

Designing SDLC, SDLLC, and QLLC Internetworks

This chapter addresses some of the special requirements for implementing routing technology within IBM System Network Architecture (SNA) environments. Internetworking within an SNA environment often involves making special accommodations for devices that were not originally designed for connection to meshed internetworks.

This chapter describes three techniques designed to enable internetworking capabilities for SNA-based network architectures:

- SDLC via STUN
- SDLLC Implementation
- QLLC Conversion

The sections that describe serial tunneling (STUN), Synchronous Data Link Control (SDLC) over the Logical Link Control, type 2 (LLC) protocol (SDLLC), and Qualified Logical Link Control (QLLC) focus on the following topics:

- Technology overview and issues
- Router technology options, implementation guidelines, and configuration examples


Note For information about IBM serial lines, refer to Appendix B, “IBM Serial Link Implementation Notes.”

SDLC via STUN

SDLC via serial tunneling (STUN) encapsulates SDLC frames into Internet Protocol (IP) packets and routes the encapsulated packets over IP-supported network media. The SDLC frame is transmitted without modification, and the information within the frame is transparent to the network. All SNA physical unit (PU) types are supported. This section focuses on the SDLC data-link protocol and its various configurations and then explains how to implement STUN.

Figure 5-1 illustrates elements of STUN configuration in an environment that includes front-end processors (FEPs) and cluster controllers.

Figure 5-1 Sample STUN network configuration.



SDLC Data Link

SDLC is the synchronous, bit-oriented protocol used by the SNA data-link control layer. As formally defined by IBM, SDLC is a line discipline for managing synchronous, code-transparent, serially transmitted bit information over a data link. Transmission exchanges can be full duplex or half duplex and can occur over switched or nonswitched links. The configuration of the link connection can be point-to-point, multidrop, or loop.

Common physical link-layer implementations are V.24 (EIA/TIA-232, formerly RS-232), V.35, and X.21. This section describes SDLC as it applies to STUN.

The SDLC data link allows a reliable exchange of information over a communication facility between SNA devices. The protocol synchronizes receivers and transmitters and detects transmission errors. It accomplishes these functions by acknowledging frame receipt and by performing a cyclic redundancy check (CRC) on the data.

Supported Data-Link Configurations

This section provides information related to router-specific hardware implementation. Table 5-1 provides a matrix of SDLC support for V.24.

Table 5-1 SDLC Support for V.24 (EIA/TIA-232)

Product Type	NRZ/NRZI	DTE/DCE	Full Duplex	Half Duplex	Maximum MTU
Cisco 7000	Both	Both	Yes	Yes	4 KB
Cisco 7010	Both	Both	Yes	Yes	4 KB
AGS+	Both	Both	Yes	Yes	4 KB
MGS	Both	Both	Yes	Yes	4 KB
Cisco 2500	Both	Both	Yes	Yes	8 KB
Cisco 4000	Both	Both	Yes	4T card only	8 KB
Cisco 4500	Both	Both	Yes	4T card only	8 KB
Cisco 3104	Both	Both	Yes	Dual serial card only	8 KB
Cisco 3204	Both	Both	Yes	Dual serial card only	8 KB


The following notes apply to the entries in Table 5-1:

- For the Cisco 7000, Cisco 4000, Cisco 4500, and Cisco 3000 products, support of data terminal equipment (DTE) or data communications equipment (DCE) functionality depends on which cable is used.
- For the AGS+ and MGS, if you are using a nonreturn to zero inverted (NRZI) applique, the systems support DCE natively. A special cable is required to support DTE mode operation. Prior to the availability of the NRZI applique, customers specifically ordered a DCE or DTE applique.
- Half-duplex support is available for the AGS+ and MGS with Software Release 9.1(7) or later. The NRZI applique, three-port SCI card, and Software Release 9.1(7) or later are all required for half-duplex support.
- Prior to software releases 8.3(6), 9.0(3), or 9.1(2), only 2-KB frame sizes were supported. When increasing maximum transmission unit (MTU) size, consider interface card buffer memory size constraints.

SDLC Frame Format

The SDLC frame format is illustrated in Figure 5-2.

Figure 5-2 SDLC frame format.



The *Flag* field starts and ends the frame and initiates and terminates error checking. When the link is idle, the router sends streaming flags to maintain link synchronization, but this is not necessary to keep the link up.

The *Addr* field contains the SDLC address of the secondary station regardless of whether the frame is coming from the primary or secondary station. The *Addr* field can contain a specific address, a group address, or a broadcast address. Routers support specific addresses and support broadcast addressing on a limited basis.

The *Control* field is a 1-byte field (for modulo 8 frames) or a 2-byte field (for modulo 128 frames). The extra byte is required for modulo 128 frames to accommodate larger send and receive *frame count fields*. The value of the *Control* field identifies three different frame formats, as shown in Table 5-2.

Table 5-2 Components of the Control Field

Format	Binary Configuration	Hex Equivalent		Command Name	Acronym
		[P/F off)	(P/F on)		
Unnumbered	000 P ¹ /F ² 0011	03	13	Unnumbered Info	UI
	000 F 0111	07	17	Request Initialization Mode	RIM
	000 P 0111	07	17	Set Initialization Mode	SIM
	000 F 1111	0F	1F	Disconnect Mode	DM
	010 F 0011	43	53	Request Disconnect	RD
	010 P 0111	43	53	Disconnect	DISC
	011 F 0011	63	73	Unnumbered Ack	UA
	100 P 0011	83	93	Set Normal Response.	SNRM
	110 P 1111	CF	DF	Set Normal Response. Mode Ex.	SNRME
	100 F 0111	87	97	Frame REJECT	FRMR
	101 P/F 1111	AF	BF	Exchange ID	XID
111 P/F 0011	E3	F3	Test	TEST	
Supervisory	RRR ³ P/F 0001	x1 ⁴	x1	Receive Ready	RR
	RRR P/F 0101	x5	x5	Receive Not Ready	RNR
	RRR P/F 1101	x9	x9	Reject	REJ
Information	RRR P/F SSS0 ⁵	xx	xx	Numbered Info Present	Transfer

1 P = Poll bit

2 F = Final bit

3 RRR = Nr (receive count)

4 x = Any single digit hexadecimal value

5 SSS = Ns (send count)

The *Info* field is a variable-length field containing a path information unit (PIU) or exchange identification (XID) information. Table 5-3 lists supported PIUs.

Table 5-3 PIU Support

PIU Type	Router Support
PIU FID0-bisync and start/stop (non-SNA)	Not supported
PIU FID1-host channel to FEP to remote FEP	Supported via STUN
PIU FID2-FEP to cluster controller (PU 2)	Supported via STUN and SDLLC
PIU FID3-FEP to SNA terminal (PU 1)	Supported via STUN
PIU FID4-FEP to FEP using virtual route	Supported via STUN
PIU FIDF-FEP to FEP (VR SNF overflow sweep)	Supported via STUN
XID 2-Contains PU parameters for PU types 1, 2, 4, and 5	PU types 2 and 4 supported via STUN and SDLLC
XID 3-APPN variable format for PU 2 and PU 2.1	Not supported

The *frame check sequence (FCS)* field is a 2-byte field that, for transmitted data, contains the result of a CRC performed on the first bit after the Flag field through the last bit of the Info field. If the frame format is unnumbered or supervisory, the CRC is performed through the last bit of the Control field. As the remote device receives the data, it performs the same CRC computation and compares the result with the contents of the FCS field. If the comparison fails to find a match, the frame is discarded and recovery procedures take place.

STUN Configuration for SDLC

The following sections provide design and implementation information for a variety of STUN-related configuration topics:

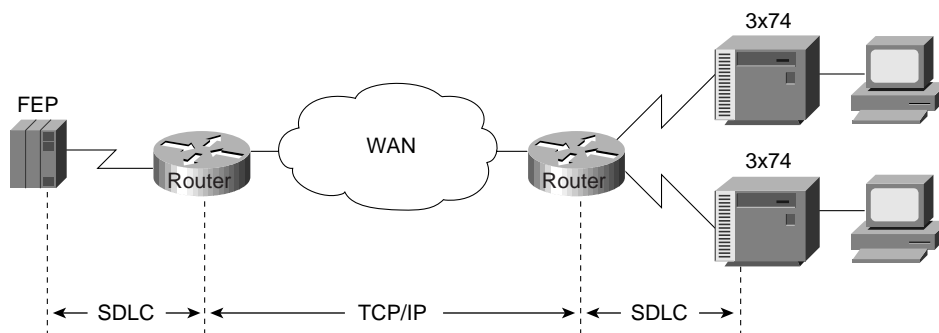
- Local Acknowledgment
- Virtual Multidrop
- SDLC Address Prioritization
- SDLC Two-Way Simultaneous Mode
- LU Address Prioritization
- Flow Control
- Transmission Groups and Class of Service Capabilities
- SNA Host Configuration Considerations for STUN

Local Acknowledgment

Local termination of SDLC sessions allows frames to be locally acknowledged by the receiving router. By locally terminating SDLC sessions, acknowledgment and keepalive traffic is prevented from traversing the backbone, and SNA sessions are preserved if the network fails.

Local acknowledgment locally terminates supervisory frames, which include receiver-ready, receiver-not-ready, and reject frames. Figure 5-3 illustrates the operation of SDLC local acknowledgment.

Figure 5-3 STUN-to-SDLC local acknowledgment.




Note Local acknowledgment requires that TCP/IP sessions be maintained between the routers to provide reliable transport.

Virtual Multidrop

Virtual multidrop exploits SDLC address mapping to allow an FEP to communicate with multiple cluster controllers. With a virtual multidrop configuration, the address of each SDLC frame is checked individually. Only addresses that match the configuration are forwarded to the specified destination, which allows an FEP to communicate with multiple 3174s from a single serial link—a multidrop link. You can also use SDLC address mapping as a security feature to restrict access based on SDLC address, as shown in Figure 5-4.

Figure 5-4 SDLCtransport in virtual multidrop environment.



The following steps are required to establish the network configuration illustrated in Figure 5-4:


- Step 1** Include a LINE definition in the Network Control Program (NCP) running in the FEP, followed by PU definitions for SDLC address 7, 9, and A. The NCP interprets these definitions as a multidrop link.
- Step 2** Use the `stun route address tcp` global configuration command to specify how to forward frames.
- Step 3** Determine whether priority queuing is required. If so, local acknowledgment is required.

SDLC Broadcast Across Virtual Multidrop Lines

The SDLC broadcast feature (introduced in Cisco IOS Software Release 10.2) allows SDLC broadcast address 0xFF to be replicated for each of the STUN peers, so that each end station receives the broadcast frame.

In Figure 5-5, the FEP views the end stations as if they were on an SDLC multidrop link. Router A duplicates any broadcast frames sent by the FEP and sends them to any downstream routers (in this example, routers B and C).

Figure 5-5 SDLC broadcast in virtual multidrop line environment.



The **sdlc virtual-multidrop** interface configuration command enables SDLC broadcast and should be used only on the router that is configured as the secondary station on the SDLC link. In addition, the **stun route address tcp** command for SDLC address 0xFF must be configured on the secondary station (in this example, Router A) for each STUN peer. A sample configuration follows:

```
stun peername xxx.xxx.xxx.xxx
stun protocol-group 1 sdlc
!
interface serial 1
encapsulation stun
stun group 1
stun sdlc-role secondary
sdhc virtual-multidrop
sdhc address 01
sdhc address 02
sdhc address 03
stun route address 01 tcp yyy.yyy.yyy.yyy local-ack
stun route address 02 tcp zzz.zzz.zzz.zzz local-ack
stun route address 03 tcp zzz.zzz.zzz.zzz local-ack
stun route address FF tcp yyy.yyy.yyy.yyy
stun route address FF tcp zzz.zzz.zzz.zzz
```

SDLC Address Prioritization


For STUN prioritization of SDLC addresses over simple serial transport connections, use the **priority-list** or the **queue-list** global configuration command and the **priority-group** interface configuration command or the **custom-queue-list** interface configuration command, respectively, on the interface that connects to the remote router (the output interface).

For STUN prioritization of SDLC addresses over TCP/IP transport connections, you must configure the **priority-list** global configuration command and use the **priority-group** interface configuration command on the interfaces that connect to the end devices (the input interfaces). Also, you must specify the **local-ack** and **priority** keywords of the **stun route address tcp** global configuration command.

SDLC Two-Way Simultaneous Mode

Two-way simultaneous mode (introduced in Cisco IOS Software Release 10.2) allows a router that is configured as a primary SDLC station to utilize a full-duplex serial line more efficiently. When two-way simultaneous mode is enabled in a multidrop environment, the router can poll a secondary station and receive data from that station while it sends data to, or receives data from, a different secondary station on the same serial line. (See Figure 5-6.)

Figure 5-6 Two-way simultaneous mode in a multidrop environment.



The **sdlc simultaneous** command enables two-way simultaneous mode in a multidrop environment. When two-way simultaneous mode is enabled for a point-to-point connection to a secondary station, the router can send data to the secondary station even if there is an outstanding poll, as long as the window size limit is not reached. The **sdlc simultaneous** command with the **single** keyword enables two-way simultaneous mode in a point-to-point link environment.

LU Address Prioritization

To prioritize logical units, use the **locaddr-priority-list** global configuration command on each router. For example:

```
locaddr-priority-list 1 02 high
locaddr-priority-list 1 03 high
locaddr-priority-list 1 04 low
```

You must also assign a priority list to the STUN priority ports using the **priority-list** global command. For example:

```
priority-list 1 protocol ip high tcp 1994
priority-list 1 protocol ip medium tcp 1990
priority-list 1 protocol ip low tcp 1991
```


The serial interfaces attached to the end systems (input interfaces) must be associated with priority lists using the **locaddr-priority** and **priority-group** interface configuration commands. The **locaddr-priority** command links the interface to a local LU priority list (specified with the **locaddr-priority-list** global configuration command). The **priority-group** command links the interface to a TCP priority list (specified with a **priority-list** global configuration command). For example:

```
interface serial 1
locaddr-priority 1
priority-group 1
```

In addition, you must specify the **local-ack** and **priority** keyword options of the **stun route address tcp** global configuration command.

The LU address prioritization feature has been enhanced to allow you to specify the PU on which an LU resides. This enhancement is important because there might be multiple PUs on a multidropped SDLC line that have the same LU address. For example, in Figure 5-7, LU02 on 3174-2 is a 3287 printer, and LU02 on 3174-1 is a 3278 terminal. Do not assign the same priority to the printer and the terminal.

Figure 5-7 LU prioritization for STUN.



As of Software Release 9.1(9), LU address prioritization for both remote source-route bridging (RSRB) and STUN solved this problem. In addition to the LU address, you can specify the SDLC address to identify a PU in a multidropped SDLC line. The syntax of the **locaddr-priority** global configuration command follows:

locaddr-priority *list lu-address sdlc secondary*

The keyword **sdlc** indicates the next byte (in hexadecimal), and *secondary* is the secondary SDLC address.

Flow Control

SDLC-level flow control is also offered with local termination. When the router detects that the TCP queue is 90 percent full, it blocks further SDLC frames until the TCP queue recedes to 80 percent full. This is accomplished by transmitting receiver-not-ready frames.

There is also a flow control protocol between STUN peers. When SDLC output queues become congested, a router can request the remotely attached router to exert back-pressure on the SDLC link.

Transmission Groups and Class of Service Capabilities


This section describes the transmission group and class of service (CoS) support that NCP-to-NCP communications provide, including the following topics:

- Typical NCP-to-NCP Communications
- NCP-to-NCP Communications over a Routed Network
- Transmission Group and CoS Support
- Transmission Group and CoS Design Guidelines and Notes

Typical NCP-to-NCP Communications

In a typical NCP-to-NCP communications arrangement, a host is channel-attached to an FEP acting as an NCP. In Figure 5-8, NCP A is the primary SDLC station, and NCP B (remote NCP) is a secondary SDLC station. The NCPs dynamically determine their relationship; the NCP with the higher subarea number becomes the primary SDLC station. NCP V5R4 and later allows you to determine which NCP is the primary and which NCP is the secondary station.

Figure 5-8 Typical NCP-to-NCP multilink transmission group communication configuration.



A *transmission group* is defined as one or more parallel SDLC links connecting adjacent PU Type 4 (NCP) nodes. Transmission groups are used to increase the reliability of the logical link connection between NCPs and to provide additional bandwidth capacity. When one link fails or is congested, data is routed on one of the other links in the group. The transmission group function is implemented at the path control (PC) layer of the NCP architectural model. The PC layer encapsulates request/response units in PIUs and sends them to the data-link control (DLC) layer for transmission.

The PC layer uses the transmission header of the PIU to route messages through the network. SNA defines different transmission header formats and identifies the different formats by Format Identification (FID) type. A transmission header of type FID 4 is used to route data between type 4 nodes that support explicit and virtual routes.

The NCP assigns a sequence order number to each link in a transmission group. In later versions of NCP, you can specify the sequence order in which an NCP should use the transmission group links; otherwise, this order is determined by the order of link activation. Deactivation and reactivation of a link cause it to become the last activated link in the transmission group, and PIU traffic will be sent on the last activated link only if all other links fail or are busy.

Traditionally, the PC layer communicates directly with the DLC layer to transmit PIUs. When sending PIUs over a multilink transmission group, a transmission group layer exists between the PC and DLC layers. The transmission group layer contains a transmit queue. When the transmission group layer gets a frame to send, it checks for the availability of a link in the transmission group in priority (activation default) order. If the transmission group layer finds an available link (that is, a link that is not down and is not busy), it assigns the PIU the next NCP sequence number and sends the frame on that link. NCP sequence numbers range from 0 to 4,095 and wrap on overflow.

When all links in the transmission group are busy, PIUs are placed in the transmission group transmit queue to await transmission. PIUs accumulate in the queue until a link becomes available. When an SDLC link becomes available, a CoS algorithm is performed on the PIUs in the transmit queue. The PIU with the highest priority is dequeued, assigned the next NCP sequence number, and sent on the available link. Sequence numbering must occur when PIUs are removed from the transmit queue because PIUs can overtake each other on the transmit queue when CoS priority processing is performed. PIUs are never preempted by other PIUs on the same SDLC link queue.

There are several reasons why PIUs might arrive at the receiving NCP out of transmission group sequence: links that operate at different speeds, PIUs on different links that have different lengths, and SDLC link transmission errors that cause retransmission. Because PIUs can arrive out of order, the receiving FEP performs resequencing by queuing incoming PIUs if their sequence number is

larger than the next expected sequence number. The algorithm is not important, but the rule is that the receiving FEP propagates PIUs by sequence number order. PIUs with a lower sequence number than expected are considered duplicates and are discarded.

Later versions of NCP deviate from the SDLC standard in their use of SDLC echo addressing. The SDLC secondary NCP sets the high-order bit of the SDLC address when sending a response. For example, the primary NCP sends frames with address 01, and the secondary NCP sends frames with address 81. This addressing scheme limits the range of SDLC addresses from 01 to 7F. Additionally, SDLC address FF is used on frames preceding and during the NCP's XID exchange. The XID and DM frames use the broadcast address.


Another deviation from the SDLC standard occurs in NCP-to-NCP communication when a host NCP loads a remote NCP. Normally, numbered information frames can be sent only after an SNRM (or SNRME), which resets the station counts to ensure that the two stations start with consistent counts. When a host NCP loads a remote NCP, the host sends a Set Initialization Mode (SIM), and the remote NCP responds with a Request Initialization Mode (RIM), which allows numbered information frames to flow. NCPs are allowed to violate the SDLC standard because these violations (echo addressing, broadcast addressing, and sending numbered information frames before SNMRE) occur only when an NCP communicates with another NCP.

NCP-to-NCP Communications over a Routed Network

There are several reasons for routers to carry NCP-to-NCP traffic. Routers allow other protocols to share high-cost bandwidth (such as leased lines) that is set aside for SNA traffic. In the traditional NCP-to-NCP communications, a leased line is required for each line in a transmission group. However, routers enable a number of leased lines to be collapsed into one leased line. In addition, routing capabilities enable routers to dynamically determine wide-area network paths, which can result in a more reliable network.

Figure 5-9 illustrates NCP-to-NCP communications in a router-based internetwork.

Figure 5-9 NCP-to-NCP communications over a routed network.



Changing from STUN pass-through to STUN local acknowledgment has the following benefits:

- Keeps SDLC poll (RR) traffic off of an overutilized WAN
- Prevents NCP timers from expiring due to network delay
- Collapses multiple WAN leased lines into one leased line
- Reduces store-and-forward delays through the router network

Consider the situation illustrated in Figure 5-9. Notice that the router attached to the SDLC primary NCP (Router A) acts as an SDLC secondary station, and vice versa for the other router (Router B). For this discussion, assume that all serial links between NCPs and routers are in the same transmission group. This means that the NCP considers the lines to be in transmission group X, and the router considers the lines to be in transmission group Y. There is no relationship between X and Y. X is used in the NCP system generation, and Y is used in the router configuration.

Transmission Group and CoS Support

The following features facilitate the support of transmission groups and CoS in Cisco routers:

- SDLC address violation allowances

Two specific instances are exempt from SDLC addressing restrictions:

- Echo addressing
- Broadcast addressing

- Remote NCP load sequence

During the load sequence, the remote NCP performs minimal SDLC functions. It cannot go into Normal Response Mode (NRM) until it is loaded. The load sequence for a remote NCP starts with a SIM/RIM exchange between NCPs, which initializes each NCP's SDLC frame count to zero. After the SIM/RIM exchange, the NCPs pass numbered information frames; this event normally does not occur until after a SNRM/UA sequence. The router's SDLC transmission group local-acknowledgment support allows loading of remote NCPs when the routers pass through all frames after a SIM/RIM sequence and before a SNRM/UA sequence. After the SNRM/UA exchange, normal local acknowledgment occurs.

- Rerouting in multilink transmission groups

When a router acknowledges an Information frame, it must ensure delivery of that frame to the receiving NCP. If, after the frame is locally acknowledged, the corresponding link in the receiving transmission group is lost, the receiving router reroutes the frame onto another SDLC link in the same transmission group.

- CoS

The sending NCP performs CoS. Each PIU is assigned a sequence number. The best service the routers can perform is to try to preserve the CoS as assigned by the sending NCP via sequence numbers. Therefore, all SNA data PIUs are treated equally with the goal to preserve PIU order. However, virtual route-pacing responses flow at SNA network priority level and do not have sequence numbers (that is, they have a sequence number of 0). The router prioritizes all SNA network priority PIUs higher than SNA data to achieve more efficient virtual route pacing.

Note The router cannot use the PIU to determine whether traffic is interactive or batch. Even if the router could make this determination, prioritizing one type of traffic over another would cause the receiving NCP to waste CPU time resequencing the PIUs. This would also degrade throughput because the receiving NCP would hold PIUs longer when resequencing.

- Flow control tuning for better CoS operation

The **tcp-queue-max** keyword of the **stun route address tcp** global configuration command allows you to tune the size of the outbound TCP queue so that when the WAN becomes congested, frames generated by an NCP can be stored in the router as well as in the NCP. When the size of the outbound TCP queue is small, back-pressure via SDLC RNRs is applied to sending NCPs sooner, causing the NCP to hold more frames. The more frames that are held by the NCP, the more frames to which the NCP's CoS algorithm is applied. The size of the outbound TCP queue should be configured to 70 or above.

Transmission Group and CoS Design Guidelines and Notes

The following guidelines should be considered when implementing transmission groups and CS:

- 1 Bandwidth of the WAN should be greater than, or equal to, the aggregate bandwidth of all the serial lines. If other protocols are also using the WAN, bandwidth of the WAN should be greater than the aggregate bandwidth of all the serial lines.
- 2 If the network delay associated with one line of an NCP transmission group is different from the network delay associated with another line in the same NCP transmission group, the receiving NCP spends additional time resequencing PIUs. This happens when one or more of the NCP transmission group lines is routed and one or more lines is directly connected between NCPs.
- 3 The Software Release 9.1 prioritizing algorithm ensures that only the highest priority traffic is guaranteed to get through. Software Release 9.21 prioritization is enhanced by the addition of *custom queuing*. Custom queuing can be used to guarantee specific bandwidth allocated to protocols with respect to bandwidth allocated to other protocols.

If you are using Software Release 9.1 and an SNA WAN as a multiprotocol backbone, give SNA traffic the highest priority and assign the next highest priority to other mission-critical protocols. In addition, make sure that your WAN bandwidth is significantly greater than your aggregate SNA serial line bandwidth so that your SNA traffic does not monopolize the WAN.

Table 5-4 lists equivalent commands for configuring priority queuing and custom queuing.

Table 5-4 Comparison of Priority Queuing and Custom Queuing Configuration Commands

Priority Queuing	Custom Queuing
priority-list 4 protocol ip high tcp 1994	queue-list 2 protocol ip 1 tcp 1994
priority-list 4 protocol ip medium tcp 1992	queue-list 2 protocol ip 2 tcp 1992
priority-list 4 protocol ip normal tcp 1991	queue-list 2 protocol ip 3 tcp 1991
priority-list 4 protocol ip low tcp 1990	queue-list 2 protocol ip 4 tcp 1990

- 4 When NCPs are directly connected, their poll-and-pause timers should be configured for maximum throughput using the NCP PAUSE statement. Configuration of this parameter depends on whether the NCP is acting as a primary or secondary SDLC station. Table 5-5 outlines the defaults and recommendations as specified in the IBM publication *Tuning and Problem Analysis for NCP SDLC Devices*.

Table 5-5 NCP PAUSE Parameter Guidelines

Pause Statement Parameter	IBM Guideline
NCP primary PAUSE	Specifies the time the NCP will wait between sending polls if it has no data to send. (Default is 0.2 seconds; 0 is recommended.)
NCP secondary PAUSE	Specifies the time that the secondary NCP will wait before returning a frame with the final bit set. (Default is 0.2 seconds; recommended to be high -0.2 to 1.0 seconds.)

Adding routers with local acknowledgment creates two SDLC sessions instead of one. The result is that the two SDLC sessions do not preserve the original characteristics of the original NCP-to-NCP SDLC session. To adapt a secondary NCP to the router environment, change its system generation PAUSE statement to a value between 0.0 and 0.1 seconds, inclusive.

SNA Host Configuration Considerations for STUN

When designing STUN-based internetworks featuring routers and IBM SNA entities, you must carefully consider the configuration of SNA nodes and routing nodes. Appendix D, “SNA Host Configuration for SDLC Networks,” provides examples of SNA host configurations that focus on two specific SNA devices:

- FEP configuration for SDLC links
- 3174 SDLC configuration example

STUN Implementation Checklist

Before implementing a serial tunneling (STUN) internetwork, make sure you are familiar with Synchronous Data Link control (SDLC). Depending on your implementation, you may need to review the “SDLC via STUN” section at the beginning of this chapter.

Use the following steps as a checklist when implementing SDLC STUN in your internetwork:

Step 1 Evaluate your current environment by answering the following questions:

- What host-attached cluster controllers or front end processors (FEPs) are being used (such as 37x5, 3172, and 3174)? The host site might be referred to as a local, core, or backbone site, or as a data center.
 - Through what media is the network connected to the host site?
- STUN: Serial connection at both ends.
- SDLLC: Token Ring at primary station and SDLC at secondary station, or Ethernet at primary stations and SDLC at secondary station.
- Reverse SDLLC: SDLC at primary station and Token Ring or Ethernet at secondary station.
- What are the link speeds for local and remote end systems?
 - What are the current SDLC line utilization measurements of those links that will attach to the router? This information will be helpful in determining the site requirements.
 - What interface types are to be used (for example, V.24 [EIA/TIA-232, formerly RS-232], V.35, X.21)?
 - What modems, data service units (DSUs), channel service units (CSUs), or modem-sharing or line-sharing devices are to be used?
 - What remote end system types are involved (for example: 3174, 3274, or AS/400)?
 - What kind of emulation requirements are involved (for example: half or full duplex, NRZ or NRZI)?
 - What are the current transaction response times? Consider peak load periods and characterize traffic patterns.
 - How many PUs are in place? How many are planned? This information is important for router utilization sizing.
 - How many LUs are in place? How many are planned? Many busy LUs attached to a PU will increase link utilization.

Step 2 Determine current host configurations. Important information includes the following:

- If the FEP is a 3745, 3725, or 3720, the Network Control Program (NCP) definition listing, especially the GROUP, LINE, PU, and LU definition statements.


- Remote controller configuration worksheets for 3x74, 5x94.
 - OS/2 Communication Manager configuration files.
 - Network topology diagram.
- Step 3** Determine what router-based IBM features will best suit your requirements:
- If remote devices are SDLC-attached PU type 2 devices, consider using SDLLC. See the following section, “SDLLC Implementation.”
 - Depending on the specific situation, STUN can be used in many instances and supports all PU types.
- Step 4** Determine what FEP-to-NCP conversion changes are required:
- Are FEP lines multidrop? Is virtual multidrop required? Refer to the “Virtual Multidrop” section earlier in this chapter.
 - Do PU addresses require changing if SDLC address prioritization is used? Refer to the “SDLC Address Prioritization” section earlier in this chapter.
 - Does the reply timeout T1 timer need to be increased to accommodate network delays if local acknowledgment is not used?
 - Does the “Retries” parameter need to be adjusted for longer elapsed retry sequences?
- Step 5** Determine how end-station controllers are configured and, if possible, configure the router to accommodate them:
- Addresses might need to be changed if you use virtual multidrop. Refer to the “Virtual Multidrop” section earlier in this chapter.
 - NRZ support might be required depending on router platform and interface used. Refer to the “Supported Data-Link Configurations” section earlier in this chapter.
 - If the controller toggles RTS (assumes half-duplex mode), refer to the “Supported Data-Link Configurations” section earlier in this chapter.

SDLLC Implementation

The SDLLC function allows serial-attached devices using the SDLC protocol to communicate with LAN-attached devices using the LLC2 protocol. The basic purpose of the SDLLC function is to consolidate the traditionally disparate SNA/SDLC networks onto a LAN-based, multiprotocol, multimedia backbone network.

Routers use the SDLLC feature to terminate SDLC sessions, to translate SDLC to the LLC2 protocol, and to forward the LLC2 traffic through remote source-route bridging (RSRB) over a point-to-point or IP network. Because a router-based IP network can use any arbitrary media, such as FDDI, Frame Relay, X.25, or leased lines, routers support SDLLC over all such media through IP encapsulation. Figure 5-10 illustrates a general SDLLC media translation internetwork arrangement.

Figure 5-10 SDLLC media translation.



Note In Figure 5-10, the Token Ring connection (Token Ring 10) could also be an Ethernet segment that connects the FEP or 3172 and router.

SDLLC Configuration

The following sections provide design and implementation information for the following SDLLC-related configuration topics:


- Local Acknowledgment
- Multidrop Access
- Router Configuration
- Encapsulation Overhead

Local Acknowledgment

Local acknowledgment of LLC2 sessions allows frames to be locally terminated by the Token Ring-attached router, which guarantees delivery to the ultimate destination through the reliable transport services of TCP/IP. Locally terminating LLC2 sessions enables packet reception to be locally acknowledged, prevents acknowledgment and keepalive traffic from traversing the backbone, and preserves SNA sessions if the network fails. The router that is performing the media translation always acknowledges the SDLC session in an SDLLC environment.

Local acknowledgment locally terminates supervisory frames, which include receiver-ready, receiver-not-ready, and reject frames. Figure 5-11 illustrates the operation of local acknowledgment.

Figure 5-11 Local acknowledgment operation.



Note Local acknowledgment requires that TCP sessions be maintained between the routers. It is not uncommon to see high router CPU utilization at idle traffic times and then decreased utilization as traffic increases. Polling overhead in the router may increase processor use.

Multidrop Access


There are two ways to configure multidrop operation for the SDLC link in an SDLLC environment. The first way is to use a line-sharing device or a modem-sharing device (MSD) to connect multiple controllers at a single site to a single SDLC port on the router. The second way is to connect multiple controllers at different sites through a multidrop service provided by a telephone company. For more information about multidrop connections, refer to Appendix B, "IBM Serial Link Implementation Notes."

Consider line speed, link utilization, and the number of controllers that will share a single line when designing a multidrop environment. In addition, consider the number of attached LUs associated with individual PUs, and determine if these LUs are being heavily used. If so, increase the bandwidth of the attached serial line. When implementing multidrop environments featuring large numbers of PUs and LUs, contact your technical support representative for specific capabilities.

Router Configuration


To configure a router for SDLLC, you need certain virtual telecommunications access method (VTAM) and NCP definition statements. Figure 5-12 illustrates the required configuration information.

Figure 5-12 Required end-to-end SDLLC information.



Consider an example of two routers that implement the SDLLC functionality in an environment that interconnects a remote site to a host channel attached to a 3174 Token Ring gateway, as shown in Figure 5-13.

Figure 5-13 SDLLC implementation with 3174 Token Ring gateway.



Note Routers also support SDLLC implementations in environments with a 3745 Token Ring gateway.

The following conditions apply to the sample network illustrated in Figure 5-13:

- The SDLC address of the 3174R is C1.
- The device called 3174L is a 3174 Token Ring that is channel attached to an IBM mainframe.

The 3174R must be defined in the configuration of the 3174L using the virtual Token Ring MAC address. This address is created in the router configuration; it includes the SDLC address as the last byte. This virtual MAC address is mapped to a host subchannel address. One host subchannel address is assigned for each downstream physical unit at host system generation time. PU and LU functions are defined to VTAM within the switched major node function. The following configuration commands are required on Router 1:

- The **sdllc traddr** interface configuration command with a virtual ring address for the 3174R. Note that the last byte must be 00 and that you must specify the appropriate SDLC address (in this case, C1) for the same last byte during the 3174L customization.

- The **sdllc partner** interface configuration command with the MAC address of the 3174L gateway and the SDLC address of the 3174R.
- The following version of the **sdllc xid** interface configuration command:

```
sdllc xid c1 00000000
```

The **sdllc xid** interface configuration command is specified with all zeros in the IDBLK/IDNUM field to establish the LLC session between Router 1 and the 3174L. All zeros in the node ID field of the XID command indicate that there is no unique node identifier in this field.

Encapsulation Overhead

Cisco routers provide several types of encapsulation solutions. Because encapsulation always incurs a certain amount of overhead, you need to assess the advantages and performance trade-offs of each encapsulation solution within the constraints of your environment.

TCP/IP encapsulation is recommended most frequently because it is very robust, provides a high quality of service, and is media independent. If SDLLC local acknowledgment is required, TCP/IP encapsulation is required. If SDLLC local acknowledgment is not required, Fast-Sequenced Transport (FST) encapsulation is highly recommended because it is less CPU intensive.

Direct High-Level Data Link Control (HDLC) encapsulation can be used only in point-to-point environments. FST and direct HDLC encapsulation are comparable in performance, but FST has more overhead, which may be an issue on low-speed serial links. TCP/IP encapsulation has the most overhead in terms of processor utilization and the WAN connection. If TCP/IP encapsulation with header compression is a requirement, use it only on link speeds of 64 Kbps or less.

Table 5-6 outlines encapsulation overhead for SDLLC and RSRB implementations.

Table 5-6 SDLLC and RSRB Encapsulation Overhead

TCP/IP	FST	TCP/IP with Header Compression	HDLC
CPU intensive	Less CPU intensive	Very CPU intensive	Least CPU intensive
4 bytes for HDLC	4 bytes for HDLC	4 bytes for HDLC	4 bytes for HDLC
20 bytes for IP	20 bytes for IP	3–8 bytes for TCP/IP	16 bytes for virtual ring
20–24 bytes for TCP	16 bytes for virtual ring	16 bytes for virtual ring	
16 bytes for virtual ring			
Total: 60–64 bytes	Total: 40 bytes	Total: 23–28 bytes	Total: 20 bytes

SDLLC Guidelines and Recommendations

The following suggestions can help improve resource response time and network performance:

- *Token Ring frame size*—Allow the Token Ring Interface Coupler (TIC) FEP to send the largest possible frame, and let the router segment the frame into multiple SDLC Information frames.
- *MAXOUT (window size)*—Change the MAXOUT value in the VTAM-switched major node for the 3174 PU. MAXOUT is IBM's terminology for *window size*. IBM recommends setting window sizes on LAN-attached devices to 1 because their tests found no performance benefit with a larger window size. The *red books*, which are published by the IBM International Systems Center, show examples with MAXOUT=1. Because the remote device is an SDLC-attached 3x74, not a Token Ring-attached device, changing MAXOUT to 7 can improve performance dramatically.

- *SDLC line speed*—Increase the line speed of the 3x74 to 19.2 Kbps (older units) or 64 Kbps (newer units) when the controller is directly attached (as opposed to being attached through a modem) to the router. Modem and communication facilities are frequently the limiting factors in determining the line speed in the prerouter configuration.
- *SDLC frame size*—Set the largest SDLC frame size to 521 on newer 3274 models (not 265, which is required for older 3274 models). *See the note that follows.*
- *Request To Send (RTS) control*—Set the 3174 for permanent RTS if the device is not connected via a multidrop service through modem-sharing devices or line-sharing devices. Modem-sharing and line-sharing connections require that RTS be toggled when the device is transmitting. Setting permanent RTS cuts down online turnaround delays and can improve link utilization by 10 percent. (However, setting permanent RTS is unlikely to achieve any perceptible response time improvements.)

Note Changing configurations of end devices, such as terminal controllers, is not recommended. The high number of devices requiring changes and the cost and unavailability associated with these changes can make these modifications onerous. Modify SDLC maximum frame size and RTS control with discretion.

SDLLC Implementation Scenarios

The following case study shows how an internetwork can evolve from a SNA-specific SDLC environment featuring 3x74 controllers and 3270 terminals to a network of PCs with client-server applications. The most important requirement for this evolution is the protection of existing SNA investment.

Assume that the original network consisted of hundreds of SDLC 3x74 controllers connected to a number of 37x5 FEPs in the data center. A disaster recovery center maintains the “mirror-image” of the data center. Many 3x74s are multidrop-connected to the host via 9.6- or 19.2-Kbps leased lines. The challenges facing the corporate MIS organization for this internetwork include the following:

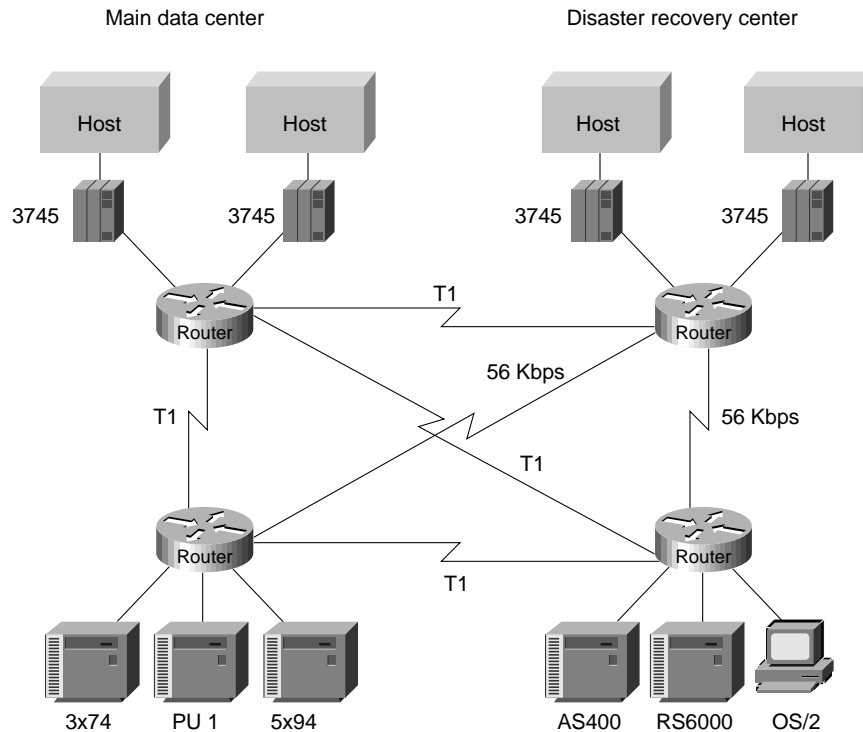
- *Reliability*—When an SDLC line goes down, all the downstream users are affected. There is no network redundancy.
- *Leased line charges*—Providing lines to multiple remote SDLC devices results in excessive service charges and must be minimized.
- *FEP CPU use*—CPU use is higher for SDLC-supported sessions than for LAN-supported sessions.
- *Maintaining VTAM and NCP*—Every move and change requires system programmers to regenerate VTAM/NCP, which increases the cost of maintaining a statistically defined network.
- *Supporting LAN-based applications*—There is a growing need to support LAN-based interconnection, both PC-to-host and PC-to-PC.
- *Availability and up time*—To maintain a competitive advantage, the organization needs to keep SNA sessions alive even if the network fails.

A phased strategy aimed at addressing these challenges would consist of three phases. Each of these phases is discussed in the following implementation examples.

Phase 1: Redundant Backbone Using STUN and Virtual Multidrop

Build a redundant backbone network with routers and high-speed E1 or T1 links in each regional office, as shown in Figure 5-14. Connect multiple SDLC devices to a router via SDLC transport with virtual multidrop. The resulting network increases reliability and minimizes leased-line charges.


Figure 5-14 Connecting multiple SDLC devices via SDLC transport with virtual multidrop.



Phase 2: Fault-Tolerant Host FEP Token Ring and SDLLC Implementation

Implement a fault-tolerant host FEP Token Ring, as shown in Figure 5-15. Connecting existing SDLC devices to the host Token Ring via SDLLC improves response time. Because SDLC devices appear as Token Ring-attached devices to the host, you do not need to regenerate NCP and reload when you are adding or changing PUs and LUs. This can be done dynamically through VTAM-switched major nodes. This implementation also reduces FEP CPU utilization.

Figure 5-15 Fault-tolerant TICs and SDLLC implementation.



Phase 3: Strategic LAN-to-WAN Implementation

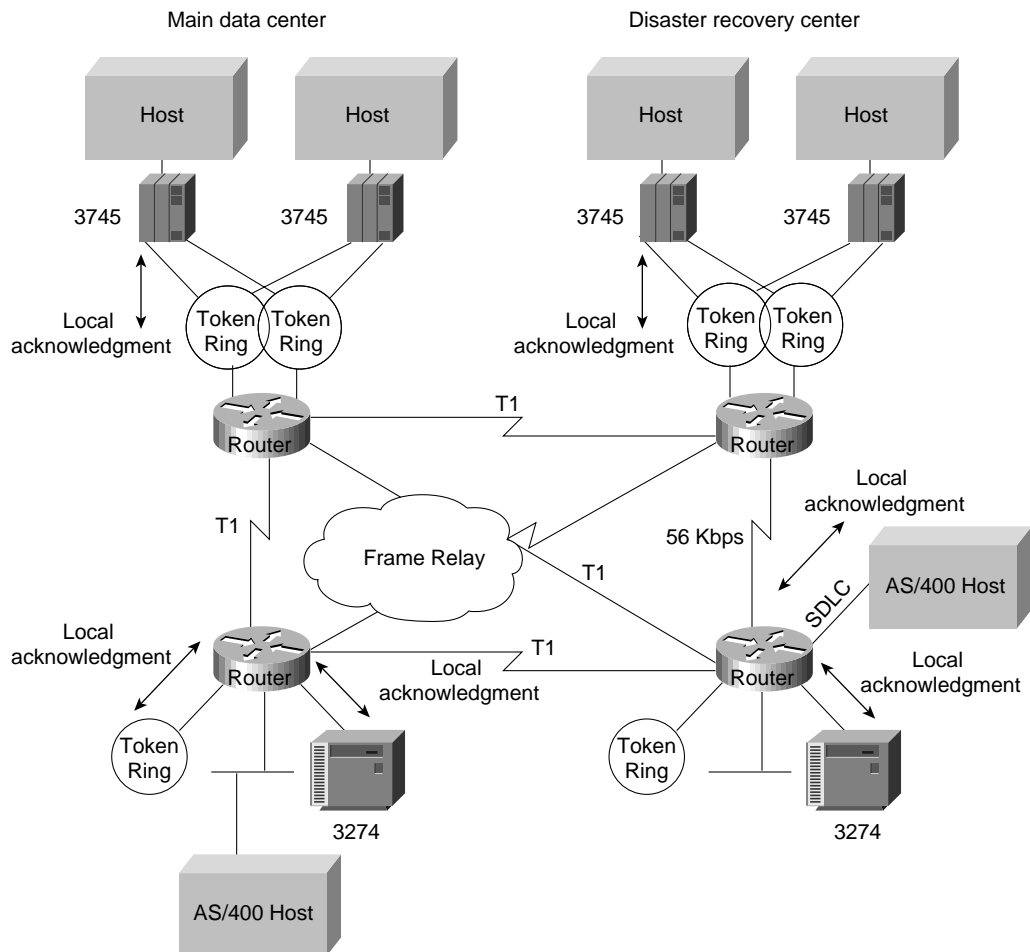
Implement LAN (both Token Ring and Ethernet) internetworks in selected locations along with alternative WAN technologies, such as Frame Relay, as shown in Figure 5-16. Connect LAN-based and remote SDLC devices to host FEP Token Ring via SDLLC, RSRB, and translational bridging, and to host FEP SDLC via reverse SDLLC (SDLC side primary). SNA session integrity is maintained through local termination of both LLC2 and SDLC traffic. These solutions provide needed support of LAN-based applications and improve availability and up time for SNA network devices.

SDLLC Implementation Checklist

Before implementing an SDLLC-based internetwork, make sure you are familiar with the information that deals with SDLC. Depending on your implementation, you might need to review the “SDLLC Configuration” and “SDLLC Implementation Scenarios” sections earlier in this chapter. In general, the following guidelines help you create a working, manageable network:

- Use a phased approach to implement your router network.
- Establish a test environment to initially bring up the routers.
- Plan a gradual cutover of end devices into the production environment.
- During the cutover period, observe the router’s behavior by using **show** commands.

Figure 5-16 Implementing alternative LAN-to-WAN technologies for an integrated solution.



Strive to create a network that has predictable behavior early in the development cycle. This strategy can prevent problems as more devices are brought online. The following is a specific SDLLC implementation checklist that you can use to identify technologies, implementations, and possible solutions for your internetwork:

Step 1 Evaluate the customer requirements for SDLLC support:

- Identify all host-attached controllers. Examples include 37x5, 3172, and 3174 devices. The host sites might be referred to as local, core, backbone, or data center sites.
- How are the host site controllers connected to the network?
- Is Token Ring already in place? Ethernet?
- What are the link speeds for remote end systems?
- What are the current line utilization measurements? Network Performance Monitor, which makes historical data available, is typically installed in the larger SNA shops.
- What interface type is required (for example: V.24 (EIA/TIA-232, formerly RS-232), V.35, or X.21)?
- What modems, data service units, channel service units, modem-sharing devices or line-sharing devices will be used?

- Is Link Problem Determination Aid (LPDA) support required? LPDA is a feature of IBM modems and data service units that reports line quality and statistics to NetView. LPDA Version 1 is not compatible with STUN and SDLLC; LAPD Version 2 may be compatible with STUN.
- What remote end-system types are expected? Examples include 3174, 3274, and AS/400.
- Will there be end-system emulation?
- What is the current transaction response time? Is subsecond response required?
- How many PUs are there? (This information will be important for router utilization sizing.)
- How many LUs are there? (Many busy LUs attached to a PU increases link utilization.)

Step 2 Determine current configuration. Important information includes the following:

- NCP system generation for 3745, 3725, and 3720 devices; in particular, note the LINE, PU, and LU definition statements.
- Local controller current worksheets for 3174 and 3172 devices.
- Remote controller configuration worksheets for 3x74 and 5x94 devices.
- OS/2 Communication Manager configuration files.
- Network topology diagram.

Step 3 Determine the SDLLC features that best suit your requirements.

Confirm that devices to be attached are SDLC PU type 2 devices. Select specific feature requirements, such as local acknowledgment and virtual multidrop.


Step 4 Determine what host conversion changes are required:

- Switched major node definitions for VTAM.
- FEP/NCP changes for Token Ring addition and SDLC link reduction.

QLLC Conversion

QLLC is a data-link protocol defined by IBM that allows SNA data to be transported across X.25 networks. With QLLC, each SDLC physical link is replaced by a single virtual circuit. Figure 5-17 illustrates a typical QLLC topology. In this topology, both ends of the connection over the X.25 network must be configured for QLLC.

Figure 5-17 Typical QLLC topology.



QLLC conversion is a feature of Cisco IOS Software Release 10.2 that causes the router to perform all of the translation required to send SNA data over an X.25 network so that IBM devices that are connected to a router do *not* have to be configured for QLLC.


QLLC conversion allows a device (typically a FEP or an AS/400) that is attached either directly to the router or through a Token Ring to communicate with a device (typically a 3174 terminal controller) that is attached to an X.25 network, as shown in Figure 5-18. In this example, only the terminal controller must be configured for QLLC and must have an X.25 interface.

Figure 5-18 Simple topology for QLLC conversion.




In some topologies, one router interface uses SDLC to communicate with the terminal controller, and another router interface uses X.25 to communicate with the remote device over the X.25 network. In Figure 5-19, the router, configured for QLLC conversion, handles SNA traffic between the terminal controller and the FEP.

Figure 5-19 Topology that uses SDLC and QLLC conversion.



QLLC conversion also supports multiple SDLC connections coming through an MSD, as shown in Figure 5-20.

Figure 5-20 QLLC conversion supports multidrop SDLC topology.



The router that is configured for QLLC conversion does not need to be on the same Token Ring as the FEP. In Figure 5-21, Router A is configured for QLLC and remote source-route bridging (RSRB), and Router B is configured for RSRB only. RSRB allows the FEP to connect to Router A. If a Token Ring connected to the X.25 network communicates with the Token Ring attached to the FEP by a protocol other than SRB, RSRB can provide connectivity.

Figure 5-21 Complex QLLC conversion topology.





Figure 5-21 shows an example using local acknowledgment, which causes the LLC2 session from the Token Ring-attached SNA device (the FEP) to be terminated at the adjacent router (Router B). A TCP session transports the data from Router B to the router attached to the X.25 network (Router A). Only Router A is configured for QLLC conversion. When enabled, local acknowledgment applies to all QLLC connections. The **source-bridge qlc-local-ack** global configuration command enables local acknowledgment and applies to all QLLC connections.

In pass-through mode, local acknowledgment is not used. Instead, the LLC2 session from the Token Ring-attached SNA device (the FEP) is terminated at the router connected to the X.25 network (Router A).

QLLC conversion also supports a configuration in which SNA end stations (3174 or equivalent) that are connected to a Token Ring reach the FEP through an X.25 connection, as shown in Figure 5-22. In this case, IBM Network Packet Switching Interface (NPSI) software is installed on the FEP.

Figure 5-22 QLLC conversion supports SNA end-station connections over Token Ring and X.25 networks.



Summary

This chapter addresses some of the special requirements for implementing routing technology within IBM System Network Architecture (SNA) environments. It describes the three techniques designed to enable internetworking capabilities for SNA-based network architectures, as follows:

- SDLC via STUN
- SDLLC Implementation
- QLLC Conversion