SNA Perspective

Volume 13, Number 9
September, 1992
ISSN 0270-7284

The single source, objective monthly newsletter covering IBM's Systems Network Architecture

In This Issue:



# Optimizing SNA over Internetworks

The implementation of multiprotocol internetworks continues to grow rapidly. Many network planners assume that these multiprotocol internetworks also will accommodate their SNA traffic requirements. While *SNA Perspective* certainly believes that SNA traffic can be transported effectively over multiprotocol internetworks, the network planner must carefully weigh the needs of the enterprise against the options available from the traditional internetworking (bridge/router) vendors and the emerging "SNA internetworking" vendors.

This article discusses alternatives available for transporting SNA traffic across a multiprotocol internetwork, highlighting the role played by "local termination" of SNA traffic in multiprotocol routers. In this article, we consider the types of SNA traffic to be transported, factors affecting transport of LAN IEEE logical link control level 2 (LLC2) data, factors affecting transport of synchronous data link control (SDLC) data, and issues to be considered when implementing local termination solutions.

*(continued on page 2)*

# SNA and the Future of X.25

Initially, SNA provided WAN connectivity via SDLC. In the early 1980s, however, IBM also began to support SNA sessions through X.25-based packet switched data networks. In the past few years, SNA sessions have been enabled across a wide variety of WAN, LAN, and interLAN/WAN protocols. These trends and changes in SNA, X.25, and other networking protocols and environments leads SNA users to consider alternatives to X.25 for SNA traffic.

This article examines four trends we believe will affect those currently using X.25 with their SNA networks. First, we consider changes in the emerging 1992 CCITT X.25 recommendations, especially addressing and speed. Then we examine the emergence of newer physical and data link technologies such as Frame Relay, SMDS, and ATM and their effect on SNA over X.25. Third, we consider the future for SNA over X.25 without NCP and NPSI when IBM eventually replaces the 3745. Finally, NPSI has used QLLC to smooth the differences between SNA/SDLC and X.25 and some internetworking vendors are looking to integrate SNA over LANs and over X.25 by supporting QLLC to LLC2 conversion, bypassing NPSI. In addition, we provide a brief overview of IBM SNA X.25 products and many ways SNA traffic is supported over X.25.

*(continued on page 11)*

*© 1992 by CSI. Reproduction in whole or in part without express written permission is prohibited.*

(continued from page 1)

## Which SNA Traffic to Transport?

Network managers typically manage a variety of LAN, WAN, and LAN/WAN networks to support SNA traffic as follows:

- LLC2 traffic for physical unit type 2 (PU 2) devices
- LLC2 traffic for node type 2.1 devices
- SDLC traffic for PU 2 devices
- SDLC traffic for 37xx-to-37xx communications

A generalized network is shown in Figure 1. The predominant method used to support LAN and LAN/WAN SNA networks is token ring LANs with source route bridges. Many sites have both SDLC and source route bridge connections. Later in this article, we will discuss options for moving SNA traffic onto a multiprotocol router–based backbone in terms of this generalized SNA network.

## The Multiprotocol Enterprise—Interconnected LANs

Multiprotocol networks have traditionally meant nonSNA networks. These multiprotocol networks evolved to meet the needs of users to communicate on a peer-to-peer basis with other computing environments (again, typically nonSNA) or with other users. These computing environments range from UNIX LAN–based technical computing systems to PC LAN–based departmental application systems.

Multiprotocol routers (supporting protocols such as TCP/IP, IPX, etc.) have become the preferred mechanism for interconnecting these peer-to-peer LAN computing environments. The substantial growth in the internetworking marketplace results from widely available products that provide reliable, cost-effective, high-performance backbone solutions. How can SNA network managers benefit from these same kinds of solutions?

Internetworking vendors are now addressing SNA connectivity and SNA transport needs in a variety of ways including:

- Support for source route and source route transparent bridging of LLC2 traffic
- Support for SDLC transport (SDLC passthrough and SDLC-to-LLC2 conversion)
- Support for local termination of SNA connections (LLC2 and/or SDLC)

## LLC2 Gaining on SDLC

The PU 2 (3x74) SDLC network continues to be the leading method, at the data link layer, for connecting SNA end users to their SNA host applications, although token ring LAN networks using the LLC2 protocol are gaining momentum rapidly. The LLC2 LAN protocol is well suited for peer-to-peer SNA traffic because it is "connection oriented," but most

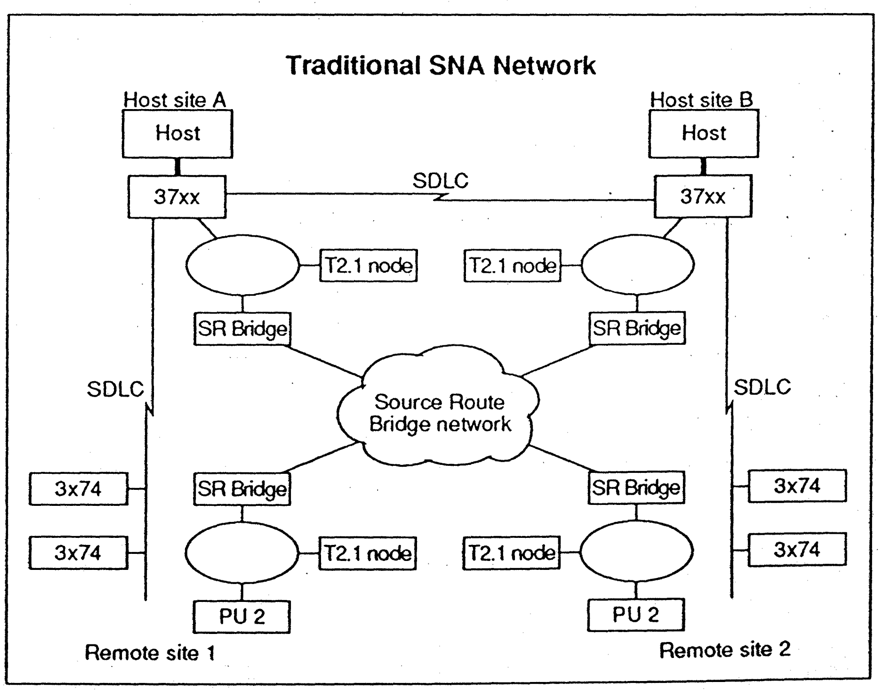


*Figure 1*

enterprises still have a relatively low volume of peer-to-peer SNA traffic. While peer-to-peer traffic will continue to grow in volume, current SNA traffic is hierarchical in nature on both the SDLC WAN network and the LLC2 LAN/WAN network. Network managers must plan how to support SDLC and LLC2 environments separately within the multiprotocol internetwork.

## Transporting SNA Traffic over LLC2

LLC2 traffic is typically transported by a source route bridge network, but any type of bridged network can be used. IBM's endorsement of the source route bridging methodology has made it the de facto standard for transporting SNA over a LAN/WAN network. Using LLC2 and source route bridging is not without its shortcomings, however.

Large networks are susceptible to source route bridge *broadcast storms*. Broadcast storms occur when many LLC2 devices attempt to establish connections at or about the same time. Each LLC2 device sends a message (TEST command with all routes broadcast) to all parts of the network in an attempt to determine the location of the device it is connecting to and to determine the path it will use through the network.

Large networks must also accommodate the volume of *keepalive* messages that must traverse the network. Keepalive messages are sent periodically to ensure that the LLC2 connection is still active.

Even with these shortcomings, LLC2 and source route bridge networks are the best understood and most reliable methodologies installed today to transport SNA LAN traffic.

Internetworking vendors have responded by providing source route bridging as one of the capabilities of a multiprotocol router. By providing source route bridging functionality, these vendors have enabled network managers to install a single LAN/WAN backbone that bridges SNA LLC2 traffic and routes most other nonSNA traffic. This integrated approach satisfies the immediate needs of many network managers by enabling a single physical network to transport all their LAN/WAN traffic.

### Response Times and Network Delays

Traditional SNA networks are noted for their deterministic qualities. Response time variations can generally be predicted reasonably accurately. However, as LLC2-based networks grow, several factors come into play that make these networks less predictable than their SDLC counterparts.

Large source route bridged networks or multiprotocol router networks are inherently not deterministic. Excessive delays can periodically result because of the bursty nature of the data traffic. These delays can cause LLC2 *timeouts* if messages are not responded to within a certain time. These timeouts terminate any active SNA sessions for that device. When LLC2 data is transported over very large router networks, the LLC2 timeout problem can easily become exacerbated as network loading and network delays become even less predictable.

The LLC2 protocol has timers that can be configured to allow network managers to tune the network and its attached devices for optimum operation. That's the good news. The bad news is that most available LLC2 products have default timer values which assume that the product will connect to other LLC2 devices on either the same LAN segment or over a relatively small LAN/WAN network. If these default values are used when connecting the product to a large, complex LLC2 network, timeouts may occur. An obvious solution to the timeout problem is to reconfigure every LLC2 device with new timer values as necessary to accommodate changes in the network. This can be a formidable and tedious task, especially when many of the devices are in remote locations.

### Local Termination to the Rescue

Internetworking vendors have introduced solutions for solving many of the LLC2 transport problems. These solutions employ a technique known as *LLC2 local termination*. (The terms local acknowledgement and local termination can be used interchangeably.) Local termination alleviates the timeout problem and lessens the bandwidth required for keepalive messages. It achieves these benefits by terminating LLC2 sessions in the router on the local LAN segment and then transporting the LLC2 data across the internetwork encapsulated within a routable protocol. However, local termination itself does not necessarily include local response to explorer packets nor termination of the route information field (RIF) count—these three features can each be implemented separately.

IBM includes local termination in its data link switching (DLS), which is a standard feature of the 6611 router (see *SNA Perspective*, August 1992, for a detailed discussion of DLS). Cisco's implementation is called LLC2 Local Acknowledgement. Other vendors have stated their intent to provide this feature. *SNA Perspective* expects additional announcements prior to the Interop show in October.

### *To Bridge or to Route SNA Traffic*

A long-standing debate in the internetworking community has centered on the question "Should I bridge or should I route my data traffic?" Popular opinion favors the use of routable protocols, where possible, to gain the most flexibility in network design and achieve the highest levels of network reliability. On the other hand, independent tests carried out by companies such as InterLAB of Sea Girt, New Jersey, indicate that bridged environments generally perform at a much higher raw performance level than routed environments. Furthermore, any bridge or bridge/router vendor's performance figures must be closely scrutinized when trying to predict how that vendor's products will perform in a particular application environment. *SNA Perspective* acknowledges that *caveat emptor*—let the buyer beware—becomes the key guideline for network managers when making any multiprotocol backbone decision.

The question for SNA traffic is "Should I bridge my LLC2 data traffic or should I encapsulate it within a routable protocol (at layer 3)?" Native SNA routing is still only possible by using the defined IBM mechanisms of the network control program (NCP) with the 37xx communication controller or by using advanced peer-to-peer networking (APPN). The local termination functions of IBM and Cisco both encapsulate the SNA data in TCP/IP to transport it.

Even though TCP/IP is used, each vendor has its own proprietary encapsulation implementation. The result is that, for local termination, IBM routers only interoperate with IBM routers, and Cisco routers with Cisco routers. Any TCP/IP router can be used as an intermediate node in the network.

## *Local Termination*

An LLC2 session connection exists between two endpoints of a network. When local termination is used, the router provides a RECEIVER READY (RR) response to the LLC2 traffic locally, thereby locally terminating the session. The router network then guarantees that data traffic will arrive at its intended destination. Figure 2 shows the logical connections for both a bridged environment and a local termination environment.

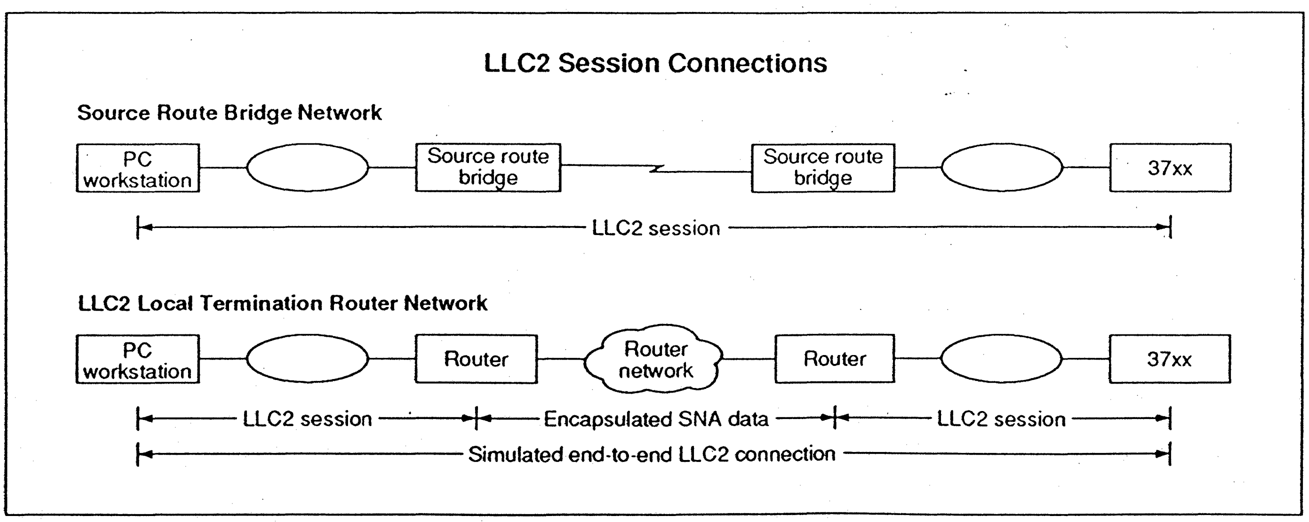


*Figure 2*

In Figure 2, note that, in the local termination case, there are actually four logical connections:

- The simulated end-to-end connection (the "circuit" in IBM terminology) between the end-user device and the 37xx communication controller (which comprises the next three items)
- The logical LLC2 connection between the end user workstation and its local router
- The logical connection between the two routers, with LLC2 data encapsulated within a routable protocol
- The logical LLC2 connection between the 37xx and its local router

Having to use three logical connections to establish each circuit seems like a lot of overhead, especially when the network may be supporting thousands of LLC2 devices. However, *SNA Perspective* believes there is an opportunity for large networks to gain benefits that should adequately outweigh any overhead incurred. Local termination:

- Eliminates LLC2 timeouts—without the need to modify default LLC2 timer settings
- Eliminates LLC2 keepalive messages traversing the network, conserving network bandwidth
- Allow vendors to take advantage of their value-added router technology capabilities

### *3270 Example*

Let's examine a simple 3270 inquiry/response application in order to better understand what happens over the network. Referring to Figure 2, assume that the workstation is a personal computer running a 3270 emulation program which appears as a single PU 2 with a single logical unit (LU) session. The normal sequence of events would be the following:

1. The user invokes the 3270 emulation program.
2. The emulation program attempts to connect to the 37xx by issuing a local TEST command to the 37xx's MAC address.
3. If the 37xx is not on the local LAN segment, then the PC will issue a broadcast TEST command to the 37xx's MAC address.
   - In an source route bridged network, this frame traverses all paths of the internetwork in an attempt to locate the 37xx. The Test Response from the 37xx is returned, acknowledging successful delivery and establishing the path through the source route bridged network for this connection. This is the process that can cause broadcast storms when many devices attempt to establish connections simultaneously.
   - In the router network using local termination, the broadcast TEST frame is intercepted by the local router and the Test Response is returned by the local router *if the location of the 37xx is known by the local router.* The process by which the location of the 37xx is known and whether or not the 37xx is active will vary from vendor to vendor. If the location of the 37xx is not known, then the router network must determine its location and find out whether it is active prior to responding.
4. Once the connection is established, the 37xx and PC continue to exchange messages throughout the duration of the LLC2 connection.
   - On an source route bridged network, each data frame and each RR travel "end-to-end."
   - On the router network utilizing local termination, RRs are generated by the local routers to acknowledge each RR or data frame generated by an end station. Guaranteed delivery of data frames is the responsibility of the router network. Figure 3 (see page 6) shows the typical flow of data for the LLC2 local termination scenario.
5. Periodically, the PC and the 37xx will issue LLC2 keepalive messages.
   - On an source route bridged network, each keepalive message (an RR) and each RR response travel end-to-end.
   - On the router network utilizing local termination, the keepalive messages are intercepted by the local router and stopped from traversing the network since the router network "knows" whether the connection is still active.

### *Congestion Control*

Two types of network congestion control can be used when implementing multiprotocol router networks with local termination:

- Flow control between routers specific to an individual TCP/IP (or other encapsulation protocol) session
- Flow control between routers specific to supporting the local termination function

Today, router vendors must address the first item in terms of its relationship to the nonSNA data on the internetwork. In addition, those vendors supporting local termination should also consider the second item. "Data Link Switching on the 6611" in *SNA Perspective,* August 1992, discusses IBM's implementation of these two mechanisms. Cisco provides the first type of flow control listed, however, and not the second type. The second type alleviates congestion more quickly, but it also throttles back all sessions, not just the congested ones.

Network managers must thoroughly understand how each of these congestion control mechanisms operate and interact. Otherwise, they will not be able to adequately plan how their SNA and nonSNA data traffic will affect each other on the multiprotocol internetwork. And, as before, each vendor's implementation will likely be unique, making the evaluation process more complex.

### *Opportunities to Excel*

Router vendors differentiate themselves by offering products that combine unique functions with a range of performance, capacity, and price options. IBM and Cisco will have functions that enhance the use of local termination, and *SNA Perspective* expects other vendors to offer solutions that provide local termination coupled with their own unique functions. Examples of the areas that are likely to be addressed include the following:

- Traffic prioritization by network address or "class of service" parameters
- Limited emulation of 37xx/NCP-type functions, such as transmission groups and class of service
- Improvements in maintaining SNA sessions despite WAN link outages on the router network
- Improved NetView visibility, via SNMP-to-NetView gateways and native NetView service access points within the router network
- Concentration of SNA traffic to conserve bandwidth on the router network
- Incorporation of APPN network node functionality tightly coupled with unique router features

The availability of these types of enhancements will be key in convincing many network managers to implement local termination solutions. These network managers will be seeking ways to tailor the

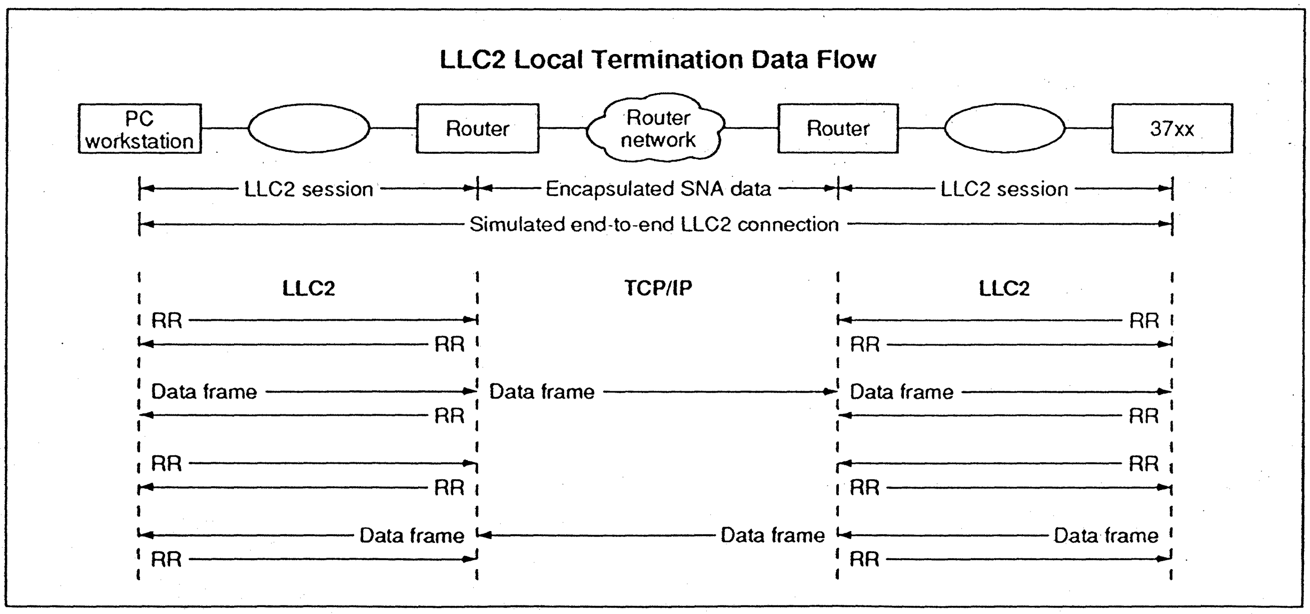


*Figure 3*

SNA aspects of their multiprotocol networks to more closely resemble the "predictable" SNA networks they manage today. Vendors providing a high level of "SNA value-add" will reap the rewards.

### Local Termination Availability

IBM's DLS function is part of the initial release of the 6611 scheduled for this month. Cisco's Local LLC2 Acknowledgement is also due to ship this month. Other vendors have stated their intent to provide similar types of functionality. Since local termination technology is still in its infancy, *SNA Perspective* recommends that users examine any vendor's local termination offering carefully to understand how it will affect their entire multiprotocol enterprise network.

## Transporting SDLC Traffic

Transporting SDLC traffic across multiprotocol networks provides significant benefits. SDLC line costs can be eliminated and end-user response times improved. Since SDLC lines make up the largest portion of today's networks, there are tremendous opportunities for both users and vendors to benefit.

There are three alternatives for transporting SDLC traffic across multiprotocol networks:

- SDLC passthrough
- Standalone SDLC-to-LLC2 conversion
- SDLC traffic conversion in routers

### SDLC Passthrough

SDLC passthrough replaces a single SDLC WAN link that connects two SDLC devices with a pair of routers operating over a multiprotocol internetwork. All of the SDLC traffic, including polling, is transmitted across the router network encapsulated within a routable protocol, typically TCP/IP. The primary motivation for using SDLC passthrough is to save the cost of a point-to-point WAN link for an SDLC device at a site where a router is located. Figure 4 illustrates a typical network scenario for using SDLC passthrough. Although this approach requires no changes to the 37xx/NCP, performance is typically less than when using a native SDLC WAN link and timeouts can occur if the router network is heavily loaded. The SDLC device is still seen by NetView, but the intervening network may not be.

SDLC passthrough is currently provided by Cisco (SDLC Tunnelling), Proteon (SDLC Relay), Wellfleet (SDLC Passthrough) and others, with additional internetworking vendors also planning to release this capability (see *SNA Perspective*, October 1991).

### Standalone SDLC-to-LLC2 Conversion

SDLC-to-LLC2 converters enable SDLC devices to connect to a LAN and appear as native LLC2

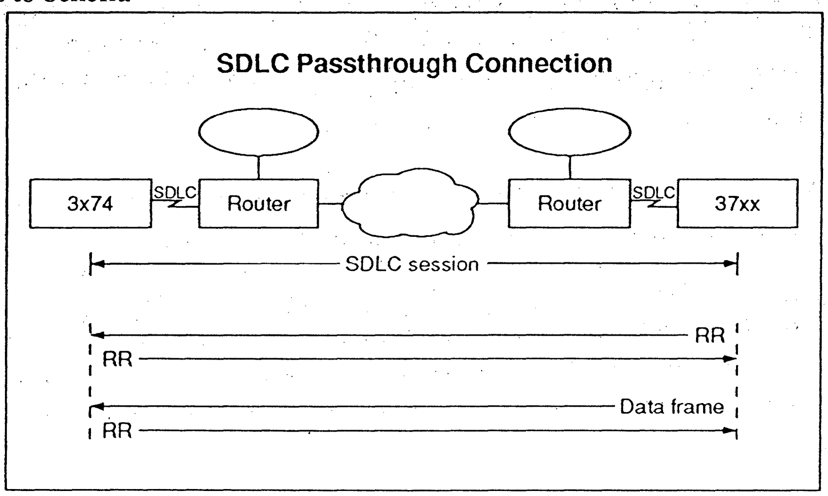


*Figure 4*

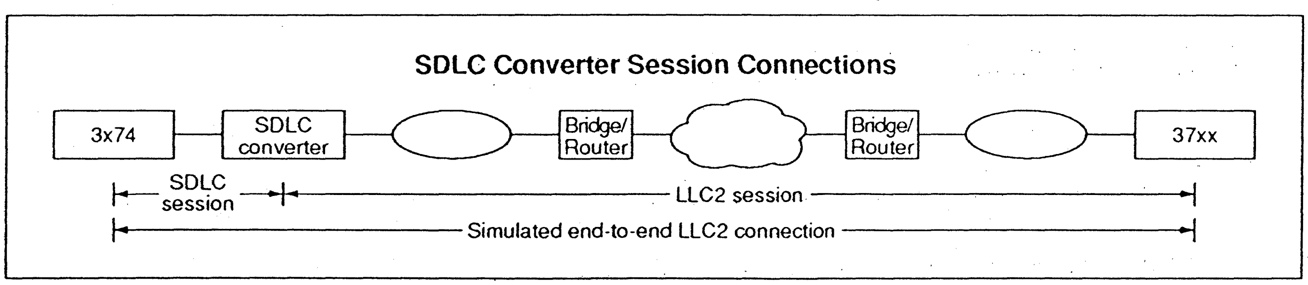


*Figure 5*

devices. These converters can be thought of as providing an alternative to upgrading a 3174-x1R model control unit to a 3174-x3R. The converter appears to attached SDLC devices as a 37xx (PU 4) and performs all the required link activation and polling functions. The SDLC devices attached to the converter then appear to the host 37xx as native LLC2 devices.

Figure 5 (see page 7) illustrates how the converter appears in a network. Figure 6 shows the typical flow of traffic using an SDLC converter.

SDLC-to-LLC2 converters provide a number of benefits over passthrough solutions:

- They save the cost of the SDLC WAN links and, in many cases, allow downsizing or elimination of 37xxs, especially at remote locations.
- They generally increase performance. Independent lab tests have shown conversion products to increase throughput between the SDLC device and the 37xx by as much as three times while passthrough degrades performance by approximately twenty percent.
- They are generally manageable directly by NetView.
- They work with any bridge/router technology, including those that utilize local termination.

SDLC-to-LLC2 converters do not require any special considerations on the LLC2 network. However, as additional LLC2 traffic is added to any network, the issues with broadcast storms and keepalive messages outlined earlier should be considered.

Standalone SDLC conversion vendors are Netlink, Ring Access, and Sync Research. These vendors have positioned themselves as SNA internetworking suppliers. Netlink OEMs its products to Apertus Technologies and Vitalink/Network Systems Corporation. Sync Research OEMs its products to McDATA. All three provide SDLC conversion to LLC2 on token ring LANs. Netlink and its OEMs currently provide Ethernet LAN support, as well. Sync Research has announced that Ethernet support will be available this month.

### SDLC Conversion in Routers

IBM and Cisco also provide SDLC conversion in their router products, both using local termination. This support is included in IBM's DLS and Cisco's offering is called SDLLC. Both of these products differ from the standalone converters in several respects:

- The SDLC conversion is performed within a router platform, with the SDLC devices attaching directly to a WAN port of the router.

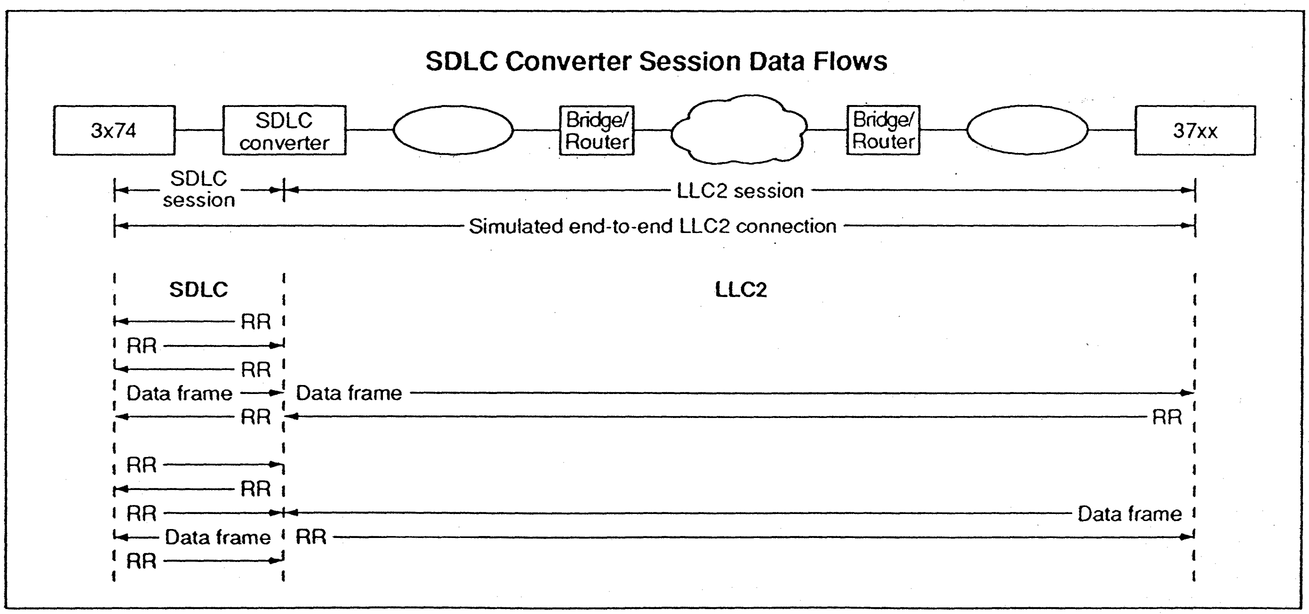


*Figure 6*

- The resulting LLC2 frame is sent as source route bridging encapsulated in TCP/IP (internal to the router) and transported across the router network. As with any source route bridge traffic on routers, the entire router network is treated as a virtual ring and counted as one hop for bridging the LLC2 frame.
- The routers are managed directly by SNMP rather than NetView. IBM uses NetView/6000 and Cisco uses NetCentral to provide a NetView service point.

The traffic flow using local termination of SDLC traffic is similar in concept to that of local termination. Figure 7 illustrates a network diagram and the typical flow of data.

### SDLC-to-SDLC Traffic with Local Termination

Local termination can also be used for SDLC-to-SDLC traffic. IBM's DLS does not support this but other vendors are developing it. As with other protocol combinations, SDLC-to-SDLC local termination removes polls and eliminates the problems with timers. In addition, it should provide a significant performance improvement over SDLC passthrough.

## Issues With Local Termination for LLC2 and SDLC

Local termination technology for both LLC2 and SDLC is still emerging, with no standards available to guide vendors or users as to the "right" way to implement these products or build networks using them. *SNA Perspective* believes that this technology will stabilize during the first half of 1993 as more vendors enter the marketplace and users determine which implementations really work well in practice. For now, users should build pilot networks to evaluate local termination solutions and compare them against using available source route bridging solutions in a multiprotocol environment. Particular areas users should examine include the following:

- How does overall network performance compare for the source route bridging solution versus the local termination solution?
- How does use of local termination for LLC2 and/or SDLC impact overall router performance?
- What value-added features are available (e.g., network tuning, traffic priority, congestion control)?

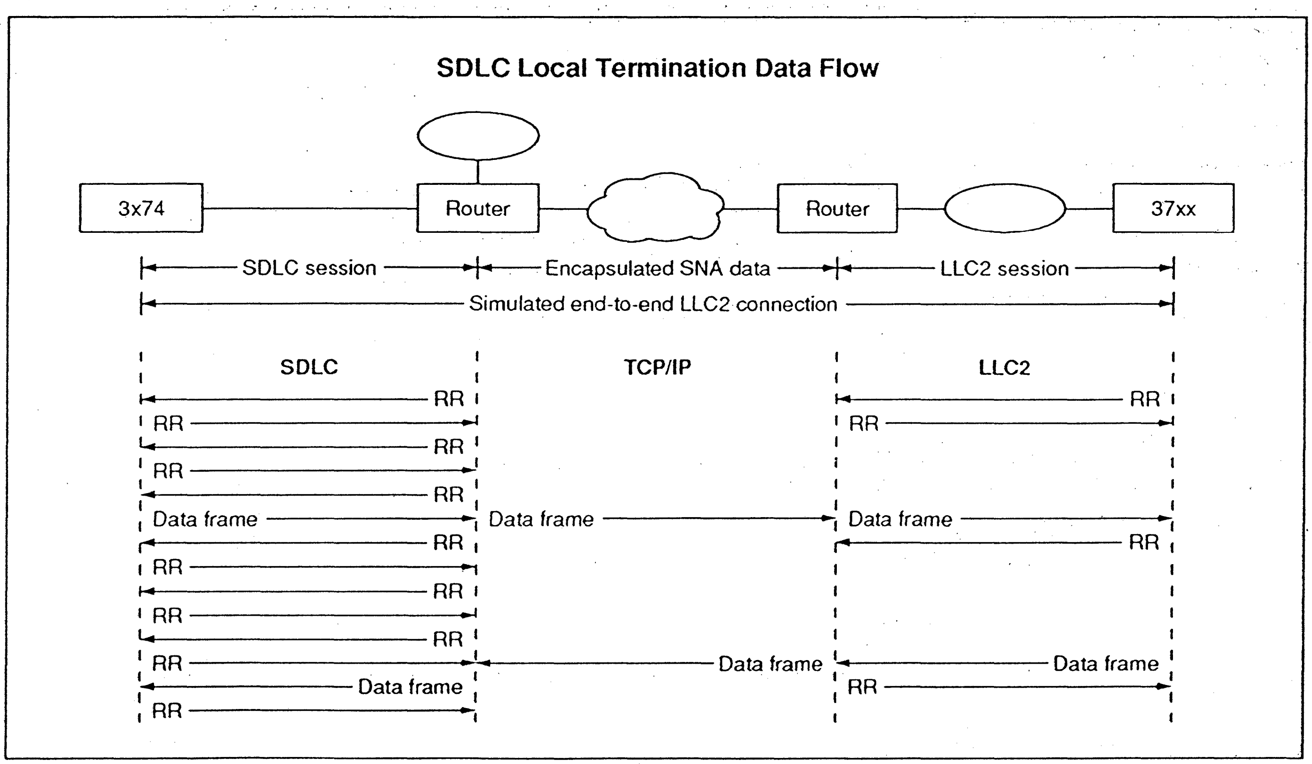


*Figure 7*

## Integrating the Network

Let's restructure the network shown in Figure 1, assuming that (1) the enterprise backbone is a multiprotocol router internetwork and (2) the objective is to have only one network connection to each site. Figure 8 shows one way to implement the multiprotocol internetwork. The primary decision points are:

- Should local termination be used instead of or in conjunction with source route bridging? As stated earlier, *SNA Perspective* recommends a close look at this new technology prior to implementing it on a network-wide basis. Even if local termination is deployed, many networks will likely benefit from using source route bridging for selected parts of the network.
- Which method of PU 2 SDLC transport should be used? *SNA Perspective* believes that standalone SDLC-to-LLC2 converters generally provide the most overall flexibility. This approach (see Remote site 1) is likely to provide the best overall cost/performance, can be implemented today with source route bridged networks, should work well with any of the local termination solutions, and eliminates lock-in to any router vendor. For sites with limited SDLC traffic and sites where some performance degradation can be tolerated, the SDLC passthrough or termination solution (see Remote site 2) may be acceptable. In these cases, however, users need to consider vendor interoperability issues.
- Should 37xx-to-37xx traffic be transported over the internetwork? Figure 8 shows an SDLC line being used to connect the 37xxs. SDLC passthrough solutions or the LLC2 transport solutions (bridging or local termination) can also be used to transport this type of traffic. *SNA Perspective* believes that users must carefully evaluate the performance of any of these internetwork approaches before implementating them.

## What About the Impact of APPN?

*SNA Perspective* views local termination of SNA traffic and the other methods of SNA transport across multiprotocol internetworks as complementary to APPN, especially through 1993. Since APPN will only accommodate peer-to-peer (node type 2.1) traffic for the near term and the majority of current SNA traffic is PU 2, these multiprotocol solutions can be viewed as safe investments in technology.

Many of the internetworking vendors who would likely implement local termination technology have already stated their intention to support APPN by using IBM licensed code for future releases of their products. This approach should enable network managers to use APPN solutions in conjunction with any other SNA solutions provided by the router vendors.

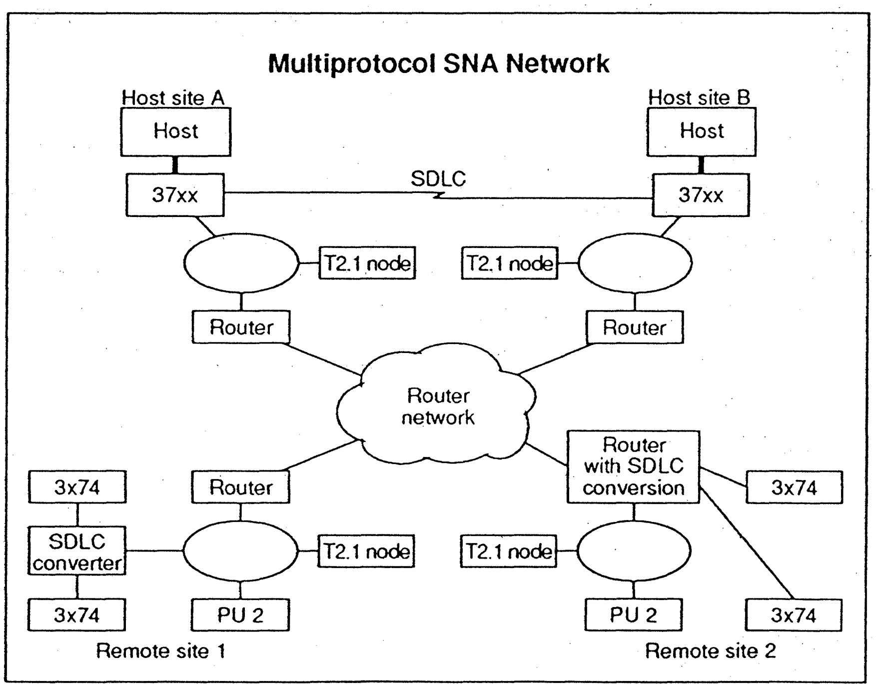


*Figure 8*

IBM's detailed plans for APPN solutions that support PU 2 devices are not clear, and probably will not be until late 1993, with delivery not likely until 1994. Network managers should be able to use any of the SDLC transport methods now and still be able to migrate those devices to APPN when desired. In fact, the SDLC transport vendors will likely have their own APPN capabilities as well.

## Summary

Implementing a multiprotocol internetwork as the enterprise backbone is a viable business decision for today's network managers. Continual advances in technology are expected over the next several years, but *SNA Perspective* believes that there is no reason to wait to begin the transition from multiple networks to a single enterprise network. Internetworking vendors and the emerging SNA internetworking vendors continue to form alliances that will speed these new SNA solutions to market.

Many options already exist for incorporating SNA as one of the many protocols that can be managed on a single multiprotocol network. *SNA Perspective* believes that these options provide credible alternatives for transporting today's LLC2 and SDLC data traffic, while remaining flexible enough to be able to implement robust local termination solutions and APPN solutions over time.

The ultimate goal of network managers is to build a manageable network that satisfies the application needs of the users—today and tomorrow. The new enterprise network must accommodate SNA as one of perhaps many protocols. As always, there is no substitute for adequate planning. Future articles in *SNA Perspective* will track the progress made in providing these SNA internetworking solutions. ■

*(continued from page1)*

## SNA over X.25

Figure 9 illustrates how IBM supports SNA sessions over the X.25 interface. The figure shows a device LU defined in a 3174 establishment controller connected in SNA session to a host LU, with an intervening PSDN. The figure could have shown several additional possible SNA over X.25 connections including between hosts, AS/400s, 5x94 controllers, System/3xs, PS/2s, and RS/6000s.

X.25 is much more popular in Europe than in the United States for several reasons including the influence of the national PTTs. Further, a much higher percentage of U.S. X.25 traffic is over public data networks than private networks. Because of these factors, a greater percentage of SNA traffic in Europe, perhaps up to 50 percent, is sent across X.25 than in the U.S., where it is probably less than 10 percent.

### X.25 Looks Like SDLC to SNA

IBM's support of SNA/X.25 provides the view to SNA path control (SNA Layer 3) that switched

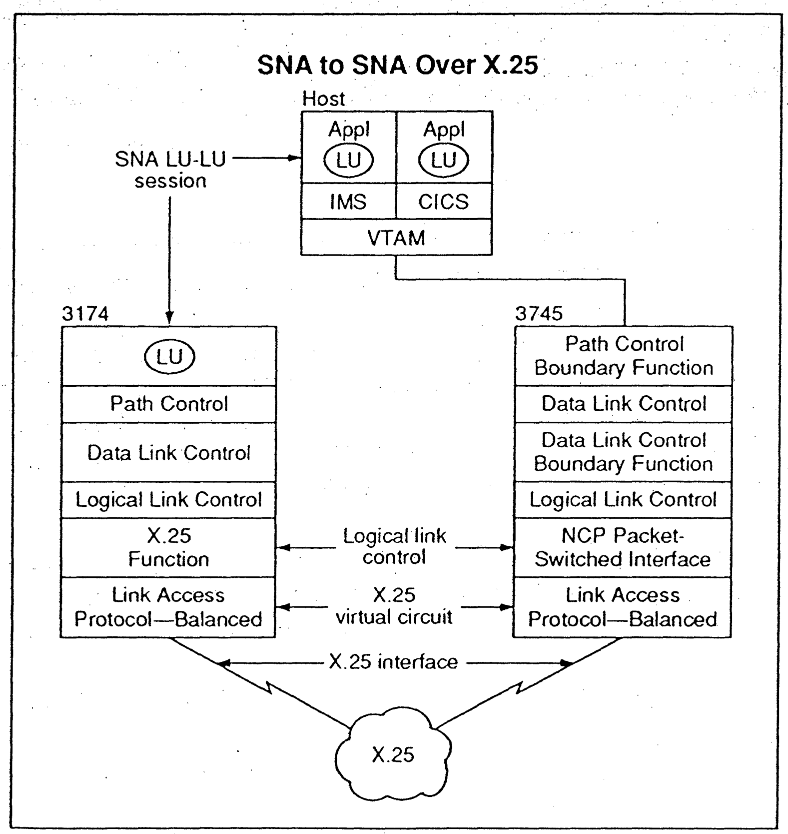


*Figure 9*

virtual circuits (SVCs) appear as switched lines and permanent virtual circuits (PVCs) appear as dedicated lines. IBM also uses logical link control (LLC) procedures to support SNA adjacent node services equivalent to SDLC functions, such as qualified logical link control (QLLC) in the 3745 and enhanced logical link control (ELLC) in the AS/400.

This reliance on SDLC procedure support across the X.25 interface leads to throughput and performance issues, particularly where the SNA/X.25 interface has been implemented via software on the 3745—Network Control Program (NCP) and NCP Packet Switching Interface (NPSI).

### *SNA Inside X.25*

Figure 10 shows how an SNA path information unit (PIU), the SNA path control layer message unit, is encapsulated in an X.25 data packet at the X.25 interface. The X.25 packet-level protocol regards the SNA PIU as user data.

## *Upcoming in X.25*

The International Telegraph and Telephone Consultative Committee (CCITT) is the standards body which develops several international "recommendations" including the X.25 family of protocols.

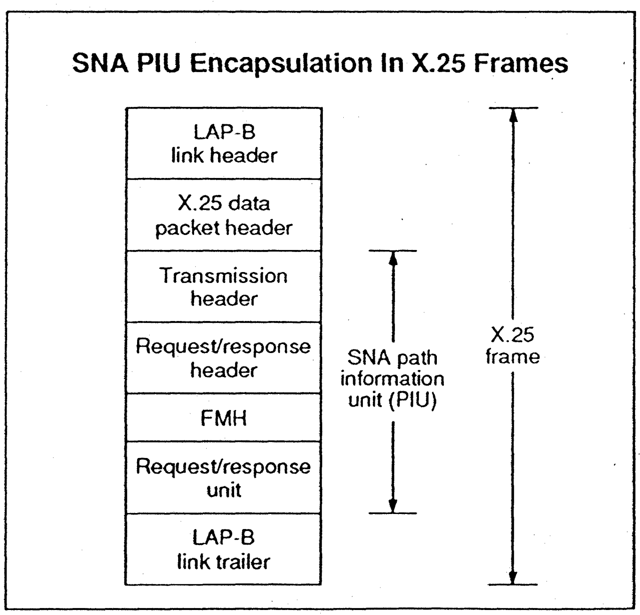


*Figure 10*

These recommendations are updated every four years. Each update is referred to by year or by the color of the cover. Some SNA-related changes in previous editions of X.25 are noted in the sidebar "X.25 Evolution." The anticipated 1992 "White Book" are still evolving at this time. However, it appears that 1992 enhancements most likely to affect SNA users relate to:

- Alternative addressing
- High-speed transmission

### *Alternative Addressing*

This is a set of alternative addressing-related facilities which enables a calling DTE to select an alternative address to identify a called DTE to establish a virtual call. Alternative addresses are those that do not conform to the formats defined in CCITT recommendations X.121 and X.301. The following alternative addresses may be supported:

- A mnemonic address according to Recommendation T.50
- An OSI network service access point (NSAP) address according to Recommendation X.213/ISO 8348 Addendum 2
- A LAN medium access control (MAC) address according to ISO 8802
- An Internet address according to Request For Comments (RFC) 877

The first two of the above represent refinements of earlier X.25 recommendations. The latter two—alternative address support for calls to LAN MAC addresses and to Internet addresses—will be of particular interest to users who want to directly interface their X.25 networks to their LANs, especially those who are already connecting their departmental LAN across multiprotocol WANs.

Definition of an Internet alternative address facility for X.25 is also important because many users have deployed TCP/IP networks in parallel with SNA and other approaches. In these cases, the 1992 Recommendation will likely allow two alternative networking options to enable X.25 DTEs to access LAN MAC addresses or Internet addresses:

- Permit a DTE to use the address block to carry any of the alternative address formats
- Allow the DTE to use the called address extension facilitiy to carry an OSI NSAP address

It is likely that definition of a standardized central directory facility to resolve alternative address formats will be part of the Recommendation X.25 (1996) study effort.

### High-Speed Transmission

The default parameters of X.25—including data link layer modulo, frame size, link window size, packet layer modulo, packet size and packet window size—are not optimized for operation over connections in which a long round-trip delay will be encountered (such as cables with long delays, or over satellite) nor for transmission rates greater than 64 Kbps. CCITT Recommendation X.25 (1992) will recommend:

- For the data link layer operating over connections with a maximum round-trip delay of 600 ms over a 64 Kbps link, use modulo-8 frame numbering with a frame size of at least 1,024 octets. For smaller frame sizes, use modulo-128 sequencing.
- For most terrestrial circuits with transmission rates of 1,920 Kbps (this speed is based on the European 1,920 Kbps H12 channel structure in PRI ISDN), the round-trip delay is observed to be approximately 1 ms, in which case modulo-8 is sufficient. Use of modulo-128 is recommended for longer round-trip delays operating at 1,920 Kbps.
- The need for running X.25 over a satellite link operating at 1,920 Kbps is not established, and is a likely study effort for the 1996 recommendation.

Though products based on these higher speed X.25 standards will not hit the market until 1994, existing subarea and Advanced Peer-to-Peer Networking (APPN) SNA sites are already experiencing bandwidth requirements exceeding the reliable 64 Kbps service provided by X.25. However, *SNA Perspective* believes that many SNA over X.25 users will wait for availability of these higher speed X.25 products rather than shift their SNA traffic onto other higher speed technologies such as Frame Relay.

## X.25 and Emerging Subnetwork Technologies

Some guidelines have been established dealing with transmission of X.25 at speeds greater than 64 Kbps. However, X.25 has traditionally supported and still primarily supports only speeds up to 64 Kbps. For a substantial part of the user base, a 64 Kbps X.25 interface will be adequate into the foreseeable future. The Integrated Services Digital Network (ISDN) basic rate interface (BRI) defines service in 64 Kbps channels and it is certainly feasible to run X.25 over BRI ISDN.

Many users and their applications, however, will require access to multimegabit and multigigabit transmission rates. This is as true in the SNA arena as in other environments. In the past few years, several technologies have emerged and are being refined that directly address requirements for higher bandwidth. These include Frame Relay, Switched Multimegabit Data Service (SMDS), ATM, Fiber Distributed Data Interface (FDDI), IEEE 802.6 Metropolitan Area Network (MAN) with Distributed Queue Dual Bus (DQDB), and Synchronous Optical NETwork (SONET).

### High Bandwidth Benefits

Users are increasingly attracted to high bandwidth technologies for two fundamental reasons:

- Client/server computing
- Transmission economics

### Client/Server Computing

A client/server computing paradigm has begun to replace host-centric networking. Users and their applications are now defined at the department level and must interconnect and share resources either with other such applications or with enterprise server applications. Personal workstation price-performance improvements are accelerating to the point where we are likely to witness the emergence of 64/64-bit supercomputers on the desktop by the end of the decade (and, perhaps, laptop, palmtop, and even wristtop—the latter likely to fulfill Dick Tracy's original requirements).

# Overview of IBM X.25 Products

IBM support for SNA through the X.25 interface is now pervasive across its major host, midrange, and workstation platforms, as summarized in Table 1 (see page 15).

IBM's endorsement of X.25 has been instrumental in the growth of the industry. Since their inception, SNA and X.25 have evolved as separate networking approaches, though they have coexisted. IBM's current pervasive support for SNA connectivity over X.25 is in direct response to major user and industry requirements. These include the user need to:

- Integrate multiple, heterogeneous networks
- Use a vendor-independent "standard" interface
- Support multiprotocol network architectures

***Integrate Multiple, Heterogeneous Networks***

Many users have installed multiple and incompatible networks at the department or cost center level in their organizations. Applications and devices in these LANs and WANs need to interconnect and share resources across the resulting sea of incoherence. SNA-to-SNA, SNA-to-nonSNA and nonSNA-to-nonSNA connections are achieved in a relatively straightforward way over a range of LAN and WAN link protocols using X.25 products from IBM and other vendors.

While X.25 was originally designed to provide a packet-mode synchronous data terminal equipment (DTE) interface to a public data network, the interface is also supported across an unprecedented number and range of private networks in more than a hundred countries. User-proximate host or terminal DTEs access data circuit-terminating equipment (DCE) node processors on the network side of the X.25 interface. In this way, the DCE serves as an entry/exit point to/from a packet switching network. Data switching equipment (DSE) is the internal switching node technology in a packet-switched data network (PSDN) and is not adjacent to the end user.

Table 1 indicates the products in which IBM has implemented X.25 DTE, DCE, and DSE functionality on its enterprise hosts (System/390/370 series); 37xx communication controllers; midrange, departmental processors (RS/6000, AS/400, System/36, System/38, System/88, Series/1 and 8100); PS/2 workstations; 3x74 and 5x94 controllers; and specialized controllers (468x, 470x, 6150, 51X0). These processor platforms, in turn, may connect to either IBM or nonIBM processors at the remote end of a PSDN through the X.25 interface.

***Provide a Vendor-Independent "Standard" Interface***

This requirement follows from the above requirement. Users demand a flexible, predictable, and reliable network interface which is not vendor-specific. X.25, developed by the CCITT, is an effective means to convey packets of various data types of various architectural origins through an end-to-end network in a reliable way.

***Support Multiprotocol Network Architectures***

The X.25 interface is used with equal frequency and success to interconnect SNA, OSI, and TCP/IP as well as a wide variety of other applications and protocols. ■

**IBM X.25 Products**

| Hardware | Software | Availability | DTE | DCE | DSE |
|---|---|---|---|---|---|
| ES/9000 ICA | VTAM[1] | January 1991 | | X | X |
| 9370 ICA | VTAM[1] | September 1987 | | X | X |
| 4361 ICA | VTAM | March 1985 | | X | |
| System/390/370 | CSFI[1] | January 1991 | | X | |
| 3745 | NPSI[1,2] | June 1988 | | X | X |
| 3720 | NPSI[1] | October 1986 | | X | X |
| 3725 | NPSI[1] | November 1983 | | X | X |
| 3705 | NPSI | September 1981 | | X | X |
| 3745/3725 | XI[1] | December 1986 | X | X | X |
| RS/6000 | AIX[1] | March 1990 | | X | |
| AS/400 (9406/9404) | OS/400[1] | August 1988 | | X | X |
| AS/400 (9402) | OS/400[1] | November 1990 | | X | X |
| System/36 | SSP | May 1984 | | X | X |
| 5363 S/36-PC | SSP | October 1987 | | X | X |
| 5364 S/36-PC | SSP | June 1987 | | X | X |
| System/38 | CPF | March 1985 | | X | |
| System/88 | OS | March 1985 | | X | X |
| Series/1 | RPS | July 1983 | | X | X |
| 8100 | DPPX | January 1986 | | X | |
| PS/2 | X25Net Switch[1] | December 1991 | X | X | X |
| PS/2 | X25Net Manager[1] | November 1991 | | | |
| PS/2 | PNA[1] | June 1990 | | X | X |
| PS/2 | OS/2 EE[1] | March 1990 | | X | |
| PS/2 | DOS[1] | June 1989 | | X | |
| 3174 | μcode[1,2] | July 1987 | | X | |
| 3274 | μcode | July 1984 | | X | |
| 5394 | μcode[1] | August 1988 | | X | |
| 5294 | μcode | May 1984 | | X | |
| 4680 | μcode[1] | September 1990 | | X | |
| 4684 | μcode[1] | September 1990 | | X | |
| 4702 | μcode | October 1985 | | X | |
| 4701 | μcode | September 1983 | | X | |
| 6150 | AIX[1] | September 1988 | | X | |
| 5150/60/70 | DOS[1] | April 1986 | | X | |

1 Supports 1984 CCITT X.25.
2 Supports 1988 CCITT X.25 and ISO 7776/8208.
AIX - Advanced Interactive Executive
CCITT - International Telegraph and Telephone Consultative Committee
CPF - Control Program Facility
CSFI - Communications Subsystem for Interconnection
DCE - Data Circuit-Terminating Equipment. Network side of the interface; also used in gateways between networks.
DOS - Disk Operating System
DPPX - Distributed Processing Programming Executive
DSE - Data Switching Equipment. Switches traffic between and among multiple user DTEs.
DTE - Data Terminal Equipment. User side of the interface.
ES - Enterprise System
ICA - Integrated Communications Adapter
ISO - International Organization for Standardization
NPSI - NCP Packet Switching Interface
OS/2 EE - Operating System/2 Extended Edition
OS/400 - Operating System/400
PNA - Programmable Network Access
RPS - Realtime Programming System
RS - RISC System
SSP - System Support Program
VTAM - Virtual Telecommunications Access Method
XI - X.25 SNA Interconnection

*Table 1*

Applications are already defined which require this level of scalable, numerically-intensive computing. Desktop and laptop 32/32-bit technology is already prevalent. Once the exclusive province of scientific and engineering processing, these technologies are now running numerically-intensive spreadsheet, database, and decision-support programs required by enterprise decision makers. Their need for interconnection, often interactively, demands reliable bandwidth far higher than 64 Kbps.

### *Transmission Economics*

Optical fiber is becoming more prevalent, particularly in the United States, leading to both economies of scale and significant bandwidth supply outstripping demand. These are both resulting in dropping prices for higher speed services. Therefore, users who can cost-justify T1 at 1.544 Mbps will soon be able to nearly as easily build a positive business case to subscribe to T3 links at 44.73 Mbps. European E1/E3 relationships are quite similar, though it will happen somewhat more slowly. Both the higher-speed X.25 standards discussed above and emerging technologies such as Frame Relay can run over these high bandwidth environments.

Following are brief descriptions of Frame Relay, SMDS, and ATM technologies with contrasts to X.25, particularly as they relate to SNA.

### *Frame Relay and X.25*

Fundamentally, Frame Relay is statistical multiplexing over a shared network. Frame Relay uses a variable-length frame to relay at Layer 2 or Layer 1, rather than at Layer 3 as X.25 is sent. Whereas X.25 networking requires a considerable degree of processing at each node in transit between the origination and destination, Frame Relay performs its switching function using the physical layer and a subset of the link layer.

Further, Frame Relay frames contain routing information fields not previously located in Layer 2 in X.25, which in effect eliminates the need for a Layer 3 in the network. Frame Relay switches examine and route frame header fields called the data link connection identifier (DLCI) at the network entry point.

Frame Relay is also designed to interconnect hierarchical as well as peer-to-peer architectures over high speed facilities and to offer bandwidth-on-demand for bursty traffic. Because of this, it is well suited to both subarea SNA and APPN. IBM is therefore strongly committed to Frame Relay. The company announced Frame Relay interfaces for the 3745 in 1991 and is expected to add Frame Relay switching support for the 3745 in the future. Therefore, SNA users should consider Frame Relay in their long-term SNA plans.

### *Fast Packet Technologies and X.25*

SMDS, IEEE 802.6 DQDB, and ATM technologies are often referred to as "fast packet" technologies in order to contrast them from X.25 packets.

**SMDS**—SMDS is a high-speed digital data service that provides packet-switched customer access. SMDS network switches will be interconnected by 45 Mbps trunks. (Plans exist to migrate SMDS into the SONET transmission hierarchy to the OC-3 rate of 155 Mbps.) SMDS uses makes use of statistical multiplexing techniques to allow multiple applications to share the same access line. Also, SMDS is a connectionless transmission scheme and requires no call setup or call takedown processes.

**802.6 MAN**—SMDS uses a cell relay architecture based on the connectionless data networking functionality of the IEEE 802.6 MAN standard. IEEE 802.6 defines segmentation of incoming data into fixed-length, 53-byte cells for packetized data transfer through network cell relay switches. The IEEE 802.6 standard also describes the use of DQDB with a dual bus to transfer packets between cell relay switches. The IEEE 802.6 Committee is articulating an isochronous SMDS specification that will support packetized voice and video communication. The SMDS migration toward handling traffic other than data and image is a move toward ATM.

**Broadband ISDN**—Broadband ISDN is a connection-oriented packet service based on the use of fixed-length ATM cells. The ATM protocol is intended for switching and transport of digital data, voice and video signals simultaneously. The ATM and IEEE 802.6 cell definitions have been matched, with each containing a 5-octet address header with 48-octet user data field (payload) for a total of 53 octets per cell. ATM network access speeds are planned to range from 155 Mbps to 2.4 Gbps.

X.25 was designed to transfer variable-length data packets over 1960s and 1970s switching and transport technologies. X.25 switches examine each arriving packet and provide confirmation of correctness back to the originating node. In contrast, SMDS, DQDB and ATM "relay" cells or packets through network nodes in a fraction of the time required for X.25 switching and in a less complex and more cost-effective way (from the perspective of the network service provider) over predominantly optical facilities with nearly error-free conditions.

### *SNA with X.25 and the New Subnets*

*SNA Perspective* believes that, while these new subnet technologies are more efficient, most SNA over X.25 users will stay with X.25 at least through the end of 1995, for two reasons—fast packet technologies will not be sufficiently shaken out until 1995 and the SNA over X.25 base must be more fully capitalized before it is replaced. Although all these technologies offer economies of scale and can support multiple protocols including SNA, they can also create problems of time-consuming convergence, variable throughput, and error recovery in connectionless networks, all of which can particularly impact SNA traffic.

IBM seems strongly committed to Frame Relay and more reluctant with respect to SMDS. This is probably because SMDS is a public service while, with Frame Relay, a user can either set up a private Relay network or attach to a public Frame Relay service. IBM, which has a strong market position in corporate backbone infrastructure, is more likely to support a network for which it can provide both interfaces (DTE) and switches (DCE).

## NCP Passage Into History and NPSI Impact

Most SNA traffic over X.25 is to hosts front-ended by NPSI which runs on 37xxs.

### No "3765," No Paris, But...

There has been speculation for some time that IBM will eventually replace the 3745 and also NCP. IBM has denied that a rumored "3765" would replace the 3745 in the near term. However, IBM has indicated that it is developing a follow-on to its oft-discussed internal prototype Paris switch technology which we believe would supersede the 3745. This would not be announced for several years, however, and we believe IBM would support the 3745 and NCP for some time after that. Until then, IBM says the 3745 and NCP remain on vigorous upgrade cycles. Several recent enhancements, for example, include Frame Relay and Ethernet support and recent IBM statements indicate additional near-term announcements.

### SNA and APPN over X.25 and Frame Relay

With regard to X.25 specifically, IBM's focus is to replace SNA with APPN and to replace SDLC and X.25 with Frame Relay in WANs. The user of SNA over X.25 considering transition to new technologies, therefore, needs to consider a migration path that may include SNA over X.25, APPN over X.25, SNA over Frame Relay, and APPN over Frame Relay.

IBM's announced strategy for APPN states that NCP will only participate in APPN as part of composite network nodes (CNNs) in conjunction with a host running Virtual Telecommunications Access Method (VTAM) (see *SNA Perspective*, April 1992). That is, NCP will not be a APPN router independent of the host. This suggests to us, given IBM's commitment to APPN, that NPSI will not be the primary vehicle to support APPN over X.25.

### IBM Likely to Support APPN over X.25

However, IBM has supported the composite LEN node (node type 2.1) over X.25 in NPSI for some time. *SNA Perspective* believes that IBM will likely introduce a new level of NPSI in 1992 and we expect this new release to support VTAM/NCP CNN APPN interfaces over X.25. However, given both the existing performance limitations of NPSI and IBM's focus on Frame Relay for APPN, this NPSI APPN-over-X.25 support is likely to be a serviceable but low-performance implementation.

On the other hand, X.25 is one of the fundamental subnetwork interfaces defined in the IBM networking blueprint (see *SNA Perspective*, August 1992). Therefore, *SNA Perspective* believes that IBM, in the long run, will support APPN over X.25 as an important mainstream protocol stack, especially in Europe, though probably with something other than NPSI.

## LLC "Glue" for SNA over X.25

For transporting SNA over X.25, IBM uses several logical link control (LLC) procedures—qualified LLC (QLLC), enhanced LLC (ELLC), and physical services header LLC (PSHLLC; the oldest LLC technique)—to act as a "bridge" between SNA and X.25. The reader should note that this is a different LLC than that used in referring to LANs. These LLC procedures are needed because X.25 does not natively provide several functions (SNA adjacent node physical services) which SNA expects SDLC to provide. These functions include:

- Operational mode selection (SNRM, SABM)
- Identification information exchange (XID)
- Link test (TEST)
- Link disconnection (DISC)

These three LLC types were developed separately for different IBM products. PSHLLC supports older 3274 cluster controllers. QLLC is supported on the majority of IBM products including NPSI on the 37xx, 3174, ES/9000, and S/88. ELLC is supported primarily on the older S/3x and current AS/400 family. Figure 11 on page 21 illustrates VC Type 2 with PSHLLC, VC Type 3 with QLLC, and VC Type 6 with ELLC.

PSHLLC is used in earlier implementations, usually where NPSI communicates with a remote IBM 5973 Network Interface Adapter (NIA) which connects over SDLC to the 3274. As indicated by the name, physical services headers are inserted in front of SNA PIUs and perform LLC functions in place of SDLC across the PSDN.

Instead of using physical services headers, QLLC uses the qualifier (Q) bit in X.25 data packets to identify unnumbered and supervisory QLLC commands and responses. QLLC uses HDLC unnumbered commands and receive ready (RR) supervisory commands and responses equivalent to their SDLC counterparts. QLLC link stations can be primary, secondary, or balanced (peer-to-peer). Where PSHLLC supported only dumb terminal-to-host connections, QLLC can support host-to-host connections.

ELLC is an enhancement developed for the S/3x and was migrated to the AS/400. ELLC provides error detection and optional BIU retransmission recovery. ELLC is user-selectable on a virtual circuit basis but the optional retransmission capability applies to the DTE/DCE interface as a whole. ELLC formats have been expanded with IBM's support of 1988 X.25 to include an adaptation of the check-sum data integrity mechanism defined in ISO 8073 Transport Protocol Specification. The check-sum detects modified or missing packets and ELLC can recover at the message level. QLLC and ELLC are not compatible so during link-level negotiation the two devices select the less functional QLLC.

IBM does not presently support ELLC in NPSI. *SNA Perspective* believes that while ELLC is functionally preferable to QLLC, it would generate unacceptable overhead if implemented in NPSI.

Table 2 (see page 18) provides an overview of QLLC, ELLC, and PSHLLC and the SDLC functions they provide. LLC functional SDLC flows are conveyed in qualified data packets. While these LLC procedures act as the "bridge" or "glue" between SNA and X.25, the need to transmit and receive a series of qualified data packets in addition to data packets containing SNA PIUs generates additional overhead.

*SNA Perspective* believes that either a direct integration of X.25 into the SNA architecture or an elegant use of X.25 to convey SNA session data, as suggested by the IBM networking blueprint, should be able to eliminate the need to send and receive SDLC sequences across the X.25 interface.

## SNA Traffic from X.25 to LAN

Many users want to directly interconnect X.25 WANs and LANs for SNA and other traffic. Frequently, these LAN MAC protocols are either token ring or Ethernet and, therefore, the LAN upper layer 2 sublayers are IEEE 802.2 LLC. It is technically feasible to directly interconnect X.25 WANs and IEEE 802.2 LLC LANs because they are both based on the same HDLC asynchronous balanced mode (ABM; Modulo-8) and ABM Extended (ABM-E; Modulo-128) functional subsets.

Sync Research of Irvine, California, supports direct interconnection of SNA traffic on X.25 WANs and the token ring using its SNAC/TRQ product. This is done through conversion between the token ring 802.2 LLC protocol and IBM's QLLC protocol used for SNA traffic on the X.25 network.

For example, SNA workstations may be connected to an X.25 network either directly though an X.25 gateway device that supports QLLC, such as IBM's Programmable Network Access, or across a LAN through an X.25/QLLC LAN gateway product, such as Eicon Technology's SNA LAN Gateway. These gateways would traditionally access an SNA host across the X.25 network through a 37xx communication controller with NPSI. However, if the 37xx provides LAN-attachment to the host, an SNAC/TRQ can be used to intercept the QLLC traffic and convert it into LLC2 traffic which will be sent across the LAN to the 37xx or even a 3172. This eliminates the need for NPSI on the 37xx, which is one benefit.

The gateways are dynamically mapped one-to-one to host destination service access point (SAP) addresses. This gateway mapping to SAPs supports host VTAM switched major node definitions, which eliminates the need to update NCP to reflect network changes, a second benefit. The SNAC/TRQ product also provides native NetView visibility through network management flows.

IBM's data link switching (DLS) on its IBM 6611 router integrates SNA/SDLC, LLC2, and NetBIOS over multiprotocol router networks. DLS currently supports SDLC-to-LLC2 conversion and transit through an IP network. Since IP Internet traffic is sent over an X.25 WAN, the 6611 DLS indirectly provides the basis for conveying SNA traffic over the IP network with underlying X.25.

**SNA/X.25 Logical Link Controls**

| SDLC function | QLLC function | ELLC function | PSH LLC function | Primary command | Secondary response | Peer-to-peer command | Peer-to-peer response |
|---|---|---|---|---|---|---|---|
| I (Information) | UDP | LI (LLC Info) | DATA | X | | X | |
| SNRM/ SABM | QSM (Q Set Mode) | LSABME | PSCONTACT | X | | X | |
| DISC (Disconnect) | QDISC | LDISC | PSDISC | X | | X | |
| XID (Exchange ID) | QXID | LXID | PSXID | X | X | X | X |
| TEST | QTEST | LTEST | PSTEST | X | X | X | X |
| UA (Unnumbered ACK) | QUA | LUA | PSCONTACT | | X | | X |
| RD (Request Disconnect) | QRD | — | PSDISC | | X | | |
| RR (Receive Ready) | QRR | LRR | — | X | X | X | X |
| RNR (Receive Not Ready) | — | LRNR | — | X | X | X | X |
| REJ (Reject) | — | LREJ | — | X | X | X | X |
| DM (Disconnect Mode) | QDM | LDM | — | | X | | X |
| FRMR (Frame Reject) | QFRMR | LPDUR | — | | X | | X |

ACK - Acknowledgement
ELLC - Enhanced Logical Link Control
L___ - Logical ___
LPDUR - LLC Protocol Data Unit Reject
PS - Physical Services
PSH LLC - Physical Services Header Logical Link Control
Q___ - Qualified ___
QLLC - Qualified Logical Link Control
SABM - Set Asynchronous Response Mode
SDLC - Synchronous Data Link Control
SNRM - Set Normal Response Mode
UDP - Unqualified Data Packet

*Table 2*

*SNA Perspective* believes that IBM is not likely to add QLLC-to-LLC conversion to the 6611 in the near term, since other protocol support is more pressing and because this would impact sales of NPSI. IBM endorsement of third-party support is more likely.

Perhaps the most intriguing feature of this X.25-to-LAN support and similar approaches is that, if implemented without the need to notify user applications of internal network differences (e.g., X.25 or nonX.25), they begin to fulfill the user requirements for transparent multiprotocol support.

## Conclusions

X.25 networks are prevalent throughout the world and IBM support for SNA over X.25 is equally ubiquitous. The major issues associated with X.25 in today's SNA user environments are discussed below.

X.25 was designed for 1970s technologies and provides a reasonable, variable-length packet service at transmission speeds up to 64 Kbps. Use of X.25 at higher speeds has been addressed by the CCITT. However, reasonable throughput and performance has not been proven at these speeds and certainly not at the higher speeds of the emerging subnet technologies.

Products from IBM and other vendors support a range of X.25 levels (1976, 1980, 1984, 1988) and therefore do not support X.25 in a consistent way. IBM products that implement X.25 at different levels do not always support upward and downward functional compatibility.

Users will increasingly require access to very high speed transmission services, especially for high-performance, lower cost desktop and laptop technologies which lead to miniature "mainframes" on LANs using complex applications interactively. Emerging technologies such as Frame Relay, ATM, SONET, and broadband ISDN will certainly address these throughput requirements and have already been shown to effectively remove Layer 3 packet processing requirements (which X.25 uses) from the underlying network delivery infrastructure.

IBM's SNA over X.25 applications have traditionally been based on the host-centric computing model and users are moving increasingly toward client/server computing.

Since IBM has said NCP will not be an APPN network node or end node (but instead will only participate with VTAM as part of a composite network node) and because of long-standing performance concerns, the long-term future of NPSI is in question by many. Alternatives to NPSI such as the Sync Research QLLC-to-LLC converter may prove popular and IBM might add X.25 support to the 3172.

*SNA Perspective* believes that IBM will support APPN over X.25 as an important mainstream protocol stack, especially in Europe, though probably with something other than NPSI.

*SNA Perspective* believes that either a direct integration of X.25 into the SNA architecture or an elegant use of X.25 to convey SNA session data, as suggested by the IBM networking blueprint, should be able to eliminate the need to send and receive SDLC sequences across the X.25 interface.

By the end of 1995, we expect that twenty percent of the existing SNA over X.25 user base in the United States will shift to higher bandwidth and bandwidth-on-demand transport technologies in lieu of X.25.

Nonetheless, *SNA Perspective* believes that the SNA/X.25 user base will remain substantial for several years, particularly outside the United States. This is due to the fact that the existing X.25 infrastructure is well established and reliable and this reliability is needed in the majority of the world that does not have pervasive fiber-based facilities and clean circuits. Further, while higher bandwidth and bandwidth-on-demand technologies are generating widespread interest, several of these are still in developmental and field trial stages and significant production shifts will occur over several years. ■

# IBM X.25 Virtual Circuit Types

Figure 11 provides a summary of how IBM accomplishes SNA-to-SNA as well as SNA-to-nonSNA connections through its X.25-supported products. The figure illustrates PSDN connections over virtual circuit types 0, 2, 3, 4, 5, and 6. These VC types are IBM-designated; they are not defined as part of the CCITT Recommendation X.25.

VC Types 0, 4, and 5 support connections between SNA and nonSNA environments. VC Types 0 and 4 are similar in that they provide access from a nonSNA X.25 DTE to an application LU in an IBM host. VC Type 0 provides access from a nonSNA, X.25 DTE into a NPSI region called Protocol Converter for Non-SNA Equipment (PCNE). PCNE emulates SNA LU Type 1 (LU1) to the host LU to enable the latter to regard the X.25 DTE as an SNA Remote Job Entry (RJE) workstation or as a 3767 printer.

VC Type 4 supports a NPSI function called General Access to X.25 Transport Extension (GATE). GATE enables a host user application called Communication and Transmission Control Program (CTCP) to monitor virtual circuits to nonSNA X.25 DTEs by processing the contents of data packets, qualified data packets, interrupt packets, call/clear packets, reset packets, and diagnostic packets. GATE CTCPs can also be used as relay programs to application subsystems such as CICS, IMS, and TSO.

VC Type 5 provides for access from nonSNA, nonX.25 DTEs into SNA host applications through the use of NPSI functions called IPAD and TPAD. IPAD is a PCNE/NPSI extension which implements a subset of X.29 for communication with TTY 33/35 and other start-stop DTEs that conform to X.28 and in turn access the PSDN over a PAD defined by X.3 (PAD parameters).

TPAD is a PCNE/NPSI extension which enables host access from remote DTEs that are nonSNA, nonX.25, and also do not conform to the CCITT Triple-X Series (X.3/X.28/X.29). That is, the remote DTE interfaces to the PSDN through a nonstandard PAD. In both the IPAD and TPAD cases, an LU Simulator functions in NPSI to present the nonSNA DTE to the host with the appearance of an SNA LU. ■

**SNA/X.25 Virtual Circuit Types**

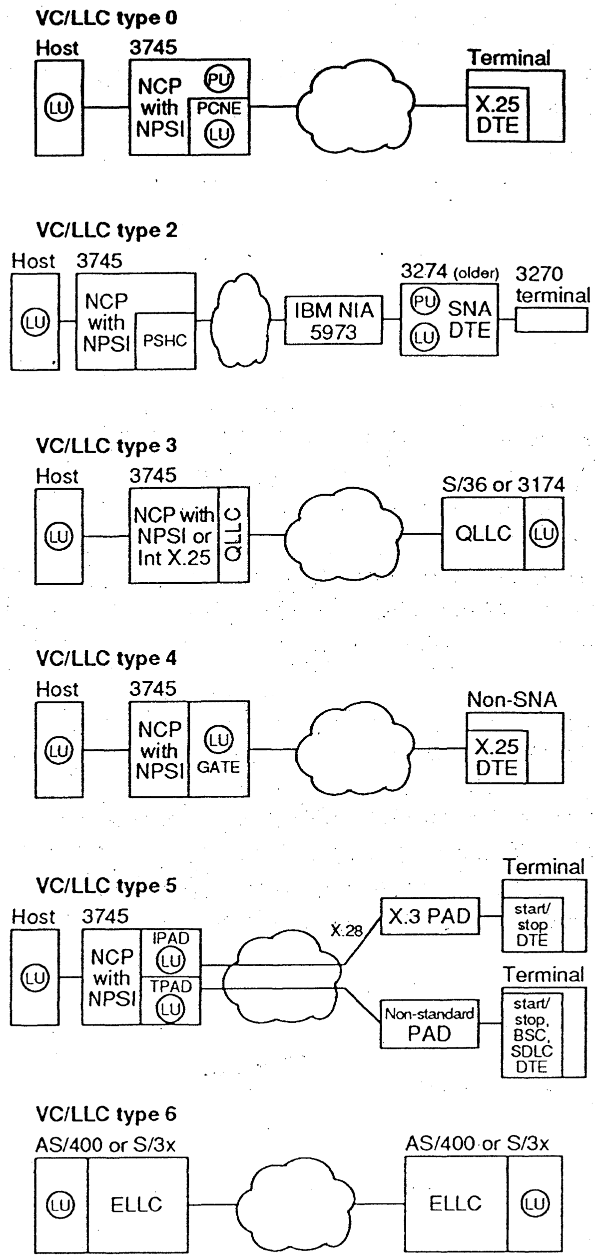


DTE - Data Terminal Equipment
ELLC - Enhanced LLC
GATE - General Access to X.25 Transport Extension
IPAD - Integrated Packet Assembler/Disassembler
Int. X.25 - Integrated X.25
PAD - Packet Assembler/Disassembler
PCNE - Protocol Converter for non-SNA Equipment
PSHC - Physical Services Header Control
QLLC - Qualified LLC
TPAD - Transparent Packet Assembler/Disassembler

*Figure 11*

## Architect's Corner

# Opportunities in SNA Transition

*by Dr. John R. Pickens*

"Captain! Captain!" cries the junior officer, throwing open the door to the ship's bridge, nearly out of breath from his dash up the ramp. "We have a problem here!" Turning to the junior officer the captain says, "Son, I thought I taught you, there are never, never problems, only opportunities in disguise." Thinking for a moment, the officer pauses, then says "Well sir, we have a serious opportunity down below."

When it comes to networked applications, we indeed seem to have serious "opportunities."

One of the inevitable side effects of technology advancement is that things that used to work don't. And things that work in the new environment don't work in the old environment. Let's be more specific with regard to SNA. As I've discussed previously on numerous occasions, the big ongoing events in SNA are the twin transitions:

- The transition from static, complex subarea technology to dynamic, simple APPN technology
- The transition from the cacophony of incompatible session types—LU 0, 1, 2, 3, etc.—to the single converged session type—LU 6.2 (APPC)

These transitions have created a problem (that is, opportunity). Old applications based on older session types only supported by the old subarea routing technology don't work on the new technology session types supported by the new APPN routing.

I'll address this inverse situation—new applications don't work on old SNA session and router types—in a future column.

Our "opportunity"? How to avoid being forced to ask the following question: "Which mutually exclusive network technology type must I choose in order to support each application?"

### *Old Applications/New Networks Opportunity*

Traditional SNA networks have experienced eighteen years of applications development. Eighteen years for IBM divisions (and users) to pick-and-choose their flavor of SNA—LU 2 for 3270 datastream, LU 1 for SNA printers, LU 0 for transaction-oriented and other specialized applications.

APPC has been (publicly) available for about eight years and APPN for six years. However, neither yet has a huge installed base. (Well, OK, the AS/400 installed base is estimated to be between 1.6 million token ring–attached, PC-based APPC LEN nodes.)

The road map in Figure 12 diagrams the state of the LU 6.2–incompatible applications base. Note the variety of applications. Hiding behind the variety is the more challenging fact that the size of this installed base is huge.

Not to despair; solutions exist at three levels—bridging and link conversion at the lower layers, transport-level gateways, and middleware interfaces to the upper layers. Of this range of solutions, some are specific to types of applications and underlying protocols while others are very general and generic.

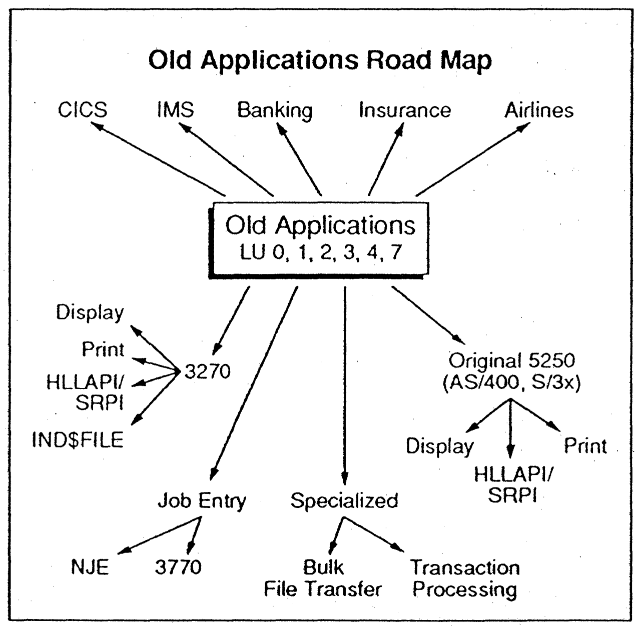


*Figure 12*

### Solution 1—Bridging

Bridging service provides the most universal broad-based solution for old applications on new networks. Why? Because the same data link layer is used for both the old and the new, and bridging operates at the data link layer—the LLC type 2 connection-oriented service for LANs and the SDLC connection-oriented service for WANs, 3270, 3770, RJE, retail, and banking can all be handled by a bridging service.

For WAN-attached devices, a conversion function is required—SDLC-to-LLC2. The remainder of the network path, beyond the conversion device, can traverse bridges. Bridging is particularly useful in networks with hierarchically-designed topologies, though the spanning tree and source routing algorithms make bridging a general solution. The weaknesses of bridges? They don't handle priority (yet). They can exhibit suboptimal behavior in the face of broadcast and multicast traffic (not used in SNA). And bridging provides switching only at the data link layer (i.e., no visibility into session layer switching, routing, and flow control).

But, presented with the requirement to migrate to new networks, the bridging function provides the most broad-based migration path for old applications migrating onto new networks.

### Solution 2—Gateway Service

Gateway service enables the support of old applications via encapsulation across new APPN network services. One gateway service has been discussed, though not yet announced, by IBM—dependent LU requester/server. This model is intended to provide a type of passthrough function for SSCP-PU and SSCP-LU session types, since neither type is capable of being routed by APPN. This enables LAN-attached nodes running old applications and old SNA protocols to be handled by APPN transport networks. After session initiation, the dependent LU data will run natively across APPN.

### Solution 3—Middleware

The first two solutions place the burden for migration on the network itself (intermediate system). The third solution places the burden on the end node (end system)—middleware.

The idea behind middleware is reasonable enough—offer the applications developer generic interfaces that keep hidden the underlying nature of the transport environment. Offer a CPI-C interface that can be mapped to underlying APPC, OSI DTP, or APPC-over-TCP services. Offer a sockets interface that can be mapped either to TCP or to APPC underlying transport services. In certain cases this solution can be mixed with gateways (solution 2). This would enable, for example, the development of a CPI-C application that originates CPI-C conversations in a node using APPC-over-SNA, passes through an SNA-to-TCP gateway, and terminates in a node using APPC-over-TCP.

Middleware is described primarily as a solution for the mixed protocol environment, but could also be used for functions like HLLAPI mapped to APPC (and then to 3270 at the remote end).

The weaknesses of middleware? Some loss of function through the mapping process. Some loss of performance. But a promising way to blend the old and the new, and the mixed protocol environment.

### Conclusion

How real are these solutions? The answer is mixed. Bridging and link conversion products are available today. Most vendors are just in the process of tidying up some missing elements such as Ethernet-to-token ring translation bridging. Gateway services are being talked about today—but are some distance from realization in products. Middleware services are beginning to become available—but some services needed for applications migration, such as HLLAPI middleware, are not really being talked about.

In an ideal world, we would not be faced with such a difficult applications migration challenge. But, in the real world, solutions have to be developed. Bridging, gateways, and middleware provide several promising migration alternatives for old applications. For SNA users and application providers, I believe the best solution is to migrate SNA applications to the new session service type—APPC. But, for now, we, like the junior officer, must learn how to identify and deal with our "opportunities." ■

## SNAP Editorial

# IBM Shows TLC, Not FUD

In August, IBM's APPC Market Enablement team hosted the first APPC/APPN Platform Developer's Conference with representatives of more than 40 peer SNA platform developer companies attending, including SNA developers, system vendors, and gateway/internetworking companies.

APPC/APPN developers have never met together before in such a forum for exchange of information and ideas. IBM was, to an unprecedented degree, open about its architectures and product direction, supportive of third-party developments, and open to feedback from developers. It was not a mutual admiration society—there were some lively, spirited discussions about topics such as delivery schedules and APPN network node licensing and fees. But most attendees were impressed with the quality of the information and the discussions. They also appreciated IBM's frankness in acknowledging many problems in the architectures and products, discussing how it will resolve them, and providing tips for handling thorny problems in the meantime.

Among the highlights: IBM is forming a group with representation from several companies to openly direct future CPI-C developments. IBM stated officially that it will support multilink transmission groups in APPN. IBM will probably publish or license border node (the APPN equivalent of subarea SNI). In addition to IBM's strategy for dependent LU support across APPN—the dependent LU requester/server model statement of direction that we discussed in the April issue—some elements within IBM are also considering full encapsulation of dependent LU traffic in LU 6.2 over APPN. Syncpoint (two-phase commit or CPI-RR) for distributed databases will be added to LU 6.2 in 1993 with developer documentation. IBM will publish and ship full-duplex APPC in 1993. IBM will support development of a standard to interface APPN to OSPF. The APPN network node licensed source code will include an SNMP MIB. IBM will provide a list of the patents related to APPN network node so that vendors who wish to develop it instead of licensing the source code can discuss specific patent licensing with IBM.

Nearly two years ago, in December 1990, we wrote, "IBM could focus on providing more support for those...who take on the challenge of implementing LU 6.2 by rewarding them for moving down this road less traveled with tools, training, and more TLC than FUD." Shortly after this, the IBM APPC Market Enablement team was established. We are not claiming credit, of course, for we were just one voice among many encouraging IBM to do what was necessary to make the new SNA successful.

*SNA Perspective* believes that the efforts of this team and the ongoing developers' dialogue started at this conference will lead to better APPC and APPN products for users—better interoperability, quicker to market, less expensive, and with more creative features—since the vendors will waste less time and energy trying to decypher IBM and its architectures. This interplay of cooperation and competition does not guarantee APPC/APPN market success, of course, but without it we believe they would probably have failed.

A note to our readers: The APPC Market Enablement team exists to serve users as well as developers with documentation, sources of products and development tools, education, and more. If you have questions or concerns about using APPC and APPN, send them a fax at (919) 254-5410. Their Internet address is appcmrkt@ralvm6.vnet.ibm.com If you have access to CompuServe, join the thousands of users on the APPC Info Exchange Forum by typing GO APPC. ■

Copyright © 1992 CSI - Communication Solutions, Incorporated, all rights reserved. Reproduction is prohibited.
• Subscription rates: U.S. - one year $350, two years $520. International - one year $385, two years $590 • *SNA Perspective* is published monthly by CSI, 2071 Hamilton Avenue, San Jose, CA 95125 • Telephone (408) 371-5790 • Fax (408) 371-5779 • Managing Editor: Louise Herndon Wells • Associate Editors: Vincent Busam, Alan Kucharski, Basil Treppa • Marketing/Development: Alisson Walsh • Circulation Coordinator: Cheryl Roberts • Contributors: Michael E. Bowman, Thomas J. Routt, Dr. John R. Pickens • Typesetting and illustration: Aaron Lyon at dSIGN
• The information and opinions within are based on the best information available, but completeness and accuracy cannot be guaranteed.