

Bridging and IBM Networking Overview

The *Bridging and IBM Networking Configuration Guide* discusses the following software components:

- Transparent and Source-Route Transparent Bridging
- Source-Route Bridging (SRB)
- Remote Source-Route Bridging (RSRB)
- Data-Link Switching Plus (DLSw+)
- Serial Tunnel and Block Serial Tunnel
- SDLC and LLC2 Parameters
- IBM Network Media Translation
- Downstream Physical Unit and SNA Service Point
- SNA Frame Relay Access Support
- Advanced Peer-to-Peer Networking
- Native Client Interface Architecture (NCIA)
- IBM Channel Interface Processor

This overview chapter gives a high-level description of each technology. For configuration information, refer to the appropriate chapter in this publication.

Note In Cisco IOS Release 11.3, all commands supported on the Cisco7500 series are also supported on the Cisco 7000 series.

Transparent and Source-Route Transparent Bridging

Cisco IOS software supports transparent bridging for Ethernet, Fiber Distributed Data Interface (FDDI), and serial media, and supports source-route transparent (SRT) bridging for Token Ring media. In addition, Cisco supports all the mandatory Management Information Base (MIB) variables specified for transparent bridging in RFC 1286.

Transparent Bridging Features

Cisco's transparent bridging software implementation has the following features:

- Complies with the IEEE 802.1D standard.
- Provides the ability to logically segment a transparently bridged network into virtual local-area networks (LANs).
- Provides two Spanning-Tree Protocols—an older bridge protocol data unit (BPDU) format that is compatible with Digital and other LAN bridges for backward compatibility and the IEEE standard bridge protocol data unit (BPDU) format. In addition to features standard with these Spanning-Tree Protocols, Cisco's proprietary software provides for multiple domains for spanning trees. The spanning-tree parameters are configurable.
- Allows frame filtering based on MAC address, protocol type, or the vendor code. Additionally, the bridging software can be configured to selectively filter local area transport (LAT) multicast service announcements.
- Provides deterministic load distribution while maintaining a loop-free spanning tree.
- Provides the ability to bridge over Asynchronous Transfer Mode (ATM), dial-on-demand routing (DDR), Fiber Distributed Data Interface (FDDI), Frame Relay, multiprotocol Link Access Procedure, Balanced (LAPB), Switched Multimegabit Data Service (SMDS), and X.25 networks.
- Provides concurrent routing and bridging, which is the ability to bridge a given protocol on some interfaces in a router and concurrently route that protocol on other interfaces in the same router.
- Provides integrated routing and bridging, which is the ability to route a given protocol between routed interfaces and bridge groups, or to route a given protocol between bridge groups.
- Provides fast-switched transparent bridging for Frame Relay encapsulated serial and High-Speed Serial Interface (HSSI) interfaces, according to the format specified in RFC 1490.
- Provides fast-switched transparent bridging for the ATM interface on the Cisco 7000, according to the format specified in RFC 1483.
- Provides for compression of LAT frames to reduce LAT traffic through the network.
- Provides both bridging and routing of virtual LANS (VLANs).

Cisco access servers and routers can be configured to serve as both multiprotocol routers and Media Access Control (MAC)-level bridges, bridging any traffic that cannot otherwise be routed. For example, a router routing the Internet Protocol (IP) can also bridge Digital's LAT protocol or NetBIOS traffic.

Cisco routers also support remote bridging over synchronous serial lines. As with frames received on all other media types, dynamic learning and configurable filtering applies to frames received on serial lines.


Transit bridging of Ethernet frames across FDDI media is also supported. The term *transit* refers to the fact that the source or destination of the frame cannot be on the FDDI media itself. This allows FDDI to act as a highly efficient backbone for the interconnection of many bridged networks. The configuration of FDDI transit bridging is identical to the configuration of transparent bridging on all other media types.

Integrated Routing and Bridging

While concurrent routing and bridging makes it possible to both route and bridge a specific protocol on separate interfaces within a router, the protocol is not switched between bridged and routed interfaces. Routed traffic is confined to the routed interfaces; bridged traffic is confined to bridged interfaces. A specified protocol may be either routed or bridged on a given interface, but not both.

Integrated routing and bridging makes it possible to route a specific protocol between routed interfaces and bridge groups, or route a specific protocol between bridge groups. Local, or unroutable, traffic can be bridged among the bridged interfaces in the same bridge group, while routable traffic can be routed to other routed interfaces or bridge groups. Figure 2 illustrates how integrated routing and bridging in a router interconnects a bridged network with a routed network.

Figure 2 IRB Interconnects a Bridged Network with a Routed Network



You can configure the Cisco IOS software to route a specific protocol between routed interfaces and bridge groups, or to route a specific protocol between bridge groups. Specifically, local or unroutable traffic is bridged among the bridged interfaces in the same bridge group, while routable traffic is routed to other routed interfaces or bridge groups. Using integrated routing and bridging, you can:


- Switch packets from a bridged interface to a routed interface
- Switch packets from a routed interface to a bridged interface
- Switch packets within the same bridge group

Bridge-Group Virtual Interface (BVI)

Because bridging operates in the data link layer and routing operates in the network layer, they follow different protocol configuration models. Taking the basic IP model as an example, all bridged interfaces would belong to the same network, while each routed interface represents a distinct network.

In integrated routing and bridging, the BVI is introduced to avoid confusing the protocol configuration model when a specific protocol is both bridged and routed in a bridge group. Figure 3 illustrates the BVI as a user-configured virtual interface residing within a router.

Figure 3 The Bridge-Group Virtual Interface in the Router



The BVI is a normal routed interface that does not support bridging, but does represent its corresponding bridge group to the routed interface. It has all the network layer attributes (such as a network layer address and filters) that apply to the corresponding bridge group. The interface number assigned to this virtual interface corresponds to the bridge group that this virtual interface represents. This number is the link between the virtual interface and the bridge group.

When you enable routing for a given protocol on the BVI, packets coming from a routed interface, but destined for a host in a bridged domain, are routed to the BVI and are forwarded to the corresponding bridged interface. All traffic routed to the BVI is forwarded to the corresponding bridge group as bridged traffic. All routable traffic received on a bridged interface is routed to other routed interfaces as if it is coming directly from the BVI.

Note Because the BVI is a virtual routed interface, it has all the network layer attributes, such as a network address and the ability to perform filtering.

To be able to receive routable packets arriving on a bridged interface, but destined for a routed interface, or to receive routed packets, the BVI must also have the appropriate addresses. MAC addresses and network addresses are assigned to the BVI as follows:

- The BVI “borrows” the MAC address of one of the bridged interfaces in the bridge group associated with the BVI.
- To route and bridge a given protocol in the same bridge group, you must configure the network-layer attributes of the protocol on the BVI. No protocol attributes should be configured on the bridged interfaces, and no bridging attributes can be configured on the BVI.

Because there can be only one BVI representing a bridge group, and the bridge group can be made up of different media types configured for several different encapsulation methods, you may need to configure the BVI with the particular encapsulation methods required to switch packets correctly.

For example, the BVI has default data link and network layer encapsulations that are the same as those available on Ethernet interfaces, but you can configure the BVI with encapsulations that are not supported on an Ethernet interface. In some cases, the default encapsulations provide appropriate results; in other cases they do not. For example, with default encapsulation, ARPA packets from the BVI are translated to SNAP when bridging IP to a Token Ring or FDDI bridged interface. But for

IPX, Novell-ether encapsulation from the BVI is translated to raw-token or raw-FDDI when bridging IPX to a Token Ring or FDDI bridged interface. Because this behavior is usually not what you want, you must configure IPX SNAP or SAP encapsulation on the BVI.

Other Considerations

- Integrated routing and bridging is not supported on cBus platforms (AGS+ and Cisco 7000 series).
- Integrated routing and bridging is supported for transparent bridging, but not for source-route bridging.
- Integrated routing and bridging is supported on all media interfaces except X.25 and ISDN bridged interfaces.
- Integrated routing and bridging supports three protocols: IP, IPX, and AppleTalk in both fast-switching and process-switching modes.
- Integrated routing and bridging and concurrent routing and bridging cannot operate at the same time.

Source-Route Transparent (SRT) Bridging Features

Cisco routers support transparent bridging on Token Ring interfaces that support SRT bridging. Both transparent and SRT bridging are supported on all Token Ring interface cards that can be configured for either 4- or 16-MB transmission speeds.

As with other media, all the features that use **bridge-group** commands can be used on Token Ring interfaces. As with other interface types, the bridge group can be configured to run either the IEEE or Digital Spanning-Tree Protocols. When configured for the IEEE Spanning-Tree Protocol, the bridge cooperates with other SRT bridges and constructs a loop-free topology across the entire extended LAN.

You can also run the Digital Spanning-Tree Protocol over Token Ring. Use it when you have other non-IEEE bridges on other media and you do not have any SRT bridges on Token Ring. In this configuration, all the Token Ring transparent bridges must be Cisco routers. This is because the Digital Spanning-Tree Protocol has not been standardized on Token Ring.

As specified by the SRT bridging specification, only packets without a routing information field (RIF) (RII = 0 in the SA field) are transparently bridged. Packets with a RIF (RII = 1) are passed to the source-route bridging module for handling. An SRT-capable Token Ring interface can have both source-route bridging and transparent bridging enabled at the same time. However, with SRT bridging, frames that did not have a RIF when they were produced by their generating host never gain a RIF, and frames that did have a RIF when they were produced never lose that RIF.

Note Because bridges running only SRT bridging never add or remove RIFs from frames, they do not integrate source-route bridging with transparent bridging. A host connected to a source-route bridge that expects RIFs can *never* communicate with a device across a bridge that does not understand RIFs. SRT bridging cannot tie-in existing source-route bridges to a transparent bridged network. To tie-in existing bridges, you must use source-route translational bridging (SR/TLB) instead. SR/TLB is described in the chapter “Configuring Source-Route Bridging.”

Bridging between Token Ring and other media requires certain packet transformations. In all cases, the MAC addresses are bit-swapped because the bit ordering on Token Ring is different from that on other media. In addition, Token Ring supports one packet format, logical link control (LLC), while Ethernet supports two formats (LLC and Ethernet).

The transformation of LLC frames between media is simple. A length field is either created (when the frame is transmitted to non-Token Ring) or removed (when the frame is transmitted to Token Ring). When an Ethernet format frame is transmitted to Token Ring, the frame is translated into an LLC-1 Subnetwork Access Protocol (SNAP) packet. The destination service access point (DSAP) value is AA, the source service access point (SSAP) value is AA, with the organizational unique identifier (OUI) value is 0000F8. Likewise, when a packet in LLC-1 format is bridged onto Ethernet media, the packet is translated into Ethernet format.



Caution Bridging between dissimilar media presents several problems that can prevent communication from occurring. These problems include bit order translation (or using MAC addresses as data), maximum transmission unit (MTU) differences, frame status differences, and multicast address usage. Some or all these problems might be present in a multimedia bridged LAN. Because of differences in the way end nodes implement Token Ring, these problems are most prevalent when bridging between Token Ring and Ethernet or between Ethernet and FDDI LANs.

Problems currently occur with the following protocols when bridged between Token Ring and other media: Novell IPX, DECnet Phase IV, AppleTalk, Banyan VINES, XNS, and IP. Further, problems can occur with the Novell IPX and XNS protocols when bridged between FDDI and other media. We recommend that these protocols be routed whenever possible.


Source-Route Bridging (SRB)

Our bridging software includes SRB capability. A source-route bridge connects multiple physical Token Rings into one logical network segment. If the network segment bridges only Token Ring media to provide connectivity, the technology is termed source-route bridging. If the network bridges Token Ring and non-Token Ring media is introduced into the bridged network segment, the technology is termed remote source-route bridging (RSRB).

Source-route bridging enables our routers to simultaneously act as a Level 3 router and a Level 2 source-route bridge. Thus, protocols such as Novell's Internetwork Packet Exchange (IPX) or Xerox Network Systems (XNS) can be routed on Token Rings, while other protocols such as Systems Network Architecture (SNA) or NetBIOS are source-route bridged.

Source-route bridging technology is a combination of bridging and routing functions. A source-route bridge can make routing decisions based on the contents of the MAC frame header. Keeping the routing function at the MAC, or Level 2, layer allows the higher-layer protocols to execute their tasks more efficiently and allows the LAN to be expanded without the knowledge of the higher-layer protocols.

As designed by IBM and the IEEE 802.5 committee, source-route bridges connect extended Token Ring LANs. A source-route bridge uses the RIF in the IEEE 802.5 MAC header of a datagram (see Figure 4) to determine which rings or Token Ring network segments the packet must transit. The source station inserts the RIF into the MAC header immediately following the source address field in every frame, giving this style of bridging its name. The destination station reverses the routing field to reach the originating station.

Figure 4 IEEE 802.5 Token Ring Frame Format

The information in a RIF is derived from explorer packets generated by the source node. These explorer packets traverse the entire source-route bridge network, gathering information on the possible paths the source node might use to send packets to the destination.

Transparent spanning-tree bridging requires time to recompute a topology in the event of a failure; source-route bridging, which maintains multiple paths, allows fast selection of alternate routes in the event of failure. Most importantly, source-route bridging allows the end stations to determine the routes the frames take.

Source-Route Bridging Features

Cisco's source-route bridging implementation has the following features:

- Provides configurable fast-switching software for source-route bridging.
- Provides for a local source-route bridge that connects two or more Token Ring networks.
- Provides *ring groups* to configure a source-route bridge with more than two network interfaces. A ring group is a collection of Token Ring interfaces in one or more routers that are collectively treated as a *virtual ring*.
- Provides two types of explorer packets to collect RIF information—an *all-routes* explorer packet, which follows all possible paths to a destination ring, and a *spanning-tree* explorer packet, which follows a statically configured limited route (spanning tree) when looking for paths.
- Provides a dynamically determined RIF cache based on the protocol. The software also allows you to add entries manually to the RIF cache.
- Provides for filtering by MAC address, link service access point (LSAP) header, and protocol type.
- Provides for filtering of NetBIOS frames either by station name or by a packet byte offset.
- Provides for translation into transparently bridged frames to allow source-route stations to communicate with nonsource-route stations (typically on Ethernet).
- Provides support for the SRB Management Information Base (MIB) variables as described in the IETF draft "Bridge MIB" document, "Definition of Managed Objects for Bridges," by E. Decker, P. Langille, A. Rijsinghani, and K. McCloghrie, June 1991. Only the SRB component of the Bridge MIB is supported.
- Provides support for the Token Ring MIB variables as described in RFC 1231, "IEEE 802.5 Token Ring MIB," by K. McCloghrie, R. Fox, and E. Decker, May 1991. Cisco implements the mandatory tables (Interface Table and Statistics Table) but not the optional table (Timer Table) of the Token Ring MIB. The Token Ring MIB has been implemented for the 4/16-Mb Token Ring cards that can be user adjusted for either 4- or 16-Mb transmission speeds (CSC-1R, CSC-2R, CSC-R16M, or CSC-C2CTR).
- Source-route bridging is supported over FDDI on Cisco 7200 series routers.
- Particle-based switching is supported (over FDDI and Token Ring) by default on Cisco 7200 series routers.

Remote Source-Route Bridging (RSRB)


In contrast to SRB, which involves bridging between Token Ring media only, RSRB is Cisco's first technique for connecting Token Ring networks over *non-Token Ring* network segments. (DLSw+ is Cisco's strategic method for providing this function.)

Cisco's RSRB software implementation includes the following features:

- Provides for multiple routers separated by non-Token Ring segments. Three options are available:
 - Encapsulate the Token Ring traffic inside IP datagrams passed over a Transmission Control Protocol (TCP) connection between two routers.
 - Use Fast-Sequenced Transport (FST) to transport RSRB packets to their peers without TCP or User Datagram Protocol (UDP) header or processor overhead.
 - Use data link layer encapsulations over a single serial line, Ethernet, Token Ring, or FDDI ring connected between two routers attached to Token Ring networks.
- Provides for configurable limits to the size of the TCP backup queue.

Figure 5 shows an RSRB topology. The virtual ring can extend across any non-Token Ring media supported by RSRB, such as serial, Ethernet, FDDI, and WANs. The type of media you select determines the way you set up RSRB.

Figure 5 Remote Source-Route Bridged Topology



Note If you bridge across Token Ring media, it is recommended that you do not use RSRB. Use SRB instead. Refer to the chapter "Configuring Source-Route Bridging."

Data-Link Switching Plus (DLSw+)

DLSw+ is a method of transporting SNA and NetBIOS. It complies with the DLSw standard documented in RFC 1795 as well as the DLSw Version 2 standard. DLSw+ is an alternative to RSRB that addresses several inherent problems that exist in RSRB, such as:

- SRB hop-count limits (SRB's limit is seven)
- Broadcast traffic (including SRB explorer frames or NetBIOS name queries)
- Unnecessary traffic (acknowledgments and keepalives)
- Data-link control timeouts

DLSw Standard

The DLSw standard, documented in RFC 1795, defines the switch-to-switch protocol between DLSw routers. The standard also defines a mechanism to terminate data-link control connections locally and multiplex the traffic from the data-link control connections to a TCP connection. The standard always calls for the transport protocol to be TCP and always requires that data-link control connections be locally terminated (the equivalent of our local acknowledgment option). The standard also requires that the SRB RIF be terminated at the DLSw router. The standard describes a means for prioritization and flow control and defines error recovery procedures that ensure data-link control connections are appropriately disabled if any part of their associated circuits breaks.

The DLSw standard does not specify when to establish TCP connections. The capabilities exchange allows compliance to the standard, but at different levels of support. The standard does not specify how to cache learned information about MAC addresses, RIFs, or NetBIOS names. It also does not describe how to track both capable or preferred DLSw partners for either backup or load-balancing purposes. The standard does not provide the specifics of media conversion, but leaves the details up to the implementation. It does not define how to map switch congestion to the flow control for data-link control. Finally, the MIB is documented under a separate RFC.

DLSw Version 2 Standard

In the Version 1 standard, the network design was required to be fully meshed so that all peers were connected to every other peer. This design created unnecessary broadcast traffic because an explorer was sent to every peer for every broadcast.

DLSw Version 2 addresses this problem and improves scalability in DLSw networks with the following enhancements:

- UDP/IP Multicast Service
- Enhanced Peer-on-Demand Routing Feature
- Expedited TCP Connection

UDP/IP Multicast Service

IP multicast service reduces the amount of network overhead because it avoids the need to maintain TCP Switch-to-Switch Protocol (SSP) connections between two DLSw peers when no circuits are available and it ensures that each broadcast results in only a single explorer over every link. Furthermore, multicast service avoids duplication and excessive bandwidth of broadcast traffic because it replicates and propagates messages only as necessary to its multicast members. This functionality is a subset of Cisco's existing DLSw+ border peering feature. In a border peer network, address resolution packets are sent via the border peer network and can take advantage of border peer caching to minimize bandwidth usage. DLSw Version 2 is for customers who run a multicast IP network and do not need the advantages of border peering.

With UDP Unicast, explorer frames and unnumbered information frames are sent via UDP rather than TCP. This feature eliminates retransmission of explorers and unnumbered information frames that could occur during congestion. Cisco's DLSw+ introduced the UDP Unicast feature prior to the release of DLSw Version 2 in Cisco IOS Release 11.2(6)F. One difference between the two enhancements is that the Release 11.2(6)F UDP Unicast feature requires that a TCP connection exist before packets are sent via UDP. Because the TCP session is up and capabilities have been exchanged, the peers know exclusive reachability information which will permit them to further reduce the explorer load on the network. DLSw Version 2, on the other hand, sends UDP/IP multicast and unicast before the TCP connection exists. Although DLSw Version 2 employs IP multicast

service when address resolution packets (CANREACH_EX, NETBIOS_NQ_ex, NETBIOS_ANQ, and DATAFRAME) are sent to multiple destinations; the response frames (ICANREACH_ex and NAME_RECOGNIZED_ex), are sent via UDP unicast.

Enhanced Peer-on-Demand Routing Feature

DLSw Version 2 now establishes TCP connections only when necessary and the TCP connections are brought down when there are no circuits to a DLSw peer for a specified amount of time. This method, known as peer-on-demand routing, was recently introduced in DLSw Version 2, but has been implemented in Cisco DLSw+ border peer technology since Cisco IOS Release 10.3.

Expedited TCP Connection

DLSw Version 2 more efficiently establishes TCP connections. Previously, DLSw created two unidirectional TCP connections and then disconnected one after the capabilities exchange took place. With DLSw Version 2, a single bidirectional TCP connection is established if the associated peer is known to support DLSw Version 2.

DLSw+ Features

DLSw+ is Cisco's version of DLSw and it supports several additional features and enhancements. DLSw+ is fully compatible with any vendors RFC 1795 implementation and the following features are available when both peers are using DLSw+:

- Peer groups and border peers.
- Backup peers.
- Promiscuous and on-demand peers.
- Scalability enhancements through explorer firewalls and location learning.
- NetBIOS dial-on-demand routing feature support.
- UDP Unicast support.
- Border Peer caching and load balancing.
- Support for LLC1 circuits
- Support for Multiple Bridge Groups
- A choice of transport options, including TCP, FST, and direct encapsulation in High-Level Data Link Control (HDLC) or Frame Relay.
- SNA type of service feature support.
- Local acknowledgment for Ethernet-attached devices and media conversion for SNA PU 2.1 and PU 2.0 devices.
- Conversion between LLC2 to SDLC between PU 4 devices.
- Local or remote media conversion between LANs and either Synchronous Data Link Control (SDLC) Protocol or QLLC.
- SNA View, Blue Maps, and IBM's LAN Network Manager support.
- MIB enhancements that allow DLSw+ plus features to be managed by the CiscoWorks Blue products, SNA Maps, and SNA View. Also, new traps alert network management stations of peer or circuit failures. For more information, refer to the current Cisco IOS release note for the location of the Cisco MIB Web site.

DLSw+ Enhancements

DLSw+ goes beyond the standard to include additional functionality in the following areas:

- Modes of Operation—Ability to determine the capabilities of the participating router and to operate according to those capabilities.
- Improved Scalability—Ability to construct IBM internetworks in a way that reduces the amount of broadcast traffic and therefore enhances their scalability.
- Improved Performance—Ability to offer higher-performance transport options when the line speeds and traffic conditions do not require local acknowledgment.
- Enhanced Availability—Overall availability of an end-to-end connection between a pair of SNA and NetBIOS end systems.

Modes of Operation

DLSw+ operates in two modes:

- Standards compliance mode—DLSw+ can automatically detect (through the DLSw+ capabilities exchange) that the participating router is manufactured by another vendor. DLSw+ then operates in DLSw standard mode.
- Enhanced mode—DLSw+ can automatically detect that the participating router is another DLSw+ router. DLSw+ then operates in enhanced mode, providing the additional features of DLSw+ to the SNA and NetBIOS end systems.

Note DLSw+ does not interoperate with prestandard implementations such as RFC 1434.

Some enhanced DLSw+ features are also available when a Cisco router is operating in standards compliance mode with another vendor's router. In particular, enhancements that are locally controlled options on a router can be accessed even though the remote router does not have DLSw+. These include reachability caching, explorer firewalls and media conversion.

Improved Scalability

One significant factor that limits the size of Token Ring internetworks is the amount of explorer traffic that traverses the WAN. DLSw+ includes the following features to reduce the number of explorers:

- Peer groups—The large Token Ring internetworks that Cisco has helped to build over the last several years have all followed a similar structure. That structure is a hierarchical grouping of routers based upon the usual flow of broadcasts through the network. A cluster of routers in a region or a division of a company is combined into a peer group.
- Border peers—Within a peer group, one or more routers are designated as border peers. When a DLSw+ router receives a test frame or NetBIOS name query, it sends a single explorer frame to its border peer. The border peer takes complete responsibility for forwarding the explorer on behalf of the peer group member. This arrangement eliminates duplicate explorers on the access links and minimizes the processing required in access routers.

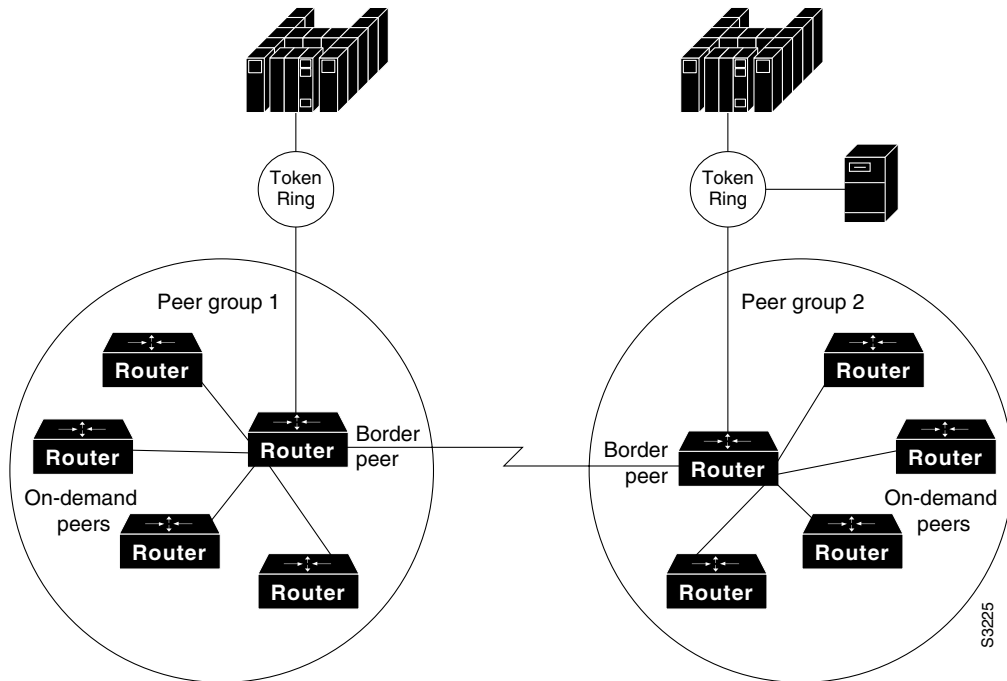
- On-demand peers—On-demand peers greatly reduce the number of peers that must be configured. As Figure 6 shows, you can use on-demand peers to establish an end-to-end circuit even though the DLSw+ routers servicing the end systems have no specific configuration information about the peers. This configuration permits casual, any-to-any connection without the burden of configuring the connection in advance. It also allows any-to-any switching in large internetworks where persistent TCP connections would not be possible.
- Explorer firewalls—An explorer firewall permits only a single explorer for a particular destination MAC address or NetBIOS name to be sent across the WAN. While an explorer is outstanding and awaiting a response from the destination, subsequent explorers for that MAC address or NetBIOS name are merely stored. When the explorer response is received at the originating DLSw+, all explorers receive an immediate local response. This eliminates the start-of-day explorer storm that many networks experience.
- NetBIOS Dial-on-Demand Routing—This feature allows you to transport NetBIOS in a dial-on-demand routing (DDR) environment by filtering NetBIOS Session Alive packets from the WAN. NetBIOS periodically sends Session Alive packets as LLC2 I-frames. These packets do not require a response and are superfluous to the function of proper data flow. Furthermore, these packets keep dial-on-demand interfaces up and this up time causes unwanted per-packet charges in DDR networks. By filtering these NetBIOS Session Alive packets, you reduce traffic on the WAN as well as some costs that are associated with Dial on Demand Routing.
- UDP Unicast Enhancement—The SSP address resolution packets are sent via UDP unicast service rather than TCP. SSP packets include: CANUREACH_EX, NETBIOS_NQ_ex, NETBIOS_ANQ, and DATAFRAME. UDP Unicast enhances the scalability of TCP peer networks because it allows DLSw+ to better control address resolution packets and unnumbered information frames during periods of congestion. Previously, these frames were carried over TCP. TCP retransmits frames that get lost or delayed in transit, and hence aggravate congestion. Because address resolution packets and unnumbered information frames are not sent on a reliable transport on the LAN, sending them reliably over the WAN is unnecessary. By using UDP for these frames, DLSw+ minimizes network congestion.

Note UDP Unicast Enhancement does not affect Fast-Sequenced Transport (FST) or direct peer encapsulations.

- Border Peer Caching—Border Peer Caching increases the scalability of DLSw+ by minimizing broadcast traffic, bandwidth requirements, and the requirement for border peers to replicate explorers. Previously, border peers did not learn reachability information from the broadcasts that they forwarded. Now, border peers learn reachability information based on relay responses and store it in their remote and group caches. After the first search when a peer is found, every subsequent explorer is forwarded to the known destination, rather than forwarded as a broadcast. As a result, the network has fewer broadcasts.
- LLC1 circuits—Support for LLC1 circuits more efficiently transports LLC1 UI traffic across a DLSw+ cloud. With LLC1 circuit support, the LLC1 UI frames are no longer subject to input queuing and are guaranteed to traverse the same path for the duration of the flow. This feature improves transportation of LLC1 UI traffic because there is no longer the chance of having a specifically routed LLC1 UI frame broadcasted to all remote peers. The circuit establishment process has not changed except that the circuit is established as soon as the specifically routed LLC1 UI frame is received and the DLSw+ knows of reachability for the destination MAC address. Furthermore, the connection remains in the CIRCUIT_ESTABLISHED state (rather than proceeding to the CONNECT state) until there is no UI frame flow for a MAC/SAP pair for 10 minutes.

- Multiple Bridge Groups—This feature allows you to assign more than one bridge group for different physical ethernet segments. See the **dlsw bridge-group** command for more information.

Figure 6 Scalability with DLSw+



Improved Performance

DLSw+ improves performance by offering higher-performance transport options in the following areas:

- Transport Connection Type
- SNA Type of Service

Transport Connection Type

The transport connection between DLSw+ routers can vary according to the needs of the network and is not necessarily tied to TCP/IP as the DLSw standard is. We support three different transport protocols between DLSw+ devices:

- TCP for transport of SNA and NetBIOS traffic across WANs where bandwidth is limited and termination of data-link control sessions is required. This transport option is required when DLSw+ is operating in standards compliance mode.
- FST for transport across IP WANs with an arbitrary topology with sufficient bandwidth to accommodate SNA and NetBIOS.
- Direct encapsulation for transport across a point-to-point connection where the benefits of an arbitrary topology are not important.

SNA Type of Service

Performance is further enhanced with the SNA Type of Service feature. Although DLSw+ priority queuing provides prioritization on the output queue of the interface, the priority characteristics are lost once the packet leaves the DLSw+ router and traverses the network. SNA Type of Service sets IP precedence in all DLSw+ frames, either setting all DLSw+ to a high precedence, or building multiple pipes for different priority traffic and setting each one appropriately. Also, when running APPN over DLSw+, APPN Class of Service (COS) characteristics are lost once the packet is delivered to the DLSw+ router for bridging over an IP network. With the new DLSw+ SNA TOS feature, however, SNA TOS works in conjunction with weighted fair queuing to reduce the response time for SNA sessions and, therefore, ensures that DLSw+ gets more bandwidth. DLSw+ traffic is prioritized and APPN COS characteristics are preserved across the network.

With the SNA Type of Service feature, DLSw+ sets the precedence bits in the IP header of outbound DLSw+ packets. When DLSw+ is used in conjunction with APPN, SNA TOS maps APPN COS to IP TOS, and preserves SNA COS across an IP backbone. If priority queuing is not configured on DLSw+, the IP precedence value of “Network” is used for all DLSw+ packets. This default value ensures more bandwidth for DLSw+ packets than for other types of packets.

Table 1 describes the various IP precedence values that map to the TCP ports.

Table 1 IP Precedence Values

TCP Ports	Priority Queue Level	IP Precedence APPN Value	IP Precedence DLSw+ Value
2065	High	Network	Network control
1981	Medium	High	Internetwork control
1982	Normal	Medium	Critical
1983	Low	Low	Flash override

When the **priority** option on the **dls w remote-peer** command is configured, DLSw+ automatically activates four TCP ports to that remote peer (ports 2065, 1981, 1982 and 1983) and assigns traffic to specific ports according to the rules defined in Table 2. Alternately, SAP prioritization, LOCADDR prioritization, or APPN COS can be used to customize how traffic is assigned to these ports.

Table 2 Port Number Priority Values

TCP Ports	DLSw+ Queue Priority	Type of Traffic (default)
2065	High	Circuit administration frames Peer keepalives Capabilities exchange
1981	Medium	None
1982	Normal	Information frames
1983	Low	Broadcast traffic

If the **priority** option is not configured in the **dls w remote peer** command, then all DLSw+ traffic defaults to port 2065 and is assigned IP precedence “network.” The default TOS settings can be changed by using policy routing based on the TCP ports that the DLSw+ router uses. For more

information on policy routing see the *Network Protocols Configuration Guide, Part 1*. For an example configuration of DLSw+ with policy routing see the **dlsw remote peer tcp** command in the *Bridging and IBM Networking Command Reference*.

Please note the following design considerations with the SNA Type of Service feature:

- APPN COS-to-IP TOS mapping occurs only if DLSw+ and APPN are running in the same router and the **priority** keyword is specified on the remote peer statement.
- SNA TOS applies only to TCP or FST encapsulation types.
- When using FST encapsulation, SNA TOS marks all DLSw+ traffic with IP precedence “network.”

Enhanced Availability

DLSw+ offers enhanced availability by caching a table of multiple paths to a given MAC address or NetBIOS name (where a path is either a remote peer or a local port). Furthermore, with the Border Peer Caching feature, border peers build an additional cache (group), which is checked before forwarding explorers for other routers.

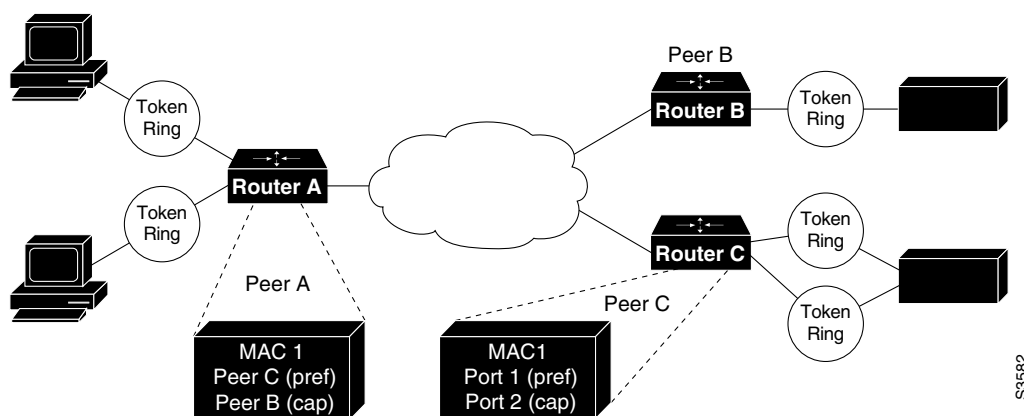
Maintaining multiple paths per destination is especially attractive in SNA networks. A common technique used in the hierarchical SNA environment is assigning the same MAC address to different Token Ring interface couplers (TICs) on the IBM front-end processors (FEPs). DLSw+ ensures that duplicate TIC addresses are found, and, if multiple DLSw+ peers can be used to reach the FEPs, they are cached.

The way that multiple capable peers are handled with DLSw+ can be biased to meet either of the following network needs:

- **Fault tolerance**—To rapidly reconnect if a data-link connection is lost. If load balancing is not enabled, the Cisco IOS software maintains a preferred path and one or more capable paths to each destination. The preferred peer is either the peer that responds first to an explorer frame or the peer with the least cost. The preferred port is always the port over which the first positive response to an explorer was received. If the preferred peer to a given destination is unavailable, the next available capable peer is promoted to the new preferred peer. No additional broadcasts are required, and recovery through an alternate peer is immediate. Maintaining multiple cache entries facilitates a timely reconnection after session outages.
- **Load balancing**—To distribute the network traffic over multiple DLSw+ peers in the network. The Cisco IOS software can be configured to perform load balancing, in which case circuits are established in round-robin fashion using the list of capable routers. When used for load balancing, this technique improves overall SNA performance.

Figure 7 shows a peer table of preferred (pref) and capable (cap) routes.

Figure 7 Enhanced Availability and Performance



In addition to supporting multiple active peers, DLSw+ supports backup peers for all encapsulation types (including direct, FST, and TCP), which are connected only when the primary peer cannot be reached.

DLSw+ Support for Other SNA Features

DLSw+ can be used as a transport for SNA features such as LNM, DSPU, SNA service point, and APPN through a Cisco IOS feature called VDLC.

LNM over DLSw+ allows DLSw+ to be used in Token Ring networks that are managed by IBM's LNM software. Using this feature, LNM can be used to manage Token Ring LANs, control access units, and Token Ring attached devices over a DLSw+ network. All management functions continue to operate as they would in a source-route bridged network or an RSRB network.

DSPU over DLSw+ allows Cisco's DSPU feature to operate in conjunction with DLSw+ in the same router. DLSw+ can be used either upstream (toward the mainframe) or downstream (away from the mainframe) of DSPU. DSPU concentration consolidates the appearance of multiple physical units (PUs) into a single PU appearance to VTAM, minimizing memory and cycles in central site resources (VTAM, NCP, and routers) and speeding network startup.


SNA service point over DLSw+ allows Cisco's SNA service point feature to be used in conjunction with DLSw+ in the same router. Using this feature, SNA service point can be configured in remote routers, and DLSw+ can provide the path for the remote service point PU to communicate with NetView. This allows full management visibility of resources from a NetView 390 console, while concurrently offering the value-added features of DLSw+ in an SNA network.

APPN over DLSw+ allows Cisco's APPN feature to be used in conjunction with DLSw+ in the same router. With this feature, DLSw+ can be used to access an APPN backbone or APPN in the data center. DLSw+ can also be used as a transport for APPN, providing nondisruptive recovery from failures and high-speed intermediate routing. The DLSw+ network can appear as a connection network to the APPN network nodes.

To use DLSw+ as a transport for other Cisco IOS SNA features requires a feature called virtual data-link control. Cisco IOS data-link users (such as LNM, DSPU, SNA service point, and APPN) write to a virtual data-link control interface. DLSw+ then reads from this interface and sends out the traffic. Similarly, DLSw+ can receive traffic destined for one of these DLUS and write it to the virtual data-link control interface, from which the appropriate DLU will read it.

In Figure 8, APPN and DLSw+ use Token Ring and Ethernet, respectively, as “real” data-link controls, and use virtual data-link control to communicate between themselves. When one of the high-layer protocols passes data to the virtual data-link control, the virtual data-link control must pass it to a higher-layer protocol; nothing leaves the virtual data-link control without going through a data-link user.

Figure 8 VDLC Interaction with Higher-Layer Protocols



The higher-layer protocols make no distinction between the virtual data link control and any other data-link control, but they do identify the virtual data link control as a destination. In the example shown in Figure 8, APPN has two ports: a physical port for Token Ring and a logical (virtual) port for the virtual data link control. In the case of the APPN virtual data link control port, when you define the APPN virtual data link control port, you can also specify the MAC address assigned to it. That means data going from DLSw+ to APPN by way of the virtual data link control is directed to the virtual data link control MAC address. The type of higher-layer protocol you use determines how the virtual data link control MAC address is assigned.

Serial Tunnel and Block Serial Tunnel

The Cisco IOS software supports serial tunnel (STUN) and block serial tunnel (BSTUN). Our BSTUN implementation enhances Cisco series 2500, 4000, 4500, 4700, 7200 routers to support devices that use the Binary Synchronous Communication (Bisync) datalink protocol and asynchronous security protocols that include Adplex, ADT Security Systems, Inc., Diebold, and asynchronous generic traffic. BSTUN implementation is also supported on the 4T network interface module (NIM) on the Cisco series 4000 and 4500. Our support of the bisync protocol enables enterprises to transport Bisync traffic and SNA multiprotocol traffic over the same network.

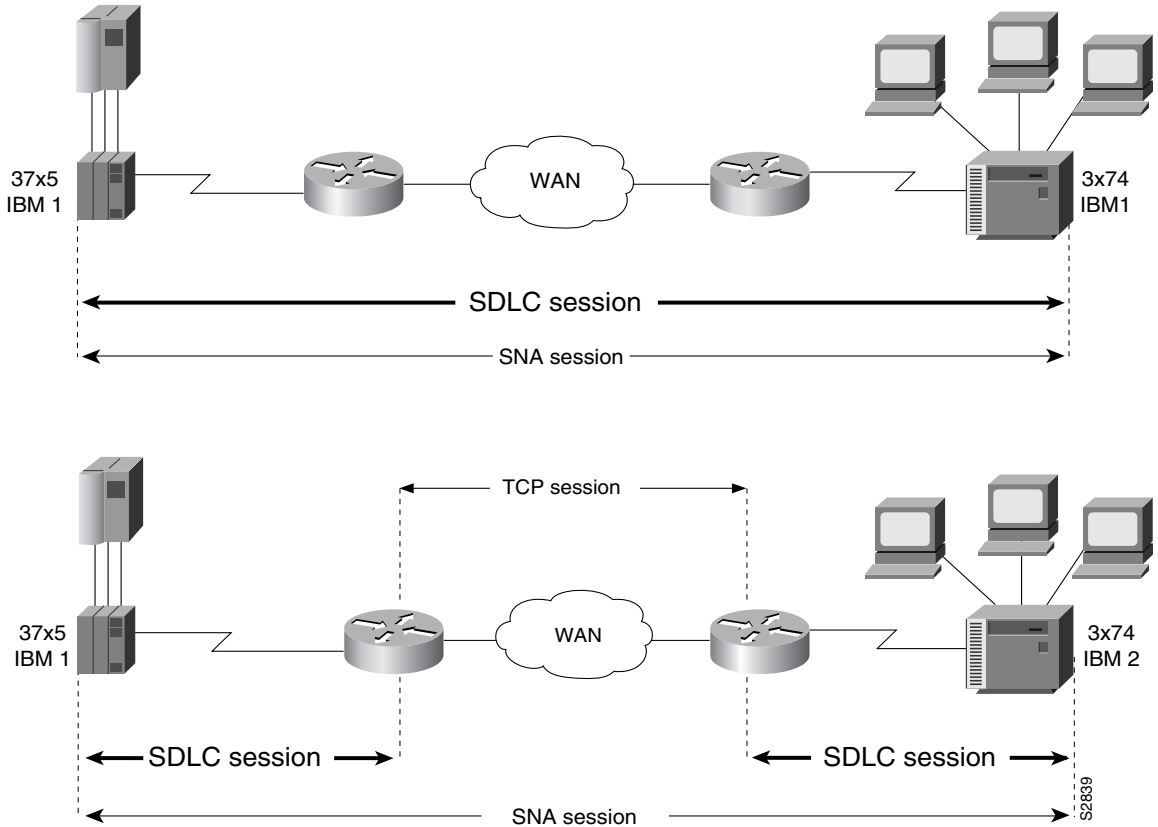
STUN Networks

STUN operates in two modes: passthrough and local acknowledgment. Figure 9 shows the difference between passthrough mode and local acknowledgment mode.

The upper half of Figure 9 shows STUN configured in passthrough mode. In passthrough mode, the routers act as a wire and the SDLC session remains between the end stations. In this mode, STUN provides a straight pass-through of all SDLC traffic, including control frames.

The lower half of Figure 9 shows STUN configured in local acknowledgment mode. In local acknowledgment mode, the routers terminate the SDLC sessions and send only data across the WAN. Control frames no longer travel the WAN backbone networks.

Figure 9 Comparison of STUN in Passthrough Mode and Local Acknowledgment Mode



Note To enable STUN local acknowledgment, you first enable the routers for STUN and configure them to appear on the network as primary or secondary SDLC nodes. TCP/IP encapsulation must be enabled. Our STUN local acknowledgment feature also provides priority queuing for TCP-encapsulated frames.

Serial Tunnel

Our STUN implementation provides the following features:

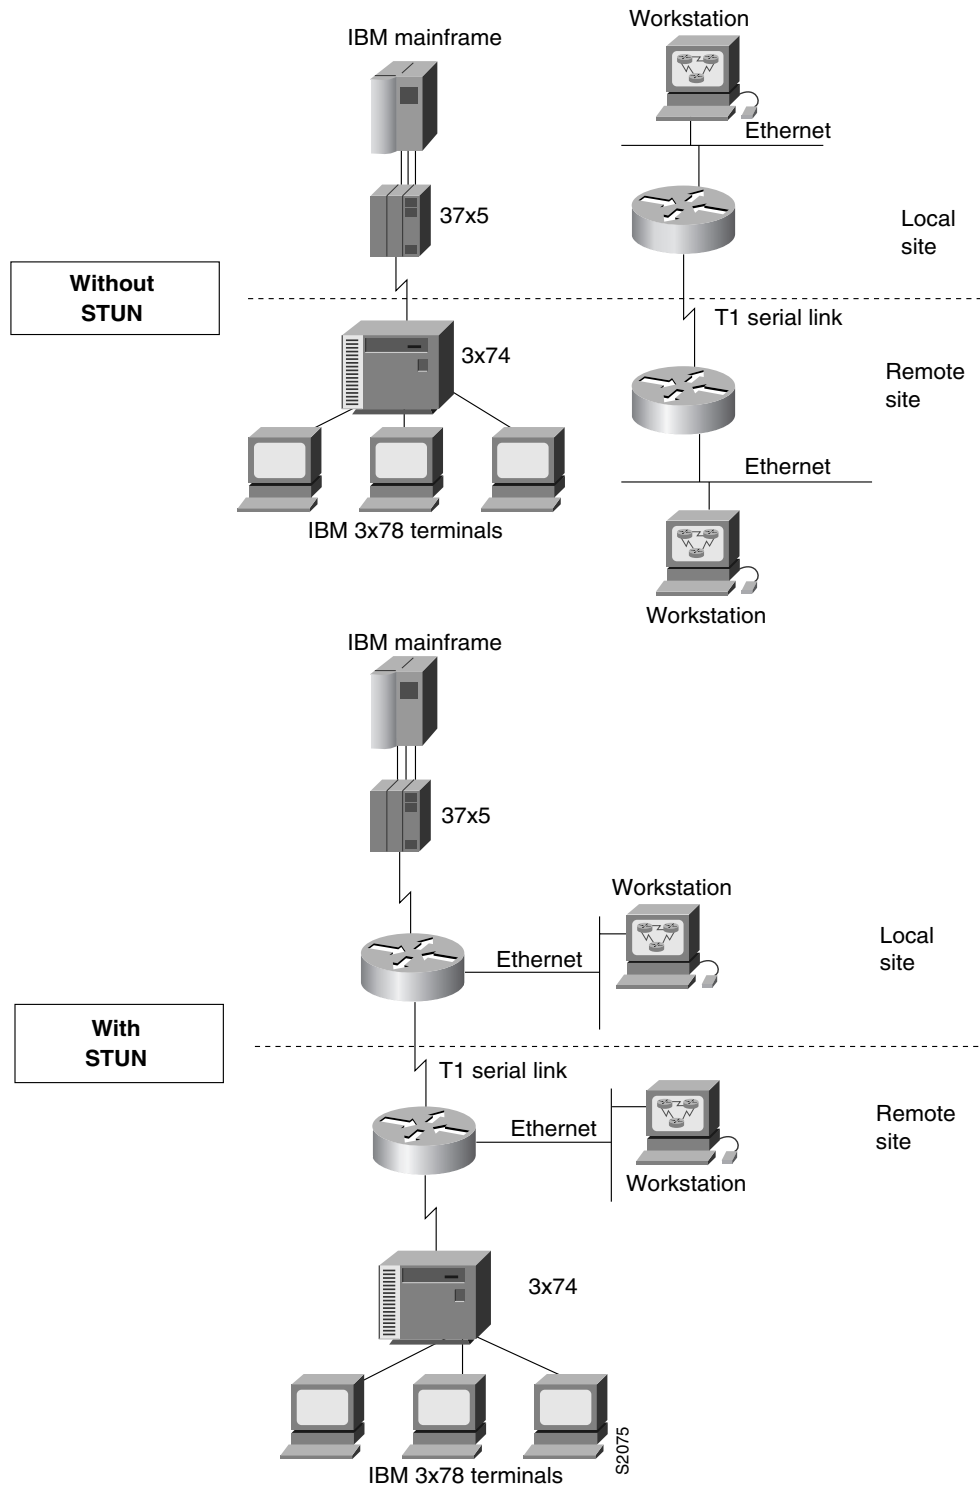
- Encapsulates SDLC frames in either the Transmission Control Protocol/Internet Protocol (TCP/IP) or the HDLC protocol.
- Allows two devices using SDLC- or HDLC-compliant protocols that are normally connected by a direct serial link to be connected through one or more Cisco routers, reducing leased-line costs.

When you replace direct serial links with routers, serial frames can be propagated over arbitrary media and topologies to another router with a STUN link to an appropriate end point. The intervening network is not restricted to STUN traffic, but rather, is multiprotocol. For example, instead of running parallel backbones for DECnet and SNA/SDLC traffic, this traffic now can be integrated into an enterprise backbone network.

- Supports local acknowledgment for direct Frame Relay connectivity between routers, without TCP/IP required.

- Allows networks with IBM mainframes and communications controllers to share data using Cisco routers and existing network links. As an SDLC function, STUN fully supports the IBM SNA and allows IBM SDLC frames to be transmitted across the network media and shared serial links. Figure 10 illustrates a typical network configuration without STUN and the same network configured with STUN.
- Encapsulates SDLC frame traffic packets and routes them over any of the supported network media (serial, FDDI, Ethernet, and Token Ring, X.25, SMDS, and T1/T3) using TCP/IP encapsulation. Because TCP/IP encapsulation is used, you can use any of the Cisco routing protocols to route the packets.
- Copies frames to destinations based on address. STUN in passthrough mode does not modify the frames in any way or participate in SDLC windowing or retransmission; these functions are left to the communicating hosts. However, STUN in local acknowledgment mode does participate in SDLC windowing and retransmission through local termination of the SDLC session.
- Ensures reliable data transmission across serial media having minimal or predictable time delays. With the advent of STUN and WAN backbones, serial links now can be separated by wide geographic distances spanning countries and continents. As a result, these serial links have time delays that are longer than SDLC allows for bidirectional communication between hosts. The STUN local acknowledgment feature addresses the problems of unpredictable time delays, multiple retransmissions, or loss of sessions.
- Allows for configuration of redundant links to provide transport paths in the event part of the network goes down.

Figure 10 IBM Network Configuration without STUN and with STUN



BSTUN Networks

The Bisync feature enables your Cisco 2500, 3600, 4000, 4500, or 7200 series router to support devices that use the Bisync datalink protocol. This protocol enables enterprises to transport Bisync traffic over the same network that supports their SNA and multiprotocol traffic, eliminating the need for separate Bisync facilities.

At the access router, traffic from the attached Bisync device is encapsulated in IP. The Bisync traffic can then be routed across arbitrary media to the host site where another router supporting Bisync will remove the IP encapsulation headers and present the Bisync traffic to the Bisync host or controller over a serial connection. HDLC can be used as an alternative encapsulation method for point-to-point links.

Block Serial Tunnel

Cisco's implementation of BSTUN provides the following features:

- Encapsulates Binary Synchronous Communications (Bisync), Adplex, ADT Security Systems, Inc., Diebold, asynchronous generic, and Monitor Dynamics Inc., traffic for transfer over router links. The tunneling of asynchronous security protocols feature (ASP) enables your Cisco 2500, 3600, 4000, 4500, or 7200 series router to support devices that use the following asynchronous security protocols:
 - adplex
 - adt-poll-select
 - adt-vari-poll
 - diebold
 - async-generic
 - mdi
- Provides a tunnel mechanism for BSTUN over Frame Relay, without using TCP/IP encapsulation.
- Supports legacy Bisync devices and host applications without modification.
- Uses standard synchronous serial interfaces on Cisco 2500 series and the 4T network interface module (NIM) on the Cisco 4000 series and Cisco 4500 series.
- Supports point-to-point, multidrop, and virtual multidrop configurations.

SDLC and LLC2 Parameters

The Logical Link Control, type 2 (LLC2) and SDLC protocols provide data-link-level support for higher-level network protocols and features such as SDLLC and RSRB with local acknowledgment. The features that are affected by LLC2 parameter settings are listed in the next section, “Cisco’s Implementation of LLC2.” The features that require SDLC configuration and use SDLC parameters are listed in the section “Cisco’s Implementation of SDLC” later in this chapter.

LLC2 and SDLC package data in frames. LLC2 and SDLC stations require acknowledgments from receiving stations after a set amount of frames have been sent before sending further data. The tasks described in this chapter modify default settings regarding the control field of the data frames. By modifying the control field parameters, you can determine the number of acknowledgments sent for frames received and the level of polling used to determine available stations. In this manner, you can set the amount of resources used for frame checking and optimize the network load.

SDLC is used as the primary SNA link-layer protocol for WAN links. SDLC defines two types of network nodes: primary and secondary. Primary nodes poll secondary nodes in a predetermined order. Secondary nodes then transmit any outgoing data. When configured as primary and secondary nodes, our routers are established as SDLC stations.

Cisco’s Implementation of LLC2

Cisco’s LLC2 implementation supports the following features:

- Local acknowledgment for RSRB

This feature is used in our implementation of RSRB as described in the chapter “Configuring Source-Route Bridging.”

Because LANs are now connected through RSRB and WAN backbones, the delays that occur are longer than LLC2 allows for bidirectional communication between hosts. Our local acknowledgment feature addresses the problem of delays, retransmissions, and loss of user sessions.

- IBM LAN Network Manager (LNM) support

Routers using 4- or 16-Mbps Token Ring interfaces configured for SRB support LNM and provide all IBM Bridge Program functions. With LNM, a router appears as an IBM source-route bridge, and can manage or monitor any connected Token Ring interface.

LNM support is described in the chapter “Configuring Source-Route Bridging.”

- SDLLC media translation

The SDLLC feature provides media translation between the serial lines running SDLC and Token Rings running LLC2. SDLLC consolidates the IBM SNA networks running SDLC into a LAN-based, multiprotocol, multimedia backbone network.

SDLLC is described in the chapter “Configuring IBM Network Media Translation.”

- ISO Connection-Mode Network Service (CMNS)

Our CMNS implementation runs X.25 packets over LLC2 so that X.25 can be extended to Ethernet, FDDI, and Token Ring media.

Cisco's Implementation of SDLC

Cisco's SDLC implementation supports the following features:

- Frame Relay access support

With our Frame Relay access support feature, a router functions as a Frame Relay Access Device (FRAD) for SDLC, Token Ring, and Ethernet-attached devices over a Frame Relay Boundary Network Node (BNN) link.

Frame Relay access support is described in the chapter "Configuring SNA Frame Relay Access Support."

- SDLLC media translation

Our SDLLC feature provides media translation between the serial lines running SDLC and Token Rings running LLC2. SDLLC consolidates the IBM SNA networks running SDLC into a LAN-based, multiprotocol, multimedia backbone network.

SDLLC is described in the chapter "Configuring IBM Network Media Translation."

- SDLC local acknowledgment

SDLC local acknowledgment is used with SDLC serial tunnel (STUN). The Transmission Control Protocol/Internet Protocol (TCP/IP) must be enabled. With local acknowledgment, STUN SDLC connections can be terminated locally at the router, eliminating the need for acknowledgments to be sent across a WAN.

SDLC local acknowledgment is described in the section "Establish the Frame Encapsulation Method" in the "Configuring STUN and BSTUN" chapter.

IBM Network Media Translation

The Cisco IOS software includes the following media translation features that enable network communications across heterogeneous media:

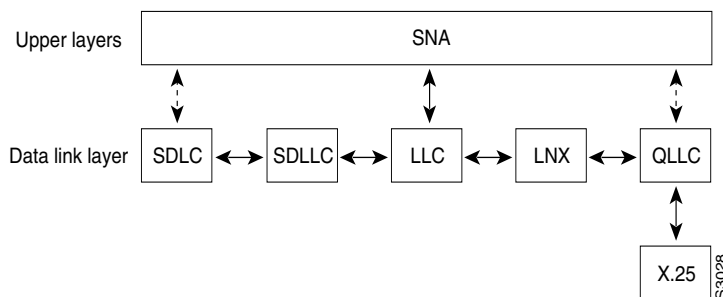
- SDLLC media translation enables a device on a Token Ring to communicate with a device on a serial link.
- QLLC conversion enables an IBM device to communicate with an X.25 network without having to install the X.25 software on local IBM equipment.

SDLLC is a Cisco Systems proprietary software feature that enables a device on a Token Ring to communicate with a device on a serial link by translating between LLC2 and SDLC at the link layer.

SNA uses SDLC and LLC2 as link-layer protocols to provide a reliable connection. The translation function between these industry-standard protocols takes place in the proprietary Cisco software.

Figure 11 illustrates how SDLLC provides data-link layer support for SNA communication.

Figure 11 SNA Data Link Layer Support



SDLLC Media Translation Features

The SDLLC software allows a physical unit (PU) 4, PU 2.1, or PU 2 SDLC device as follows:

- SDLLC with direct connection—A 37x5 FEP on a Token Ring and the 3x74 cluster controller connected to a serial line are each connected to an interface on the same router configured with SDLLC.
- SDLLC with RSRB—A 37x5 FEP on a Token Ring and a 3x74 cluster controller connected to a serial line are connected to different routers. Only the device to which the 3x74 is connected is configured with SDLLC. The routers communicate via RSRB using direct encapsulation, RSRB over an FST connection, or RSRB over a TCP connection.
- SDLLC with RSRB and local acknowledgment—A 37x5 FEP on a Token Ring and a 3x74 cluster controller connected to a serial line are connected to different routers. Only the device to which the 3x74 is connected is configured with SDLLC. The routers communicate via RSRB over a TCP connection that has local acknowledgment enabled.

In all these topologies, each IBM end node (the FEP and cluster controller) has no indication that its counterpart is connected to a different medium running a different protocol. The 37x5 FEP responds as if the 3x74 cluster controller were communicating over a Token Ring, whereas the 3x74 responds as though the 37x5 FEP were communicating over a serial line. That is, the SDLLC software makes translation between the two media transparent to the end nodes.

Virtual Token Ring Concept

Central to Cisco’s SDLLC feature is the concept of a virtual Token Ring device residing on a virtual Token Ring. Because the Token Ring device expects the node with which it is communicating also to be on a Token Ring, each SDLLC device on a serial line must be assigned an SDLLC virtual token ring address (SDLLC VTRA). Like real Token Ring addresses, SDLLC VTRAs must be unique across the network.

In addition to the SDLLC VTRA, an SDLLC virtual ring number (SDLLC VRN) must be assigned to each SDLLC device on a serial line. (The SDLLC VRN differs from the virtual ring group numbers that are used to configure RSRB and multipoint bridging.)

As part of its virtual telecommunications access method (VTAM) configuration, the IBM node on the Token Ring has knowledge of the SDLLC VTRA of the serial device with which it communicates. The SDLC VTRA and the SDLLC VRN are a part of the SDLLC configuration for the router’s serial interface. When the Token Ring host sends out explorer packets with the SDLLC VTRA as the destination address in the MAC headers, the router configured with that SDLLC VTRA

intercepts the frame, fills in the SDLLC VRNA and the bridge number in the RIF, then sends the response back to the Token Ring host. A route is then established between the Token Ring host and the router. After the Cisco IOS software performs the appropriate frame conversion, the system uses this route to forward frames to the serial device.

Resolving Differences in LLC2 and SDLC Frame Size

IBM nodes on Token Ring media normally use frame sizes greater than 1 KB, whereas the IBM nodes on serial lines normally limit frame sizes to 265 or 521 bytes. To reduce traffic on backbone networks and provide better performance, Token Ring nodes should send frames that are as large as possible. As part of the SDLLC configuration on the serial interface, the largest frame size the two media can support should be selected. The Cisco IOS software can fragment the frames it receives from the Token Ring device before forwarding them to the SDLC device, but it does not assemble the frames it receives from the serial device before forwarding them to the Token Ring device.

Maintaining a Dynamic RIF Cache

SDLLC maintains a dynamic RIF cache and caches the entire RIF; that is, the RIF from the source station to destination station. The cached entry is based on the best path at the time the session begins. SDLLC uses the RIF cache to maintain the LLC2 session between the router and the host FEP. SDLLC does not age these RIF entries. Instead, SDLLC places an entry in the RIF cache for a session when the session begins and flushes the cache when the session terminates. You cannot flush these RIFs because if you flush the RIF entries randomly, the Cisco IOS software cannot maintain the LLC2 session to the host FEP.

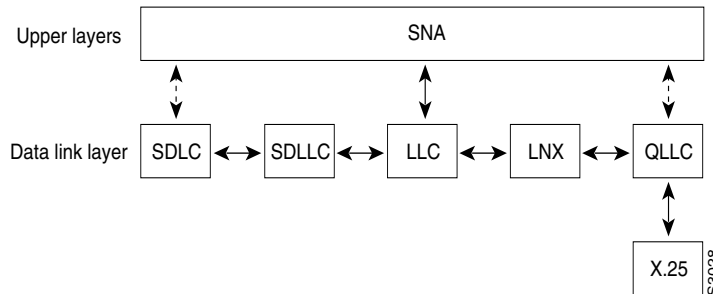
Other Considerations

- As part of Cisco's SDLC implementation, only modulus 8 Normal Response Mode (NRM) sessions are maintained for the SDLC session.
- SDLC sessions are always locally acknowledged. LLC2 sessions can be optionally configured for local acknowledgment.
- SDLLC does not apply to SNA subarea networks, such as 37x5 FEP-to 37x5 FEP communication.
- Parameters such as the maximum number of information frames (I-frames) outstanding before acknowledgment, frequency of polls, and response time to poll frames can be modified per interface. If local acknowledgment is not enabled, these parameters are modified on the SDLC interface; if local acknowledgment is enabled, these parameters are modified on the Token Ring interface.
- Local acknowledgment only applies when the remote peer is defined for RSRB using Internet Protocol (IP) encapsulation over a TCP connection. If no local acknowledgment is used, the remote peer can be defined for RSRB using direct encapsulation, RSRB using IP encapsulation over an FST connection, or RSRB using IP encapsulation over a TCP connection.

QLLC Conversion


Qualified Logical Link Control (QLLC) is a data link protocol defined by IBM that allows SNA data to be transported across X.25 networks. (Although IBM has defined other protocols for transporting SNA traffic over an X.25 network, QLLC is the most widely used.) Figure 12 illustrates how QLLC conversion provides data link layer support for SNA communication.

Figure 12 SNA Data Link Layer Support




As shown in Figure 13, any devices in the SNA communication path that use X.25, whether end systems or intermediate systems, require a QLLC implementation.

Figure 13 SNA Devices Running QLLC




As shown in Figure 14, the QLLC conversion feature eliminates the need to install the X.25 software on local IBM equipment. A device attached locally to a Token Ring network can communicate through a router running the QLLC Conversion feature with a remote device attached to an X.25 network using QLLC. Typically, the locally attached device is a front-end processor (FEP), an AS 400, or a PS/2, and the remote device is a terminal controller or a PS/2. In this case, only the remote device needs an X.25 interface and the FEP can communicate with the terminal controller as if it were directly attached via a Token Ring network.

Figure 14 Router Running QLLC Conversion Feature



More elaborate configurations are possible. The router that implements QLLC conversion need not be on the same Token Ring network as the FEP. As shown in Figure 15, QLLC/LLC2 conversion is possible even when an intermediate IP WAN exists between the router connected to the X.25 network and the router connected to the Token Ring.

Figure 15 QLLC Conversion Running on Router with Intermediate IP Network

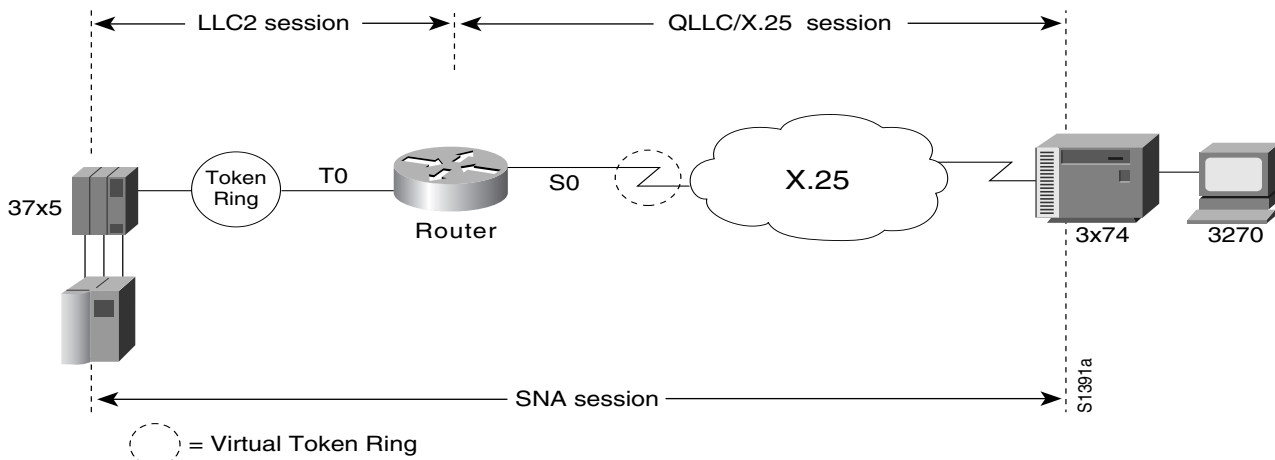


Cisco's Implementation of QLLC Conversion

SNA uses QLLC and X.25 as link-layer protocols to provide a reliable connection. QLLC itself processes QLLC control packets. In a Token Ring environment, SNA uses LLC to provide a reliable connection. The LAN-to-X.25 (LNX) software provides a QLLC conversion function to translate between LLC and QLLC.

Figure 16 shows the simplest QLLC conversion topology: a single Token Ring device (for example, a 37x5 FEP) communicates with a single remote X.25 device (in this case a 3x74 cluster controller). In this example, a router connects the Token Ring network to the X.25 network.

Figure 16 QLLC Conversion between a Single 37x5 and a Single 3x74

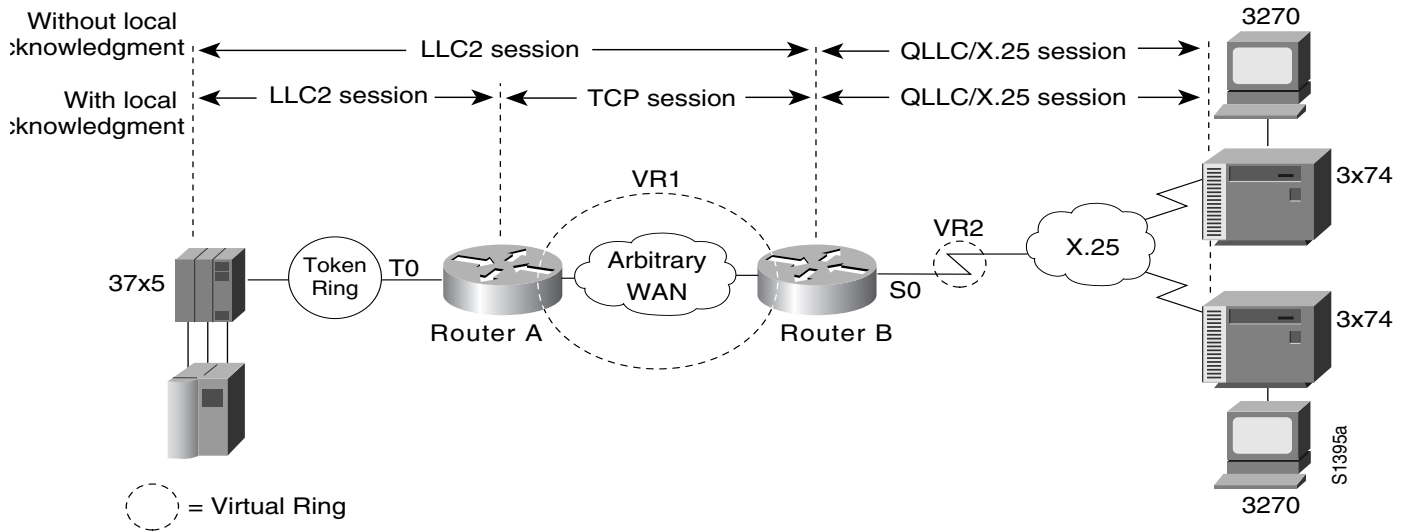


In Figure 16, each IBM end node has no indication that its counterpart is connected to a different medium running a different protocol. The 37x5 FEP responds as if the 3x74 cluster controller were communicating over a Token Ring, whereas the 3x74 responds as though the 37x5 FEP were communicating over an X.25 network. This is accomplished by configuring the router's X.25 interface as a virtual Token Ring, so that the X.25 virtual circuit appears to the Token Ring device (and to the router itself) as if it were a Token Ring to which the remote X.25 device is attached.

Also in this figure, the LLC2 connection extends from the 37x5 FEP across the Token Ring network to the router. The QLLC/X.25 session extends from the router across the X.25 network to the 3x74 cluster controller. Only the SNA session extends across the Token Ring and X.25 networks to provide an end-to-end connection from the 37x5 FEP to the 3x74 cluster controller.

As Figure 17 shows, a router need not directly connect the two IBM end nodes; instead, some type of backbone WAN can connect them. Here, RSRB transports packets between Router A and Router B, while Router B performs all conversion between the LLC2 and X.25 protocols. Only the router attached to the serial line (Router B) needs to be configured for QLLC conversion. Both Router A and Router B are configured for normal RSRB.

Figure 17 QLLC Conversion between a Single 37x5 and Multiple 3x74s across an Arbitrary WAN



How communication sessions are established over the communication link varies depending on whether or not LLC2 local acknowledgment has been configured on Router A's Token Ring interface. In both cases, the SNA session extends end-to-end and the QLLC/X.25 session extends from Router B to the 3x74 cluster controller. If LLC2 local acknowledgment has not been configured, the LLC2 session extends from the 37x5 FEP across the Token Ring network and the arbitrary WAN to Router B. In contrast, when LLC2 local acknowledgment has been configured, the LLC2 session extends from the 37x5 FEP to Router A, where it is locally terminated. A TCP session is then used across the arbitrary WAN to Router B.

Comparing QLLC Conversion to SDLLC

Although the procedures you use to configure QLLC are similar to those used to configure SDLLC, there are structural and philosophical differences between the point-to-point links that SDLC uses and the multiplexed virtual circuits that X.25 uses.

The most significant structural difference between QLLC conversion and SDLLC is the addressing. To allow a device to use LLC2 to transfer data, both SDLLC and QLLC provide virtual MAC addresses. In SDLLC, the actual MAC address is built by combining the defined virtual MAC (whose last byte is 0x00) with the secondary address used on the SDLC link; in this way, SDLLC supports multidrop. In QLLC conversion, multidrop is meaningless, so the virtual MAC address represents just one session and is defined as part of the X.25 configuration. Because one physical X.25 interface can support many simultaneous connections for many different remote devices, you only need one physical link to the X.25 network. The different connections on different virtual circuits all use the same physical link.

The most significant difference between QLLC conversion and SDLLC is the fact that a typical SDLC/SDLLC operation uses a leased line. In SDLC, dial-up connections are possible, but the maximum data rate is limited. In QLLC, both switched virtual circuits (SVCs) and permanent virtual circuits (PVCs) are available, but the favored use is SVC. While the router maintains a permanent connection to the X.25 network, a remote device can use each SVC for some bounded period of time and then relinquish it for use by another device. Using a PVC is very much like using a leased line.

Table 3 shows how the QLLC commands correspond to the SDLLC commands.

Table 3 QLLC and SDLLC Command Comparison

QLLC Command	Analogous SDLLC Command
qllc largest-packet	sdllc ring-largest-frame, sdllc sdlc-largest-frame
qllc partner	sdllc partner
qllc sap	sdllc sap
qllc srb, x25 map qllc, x25 pvc qllc	sdllc traddr
qllc xid	sdllc xid
source-bridge qllc-local-ack	source-bridge sdllc-local-ack

Other Implementation Considerations

Consider the following when implementing QLLC conversion:

- To use the QLLC conversion feature, a router must have a physical link to an X.25 public data network (PDN). It must also have an SRB/RSRB path to an IBM FEP. This link could be a Token Ring or Ethernet interface, or even FDDI, if RSRB is being used.
- QLLC conversion can run on any router with at least one serial interface configured for X.25 communication and at least one other interface configured for SRB or RSRB.
- QLLC conversion security depends upon access control in SRB/RSRB and X.25 and upon XID validation.

You can configure DLSw+ for QLLC connectivity, which enables the following scenarios:

- Remote LAN attached devices (physical units) or SDLC-attached devices can access a front-end processor (FEP) or an AS/400 over an X.25 network.
- Remote X.25-attached SNA devices can access a FEP or an AS/400 over a Token Ring or over SDLC.

For information on configuring DLSw+ for QLLC conversion, refer to the “Configuring DLSw+” chapter.

You can configure DSPUs for QLLC. For more information on this configuration, refer to the “Configuring DSPU and SNA Service Point Support” chapter.

Downstream Physical Unit and SNA Service Point

Downstream physical unit (DSPU) is a software feature that enables the router to function as a physical unit (PU) concentrator for SNA PU type 2 nodes. PU concentration at the device simplifies the task of PU definition at the upstream host while providing additional flexibility and mobility for downstream PU devices.

The DSPU feature allows you to define downstream PU type 2 devices in the Cisco IOS software. DSPU reduces the complexity of host configuration by letting you replace multiple PU definitions that represent each downstream device with one PU definition that represents the router.

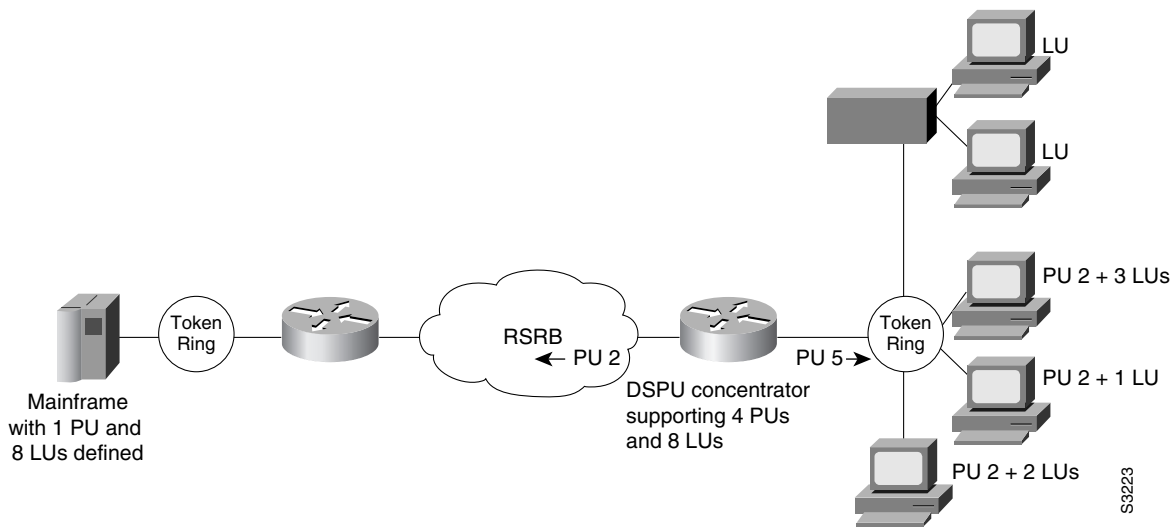
Because you define the downstream PUs at the router rather than the host, you isolate the host from changes in the downstream network topology. Therefore you can insert and remove downstream PUs from the network without making any changes on the host.

The concentration of downstream PUs at the router also reduces network traffic on the WAN by limiting the number of sessions that must be established and maintained with the host. The termination of downstream sessions at the router ensures that idle session traffic does not appear on the WAN.

Our SNA Service Point support in the Cisco IOS software assumes that NetView or an equivalent product is available at the SNA host. The user interacts with the network management feature in the router and at the SNA host. In the Cisco IOS software, you can configure the host connection and show the status of this connection. At the SNA host, you can use the NetView operator's console to view alerts and to send and receive Cisco syntax commands to the Cisco device.

Figure 18 shows a router functioning as a DSPU concentrator.

Figure 18 Router Acting as a DSPU Concentrator

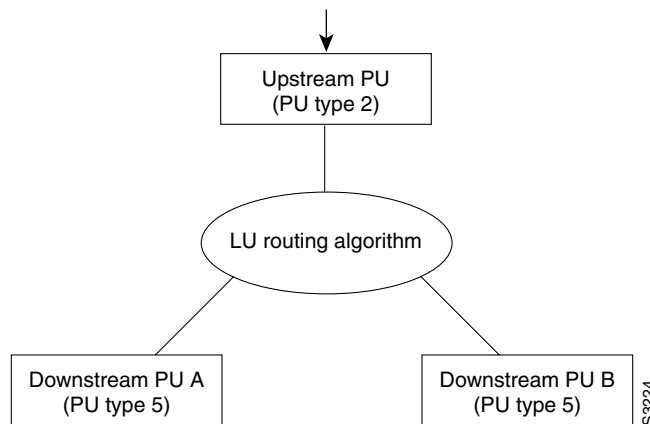


Typically, a router establishes one or more upstream connections with one or more hosts and many downstream connections with PU type 2 devices. From an SNA perspective, the router appears as a PU type 2 device to the upstream host and assumes the role of a system services control point (SSCP) appearing as a PU type 5 device to its downstream PUs.

The SSCP sessions established between the router and its upstream host are completely independent of the SSCP sessions established between the router and its downstream PUs. SNA traffic is routed at a logical unit (LU) level using a routing algorithm that maps downstream LUs onto upstream LUs.

Figure 19 illustrates the SNA perspective of DSPU.

Figure 19 SNA Perspective of DSPU



S3224

SNA Frame Relay Access Support

Using Frame Relay Access Support (FRAS) the Cisco IOS software allows branch SNA devices to connect directly to a central site front-end processor over a Frame Relay network. FRAS converts LAN or SDLC protocols to a Frame Relay format understood by the Network Control Program (NCP) that runs in a front-end processor. The Cisco IOS software and the NCP support two frame formats:

- RFC 1490 routed format for LLC2, specified in the FRF.3 Agreement from the Frame Relay Forum and known in NCP literature as Frame Relay Boundary Network Node (BNN) support. Support for this feature requires NCP 7.1 or higher.
- RFC 1490 802.5 source-route bridged format, known in NCP literature as Frame Relay Boundary Access Node (BAN) support. Support for this feature requires NCP 7.3 or higher.

Management service point support in FRAS allows the SNA network management application, NetView, to manage Cisco routers over the Frame Relay network as if it were an SNA downstream PU.

FRAS provides dial backup over RSRB in case the Frame Relay network is down. While the backup Public Switched Telephone Network (PSTN) is being used, the Frame Relay connection is tried periodically. As soon as the Frame Relay network is up, it will be used at once.

RFC 1490 Routed Format for LLC2 (BNN)

RFC 1490 specifies a standard method of encapsulating multiprotocol traffic with data link (Level 2 of the OSI model) framing. The encapsulation for SNA data is specified in the FRF.3 Agreement.

The Frame Relay encapsulation method is based on the RFC 1490 frame format for “user-defined” protocols using Q.933 NLPID, as illustrated in Figure 20.

Figure 20 Frame Relay Encapsulation Based on RFC 1490

DLCI Q.922 Address	Control 0x30	NLPID Q.933 0x08	L2 Protocol ID 0x4c (802.2) 0x08	L3 Protocol ID	DSAP SSAP	Control	F C S
--------------------------	-----------------	------------------------	--	-------------------	--------------	---------	-------------

S3217

Note The protocol ID for SNA subarea FID4 is 0x81. The protocol ID for SNA subarea FID2 is 0x82. The protocol ID for APPN FID2 is 0x83.

Frame Relay access support allows the router acting as a FRAD to take advantage of the SNA BNN support for Frame Relay provided by ACF/NCP 7.1 and OS/400 V2R3. Downstream PU 2.0 and PU 2.1 devices can be attached to the router through SDLC, Token Ring, or Ethernet links. The router acting as a FRAD is connected to the Network Control Program (NCP) or AS/400 through a public or private Frame Relay network, as illustrated in Figure 21.

Figure 21 SNA BNN Support for Frame Relay

The frame format that communicates across the Frame Relay BNN link is defined in RFC 1490 for routed SNA traffic. From the perspective of the SNA host (for example an NCP or AS/400), the Frame Relay connection is defined as a switched resource similar to a Token Ring BNN link. Because the frame format does not include link addresses to allow the NCP to distinguish among SNA devices on the same permanent virtual circuit, Cisco supports a feature known as service access point (SAP) multiplexing. SAP multiplexing allows you to configure unique LLC2 SAPs for each downstream SNA device so that they can share a single permanent virtual circuit to the front-end processor.

The Cisco IOS software is responsible for terminating the local data link control frames (such as SDLC and Token Ring frames) and for modifying the data link control frames to 802.2 compliant LLC frames. The LLC provides a reliable connection-oriented link layer transport required by SNA. (For example, 802.2 LLC is used to provide link layer acknowledgment, sequencing, and flow control.)

The Cisco IOS software encapsulates these 802.2 LLC frames according to the RFC 1490 format for SNA traffic. The frames are then forwarded to the SNA host on a Frame Relay permanent virtual circuit (PVC). In the reverse direction, the software is responsible for de-encapsulating the data from the Frame Relay PVC, and for generating and transmitting the appropriate local data link control frames to the downstream devices.

RFC 1490 Bridged Format for LLC2 (BAN)

BAN provides functionality similar to BNN except that it uses a bridged frame format, as illustrated in Figure 22.

Figure 22 RFC 1490 Bridged Frame Format

Q.922 Address			
Control	0x03	pad	0x00
NLPID	SNAP 0x80	OUI	00x0
OUI 0x80-C2 (bridged)			
PID 0x00-09			
pad 0x00		Frame Control	
Destination/Source MAC (12 bytes)			
DSAP		SSAP	
Control			
SNA Data			
PCS			

H7115

Because it includes the MAC header information in every frame, BAN supports multiple SNA devices sharing a single permanent virtual circuit without requiring SAP multiplexing. BAN also supports load balancing across duplicate data-link connection identifiers to the same or different front-end processors at the data center to enhance overall availability. BAN works for devices attached by either Token Ring or Ethernet.

Advanced Peer-to-Peer Networking

Advanced Peer-to-Peer Networking (APPN) is the second generation of SNA. APPN provides support for client/server applications and offers more dynamics than traditional hierarchical SNA, such as dynamic directory and routing services.

Cisco's APPN implementation includes the following features:

- Implements High Performance Routing (HPR), the next step in the evolution of SNA networking. HPR replaced the APPN routing technique called intermediate session routing (ISR) with two elements that provide significant performance improvements:
 - Rapid Transport Protocol (RTP) provides functions including error recovery, packet resequencing, segmentation, selective retransmissions, flow control and congestion control. RTP incorporates a new congestion avoidance algorithm, Adaptive Rate-Based (ARB) congestion control.
 - Automatic Network Routing provides a low-level routing mechanism that minimizes cycles and storage requirements for routing packets through intermediate nodes.

- Supports a wide variety of media: Token Ring, Ethernet, FDDI, Frame Relay, QLLC over X.25, and SDLC. APPN also enables interoperability with local and remote source route bridging via APPN ports that can connect directly to source route bridge ring groups.
- Supports APPN connection networks over RSRB and DLSw+.
- Enables existing APPN products, such as IBM's Communications Manager/2 (CM/2), Application System/400 (AS/400), System/36 (S/36), and Advanced Communication Facility/Virtual Telecommunications Access Method (ACF/VTAM), to connect as network nodes (NNs) or end nodes (ENs) over a network of Cisco routers.
- Implements APPN across the channel. Cisco routers with Channel Interface Processors (CIPs) can replace FEPPs and other channel-attached controllers, enabling ACF/VTAM definition either as an EN or NN that uses the Cisco network for session routing. Both APPN and the CIP offer the ability to communicate through the source route bridging mechanism in the Cisco router.
- Implements the Dependent Logical Unit Requester (DLUR) feature. Enables Dependent Logical Units (DLUs) that support only legacy applications (such as those using 3270 data streams) to traverse the APPN network to access those applications.
- Allows prioritization of traffic and bandwidth based on user selected criteria.
- Implements scalability enhancement features, including Locate Throttling and Negative Caching, which allow you to conserve network resources.
- Supports SNMP management via APPN MIB (RFC 1593), an informational RFC provided by IBM.
- Supports a new MIB definition approved by the APPN Implementors Workshop (AIW). The new MIB provides better manageability of APPN network nodes across implementations and adds objects for supporting connection networks. Cisco supports both the current and new MIBs to allow for migration of our application customers from the current version which supports RFC 1593 to a new version for this new MIB.

The following section describes some of the components of an APPN network and reviews basic SNA terminology. The section identifies and compares node types and compares APPN with subarea SNA.

SNA Terminology

The basic component of an SNA network, subarea or APPN, is the network addressable unit (NAU). A NAU is assigned a unique eight-character name and an eight-character network identifier. Examples of NAUs are LUs, PUs, control points (CPs) and system services control points (SSCPs).

Logical Unit

A logical unit (LU) is an interface that enables end users to gain access to network resources and communicate with each other. Examples of LUs are printers, terminals, and applications. LUs communicate with each other via LU-LU sessions. The LU-LU session is the basis of communication in SNA; all end user data traffic communicates through this session type.

To participate in an SNA network, a DLU requires the services of a VTAM host acting as an SSCP. The SSCP must establish connections with each DLU before the DLU is able to participate in the network. Dependent LUs, such as 3270 terminals, are only capable of maintaining a single LU-LU session at any one time.

An Independent LU (ILU) does not require the services of an SSCP to participate in an SNA network. In addition, an ILU can establish sessions to more than one partner in the network, and can have multiple parallel sessions with the same partner LU. Applications implementing Advanced Program-to-Program Communications (APPC) are examples of independent LUs.

Physical Unit

SNA defines a physical unit (PU) as the representation of the physical device. The PU manages and monitors the resources (such as attached links and adjacent link stations) associated with an SNA node.

Physical Unit Type 2 (PU2), the legacy physical unit, only supports dependent LUs, and requires the services from a VTAM host to perform network functions.

Physical Unit Type 2.1 (PU2.1), also known as a type 2.1 node, offers peer node capabilities in an SNA environment. A PU2.1 can support both dependent and independent LUs. In addition, a PU2.1 can support a control point, which is central to APPN networking.

APPN extends the PU T2.1 architecture to provide dynamic discovery and definition of resources and routing capabilities for large, complex networks.

Control Point

A control point (CP) identifies the networking components of a PU Type 2.1 node. In APPN, CPs are able to communicate with logically adjacent CPs by way of CP-CP sessions. Almost all APPN functions, including searches for network resources and discovery of network topology, use CP-CP sessions as the means of communication between nodes.

APPN Node Types

In an APPN network, different node types distinguish different levels of networking capabilities. This section describes APPN node types.

Low-Entry Networking Node

A Low-Entry Networking (LEN) node, sometimes called a LEN end node, is a PU 2.1 without APPN enhancements. The following are some of the characteristics of a LEN node:

- A LEN node does not support CP-CP sessions. The CP does not communicate with other nodes.
- The LEN node participates in the APPN network by using the services of an adjacent Network Node (NN).
- Destination LUs in an attached APPN network must be defined to the LEN node, and destination LUs on the LEN must be defined on the attached NN.

LEN nodes predate APPN, but are able to interoperate with an APPN network. Because there is no CP-CP session between a LEN node and its NN, resources at the LEN node must be defined at the NN, reducing dynamic resource discovery capabilities.

End Node

An EN is a PU 2.1 that includes the APPN support necessary to gain full access to an APPN network, but is not capable of performing in an intermediate role when routing APPN sessions. The following are the characteristics of an EN:

- Contains a CP to manage its resources.
- Uses the services of an adjacent NN to access the network.
- Provides partial network services.
- May connect to multiple NNs.
- Is always a session end point.

Because an EN lacks full APPN routing capability, it might be thought of as an application host or point of user access. The EN establishes CP-CP sessions with its NN server, so topology and directory information can be exchanged dynamically, eliminating the need to define resources on the NN. The EN may connect to multiple network nodes and LU-LU sessions can be established through any of the connected network nodes, although only one network node can be its network node server at any one time.

Network Node

An NN implements the APPN extensions to the PU 2.1 architecture that allow it to provide intermediate routing services to LEN nodes and ENs. An NN contains a CP to manage its own resources and the ENs and LEN nodes in its domain.

An NN provides the following network services:

- Connectivity
- Location of resources in the network
- Route selection
- Intermediate session routing
- Data transport
- Network management

An NN may be a session end point or an intermediate system.

An NN Server is a network node that provides resource location and route selection services for the LEN nodes, ENs, and LUs it serves. The nodes served, ENs or LEN nodes, are defined as being in the network node server's domain.

APPN Components

This section describes the basic components of APPN and how they interact to provide APPN networking functions.

APPN Connections

The configuration services (CS) component of an APPN node manages local interfaces and connections to the APPN network. CS controls the ports and link stations on the node. A port defines a connection to a transport media accessible to APPN, while a link station identifies the addressing information and characteristics of a connection with another node.

APPN Network: Connection Phase

When two APPN nodes connect, the following activities occur:

- An APPN port is activated, allowing the node to access an interface and its corresponding transport media.
- A link station is activated, initiating a connection-oriented link with the partner node.
- If the CP is APPN capable, and there is no CP-CP session already between these two nodes over a different link, a pair of CP-CP sessions may be established between adjacent CPs. EN-NN and NN-NN CP-CP sessions can be established. Although EN-EN APPN connections can be defined, an EN will not establish CP-CP sessions with another EN. LEN nodes are not capable of establishing CP-CP sessions.

The entire connection phase may be configured to occur automatically, or the connection can be initiated manually via EXEC commands.

A link is defined as both the link stations within the two nodes it connects and the link connection between the nodes. A link station is the hardware or software within a node that enables the node to attach to, and provide control over, a link connection. A link connection is the physical medium over which data is transmitted. In legacy SNA, a transmission group may consist of one or more links between two nodes, but in the APPN architecture, transmission groups are limited to a single link. It is therefore common in APPN to use the terms link and transmission groups interchangeably.

Details of Link Activation

The connection phase in APPN begins with link activation, which initiates communication between nodes. It is independent of the DLC chosen, and may not be required for some DLC types. For switched connections, the connection phase is similar to “dial” and “answer” procedures. In an X.25 network, the connection phase would be the establishment of a virtual circuit. When the connect phase is complete, the two nodes can exchange and establish node characteristics through exchange identification (XID).

The exchange of XIDs allows a node to determine whether the adjacent station is active and to verify the identity of the adjacent node. Node identification fields, including the CP name, will be exchanged. This information exchange allows, for example, a node to correlate an incoming connection with a link station definition on the node.

During XID exchange, primary and secondary link stations are determined. The two link stations compare the XID node identifier values (block number plus ID number). If the link stations are defined as negotiable, then the higher node ID becomes primary. If the nodes have the same node ID, each generates a random number and the node with the higher random number becomes primary. The result of the primary-secondary role negotiation determines which node will send the mode-setting command—Set Normal Response Mode (SNRM) for SDLC, Set Asynchronous Balanced mode (SABM) for X.25, and so on.

After the link activation XID exchange is complete, CS creates a new path control and instructs the Address Space Manager (ASM) to activate a new address space. Then CS notifies Topology and Routing Services (TRS) that a transmission group (TG) has become active so the APPN topology can be updated. Finally, if a CP-CP session is to be activated, CS notifies SS. SS then activates the CP-CP session.

CP-CP Session Activation

After a link is established, the nodes determine whether CP-CP sessions should be established. Between network nodes, CP-CP sessions are normally activated on the first link to become active between the nodes. An EN determines which of the NNs will be used for a CP-CP session. The EN

indicates which NN is the server by sending a request to activate a session, known as a BIND, to the CP on the adjacent NN. The NN accepts by sending BIND for a second LU 6.2 session, completing the CP-CP session pair. Each node uses one session to send CP-CP communication data, while the other is reserved for receiving data from the partner node.

CPSVCMG is the class of service (COS) used for CP-CP sessions. It indicates a transmission priority of network, which is the highest of the four transmission priorities in APPN.

When the CP-CP session is established, the CPs exchange capabilities, and, in the case of EN to NN CP-CP sessions, register the local LUs.

APPN Topology

APPN maintains a map of the APPN network as known to a particular node. TRS is the APPN component responsible for maintaining the topology database.

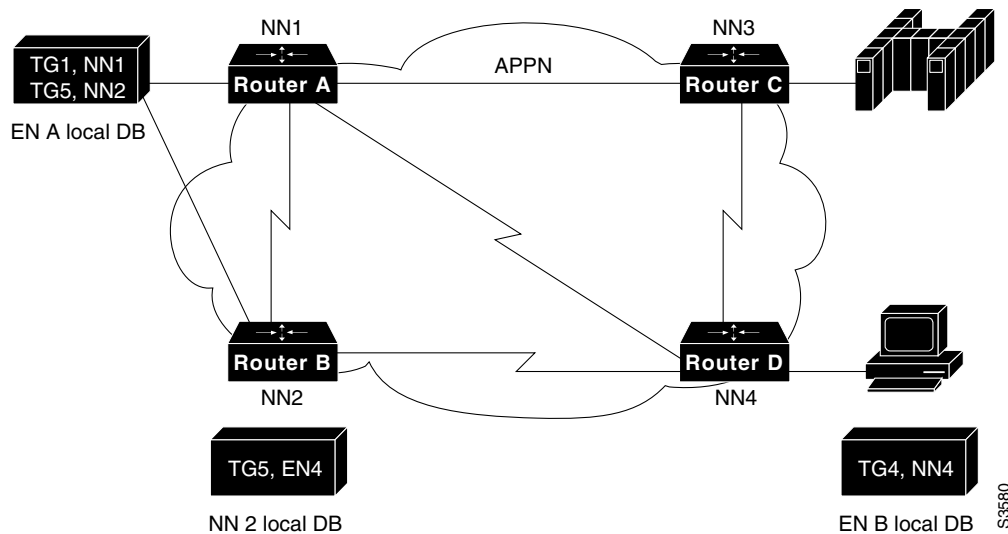
There are two types of information in a topology database: local topology and network topology.

- Local topology, found in both ENs and NNs, describes the local links attached to a particular node, and the partner control points to which it is connected.
- Network topology, found only in NNs, describes all the NNs in the network and the links that interconnect them.

An EN uses the information in its local topology database to send local TG information to the NN server on APPN search requests and replies. The TG information is passed to the NN server when a route is requested, so the NN can select the best TG from the EN to one of its adjacent NNs. In a NN, the local topology database includes information about the attached ENs and TGs.

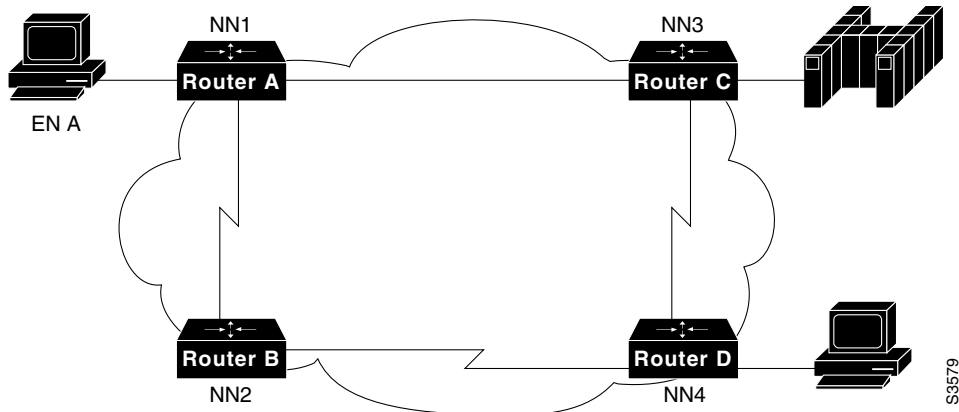
Figure 23 shows the local topology as it is known to EN A, NN2, and EN B.

Figure 23 APPN Local Topology



Network topology contains information on all network nodes in the APPN network as well as the TGs interconnecting them. Every NN maintains a fully replicated copy of the network topology database. Figure 24 illustrates an APPN network topology.

Figure 24 Network Topology



The network topology database is built from information about the local NN and its TGs to other NNs, and from Topology Database Updates (TDUs) received from adjacent NNs. TDUs are exchanged whenever NNs establish CP-CP sessions with each other. As updates occur in NNs or TGs between NNs, the owning NN sends a TDU to its adjacent NNs, which propagates the TDU to its adjacent NNs until the network topology database is again replicated.


For each NN, certain properties are specified in the topology database. Each node and TG in a network topology database is assigned a unique resource sequence number. These numbers are incremented to the next even value whenever the owning network node creates a TDU for that resource. TG database entries include whether CP-CP sessions are supported, and include the TG weight, TG status (operative or inoperative), TG number, partner node information, and TG characteristics such as cost per byte, cost per connect time, and security level.

TDUs provide updated information to NNs about the node itself and information about all locally owned TGs to other network nodes. TDUs can be triggered by changes in node or TG characteristics.

TRS broadcasts TDUs containing local node information every five days to prevent other network nodes from discarding valid information, which occurs after 15 days with no update. This 15-day cleanup of the database is called garbage collection.


In Figure 25, NN3 adds itself to the network. NN3 forwards information about itself to NN1. NN1 forwards the information to NN2 and NN4. NN2 and NN4 then forward a second TDU to each other. Because they will have the same RSN and information, the second TDUs will be discarded.

Figure 25 Topology Database Update: Adding a Node



In Figure 26, TG6 is activated between NN4 and NN3. TDUs are exchanged between the two nodes, so each node can build a new topology database. New information is propagated to adjacent nodes once NN4 and NN3 have updated topology databases. TG6 could be active, but with no CP-CP session established. The TG will still be included in the network topology and forwarded, so that path can be used for sessions.

Figure 26 Topology Database Update: Adding a Transmission Group



APPN Directory and APPN Searches

Directory services is the APPN component responsible for managing the directory database and searching for network resources throughout the APPN network. The directory database should not be confused with the topology database. The directory database maintains information about resource names and their owners, while the topology database maintains a network map of NNs and TGs. Directory services locates a session partner. Topology service computes an optimal route to the session partner when it has been located.

At a minimum, each APPN network resource must be defined on the node where the resource exists. When the resource is defined, it can be found through network searches. Optionally, the resource location may be defined in other nodes to optimize directory services search logic.

Registered directory entries are created dynamically in the local directory of the NN. These entries represent local LUs of an adjacent EN for which the NN is providing network node server function.

Cached directory entries are added dynamically and are based on the results of previous searches. A cache entry allows the node to attempt a directed search straight to the destination resource. If a cache entry does not exist, or the entry is incorrect, a broadcast search is issued to query each NN in the network for the destination resource. Each cache entry is verified before use to make sure the resource is still available. If the resource is not found, a broadcast search is attempted.

Some implementations, including Cisco's, support safe store of the directory cache. The cache entries in a network node's directory database are periodically written to a permanent storage medium. Safe store permits faster access (after a network node failure or initial power-on) by eliminating network broadcast searches for safe-stored resources.

The central resource registration feature implements the registration requester function of the APPN options set 1107, Central Resource Registration of LUs.

An EN registers its local resources at its NN server. The NN server, in turn, registers those resources at a central directory server. This feature significantly reduces the number of network broadcast searches in an APPN network. If every resource is registered, then all network nodes can query the central directory server, which eliminates the need for broadcast searches.

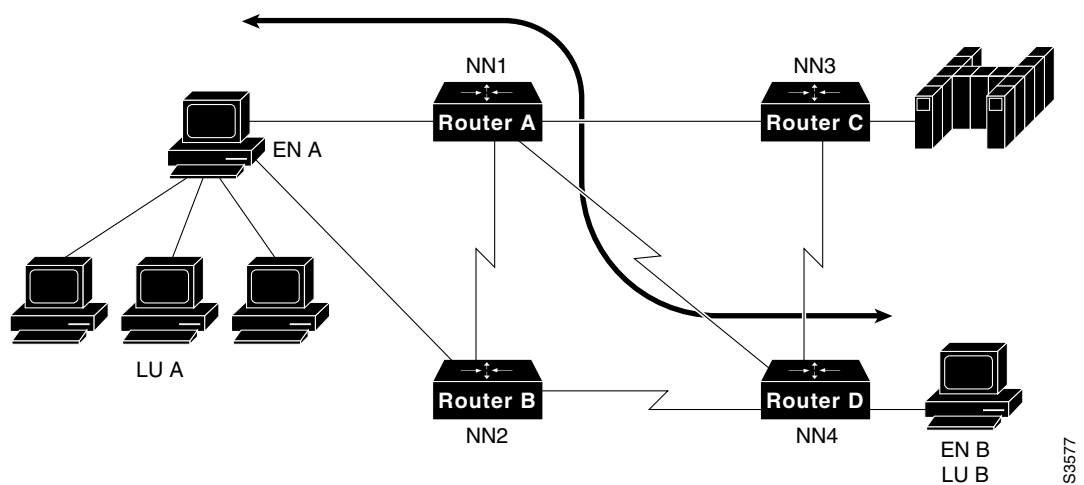
APPN Route Selection

Each APPN session is assigned a route on which the data path for this session will travel. APPN TRS is responsible for computing a route for an LU-LU session. A route is an ordered sequence of nodes and TGs that represents a path from an origin node to a destination node. TRS uses the topology database and its COS definitions to obtain the information necessary to perform a route computation.

While multiple routes may be available between an origin node and a destination node, the best route is selected (see Figure 27). The lowest-cost path that provides the desired level of service is selected.

When a session is requested, the directory services function locates the target resource, and TRS selects a route.

Figure 27 An APPN Network: Selecting the Best Path



When the origin LU requests a session, it either specifies a mode or uses a default mode for the session. The mode determines a COS for this session request. Acceptable COS characteristics are compared to node and TG characteristics to select the best route.

A COS table entry specifies transmission priority, COS name and one or multiple rows of COS characteristics that are acceptable for that COS. Note that traffic on sessions with the same COS can only flow at the same transmission priority. You need a separately named entry to achieve a different priority on the same route.

TG characteristics include link speed, cost per connect time, cost per byte, security class (7 levels), propagation delay, and three user-defined fields.

Node properties include route additional resistance (RAR) and a congestion indicator.

The COS table can contain multiple entries meeting the criteria. Some entries are more acceptable than others. The COS table lets you assign weight to each entry to differentiate the value of each entry. For each possible route, compare the characteristics of the component nodes and transmission groups to the ranges of acceptable characteristics as defined in the COS table for that COS. For each NN or TG the characteristics are compared to see if they are within the range of tolerance. If so, a weight is assigned to each node and each TG and is added to the total weight for that particular route.

When all routes are weighed, the one with the lowest weight is selected. Ties are broken by random selection.

Route selection is a complex process and, where multiple routes are available, may involve considerable overhead. To reduce this overhead, TRS uses routing trees to store the best route path to each node in the network for a specific COS. You may configure the number of routing trees the node will maintain at any one time.

Connection Networks

Connection networks provide extensions to the APPN route calculation algorithm to simplify definitions and enhance connectivity on shared transport media. If you specify a node as a member of a connection network, that node can establish a direct connection, when needed, between itself and any other member of the connection network.

For example, on one Token Ring, there may be 50 end nodes with links and a CP-CP session active to the same network node. The end nodes may communicate with each other through the network node, but this data path, which routes data through the network node then back on the same media to the target end node, is an inefficient use of both the network node and the transmission media.

Alternatively, APPN links can be configured between each end node in a mesh fashion, but this would require over 1000 link definitions. Instead, each end node can be configured as a member of the same connection network. This allows APPN route calculation to calculate a route through the connection network between the two end nodes. The resulting data path is a direct connection between end nodes without a corresponding link definition on either node.

APPN Intermediate Session Routing

After a route is selected, a BIND session command flows from the origin LU to the destination LU over the specified route. The BIND includes:

- A route selection control vector (RSCV), the route for this session.
- A fully qualified procedure correlation identifier (FQPCID), which uniquely identifies the session.
- A transmission priority that indicates the relative priority for frames traveling on this session relative to frames on other sessions. There are four priority levels: network, high, medium, and low. The transmission priority is assigned based on the session COS. In SNA, transmission priority provides the mechanism for delivering high priority traffic, such as interactive traffic, ahead of lower priority traffic, such as batch jobs and file transfers.

- A pacing indicator that indicates the maximum window size (the maximum number of messages that can be transmitted prior to receiving a SNA response). It also indicates whether fixed window size or adaptive pacing is supported.
- A segmenting indicator that identifies the capabilities of adjacent nodes to perform segmentation of data frames received on this session.

As the BIND passes through each NN, the NN records the inbound session identifier or local form session identifier (LFSID), and assigns a new session identifier to be used on the outbound port. It also builds a “session connector” that connects the inbound session identifier with the outbound session identifier, so that it knows how to forward subsequent packets on this session. The session connector also stores information carried in the BIND, such as segmentation values and transmission priority, so that other components, such as path control, will be able to use the information on session traffic.

No subsequent data packets contain any routing information—they contain only the local form session identifier. Packets are forwarded based on information in session connectors, including port and outbound LFSID. LFSIDs are swapped at each NN because they have only local significance.

Dependent Logical Unit Requester/Server

A dependent LU (DLU) is an LU that depends on the SSCP in VTAM to provide services for establishing sessions. Common DLU types are DLU 0, 1, 2, and 3.

Dependent LUs originated in subarea SNA. In early APPN environments, only LU 6.2 independent LUs were supported. The PU supporting dependent LUs still required a logical link directly to a subarea network boundary function (VTAM or NCP) in order to operate. Dependent LU Requester (DLUR) provides the extensions to APPN necessary to allow dependent LUs to interoperate in an APPN network and removes the restriction that dependent LUs must have a direct connection to a subarea network boundary function.

DLUR is the client half of Dependent LU Requester/Server (DLUR/DLUS). The Dependent LU Server (DLUS) is currently implemented in IBM’s VTAM version 4.2.

There are three main concepts in DLUR/DLUS:

- Control flows to the DLUs from VTAM, such as active logical unit (ACTLU) and active physical unit (ACTPU), and session requests flow on a pair of LU 6.2 session between the DLUR and the DLUS.
- Upon receiving a session request from the LU over control session, the VTAM acting as DLUS notifies the VTAM owning the target application that a session is requested.
- The application initiates the session, and its serving network node selects the best route over the APPN network to the DLUR and destination LU.

The DLUR function resides on an EN or NN that owns the dependent LUs. In addition, the DLUR exists in a network node that offers its DLUR services to PU type 2.0 and PU type 2.1 nodes which connect to the DLUR. This arrangement consolidates the LU6.2 control sessions (only one pair is needed for all downstream PUs that the DLUR serves). In addition, this arrangement provides considerable cost savings over upgrading each device that owns dependent LUs to be APPN and DLUR capable.

The DLUR function does not have to be active in all APPN nodes—only in those nodes with DLUs directly attached. Once encapsulated in the LU 6.2 session, the DLUR control traffic (encapsulated SSCP-PU and SSCP-LU control flows) looks like regular APPN traffic.

By providing DLUS support in VTAM, SSCP support is extended to LUs residing on nodes that are nonadjacent to the VTAM or NCP boundary function. Traditional SSCP-PU and SSCP-LU session flows are multiplexed in LU 6.2 sessions between the DLUR and DLUS.

The benefits of using the DLUS/DLUR function include:

- The ability to use dynamic definition of the dependent LUs in VTAM, providing the served PU supports it, to reduce system definition and allow nodes to be moved in the network. All DLUR attached resources appear as switched resources that are not associated with a particular communications line or port.
- The SSCP owns resources that are non-adjacent to a VTAM or NCP.
- Dependent LUs are fully visible as VTAM resources for network management purposes.
- APPN routing can be used to select an optimal route for the session path, even though the secondary logical unit (SLU) is not APPN-capable. A key advantage of DLUR over SNA gateway functions is that the LU-LU session path from the application host to the dependent LU can differ from the path between the DLUR to the DLUS. SNA gateway and PU concentration devices are restricted to having their LU-LU session data paths identical to the path over which the SSCP-PU and SSCP-LU traffic flows.
- Because DLUR offers the ability to separate the control functions for dependent LUs from the session traffic, you can interrupt and recover, or even change, DLUS nodes without interrupting currently active LU-LU sessions to the dependent LUs.
- No changes to end user systems or applications are required.
- Network consolidation—combining the dependent LU traffic with other multiprotocol traffic—provides cost savings.

Native Client Interface Architecture (NCIA)

NCIA is a new software architecture introduced by Cisco Systems to make accessing IBM SNA applications over routed internetworks more scalable and flexible. NCIA is a component of the Cisco IOS software. The architecture is intended to combine the benefits of the native SNA interface at end stations and mainframes with those of TCP/IP across the network backbone.

NCIA extends the use of the TCP/IP protocol all the way to the SNA end station. Because of the wide range of media supported by TCP/IP, including dial-up telephone lines for remotely located users, NCIA makes multiprotocol access to corporate backbone networks much more flexible for SNA users.

NCIA allows SNA end stations such as PCs or workstations to encapsulate SNA traffic in TCP/IP, rather than requiring the traffic to travel through routers. The first phase of NCIA (NCIA I), used Cisco remote source-route bridging (RSRB) encapsulation. The current phase (NCIA Server) uses a new client/server model. NCIA Server is not backward compatible to NCIA I.

NCIA I

Using Cisco's RSRB technology, NCIA I encapsulates the Token Ring traffic inside IP datagrams passed over a TCP connection between a router and a client. A virtual ring is created to allow the router to interconnect any client. The virtual ring acts as a logical Token Ring in the router, so that all the Token Rings connected to the router are treated as if they are all on the same Token Ring. The virtual ring is called a ring group. The ring group number is used just like a physical ring number and shows up in any route descriptors contained in packets being bridged. A ring group must be assigned a ring number that is unique throughout the network.

An NCIA I client acts as both an RSRB router and an end station. It must have a “fake” ring number and a “fake” bridge number so that it looks like an end station sitting on a real Token Ring. The fake ring and bridge numbers are visible to both the RSRB router and the NCIA client. The client must also have an LLC2 so that it can handle the LLC2 sessions.

NCIA Server

The NCIA Server feature extends the scalability of NCIA I, enhances its functionality, and provides support for both the installed base of RSRB routers and the growing number of DLSw+ routers. The NCIA Server feature includes the following enhancements:

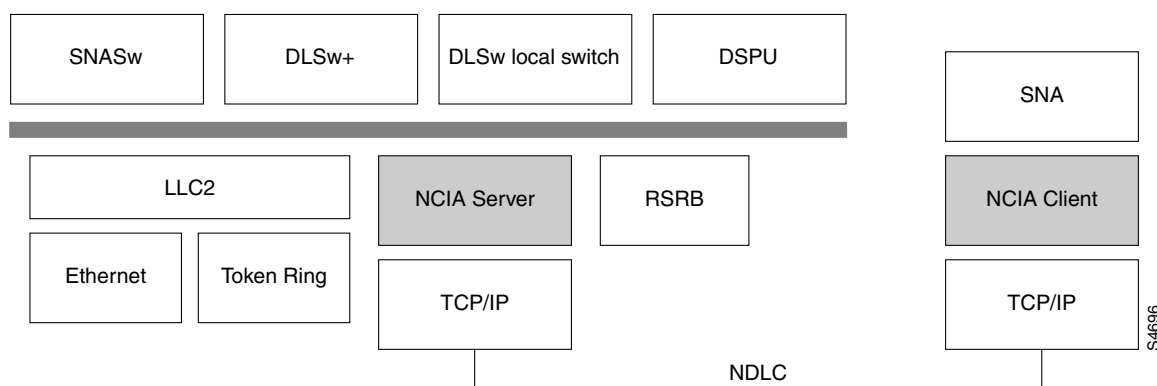
- No need to configure a ring number on the client.
- No need to configure each client on the router.
- MAC address can be dynamically assigned by the NCIA server running on the router.
- SNA is directly on top of TCP/IP; LLC2 is no longer required at end station.
- A client is a true end station, not a router peer.
- NCIA Server communicates with other components in router, such as RSRB, APPN, DLSw+, and DSPU.
- Supports both connect-in and connect-out.
- NCIA client/server model is independent of the upstream implementation.
- Efficient protocol between client and server.

NCIA Client/Server Model

The NCIA Server feature uses a client/server model (see Figure 28), where the NCIA server is a software module on a Cisco router and the NCIA client is a PC or workstation. The NCIA server performs two major functions:

- Establishes TCP to NCIA Data Link Control (NDLC) sessions with clients for the purpose of sending and receiving data.
- Uses the Cisco link services interface (CLSI) to communicate with other software modules in the router, such as APPN, DLSw+, and DSPU, and acts as the data intermediary between them and NCIA clients. The NCIA server’s role as an intermediary is transparent to the client.

Figure 28 NCIA Server Client/Server Model



NCIA Data Link Control (NDLC) is the protocol used between clients and servers. NDLC serves two purposes:

- Establishes the peer connection.
- Establishes the circuit between the client and the server.

The peer session must be established before an end-to-end circuit can be set up. During the set up period for the peer session, the MAC address representing a client is defined. The MAC address can be defined by the client or by the server when the client does not have a MAC address.

The NCIA Server feature supports connect-in and connect-out (from the server's perspective), but connect-out is not supported if the client station does not listen for the incoming connection. For a server to connect-out, clients must connect to the server first. After registering itself by providing its own MAC address, the client can then optionally disconnect from the server. When a server receives an explorer, and its destination MAC address is registered, an NCIA server will connect to that client if it is not connected. For NetBIOS explorers (addressed to functional address 0xC0000000080), the TCP session must remain up so that the server can broadcast the explorers to the client. If the TCP session is down, the server will not send the NetBIOS explorers to a client, even when the client is registered.

After the peer session has been established, the NDLC protocol establishes the circuit between the client and server. This circuit is used to transfer end-user data between the client and the server. Because the client and its target station are not on the same transport, they cannot form a direct, end-to-end circuit. Each client must form a circuit between the client and server, and the server must form another circuit between the server and the target station. The server links those two circuits to form an end-to-end circuit. The server acts as a mediator between the client and the target station so that packets can be transferred between them.

In the NCIA server only peer keepalive is maintained. There is no keepalive at circuit level.

The NCIA server acts as a data-link provider, like Token Ring or Ethernet, in the router. It uses CLSI to communicate with other software modules, just as other data-link providers do. The network administrator configures the router to communicate with specific modules. For data-link users, such as APPN, DLSw+, and DSPU, the NCIA server can interface to them directly. For other data-link providers, the NCIA server must go through a DLSw+ local peer to communicate with them. The DLSw+ local peer passes packets back and forth among different data-link providers.


Advantages of the Client/Server Model

The client/server model used in the NCIA Server feature extends the scalability of NCIA. In addition, it provides support for both the installed base of RSRB routers and the growing number of DLSw+ routers.

Extended Scalability

The client/server model minimizes the number of central site RSRB or DLSw+ peer connections required to support a large network of NCIA clients (see Figure 29). Rather than each client having a peer connection to a central site router, the clients attach to an IP backbone through an NCIA server that, in turn, has a single peer connection to a central site router. This scheme can greatly reduce the number of central site peer connections required. For example, in a network with 1000 clients and 10 NCIA servers, there would be only 10 central site peer connections. Note that there would still be 1000 LLC2 connections that must be locally acknowledged at the central site router, but this can easily be handled in a single central site router. When the number of LLC2 connections (or the number of clients) is in the tens of thousands, NCIA servers can take advantage of downstream PU concentration to minimize the number of LLC2 connections that must be supported by the central site routers.

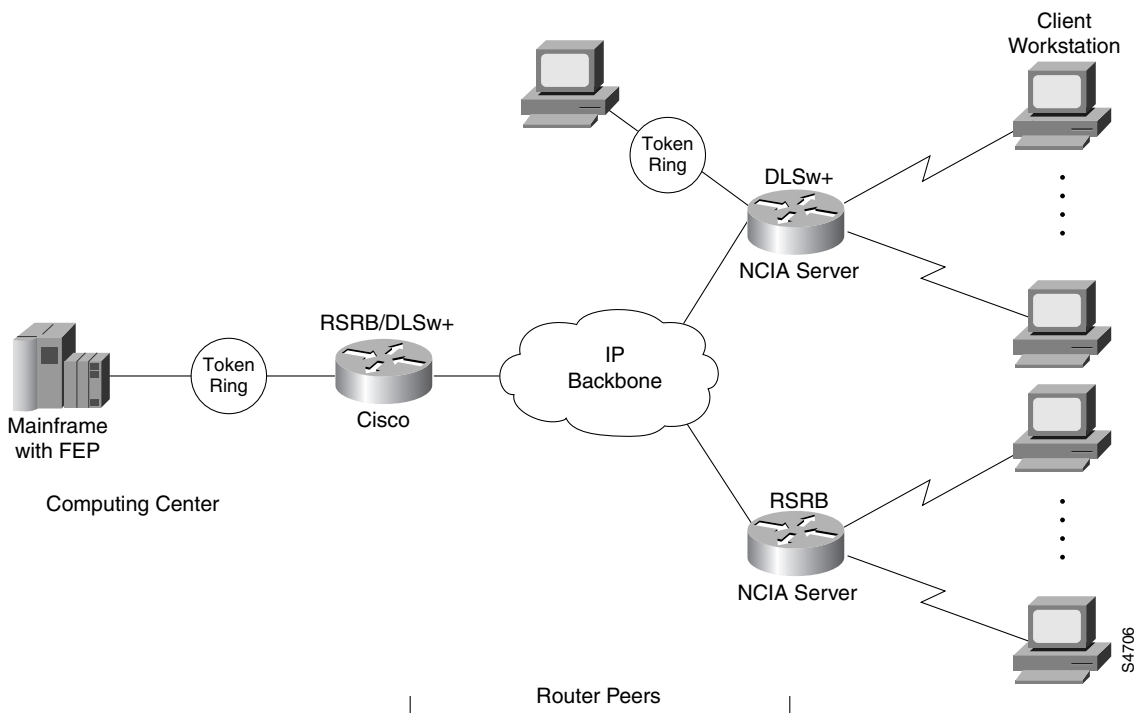
Figure 29 NCIA Server Provides Extended Scalability to Support Large Networks



Migration Support

Using a client/server model allows the NCIA Server feature to be independent of the upstream implementation, allowing it to be implemented in a network that is still using RSRB, as well as in a DLSw+ network. It also greatly simplifies migration from RSRB to DLSw+, because it requires no changes at the client. A single NCIA server can support either approach (but not both). As Figure 30 illustrates, a central site router can support RSRB and DLSw+ concurrently, allowing a portion of the NCIA servers to communicate using RSRB and another portion to communicate using DLSw+.

Figure 30 NCIA Server Provides Independence from the Upstream Network Implementation



IBM Channel Interface Processor

The Cisco 7000 series supports the Cisco IOS software mainframe CIP application, which in turn supports the IBM channel attach feature.

IBM (and IBM-compatible) mainframe hosts are connected to each other and to communication controllers through high-performance communication subsystems called mainframe channels. Cisco supports IBM channel attachment technologies, including the fiber-optic Enterprise Systems Connection (ESCON) channel introduced on the ES/9000 mainframe and the parallel bus-and-tag channel supported on System 370 and later mainframes.

The Cisco 7000 series configured with the CIP (and other interface processors) is an ideal connectivity hub for large corporate networks, and provides the following routing services between mainframes and LANs:

- Replaces an IBM 3172 interconnect controller, which enables mainframe and peripheral access from LAN-based workstations.
- Simplifies the network because the number of network devices is reduced, especially in situations where a router can replace an IBM 3172.
- Ensures 100 percent IBM compatibility with the Cisco ESCON Channel Adapter (ECA) which uses the IBM ESCON chipset.

Note Refer to the Cisco 7000 hardware installation guide for more information on installing the CIP processor.

Common Link Access for Workstations (CLAW) Support

Cisco has implemented Common Link Access for Workstations (CLAW) support in the CIP, which is a link-level protocol used by channel-attached RISC System/6000 series systems and by IBM 3172 devices running TCP/IP offload. The CLAW protocol improves channel efficiency and allows the CIP to provide the functionality of an IBM 3172 in TCP/IP environments and support direct channel attachment. The output from TCP/IP mainframe processing is a series of IP datagrams that the router can switch without modifications.

TCP Offload Support

Cisco has implemented offload processing support for TCP/IP. Like the offload feature of the IBM 3172 Model 3, the TCP offload feature on the CIP is designed to remove processing cycles from the mainframe by executing the TCP protocol on the CIP card. But while the IBM 3172-3 executes TCP in an OS/2 environment, the CIP utilizes the MIPS processor and high-speed channel software to deliver vastly improved performance and scalability. The TCP/IP protocol suite runs on the CIP board and delivers routable IP frames to the Cisco 7000 series router.

CIP Systems Network Architecture Support


CIP Systems Network Architecture (CSNA) support in a Cisco 7000 series router provides mainframe connectivity to SNA network nodes. The CIP supports both ECA and Parallel Channel Adapter (PCA) connections to an IBM mainframe using SNA network features. The CSNA feature provides an SNA LAN gateway to VTAM using a high-speed channel connection.

The CSNA feature also allows you to replace currently installed IBM 3172 interconnect controllers with a Cisco 7000 series router and experience no loss of functionality. You will, in fact, gain functionality with minimal or no changes to VTAM or site configuration.

TN3270 Server Support

The Cisco implementation of TN3270 server for CIP and CSNA provides mapping between an SNA 3270 host and a TN3270 client connected to a TCP/IP network as shown in Figure 31.

Figure 31 TN3270 Implementation



Functionally, it is useful to view the TN3270 server from two different perspectives: SNA functions and Telnet Server functions.

- **SNA Functions**

From the perspective of an SNA 3270 host connected to the CIP, the TN3270 Server is an SNA device that supports multiple physical units (PUs), with each PU supporting up to 255 Type 1, 2, or 3 logical units (LUs). The SNA host is unaware of the existence of TCP/IP extension on the implementation of these LUs.

The LUs implemented by TN3270 server are dependent LUs. To route these dependent LU sessions to multiple VTAM hosts connected to the server in the CIP card, rather than routing in the virtual telecommunications access method (VTAM) hosts, the TN3270 server implements an SNA session switch with end node Dependent LU Requester (DLUR) function. Using the DLUR is optional, so that the TN3270 server can be used with VTAM versions prior to version 4.2, which provide no APPN support.

SNA session switch allows you to eliminate SNA subarea routing between hosts of TN3270 traffic by establishing APPN links with the primary LU hosts directly.

- **Telnet Server Functions**

From the perspective of a TN3270 client, the TN3270 server is a Telnet server that can support approximately 8000 concurrent Telnet sessions. The server on the CIP card supports Telnet connection negotiation and data format as specified in RFC 1576 (referred to as “traditional TN3270”) and RFC 1647 (referred to as “TN3270E”).

For more information on TN3270 on the CIP, refer to the “Configuring IBM Channel Attach” chapter.

IBM Channel Attach Hardware Requirements

Support for IBM channel attach requires the following hardware:

- A Cisco 7000 series router with one available card slot.
- A CIP with one or two adapter cards (ECA, PCA, or a combination).
- Cables for interconnecting the adapter cards to the mainframe or ESCON director switch.

IBM Channel Attach Host Software Requirements

Your mainframe host software must meet the following minimum requirements:

- IBM TCP/IP for VM Version 2 Release 2, with program temporary fix (PTF) enhancements for RISC System/6000 series ESCON support.
- IBM TCP/IP for MVS Version 2 Release 2.1, with program temporary fix (PTF) enhancements for RISC System/6000 series ESCON support.