

Understanding Logical Link Control

Document ID: 12247

Introduction

Prerequisites

Requirements

Components Used

Conventions

Background Information

Implementations of LLC

Basic Information You Must Know in Order to Troubleshoot

LLC Frame Formats

DSAP Field

SSAP Field

Control Field

LLC Control Field Summary

LLC2 Modes and Session Establishment

Asynchronous Balanced Mode Extended (ABME)

Asynchronous Disconnect Mode (ADM)

LLC2 Asynchronous Balanced Mode Operation

LLC2 Tunable Parameters

Examples of LLC2 Parameter Configurations

LLC2 Error Conditions

Related Information

Introduction

IEEE standard 802.2 defines Logical Link Control (LLC) as a data link control layer used on 802.3, 802.5, and other networks. IBM originally designed LLC as a sublayer in the IBM Token Ring architecture.

Prerequisites

Requirements

Cisco recommends that you have knowledge of these topics:

- A basic understanding of LLC

Components Used

This document is not restricted to specific software and hardware versions.

The information in this document was created from the devices in a specific lab environment. All of the devices used in this document started with a cleared (default) configuration. If your network is live, make sure that you understand the potential impact of any command.

Conventions

Refer to Cisco Technical Tips Conventions for more information on document conventions.

Background Information

The LLC layer provides **connectionless** and **connection-oriented** data transfer.

Connectionless data transfer is commonly referred to as LLC type 1, or LLC1. Connectionless service does not require you to establish data links or link stations. After a Service Access Point (SAP) has been enabled, the SAP can send and receive information to and from a remote SAP that also uses connectionless service. Connectionless service does not have any mode setting commands (such as SABME) and does not require that state information is maintained.

Connection-oriented data transfer is referred to as LLC type 2, or LLC2. Connection-oriented service requires the establishment of link stations. When the link station is established, a mode setting command is necessary. Thereafter, each link station is responsible to maintain link state information.

Implementations of LLC

LLC2 is implemented whenever Systems Network Architecture (SNA) runs over a LAN or virtual LAN. LLC2 is also directly encapsulated into Frame Relay. Sometimes the router simply passes LLC2 frames and sometimes the router implements an LLC2 linkstation. NetBIOS also uses LLC. NetBIOS uses LLC1 to locate a resource. LLC2 connection-oriented sessions are then established.

The router implements an LLC2 stack when any of these features are enabled:

- Data-Link Switching (DLSw) (connection to LAN)
- Remote Source-Route Bridging (RSRB) with local ACK
- Channel Interface Processor (CIP)
- Advanced Peer-to-Peer Networking (SNASwitching (SNASw))
- Synchronous Data Link Control (SDLC) to LCC Conversion (SDLLC)

Basic Information You Must Know in Order to Troubleshoot

A basic knowledge of LLC is enough to isolate and resolve most problems. Because there are no link states or sessions to maintain, problems are rare in LLC1.

In LLC2, two categories of problems can occur:

1. Sessions that do not establish
2. Established sessions that intermittently fail

In order to solve these issues you need to know about these topics:

- LLC Frame Formats
- LLC2 Modes and Session Establishment
- LLC2 Asynchronous Balanced Mode Operation
- LLC2 Error Conditions

LLC Frame Formats

This section provides information on LLC frame formats.



Destination Service Access Point (1 byte)		Control Field – Unnumbered (1 byte)		
dddd ddx xxxx xxlx xxxx xxxl	Dest. Addr. IEEE Defined Group Address	CCCC CC11 000F 1111 010P 0011 011F 0011 011P 1111 100F 0111 101z 1111 111z 0011	xx-xx 0F-1F 43-53 63-73 6F-7F 87-97 AF-BF	Unnumbered format Disconnect Mode Disconnect Unnumbered Ack. SABME Frame Reject XID
Source Service Access Point (1 byte)		E3-F3 Test		
		Control Field – Supervisory (2 bytes)		
ssss ssxx xxxx xxlx xxxx xxxl	Source Address IEEE defined Response LPDU	CCCC CC01 0000 0001 0000 0101 0000 1001	xx-xx 01-xx 05-xx 09-xx	Supervisory Format Receiver Ready Receiver Not Ready Reject
		Control Field – Information frames (2 bytes)		
		ssss sss0 xxxx		Information format
<p>P = Poll bit set to "1"</p> <p>F = Final bit set to "1"</p> <p>Z = Poll/Final bit set to either "0" or "1"</p>				

An LLC frame is called an LLC Protocol Data Unit (LPDU), and is formatted as shown here:

DSAP (1 byte)-SSAP (1 byte)-Control Field (1 or 2 bytes)-Information Field (0 or more bytes)

DSAP Field

The Destination Service Access Point (DSAP) identifies the SAP for which the LPDU is intended. The DSAP consists of six address bits, a user bit (U) and an Individual/Group (I/G) bit, organized as shown here:

D-D-D-D-D-D-I/G

The U bit indicates whether the address is defined by the IEEE (1) or user-defined (0). The I/G bit indicates whether the SAP is a group address (1) or individual address (0). For our purposes, neither of these bits are too important. All you really need to know is that the DSAP is the destination of the LPDU. Some common ones appear over and over.

SSAP Field

The Source Service Access Point (SSAP) identifies the SAP which originated the LPDU. The SSAP consists of six address bits, a user bit (U) and a Command/Response (C/R) bit, organized as shown here:

S-S-S-S-S-U-C/R

The U bit indicates whether the address is defined by the IEEE (1) or user-defined (0). The C/R bit indicates whether the LPDU is a command or response. When LPDU frames are received, the C/R bit is not considered part of the SSAP. Therefore, the SSAP is normally considered to be only the leftmost seven bits.

Control Field

The LPDU control field contains command, response, and sequence number information. You need to know how to decode the control field in order to determine what happens on a particular LLC2 session. However, decoding information is readily available.

There are three typws of frames:

- I Frames
- Supervisory Frames
- Unnumbered Frames

Although each type has a different format for the control field, you can easily distinguish them through an examination of two bits in the control field.

X-X-X-X-X-X-X-0 = I Frame

X-X-X-X-X-X-X-0-1 = Supervisory Frame

X-X-X-X-X-X-1-1 = Unnumbered frame

The next few sections explain each type of control field.

I Frame

I frames enable you to transfer sequentially-numbered LPDUs that contain information (connection-oriented) between link stations. The format of the I frame contains an NS and NR count. The NS count is the sequence number (modulo 128) of the LPDU currently in transmission. The NR count is the sequence number of the next LPDU I frame that the sender expects to receive. To help you later, remember that NR means "next receive."

NS-NS-NS-NS-NS-NS-NS-NS-0-NR-NR-NR-NR-NR-NR-NR-P/F

The P/F bit is called the P bit in command LPDUs and the F bit in response LPDUs. The P/F bit is set in command LPDUs to request that the remote link station send a response with this bit set. Only one response must be received with the F bit set for every command sent with the P bit set. There are some other details about the use of the P/F bit in relation to error recovery, but that is the general rule.

Supervisory Frame

Supervisory frames perform supervisory control functions, for example, to acknowledge I Frames (RR), to request retransmission of I frames (REJ), and to request temporary suspension (RNR) of I frames. Supervisory frames do not contain an information field. Therefore, supervisory frames do not affect the NS in the sending station, and so do not contain an NS field. Here is the format of a supervisory frame:

0-0-0-0-S-S-0-1-NR-NR-NR-NR-NR-NR-NR-P/F

The "S" bits indicate the type of supervisory frame.

- B'00' = Receiver Ready

A station uses the RR to indicate that the station is ready to receive, and contains the NR count of the next I frame that is due to arrive. When a station sends an RR frame, the station acknowledges receipt of numbered I frames from the remote station of up to NR - 1.

- B'01' = Receiver Not Ready

A station uses the RNR to indicate that the station is temporarily not ready to receive. RNR also contains the NR count which follows the same rules RR. Transient periods of RNRs are not always indicative of a network problem. If RNRs are persistent, look for congestion in the end station.

- B'10'=Reject

A station uses the REJ to request the retransmission of I frame LPDUs starting with the number indicated in the NR count. REJ is not indicative of a serious problem (which means it is recoverable). If you see many REJ commands, look for missing (dropped) I frames in the opposite direction. Do not confuse an REJ with a Frame reject (FRMR). A FRMR is an unnumbered frame and is always indicative of a serious problem.

Unnumbered Frames

Unnumbered frames provide link control functions, for example, mode setting commands and responses. In some cases, unnumbered information frames can also be sent. Unnumbered frames are only one byte in length. They do not contain fields for NR or NRS counts. Here is the format of an unnumbered frame:

M-M-M-P/F-M-M-1-1

The "M" bits indicate the type of unnumbered frame.

- B'00011'=DM Response (0x1F)

A link station sends a DM response to report that it is in asynchronous disconnect mode. This means is that the link is not active. If a link station was active and suddenly begins to send DMs, the link station was probably reset.

- B'01000'=DISC Command (0x53)

A link station sends a DISC to terminate asynchronous balanced mode. The DISC command informs the remote link station that it suspends operation. The correct response to a DISC command is a UA (if the station is in ABM), or a DM (if the station is in ADM).

- B'01100'=UA Response(0x73)

A link station sends an UA in response to the SABME and DISC commands.

- B'01111'=SABME Command(0x7F)

A link station sends a SABME to initiate data transfer in asynchronous balanced mode. The correct response to a SABME is a UA. When a station receives a SABME command, the station resets the NR and NS counts to zero. The sending station does likewise when it receives the UA response.

- B'10001'=FRMR Response(0x87)

A link station sends a Frame Reject response to report an error in an incoming LPDU from the other link station. When you see a FRMR, the station that sends the FRMR has detected an unrecoverable error. It is not the cause of the error. Any frames that arrive after the FRMR error has occurred are ignored until a DISC or SABME is received.

A FRMR response contains information about the cause of the FRMR condition.

Bytes 0 and 1 contain the contents of the control field of the LPDU which caused the frame reject. Bytes 2 and 3 contain the NS and NR counts, respectively. Byte 4 contains several bits that identify the type of error as shown here:

0-0-0-V-Z-Y-W-X

The V bit indicates that the NS number carried by the control field in bytes 0 and 1 is invalid. An NS is invalid if greater than or equal to the last NS plus the maximum receive window size. When this condition occurs, the link station sends a REJ LPDU, not a FRMR response.

The Z bit indicates that the NR that the control field carries indicated in bytes 0 and 1 does not refer to either the next I frame or an I frame that has already been transmitted but not acknowledged.

Note: It is all right to receive the the same NR count multiple times.

The NR count is only invalid if the count references an I frame that has already been acknowledged or if the count skips ahead to one that has not been transmitted yet. The former is the most common case of this type of error. When you see this type of error, it usually means frames were received out of sequence, and you should look for the network that delivers frames out of order. It is possible that the sending link station transmitted them out of order, but very unlikely.

The Y bit indicates that the length of the I field in the received LPDU exceeded the available buffer capacity. If this situation occurs, look for problems in the end stations, not the network.

The X bit indicates that the LPDU contained an I field when it must not have, or a FRMR response was received that did not contain 5 bytes. This appears to be an end station problem, not a network problem.


The W bit indicates that an unsupported LPDU was received. This is an end station problem.

- B'10111' XID Command or Response

A link station uses the XID command to convey characteristics of the sending node and to cause the remote link station to respond with an XID response. Link stations can send and receive XIDs in various formats, including SNA formats.

- B'11100' TEST Command or Response

A link station sends the TEST command to cause the remote link station to respond with a TEST response as soon as possible. The TEST command is generally used for path discovery in a source–route bridging environment.



LLC Control Field Summary

Value	Unnumbered Frames
0x0F or 0x1F	Disconnect Mode (DM) Response
0x43 or 0x53	Disconnect (DISC) Command
0x63 or 0x73	Unnumbered Acknowledgment (UA) Response
0x6F or 0x7F	Set Asynchronous Balanced Mode (SABME) Command

0x87 or 0x97	Frame Reject (FRMR) Response
0xAF or 0xBF	Exchange Id (XID) Command or Response
0xE3 or 0xF3	Test (TEST) Command or Response

Value	Supervisory Frames
0x01	Receiver Ready (RR)
0x05	Receiver Not Ready (RNR)
0x09	Reject (REJ)

Value	Information Frames
0bnnnnnnn0	Information Frame (INFO)

LLC2 Modes and Session Establishment

There are two modes of LLC2 operation:

- Asynchronous Balanced Mode Extended
- Asynchronous Disconnect Mode

Asynchronous Balanced Mode Extended (ABME)

ABME is a balanced mode of operation between two link stations. Balanced mode refers to the fact that either station can send commands at any time, independently of the other link station. Contrast this with SDLC, which operates in unbalanced mode. In unbalanced mode, the secondary station must wait to be polled by the primary before it can send data. As a result of balanced mode operation, polling does not occur on LLC2 circuits in the traditional sense. A station does send keepalives to maintain the session, but it is not necessary to send these frequently for optimal performance as in SDLC. For this reason, the keepalive timer is typically 10 seconds or greater. It is important to note that the end stations can increase this keepalive timer to reduce overhead. An increase the keepalive timer has no negative effect on throughput or response time.

A station enters ABME after the station sends or receives a UA to to SABME command. When in ABME, the station can send and receive numbered information frames.

Asynchronous Disconnect Mode (ADM)

Before and after a station terminates ABME, the station is in Asynchronous Disconnect Mode. In ADM, the link is logically disconnected; therefore, no I frames or supervisory frames can be sent. A station can enter ADM under these conditions:

- Receipt of a DISC command
- The linkstation is activated
- Receipt of a DM response
- Retry limit is is exhausted

Here is an example of a link station activation sequence:

```
To1 4000.0840.0001 8800.5a94.7d94 SABME F0F07F
```

```
To1 4000.0840.0001 8800.5a94.7d94 UA F0F173
To 1 4000.0840.00018800.5a94.7d94 RR nr=0 F0F001
To1 4000.0840.0001 8800.5a94.7d94 INFO nr=0 ns=0 F0F00000 ...
To1 4000.0840.0001 8800.5a94.7d94 RR nr=1 F0F101
To1 4000.0840.0001 8800.5a94.7d94 INFO nr=1 ns=1 F0F00202 ...
To1 4000.0840.0001 8800.5a94.7d94 RR nr=2 F0F101
To1 4000.0840.0001 8800.5a94.7d94 INFO nr=2 ns=2 F0F00000 ...
```

LLC2 Asynchronous Balanced Mode Operation

Stations that operate in ASBM do not have a strict sense of primary or secondary stations. Stations do not need to poll or be polled to transfer data. Stations can transmit data to any station asynchronously. Stations have peer-to-peer relationships.

Even though there is no strict sense of primary and secondary, a sending station requires a link level response called an acknowledgment from the receiving station for each numbered information frame sent. A station can continue to transmit I frames to another station until the number of unacknowledged frames reaches a limit. This number is called the "window size," and typically defaults to 7. You can increase the window size on circuits where there is a lot of latency to avoid the need for the sending station to stop and wait for a response. This is usually not necessary, especially in situations where LLC is locally acknowledged. When a sending station reaches the sending window, the station sets the poll bit to force the receiving station to send a response. In the router, the window size is called the llc2 local-window.

A receiving station has the option to withhold acknowledgments until either a certain number of I frames arrive or a timer expires. These parameters are called N3 and T2, respectively. This way, multiple frames can be acknowledged with one RR frame, or an acknowledgment can be sent on top of on an I frame. Cisco calls the N3 counter llc2 ack-max. The default value of three indicates that the router withholds an acknowledgment until the router receives three I frames, or until the T2 timer, or the llc2 ack-delay-time, expires.

Modification of these parameters on a station without consideration of the partner station can affect response time and throughput. For example, consider what would happen if the sending station local-window is set to 5 and the receiving station has values of 7 for ack-max and 500 milliseconds for ack-delay-time.

In this case, the sending station sends five frames, then wait for an acknowledgment before continuing. Because the receiver withholds acknowledgments until seven frames are received, it will not send an acknowledgment until the 500 millisecond delay time expires. You can dramatically improve performance if you lower the ack-max value on the receiving station.

Another common LLC2 parameter is called the Ti timer. The router calls this the llc2 idle-time. The purpose of the Ti timer is to keep the LLC2 session active during periods when no I frames are being transmitted. You cannot improve throughput and performance if you lower this value. When the Ti timer expires, an RR frame is sent with the poll bit on to cause a response from the receiver. If the station does not respond, the station is retried after llc2 tpf-time until the number of retries defined by llc2 n2 has expired. At that time, the session is torn down.

Increase the idle time to reduce the amount of overhead on an LLC2 circuit and you can adjust this as an alternative to local ack. Consider an example where 200 DSPUs are connected to an NCP. Each of the PUs

maintains an independent LLC2 session. If each one sends a keepalive every ten seconds, there are 20 frames of overhead every second. If you increase the idle time to 30 seconds, the amount of overhead frames reduces to 6.67 frames per second. The drawback to this approach is that stations take longer to discover that their partner is unreachable. But depending on your situation, this could be a good thing.

LLC2 Tunable Parameters

Command	Default	Description
llc2 ack-delay-time>/b> msec	100	The amount of time to wait for a response before sending an acknowledgment when the ack-max value has not been reached.
llc2 ack-max count	3 frames	The number of frames to receive before sending an acknowledgement.
llc2 idle-time msec	10000	The amount of time between polls during idle periods of time.
llc2 local-window count	7	The number of frames to send before waiting for a response.
llc2 n2 count	8 retries	The number of times unacknowledged I frames or polls are sent without receiving a reply before termination the session.
llc2 t1-time msec	1000	The amount of time to wait for a response before resending I frames. This time needs to be large enough to accommodate the round-trip delay.
llc2 tbusy-time msec	9600	The amount of time to wait before polling a station that has sent a RNR. Change the value only to increase the value for stations that have unusually long, busy periods before they clear their status.
llc2 tpf-time msec	1000	The amount of time to wait for a final response before resending the poll frame.
llc2 trej-time msec	3200	The amount of time to wait for a correct frame after sending a REJ.


You can use the **show llc** command to see the values of these parameters:

```
ibu-7206#sh llc
LLC2 Connections: total of 1 connections
TokenRing3/0 DTE: 4001.68ff.0000 4000.0000.0001 04 04 state NORMAL
```

V(S)=5, V(R)=5, Last N(R)=5, Local window=8, Remote Window=127
 akmax=3, n2=8, Next timer in 8076
 xid-retry timer 0/60000 ack timer 0/1000
 p timer 0/1000 idle timer 8076/10000
 rej timer 0/3200 busy timer 0/9600
 akdelay timer 0/100 txQ count 0/2000

Examples of LLC2 Parameter Configurations

In a typical DLSw+ network with a Token Ring LAN at either end, the configuration of LLC2 parameters are made on the outgoing token ring interface.




There are two separate LLC2 sessions. Therefore, configure LLC2 parameters as shown here:

```
hostname dlsw1
!
source-bridge ring-group 100
!
dlsw local-peer ...
dlsw remote-peer ...
!
interface token-ring 0
source-bridge 10 1 100
llc2 tpf-timer 2000
llc2 n2 20

hostname dlsw2
!
source-bridge ring-group 100
!
dlsw local-peer ...
dlsw remote-peer ...
!
interface token-ring 0
source-bridge 20 1 100
llc2 tpf-timer 2000
llc2 n2 20
```

Note: These skimmed configurations show only relevant LLC2 parameter configurations.

LLC2 parameter configurations must match the LLC2 parameters to the FEP (to DLSw1 router) and PC (to DLSw2 router). When the central site DLSw+ peer is on a CIP router, the configuration is slightly different.




The remote DLSw+ router configuration remains unchanged. However, the LLC2 session at the central site is between the CIP and the LLC2 stack in IOS. The CIP represents the Mainframe, and the LLC2 parameters from the Mainframe out towards IOS are configured under the adapters on the LAN Token Ring (on the CIP). The LLC2 parameters from IOS towards the Mainframe are configured on the outgoing interface. That is, interface channel x/2 (for CIP) and interface channel x/0 (for xCPA). For example:

```
hostname dlsw1
!
source-bridge ring-group 100
!
dlsw local-peer ...
dlsw remote-peer ...
!
interface channel 0/2
llc2 tpf-timer 2000
llc2 n2 20
lan tokenring 0
source-bridge 10 1 100
adapter 0 4000.7513.0000
llc2 tpf-timer 2000
llc2 n2 20
```

Note: These skimmed configurations show only relevant LLC2 parameter configurations.


If the CIP router connects over the LAN to a local station, you need only the LLC2 parameters under the CIP adapters. The LLC2 parameters would then be matched to those of the PC. Any LLC2 parameters under interface channel 0/2 are ineffective.



```
hostname rtr1
!
source-bridge ring-group 100
!
interface channel 0/2
lan tokenring 0
source-bridge 10 1 100
adapter 0 4000.7513.0000
llc2 tpf-timer 2000
llc2 n2 20
```

Note: These skimmed configurations show only relevant LLC2 parameter configurations.

If devices connect into DLSw+ through bridge-groups, LLC2 parameters are configured on the DLSw+ bridge-group statement as shown here:




```
hostname dls2
!
dls local-peer ...
dls remote-peer
dls bridge-group 1 llc2 tpf-timer 2500 n2 20
!
interface ethernet 0
bridge-group 1
bridge 1 protocol ieee
```

Note: These skimmed configurations show only relevant LLC2 parameter configurations.

Note: Although you can configure LLC2 under the ethernet 0 interface, these parameters have no effect. DLSw bridge-group LLC2 was new in Cisco IOS Software Release 11.3.

When the router is configured as an end-station (for example, SNASw and DSPU), you must configure the LLC2 parameters on the outgoing interface. Note that not all virtual interfaces support configuration of LLC2 parameters. For example:



Note: These skimmed configurations show only relevant LLC2 parameter configurations.

```
hostname snasw1
!
Interface fastethernet 0/0
llc2 tpf-timer 2000
llc2 n2 20
!
snasw cpname neta.snasw1
snasw port FASTETH0 FastEthernet0/0 conntype nohpr
```

LLC2 Error Conditions

Some errors on LLC2 sessions are normal and recoverable, for example, occasional missed frames or frames out of order. These usually result in a REJ and retransmitted frames. Excessive retransmits are not normal, and you must identify the cause and resolve the issue. Excessive retransmits can occur due to drops, multiple paths, congestion and excessive latency.

Some errors in LLC2 are unrecoverable and always result in a session outage. These errors are exclusively protocol violations. They can occur when stations send undefined control fields or other erroneous information. These protocol violations can cause a station to send a FRMR response. The station that sends the FRMR response is most likely not the violator, but merely the messenger. The FRMR contains information that identifies why the FRMR is sent, which is most commonly when a station requests another station to resend an I frame that it has already acknowledged. Because the station cannot retransmit the frame (because it has already discarded it), it has no choice but to terminate the LLC session. When this type of error occurs, the most likely cause is that frames are out of order.

Related Information

- [Technical Support & Documentation – Cisco Systems](#)

[Contacts & Feedback](#) | [Help](#) | [Site Map](#)

© 2008 – 2009 Cisco Systems, Inc. All rights reserved. [Terms & Conditions](#) | [Privacy Statement](#) | [Cookie Policy](#) | [Trademarks of Cisco Systems, Inc.](#)

Updated: Sep 09, 2005

Document ID: 12247
