\_NACK header

NACK is sent when a hole in the stream of received reliable packets is detected. If a sequence of packets is missing all are NACKed by the same NACK packet. A timer ensures that NACKs are sent as long as there are out-of-order packets in the receive queue.

0										1										2											3	
0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	
0								1								2								3								
Next				Reserved				Count								Res				Seqno												

Table 25 Nack Header

Field	Definition
Next	Next header, the protocol number of the following Enea LINX header or 1111b if last header.
Reserved	Reserved, must be 0.
Count	Number of NACKed seqnos.
Res	Reserved, must be 0.
Seqno	First NACKed seqno, the rest are assumed to follow consecutively seqno + 1, seqno + 2, ... , seqno + count - 1.

Table 26 nack Header Description

### 5.1.5 Enea LINX Connection Supervision Protocol

The purpose of the Connection Supervision function is to detect crashes on the B-side. If B has been silent for some time A sends a few rapid pings and, if B doesn't respond, A decides that B has crashed, notifies higher layers, and initiates teardown of the connection. There is no separate PDU for the Connection Supervision function, the ACKR header from the reliability protocol doubles as ping PDU.

#### 5.1.5.1 Connection Supervision protocol description

##### A-side

When a connection has been established, the Ethernet Connection Manager initializes the Connection Supervision function to state PASSIVE and starts a timer which will fire at regular interval as long as the link is up.

1. When the timer fires, the Ethernet Connection Manager checks if packets have been received from B since the last time the timer was active.
  - a. If yes, clear packets counter, set Connection Supervision state to PASSIVE and go to step 2.
  - b. If no, check if connection supervision is already in active state.
    - i. If no, enter state Active and send an ACKR to B.
    - ii. If yes, check if the ping limit has been reached.
      1. If yes, the connection is down, notify higher layers and go to state Disconnected.
      2. If no, resend ACKR and increment the ping counter.
2. Restart the timer.
3. If the connection is closed, stop the send timer.

##### B-side

1. When an ACKR is received, reply to the sender with an empty ACK [RN,SN<sub>max</sub>-1].

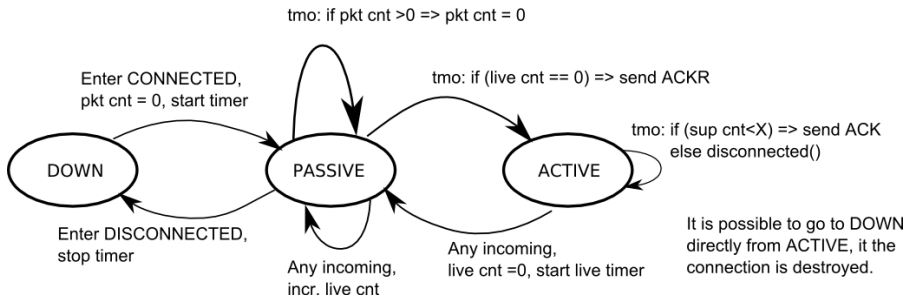


Figure 5 State diagram for Connection Supervision

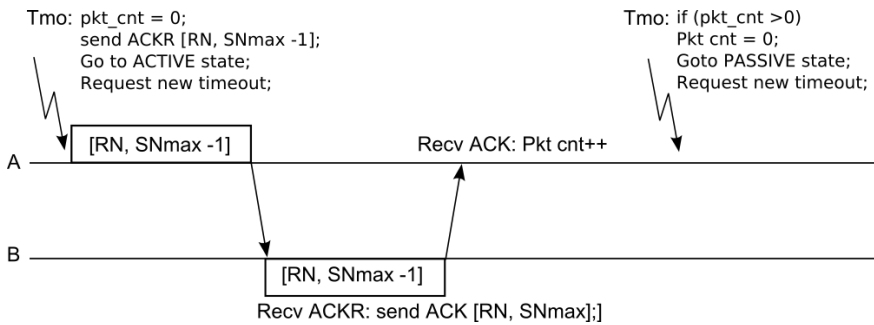


Figure 6 Connection supervision: the peer replies.

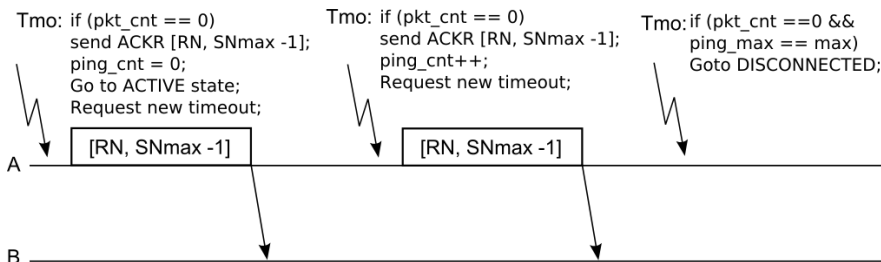


Figure 7 Connection supervision: the peer is down.

## 5.2 LINX Discovery Daemon

The LINX Discovery Daemon, **linxdisc**, automatically discovers LINX network topology and creates links to other LINX hosts. Linxdisc finds LINX by periodically broadcasting and listen for LINX advertisement messages on available Ethernet interfaces. When advertisements from other linxdiscs arrive linxdisc examines the messages and if it passes a number of controls a LINX connection is created to the sending host.

### 5.2.1 Linxdisc protocol

#### A-side

At startup read configuration file, build a list of all active interfaces, extract hostname and cluster name.

1. Periodically broadcast discovery messages on interfaces allowed by the configuration.

#### B-side

1. When a discovery message arrives, linxdisc performs the following steps to determine whether to connect to the sender:
  - a. Not one of my own messages, i.e. the sender address is the same as one of my interfaces.
  - b. Cluster ID is same as configuration.
  - c. Unique connection name, the name must differ from my name and the name must differ from names to which there already exist connections.
  - d. Check that the configuration doesn't forbid connections to this host.
2. If all tests pass create a connection to the peer using the advertised name and store information about the connection for later.
3. When terminated, linxdisc closes all connections it has created before exiting.
4. If the configuration is changed, closes connection not allowed by the new configuration, and updates advertisements messages to reflect the new configuration.

0										1									2											3																			
0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1																		
										0										1										2										3									
Type																																																	
Linklen																																																	
Netlen																																																	
Addr																																																	
Strings																																																	

Table 27 Linxdisc Discovery Message

Field	Description
Type	Linxdisc message type, currently only advertisements (2) is used.
Linklen	Length of linxname, c.f. below.
Netlen	Length of netname, c.f. below.
Addr	MAC-address to use when connecting to the sending host.
Strings	Carries two null-terminated strings, Linkname and Netname (cluster id). Link name is a unique identifier for the sending host ment to be used as the name of the link created by the receiving linxdisc. Netname is a unique cluster id string identifying a group of hosts that will form a <b>LINX cluster</b> .

Table 28 Linxdisc Discovery Message Definition