Volume 13, Number 12
December, 1992
ISSN 0270-7284

SNA Perspective

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

APPI: The Product and the Protest

In the past few months, an interesting twist has enlivened the world of Advanced Peer-to-Peer Networking (APPN). It has been generated not by IBM, though, but by several networking vendors led by Cisco Systems under the banner of Advanced Peer-to-Peer Internetworking (APPI). APPI is proffered as an alternative to APPN, supporting APPN end nodes but running over TCP/IP.

Consider an analogy of two adjacent kingdoms, ruled by subarea SNA and TCP/IP. As subarea SNA ages, APPN emerges as the young heir apparent. TCP/IP, in the thriving border kingdom, offers a union by marriage (APPI), with the stipulation that TCP/IP's home be the new capital of the combined kingdom. APPN prefers an alliance but must reckon with TCP/IP's influence. APPN must negotiate the most advantageous terms for the kingdom. But an alliance there must be; the alternative might be war.

SNA users should understand the motivations behind APPI, what it is technically, and what it represents for their networks. This article discusses these issues as well as recent IBM moves in the APPN internetworking arena. In addition to its technical proposal, we consider APPI both a voice of protest and, in a way, a vote of confidence for APPN.

(continued on page 2)

Sprucing Up Your 3270 Controller

The 1980s were a decade of significant growth for SNA and for the 3270 display system that supported SNA access for terminals and personal computers. This growth slowed in the late 1980s and shipments have actually declined for the last few years. PC-based 3270 LAN gateways and other mainframe access options often offer more price-competitive solutions.

But existing 3270 systems represent a significant investment in controllers, workstations, PC adapters, software, cabling, and user experience. About 850,000 3270 controllers are installed worldwide, more than half purchased within the last five years. Although new purchases are declining, users are considering ways to expand the functionality of their existing 3270 controllers. This article discusses current networking options including TCP/IP, AS/400 5250, DEC LAT, Ethernet, APPN, ISDN, ESCON, and SDLC converters for sprucing up your 3270 controller.

(continued on page 14)

In This Issue:

(continued from page 1)

This article is not intended to be a technical comparison of APPN and TCP/IP as alternative migration paths for subarea SNA users. A future article or series of articles in *SNA Perspective* will address this important topic. This article focuses instead on the issues raised by APPI, which are primarily marketing rather than technical.

To understand APPI, some knowledge of APPN is necessary. We include here a sidebar A Brief Review of Subarea SNA and APPN. For more details on APPN, see "APPN Insights and Design Clues" in *SNA Perspective*, April 1992.

Three Reasons for APPI

Three categories of issues are cited by APPI proponents:

- technical
- industry (vendors)
- marketplace (users)

The technical issues relate to adaptive routing, media support, multiple backbone protocols, and APPN's limited track record. The industry issues center around two issues: first, the APPN specifications are proprietary and implementers must pay IBM for specifications, source code, and/or patent licenses; second, any enhancements are available first to IBM and its selected allies. Marketplace issues concern integration of subarea SNA/APPN with existing router-based internetworks, and, again, IBM's ownership of APPN. Some of these issues have been mitigated to some extent by recent IBM announcements, as we shall see.

For a full discussion of these issues, see the section The Case for APPI on page 9.

The APPI Forum

In August 1992, Cisco Systems announced the APPI development project which will be supported by the APPI Forum. Under Cisco's leadership, the APPI Forum has been organized under the auspices of the Interop Company, which has sponsored several other forums.

Principal memberships in the APPI forum cost $8,000 while auditing members, such as users and consultants, can join for $1,500. Principal members can participate in the working committee meetings and have a vote. Auditing members get all technical documents, but can only attend the general meetings. Membership is open to any vendor, user, or other industry participant.

The APPI Forum formation meeting was held in October, 1992, in San Francisco, California, during Interop. The first committee meetings are scheduled during ComNet in Washington, D.C., in February 1993, followed by another general meeting at Interop Spring '93 in Washington, D.C. Work will proceed via electronic mail between these meetings. The APPI Forum has indicated that it hopes to have a complete specification approved and available by June 1993. *SNA Perspective* considers this quite an ambitious timeframe.

If this schedule is met, the first products based on these specifications could be available by late 1993. There has been no public discussion of a timeframe for future releases of APPI.

APPI Forum Members (as of November 1992)

Principal

Alcatel *	Infonet *
Arkhon	McData *
British Telecom, UK *	MCI
Cabletron *	Netrix *
Cascade Communications *	Proteon *
CompuServe	Rabbit
Cisco Systems *	Sprint
Data Connection, Ltd.	SunConnect *
Digital Equipment Corporation *	SynOptics *
Hewlett-Packard *	

Auditing

British Telecom, North America
Computer Communications, Inc. (CCI)
Eicon Technology
Proginet, Corp.
SourceCom Corp.

* Founding member

Table 1

A Brief Review of Subarea SNA and APPN

Advanced Peer-to-Peer Networking (APPN) is an IBM networking architecture that is an evolutionary extension of the company's Systems Networking Architecture (SNA). Subarea SNA was predicated on a hierarchical network, in which selected mainframe-based nodes controlled the network and all other nodes.

Subarea SNA

Subarea SNA uses several hierarchical levels or types of nodes called physical units (PUs). PUs are implemented at roughly layers 3 and 4 of the OSI reference model. PU 5 is implemented as the System Services Control Point (SSCP) on the host and controls all the nodes in its domain within an SNA network. PU 4 is implemented in the Network Control Program (NCP) of a 37xx communication controller. Traffic between PU 4s and between a PU 4 and a PU 5 is called SNA subarea traffic.

PU 2 is called a peripheral node and is implemented in 3270 controllers, gateways, and protocol converters. A PU 2 must logically connect to a PU 4 or a PU 5. The PU 4 or PU 5 that a PU 2 connects to is called its boundary function. The traffic between them is called local or peripheral SNA traffic. This boundary function converts its local (nonroutable) traffic into subarea (routable) traffic.

APPN

APPN is based on one node type, type 2.1. An APPN end node (EN) or network node (NN) also contains a control point (CP). An earlier type 2.1 node, the low-entry networking node, had no control point and thus lacked much of the flexibility of end nodes and network nodes. Rather than relying on a central SSCP, each APPN network node keeps network topologies and makes routing decisions.

In internetworking terminology, end nodes would be called end systems and network nodes would be called intermediate systems. In addition to end node capabilities, a network node can perform intermediate routing and distributed directory services. ■

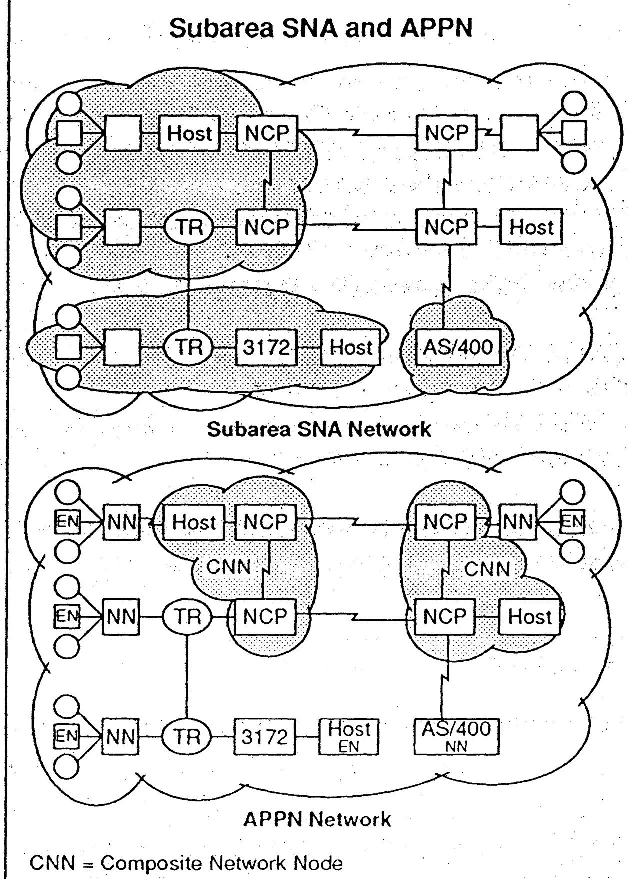

Figure 1

A Chair for IBM?

IBM was invited to join the APPI Forum but has declined to participate at this point, indicating that it does not yet see a significant benefit to users in APPI instead of, or in addition to, APPN.

The APPI Concept

The APPI Forum specifications will probably be developed in several phases. The first iteration will use an APPI node which will take information from APPN end nodes and route it as IP traffic instead of APPN. A future release will add more APPN properties such as flow control, class of service, and network management. Eventually, the APPI Forum is expected to implement an APPI node that will interact with APPN network nodes.

To understand how APPI will work, it is important to understand that APPN is not directly comparable to TCP/IP. See the sidebar TCP/IP and APPN Routing Protocols on page 5.

The communication between an APPN end node and network node (EN-NN) could be considered a *network access* protocol, since it is the point of access to the APPN network and is usually local.

The communication between network nodes (NN-NN) could be called a *network transport* protocol, since it is the process by which the traffic is transported across the network. Technically, it should be possible to support access protocols only to route APPN traffic over a multiprotocol transport, which is what APPI proposes. The NN-NN communication would be replaced by features in TCP/IP.

End Node to Open Network Node

An APPI node will be called an open network node (ONN). The ONN would *appear* to any APPN end node as an APPN network node.

Specifications for this ONN can be directly developed from the *Type 2.1 Node Reference* which has been published by IBM (IBM document SC30-3422-2). It contains the specifications for the end node and the EN-NN communication.

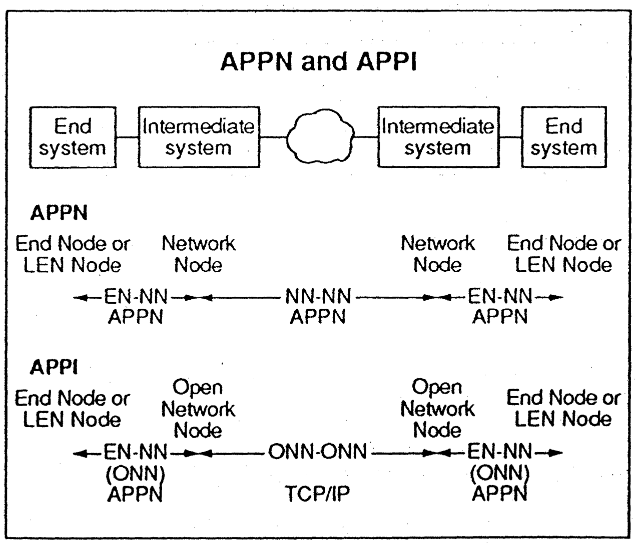

Figure 2

The APPN end nodes would register their resources to an ONN, as they would with an APPN network node, through the usual APPC general data stream (GDS) variables. APPI ONNs would also support the older LEN nodes, but these nodes and their resources would have to be statically configured in the ONN or a distributed directory server. An ONN could also support a network node by treating it as a LEN node and configuring a table of its resources.

Connection Networks. *SNA Perspective* believes that APPI would probably use *connection networks* as the basis for its end node support. A connection network is an APPN feature that allows an end node to specify its connection to a LAN or a bridge/router internetwork as a virtual routing node. In this way, two end nodes specifying the same connection network can connect directly with a single APPN link with the underlying network transparent to APPN. A connection network can thus be used instead of APPN hop-by-hop routing. Connection network support will be included in the APPN network node source code and specifications that will be available in first quarter 1993.

Instead of Network Node

Between ONNs, APPN is completely replaced.

- The traffic will be in IP format instead of APPN.
- The routing protocol can be any that routes IP traffic, such as OSPF, RIP, IGRP, or integrated IS-IS, rather than the APPN network node link state.
- The directory service will be a new APPI distributed directory service instead of the APPN distributed directory service.

A router with ONN would be needed only at points where APPN end nodes access the router network. Inside the network, since the traffic is TCP/IP, internal routers do not need ONN.

Directory Services

At the APPI Forum formation meeting, Cisco discussed the proposal it will make to the APPI Forum for directory services. Each ONN will be a distributed directory client (DDC). Some ONNs will also contain a distributed directory server (DDS), which will have a view of the network. Taken together, the

TCP/IP and APPN Routing Protocols

Internet Protocol

TCP/IP is an umbrella term for several protocols. TCP/IP includes two transport-layer protocols: transmission control protocol (TCP) is a connection-oriented, reliable layer-four protocol and the user datagram protocol (UDP) is a connectionless, best-effort delivery layer-four protocol.

A routing protocol, in internetworking parlance, is one that can interpret a network-level (layer three) address, has a routing table or can discover the location of the destination, and determine the appropriate route to the destination.

The internet protocol (IP) is a local routing protocol. It can do local routing within a network boundary, but not across internetworks. The TCP/IP standard internetwork routing protocol, routing information protocol (RIP), offers minimal functionality so most routers include one or more other protocols for internetworking IP traffic.

These include Cisco's proprietary Interior Gateway Routing Protocol (IGRP), the Internet and OSI integrated intermediate system to intermediate system (integrated IS-IS), and the Internet standard open shortest path first (OSPF). These routing protocols are incompatible with each other—a router that only supports IGRP, for example, cannot communicate with a router that only supports OSPF, even though both may be routing IP packets.

APPN

APPN is also an umbrella term for several elements that are usually referred to separately in internetworking parlance. It is a network protocol, a transport protocol, a routing protocol, and includes an integrated directory services. (Technically, the path control layer could be considered separate from APPN, but for purposes of comparison, it is helpful to consider it under the APPN umbrella.)

The approximate equivalent to TCP and IP in APPN is called intermediate session routing (ISR). ISR uses a connection-oriented network protocol, while IP is connectionless. ISR uses both layers three and four at each intermediate router, while TCP/IP uses only the IP layer at intermediate nodes.

IBM has begun to discuss APPN+, a forthcoming version of APPN. The alternative to ISR in APPN+ is called high-performance routing (HPR). The HPR layer three protocol is connectionless. With HPR, the intermediate nodes will only use layer three. HPR will also support adaptive routing. ■

DDS nodes would make up the distributed database. There are also plans for an optional central directory server.

The APPI directory services proposed by Cisco will probably be based on extensions to the TCP/IP domain name service. To meet the goal of APPI being based on standards, the APPI Forum would need to work with the Internet Engineering Task Force (IETF) working group on domain name service to standardize these extensions. Any APPI Forum member could propose another distributed directory service. These could be based, for example, on the Open Software Foundation (OSF) Distributed Computing Environment (DCE) directory service or the OSI/CCITT X.500 directory service.

No Patents To Be Used—At First

One goal of the APPI Forum is to not implement any features of APPN that are covered by IBM patents. These patents include such features as adaptive session level pacing, transmission group number negotiation, and distributed database locate requests between both EN-NN and NN-NN. *SNA Perspective* sees this as a drawback to APPI, since IBM probably patented what it felt were the most valuable elements of its design.

Further, when an APPI node is developed in a future release of APPI that can communicate with an APPN network node, it may have to implement features that are covered by those patents. Therefore, an implementer may need to license at least the patents from IBM.

The concepts covered by these patents are not new. Adaptive session level pacing is just one type of variable sliding window protocol, for example, and APPN locate requests are only one way of locating resources in a distributed network. Therefore, the APPI Forum would not be starting from scratch to develop an APPI implementation that does not infringe on any patents. But it is still a formidable task and raises issues for future interoperability.

APPI Reduces Number of Routing Protocols

APPI offers the advantage of having one less protocol on the backbone. Rather than offering coexistence for APPN and TCP/IP networks, APPI envisions a corporate network supporting end nodes but not supporting APPN as a network. Also, APPI will adapt several of the technical benefits of APPN over TCP/IP, over time, to run on TCP/IP networks.

APPI Cannot Recognize APPN Networks

Based on the initial APPI proposal, APPI networks cannot communicate with APPN networks, as shown in Figure 3. This is because ONNs are not actually network nodes or even APPN nodes at all. ONNs cannot communicate in any way with network nodes. Therefore, sessions cannot flow from an end node supported by an ONN and an end node supported by a network node.

Eventually, the APPI Forum is expected to develop a node which can communicate with APPN network nodes. But, until that time, the two networks would be incompatible. APPI could access certain network nodes that have selected type 2.1 links defined for LEN communication, and develop enhanced communication with LEN nodes that does not involve control points. An ONN might also communicate with an end node through the same LAN as its network node server. But how and when these would be done is unclear.

The main impact of this incompatibility with network nodes, in the short run, would be lack of support in APPI for the installed base of AS/400s that comprise the largest number of APPN networks. (AS/400s can be configured as end nodes, but must be network nodes to use PC Support/400.) APPI would have the same challenge with IBM 3174s, which can be configured as network nodes but not as end nodes. Table 2 (see page 7) shows the current and planned APPN products from IBM and other vendors.

APPI Cannot Support Hosts through 3745

With VTAM 4.1, a host can be defined as either an end node or a network node. If not intended to be involved in routing, the host can be configured as an end node, and APPI can support it.

However, if a host owns any 3745s, it must be configured as part of a composite LEN node or composite network node. Therefore, a host that connects to a LAN through a 3745 (which includes most LAN-attached hosts today) cannot be an end node. APPI could support such a host if it were configured as a composite LEN node, but only at a significant loss of APPN functionality.

VTAM hosts configured as end nodes and connected to APPN through a 3172 or similar LAN-mainframe gateway or through an integrated communications adapter (ICA) could be supported by APPI.

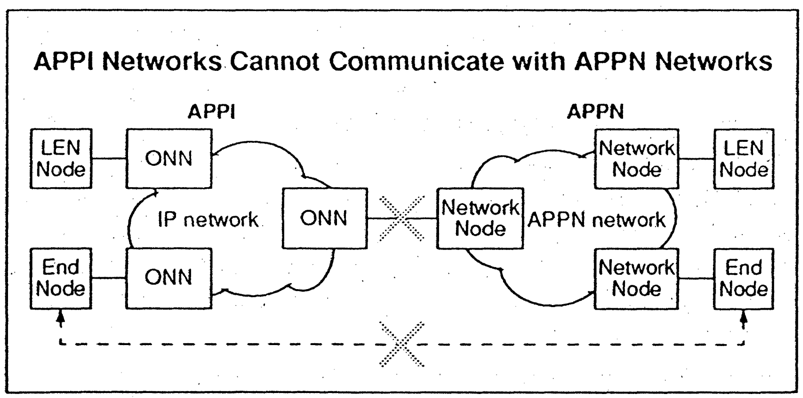

Figure 3

No Dependent LU Support

The APPI Forum has, as one of its goals, support for dependent LU traffic. However, the APPI ONN, in its initial specifications, will not be able to support dependent LU traffic sent through an APPN node.

In the November 1992 issue of *SNA Perspective*, we discussed at length existing and expected support from IBM for dependent LUs over APPN. VTAM 4.1 can support dependent LUs that are logically adjacent to it or its 3745s. However, it must be configured as a network node or, if configured as an end node, must have a VTAM 4.1 node as its network node.

State of APPN

Implementation	LEN	EN	NN
IBM			
AS/400	✔	✔	✔
System/36	✔	no	✔
OS/2 Communication Mgr	✔	✔	✔
3174 w/ Config Support C	no	no	✔
VTAM/NCP	✔	1993	1993
6611	n/a	n/a	1993
RS/6000	✔	1993	1993
OEM	no	yes?	1993
NS/DOS	late 1992	n/a	n/a
Non-IBM			
System Strategies, Inc. (OEM)	✔	late 1992	?
Data Connections Ltd. (OEM)	✔	1993	1993
Novell - Netware for SAA	✔	SOD	1993
Eicon - SNA LAN Gateway	✔	n/a	n/a
DCA - Select Comm Server	✔	n/a	n/a
3Com Corporation	no	no	1993
Network Equipment Tech.	no	no	?
Apple Computer	✔	SOD	?

Source: APPI Forum and *SNA Perspective*

Table 2

Since ONNs cannot support network nodes, they cannot support these VTAM hosts. Further, since VTAM 4.1 dependent LU support sends the dependent LU traffic natively on APPN and not encapsulated in APPC sessions, it would be difficult for APPI to support this.

The licensed APPN source code and specification to be published in first quarter 1993 do not contain dependent LU support either. *SNA Perspective* expects that IBM will eventually provide this support and certain other incremental network node capabilities, such as central directory server and border node, as separate options for developers. We do not expect them to be available until sometime in 1994. (Border node, to connect two APPN networks, is not even included in VTAM 4.1 but has a statement of direction for a future release.)

APPN Across TCP/IP from IBM

IBM actually has several approaches to integrating APPN and TCP/IP. However, none of them are shipping today, only one has been announced (in October), and the others can only be discerned from an examination of IBM networking strategy. Understandably, several vendors and users are quite interested in IBM's plans for APPN and TCP/IP and frustrated over the veil of secrecy around such an important topic.

An Encapsulation Approach

At Interop Fall '92 in October, IBM demonstrated the ability for its 6611 router to route APPN traffic

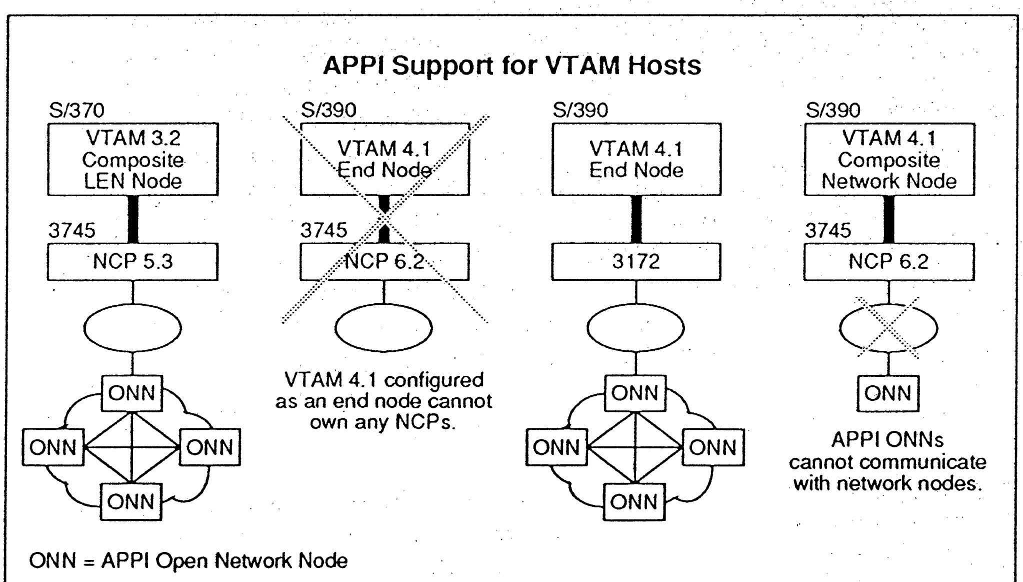

Figure 4

encapsulated in IP. This capability will be available in first quarter 1993. It is interesting to note that, for the 6611, APPN over TCP/IP will be available before native APPN routing, which is not expected to ship until late 1993.

Figure 5 shows two IBM 6611 routers equipped with TCP/IP, data link switching, and APPN and illustrates communication between two APPN end nodes on token rings. This support can use an APPN connection network, which is described above in the section The APPI Concept and shown in Figure 5. Standard APPN ISR hop-by-hop routing can also be used, in which case the data would travel between the routers on a sockets session rather than over data link switching.

Data link switching is IBM's procedure for supporting SNA, APPN, and NetBIOS traffic on a router network (see "Data Link Switching on the IBM 6611," August 1992, *SNA Perspective*). IBM will submit a Request for Comments (RFC) to the IETF on data link switching and APPN over sockets so other vendors can implement them.

The network nodes in this implementation communicate across TCP/IP over a sockets interface. The IBM network node source code and specifications that will ship in first quarter 1993 will include these interfaces for APPN over token ring and sockets. No other interfaces will be provided, though the code can be used to develop interfaces over Ethernet, SDLC, PPP, frame relay, and others.

This support differs significantly from APPI. This implementation encapsulates APPN protocols and data over TCP/IP, essentially treating TCP/IP as a reliable data link. APPI uses TCP/IP protocols instead of APPN protocols to send data between APPN nodes.

The Blueprint Approach

SNA Perspective believes that this encapsulation approach is an interim solution for IBM and that its strategic solution for integrating APPN and TCP/IP will be based on its networking blueprint. In March, IBM unveiled this networking blueprint (see "Blueprint to Integrate the Architectures," August 1992, *SNA Perspective*) which is IBM's framework for multiprotocol integration.

An important element of the blueprint is the common transport semantics, an interface that rides on top of the transport layer. Common transport semantics will be based in large part on IBM's multiprotocol transport networking (MPTN). MPTN is an extension to the X/Open transport interface (XTI). XTI allows OSI applications to communicate over TCP/IP and TCP/IP applications to run over OSI. MPTN extends this concept to SNA and NetBIOS. IBM published MPTN in 1991 and has proposed it to X/Open for adoption as a industry specification.

MPTN can be used in two ways—as a server or a gateway. First, it can allow traffic from an application that expects a certain

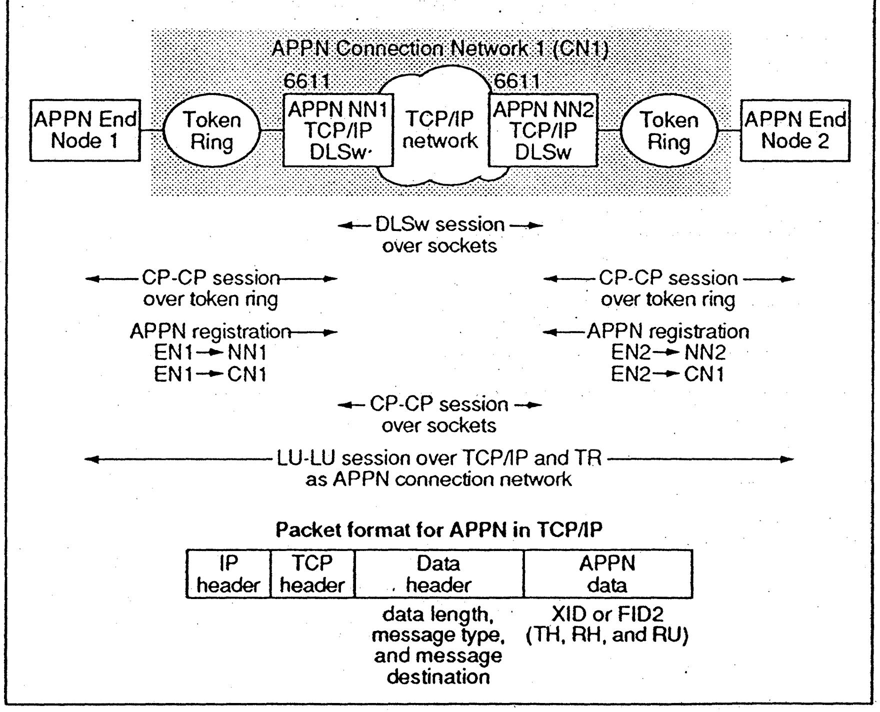

IBM 6611 APPN over TCP/IP with Sockets and Data Link Switching (DLSw)

Figure 5

transport to run over an alternative transport. IBM has discussed as a statement of direction two products that will use MPTN in this way: support for CPI-C applications over TCP/IP, and support for sockets applications over SNA/APPN (informally referred to as SNAckets). Second, MPTN can be used for a standard transport gateway between different transport and routing stacks. *SNA Perspective* expects that IBM will develop an MPTN-based gateway for communication between APPN and TCP/IP networks. The same gateways could be used to support APPC sessions across TCP/IP and sockets applications across APPN.

MPTN differs from APPI by using two transport stacks in each network gateway, while APPI supports APPN EN-NN CP-CP services and EN-NN APPN XIDs but not APPN transport.

The Case for APPI: Technical Issues

As stated earlier, the APPI Forum presents its case in three categories: technical, industry, and marketplace issues.

Some of the APPI Forum technical issues relate to a comparison of APPN and TCP/IP; that is, whether TCP/IP is a better technology than APPN. In the sidebar APPN Pros and Cons on page 10, we touch briefly on several differences. This comparison will be developed at length in a future issue of *SNA Perspective* from the point of view of a subarea SNA user planning a migration strategy.

SNA Perspective believes, however, that a strict technical comparison is not a primary concern for the APPI Forum. The real technical issue is whether it is better to have one backbone protocol, such as TCP/IP, or several protocols running on a shared backbone, or several networks with different internal protocols able to internetwork with each other.

The APPI Forum believes that a single protocol backbone is better and that TCP/IP is the best candidate, not primarily because of technical capabilities but because of its enormous installed base and because it is an openly developed, public domain standard. APPI's long-term direction appears to involve taking the best features of APPN and recreating them in TCP/IP. The best features, of course, may include the ones IBM chose to patent, which brings us to the industry issues.

Industry Issues

The industry issues raised by the APPI Forum relate to effects on other vendors in the market. The two main stated issues are fairness and openness. We also address the underlying concern of price versus risk. Because openness is also a primary issue for users, it is discussed below under User Marketplace Issues.

There are five main concerns about fairness:

- Early access for IBM and a selected group of vendors.
- Requirement to license code rather than being able to build to a published specification. (This was true when the APPI Forum was formed in August, but was obviated by IBM's decision in October to publish the specification.)
- Price of license fees and royalties.
- Secrecy regarding license fees and royalties, leading to concerns about favoritism and a level playing field.
- Uncertainty regarding possible patent license fees for both APPN end node and network node.

Early Access

Several APPI members are particularly upset that IBM selected a few vendors—3Com, Network Equipment Technology, and Novell—to assist in beta testing of the APPN source code. This early participation gave these few vendors a full year of access to the code before any other industry participants, since the testing began in March 1992 and IBM plans to make the source code available publicly in first quarter 1993. Further, since IBM is the sole APPN architect, it will always have this timing advantage over other vendors.

APPN Pros and Cons

APPN offers several technical advantages over TCP/IP on multiprotocol routers. First, APPN has integrated a distributed directory service with the routing protocol, and its directory service is more functional than the TCP/IP domain name service. Second, APPN supports class of service (COS). Third, it offers a larger address space. Fourth, APPN can handle much larger packets, depending on the capability of the subnetwork. Fifth, APPN supports congestion control through adaptive pacing.

On the other hand, there are some disadvantages. First, APPN offers no multiprotocol support at this time. Second, several analysts consider current implementations of APPN to offer poor price/performance. Third, it is currently implemented on token ring, SDLC, and X.25, while IP runs over a much wider range of link types, including T1/E1, T3, ISDN, Ethernet, SMDS, FDDI, and frame relay.

Fourth, APPN ISR has no adaptive routing—it uses session-level routing rather than packet-level routing, so all traffic for a given session follows the same route even if traffic conditions change and the session is lost if any link on an APPN route goes down while TCP/IP has automatic reroutes. Fifth, APPN has a more limited deployment track record than TCP/IP—although IBM claims it offers greater scalability, the largest APPN network today is smaller than the largest TCP/IP multiprotocol router network.

It is important to note that both TCP/IP and APPN are rapidly evolving, so these comparisons are only temporary. For example, the IETF is considering five proposals to expand IP's addressing, which is expected to be exhausted by 1995. Also, APPN+ HPR will have automatic reroutes and APPN is being implemented on several more link types. ■

No Published Specification

At the time the APPI Forum was formed in August, IBM was not planning to publish the network node specifications. Instead, it was going to make network node available only through license of the source code. Many vendors wanted the option to develop APPN network node themselves. Based on significant industry pressure from several quarters, including but not limited to the APPI Forum, IBM changed its mind in late October and announced that it would publish the specifications at the same time as the licensed code is available—first quarter 1993.

License Fees and Royalties

IBM has stated since March that APPN network node source code license fees and royalties would be set on a case-by-case basis. Anyone discussing APPN network node licensing and royalties with IBM is held to very strict nondisclosure. The company acknowledged in October that the "list price" of the source code license is $400,000, but said that the actual fee paid would depend on negotiations regarding products, cross-licenses, or other elements the licensee may bring to the table.

SNA Perspective expects that the per-unit royalty fee will probably be a percentage of either the cost of the entire system or the price charged by the vendor just for APPN. The concerns are common with most royalty structures, such as how to fairly set a royalty based on the cost of a system of which this code is only a small part, and the problem of having to inform a competitor of one's shipment numbers through royalty payments. Pricing for APPN on the 6611 router is not available, but IBM charges $129 to upgrade its 3174 controller to code that supports APPN. Vendors indicate that the royalty fee structure IBM is discussing would require them to charge a much higher and therefore uncompetitive price. The multiprotocol router vendors package their protocols differently, which is why they have different levels of concern about this (see sidebar Multiprotocol Router Protocol Pricing on page 12).

Secrecy Regarding License Fees and Royalties

IBM's confidential discussions led to concerns about favoritism and a level playing field—if some vendors are able to pay a lower fee or royalties, they can price their products more competitively.

Uncertainty Regarding Patents

IBM has been granted several patents on different aspects of APPN, that affect the end node, the network node, or both. Implementers therefore have to sign a patent license agreement with IBM for them or amend existing license agreements. IBM has not stated whether it will require patent license fees for APPN end node or network node products. Since IBM published end node in 1991, some vendors have implemented it in their products and IBM has apparently not charged them patent fees. However, it legally retains the right to do so at some point. The uncertainly regarding whether IBM will charge for network node patents makes it difficult for vendors to consider the make-versus-buy decision regarding network node.

Price versus Risk

SNA Perspective believes that the $400,000 price tag is reasonable for 120,000 lines of clean code, documentation, and support. Our research indicates that it would take from four to ten person-years to develop the code internally. With cost of at least $100,000 per person-year, the cost of the source code license is attractive.

On the other hand, pricing must take into account the competition. TCP/IP source code is widely available, as are experienced TCP/IP programmers, and the price for quality, richly-featured TCP/IP source code is much less than $400,000.

In addition, all the router and bridging protocols on a Cisco router together are in the range of a few hundred thousand lines of code so, while $400,000 may be "appropriate" for 120,000 lines of code, APPN would make up a significant percent of a vendor's code.

SNA Perspective believes that the unspoken but important aspect behind the pricing concerns is the risk inherent in investing in APPN. SNA certainly has the largest installed base of networking worldwide. But TCP/IP is reaching the same range. It is also growing much more quickly, while analysts say SNA is growing slowly, staying the same, or even shrinking. TCP/IP is even hitting the mainframe market (see the three-part series "Integrating TCP/IP into SNA," *SNA Perspective*, May, June, and July 1992). More than ten percent of IBM mainframes have TCP/IP installed today and that number could reach twenty-five percent within a year. Probably a third have offboard TCP/IP access to the host.

APPN is the architectural successor to subarea SNA and will offer a smoother migration path than TCP/IP. However, though it has been discussed since 1982 and the first APPN product appeared in 1986, APPN will not ship for the mainframe until sometime in the first half of 1993. In the intervening years, many users have choosen TCP/IP to provide the flexibility, dynamism, and peer support IBM had been promising but not delivering. APPN might be the natural child of SNA, but TCP/IP is also being "adopted" and stands to inherit a significant portion of the SNA estate.

Faced with this situation and having been burned by the bright promise and dim reality of OSI (which was an openly developed and publicly owned standard), vendors are understandably cautious about investing in APPN. Even if they were given APPN source code at no charge, they still might think twice about the cost of training, porting, supporting, and marketing for APPN.

User Marketplace Issues

The two main user issues raised by APPI are integration and openness.

Integration

Users would prefer to have as few protocols as possible on their backbone network. They also want the best networking support. And they also want to maintain their existing investments. Tradeoffs must be made.

Multiprotocol Router Protocol Pricing: What and How

Among the factors in choosing a multiprotocol router are not only the router capabilities and the types of protocols offered, but how the software is provided. Different vendors price their products in different ways.

Market leader Cisco Systems sells products that come complete with all current routing protocols, but bridging, packet switching, and DDN are sold separately.

Wellfleet handles the same situation differently. While the company provides bridging with its products along with one routing protocol, each additional protocol is priced as an option.

3Com approaches the issue in yet another way. The customer *must* purchase software for the router, and can choose between buying the local bridging/routing software or the local and remote version.

A fourth packaging system is used by companies such as ACC and Network Systems/ Vitalink. These companies bundle all their software, bridging, routing, packet switching, etc. with the router.

Each of these approaches has its merits and liabilities. If a company wishes to handle all traffic via routing, why should it purchase bridging? If it needs only certain protocols, why purchase all of them? If only local operations are being used, why pay for the unused remote capabilities? If a company may have many protocols running on its network over time, why not buy bundled software with all available protocols? ■

IBM is proposing TCP/IP and APPN coexistence, whether in neighboring networks or sharing access across the same network. APPI proposes TCP/IP instead of APPN routing, supporting APPN only as end nodes at the periphery of the network.

In theory, this sounds attractive. But most APPN end nodes will not be found at the periphery of a TCP/IP network. They will be upgraded subarea SNA nodes at the periphery of an existing SNA network. Also, many IBM APPN nodes can only be configured as network nodes, such as the 3174, many AS/400s, and VTAMs with 3745s and dependent LU support.

Not supporting these in APPI makes it difficult to claim integration. But neither IBM's integration products nor APPI's actual specifications are available yet to see how they address the need for integration.

Openness

Users are increasingly insistent on using standards, particularly on the backbone where IBM wants APPN to be. But, as the industry discovered with OSI, users do not always put their money where their mouth is.

IBM considers industry-standard multivendor protocols such as OSI and TCP/IP to be very important, has widely implemented them on its products, and is actively marketing them. On the other hand, it considers APPN to be an IBM architecture primarily for connecting IBM systems, a migration path for its subarea SNA networks.

Although IBM has published the APPN end node specifications and has said it will publish the network node specifications, it still owns APPN. Implementers can be required to pay patent license fees up front or for each copy even if they develop their own code. IBM also owns the development—it will develop the features and products it believes it can sell to the most users and will then publish the specification for these new features.

There are several levels of openness: published interfaces, source code licensing, published specifications, free or nominal fee patent rights, published development plans, open industry development and participation, and open ownership.

IBM has taken several significant steps toward openness in the last year. It has published several APPN-related interfaces and protocols. It published APPN end node. It proposed MPTN to X/Open. It has held an APPC/APPN Developer's Conference. It is licensing APPN network node to vendors. It has decided to publish network node.

The company is struggling to decide the appropriate point to stop. This is the hundred million dollar decision. Can it make back its investment by owning a larger percentage of a smaller pie or can it make it by opening further, thus possibly creating a larger pie of which it might have smaller piece?

Conclusions

Technical Issues

On the surface, the APPI concept seems to promise the best of both worlds—allowing a smooth migration from SNA to APPN inside the end systems and access to TCP/IP at the network threshold. However, its current design has significant limitations. There are several other ways to integrate APPN and TCP/IP.

There is not yet enough firm information to make a detailed technical analysis of the strengths and weaknesses of APPI. We believe, in fact, that the actual APPI specification will probably differ significantly from its original concept, hopefully addressing some of these concerns.

Industry Issues

SNA Perspective considers the movement for APPI to be primarily a voice of protest. We believe that several members have joined the APPI Forum not so much in appreciation for the technical concept but to join together to share information and express concerns regarding the future of APPN.

APPI is also, in a way, a vote of confidence for APPN. If the members of the APPI Forum thought that APPN was without merit, they would have ignored it. The APPI movement may turn out, in the long run, to have been a boon to APPN, because it has raised existing concerns to a high level quickly rather than allowing them to fester and hamper APPN growth and industry participation.

Marketplace Issues

The enormous installed base of TCP/IP, the significant growth of TCP/IP in and to the mainframe, and the market resistance to proprietary protocols give TCP/IP significant momentum that IBM must counter for APPN to be successful. The existing subarea SNA market is not a set of users without alternatives.

This means that IBM must make clear the benefits of APPN compared to TCP/IP, not just compared to subarea SNA. It must also clarify its direction as far in the future as the market can see in the TCP/IP standards development process.

Integration

Returning to our analogy of the two kingdoms at the beginning of this article, we consider APPI to be but one, albeit the loudest, of many emerging proposals for an alliance between the two realms.

We must also amend the analogy to note that both kingdoms, especially that of subarea SNA and APPN, are experiencing an increasing surge of democracy. It is the citizens, the users, who will vote with their dollars. ■

(continued from page 1)

The 3270 controller has come a long way from simply supporting "dumb" terminals, as shown in Figure 6. Most users know that, in addition to connecting 3270 terminals and printers, IBM and other vendors offer support on their 3270 controllers for PCs and ASCII terminals and hosts, and connection to token ring and X.25 as well as SDLC.

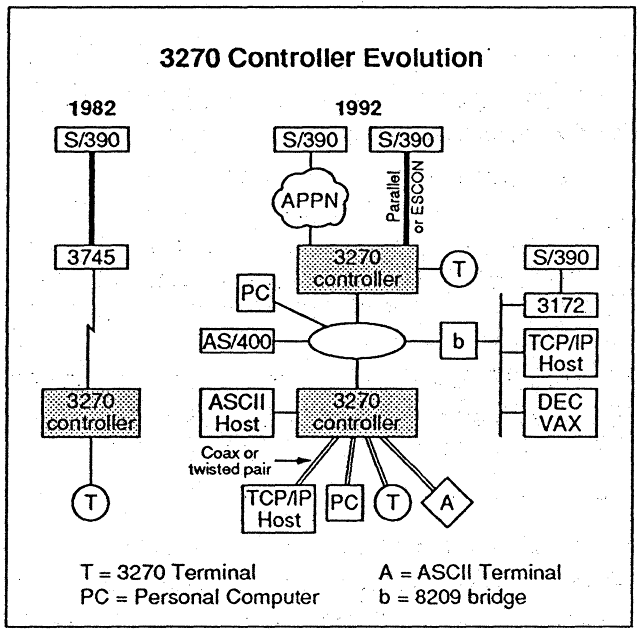

Figure 6

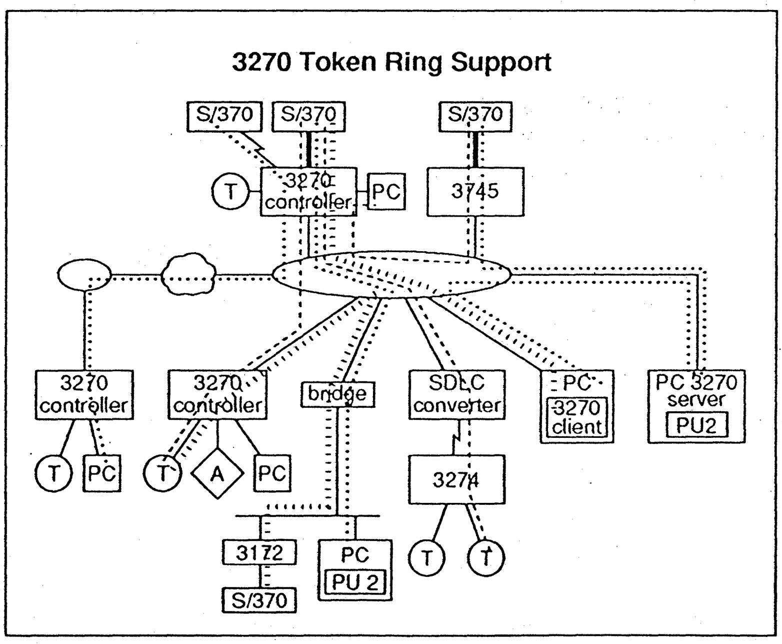

Figure 7

But many users are unaware that their controllers can be upgraded to access multiple SNA hosts, Digital Equipment Corporation (DEC) systems through local area transport (LAT) protocols, TCP/IP systems through telnet and tn3270, and AS/400 systems through 5250 emulation. Controllers can also be enhanced to connect to APPN, Ethernet, and ISDN networks and the ESCON channel. Even older controllers that do not support LAN adapters can access LANs through SDLC converters.

These networking features are provided on 3270 controllers in addition to enhanced functionality such as local format storage, dynamic definition of dependent LUs, multiple logical terminals, split screen, and network management features.

This article focuses on 3270 controller networking support in six areas:

- Token ring
- PC support
- TCP/IP
- Ethernet
- SDLC passthrough and conversion
- APPN

Token Ring

For many years, IBM and other vendors have offered token ring support in both gateway and downstream configurations (see Figure 7).

Gateway support allows physical unit (PU) 2 traffic to flow through the 3174. Without gateway support, PU 2 traffic cannot go through another PU 2 on the way to its boundary function (PU 4 or 5). Most 3270 controllers in gateway mode are limited to the gateway function—they usually cannot support directly-attached terminals and PCs to access hosts back across the LAN. (3174 peer communications allows this for PCs.) A gateway can be defined to support up to 250 downstream nodes.

The downstream PU (DSPU) support allows the device to access a number of hosts through the LAN internet. The actual number of hosts varies by vendor; the 3174 allows support to up to eight hosts across the LAN. These nodes must be preconfigured in the gateway's definition as downstream nodes.

PCs on 3270 Controllers

Although many consider the 3270 to be primarily a terminal controller, about half of the displays on 3270 controllers today are PCs. This impacts the features being added to controllers. For example, the 3270 protocol has been enhanced to support file transfer, APPC support, and other intelligent workstation capabilities. One of the PC-related features offered by IBM is called peer communications.

Figure 8 illustrates the peer communications feature available on the IBM 3174. Peer communications provides three capabilities:

- A "virtual LAN" for coax-attached PCs
- A bridge from the virtual LAN to a token ring
- Host access for PCs with TCP/IP

Peer communications requires configuration support C and is offered as a no-charge option. Peer communications can be used outside a LAN environment, but still requires the token ring adapter. Further, peer communications does not provide a LAN network operating system; the coax-attached PCs would need to have whatever client software is required to access a LAN server.

TCP/IP

TCP/IP support is emerging for 3270 controllers as vendors seek to broaden their range of support beyond the capabilities of competing PC LAN gateways. 3270 controllers have long offered basic ASCII terminal and host support through asynchronous communication adapters. Further, most allow 3270 terminals to access ASCII hosts and provide 3270 terminal emulation for attached ASCII terminals. But this ASCII support is limited to character and line mode communication.

Using 3174 peer communications, all TCP/IP capabilities from the PC TCP/IP software (whether based on DOS or OS/2) can pass through the 3174 to TCP/IP hosts across the LAN, as shown in Figure 9 on page 16. However, the TCP/IP support can flow only over the token ring adapter. TCP/IP traffic cannot pass upstream through the 3174's host channel, remote SDLC, or X.25 connection to a host with TCP/IP software. Further, since the IBM 3745 token ring adapter only supports SNA traffic, the user cannot access TCP/IP on mainframes through a 3745 on a token ring.

Peer communications does not support terminals. A completely separate 3174 feature called TCP/IP telnet support is available as a request for price quotation (RPQ) feature, which provides telnet support for both 3270 and ASCII terminals. This support can also be accessed by a PC acting as a terminal. Through the token ring connection, these terminals can access TCP/IP hosts on the ring or across a bridge/router network. IBM made a statement of

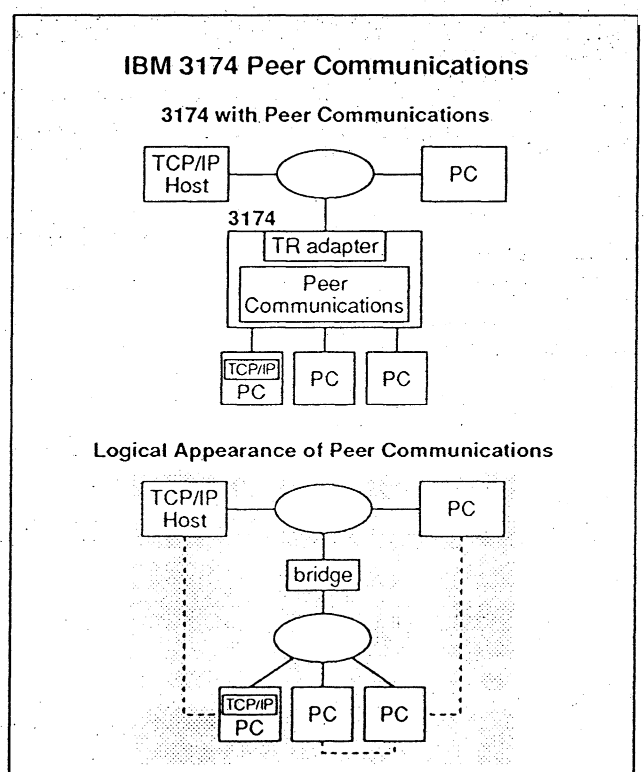

Figure 8

direction in September that it would add tn3270 support to this feature, allowing access to 3270 applications as a 3270 terminal across a TCP/IP network.

Ethernet

Vendors such as IDEA Courier and McDATA have Ethernet adapters for their controllers. However, they do not offer TCP/IP support yet. Instead, they chose to first support access to DEC hosts though the LAT protocol (see Figure 10). *SNA Perspective* expects that TCP/IP support is probably in their plans.

IBM made a statement of direction in September 1992 to provide an Ethernet adapter for the 3174.

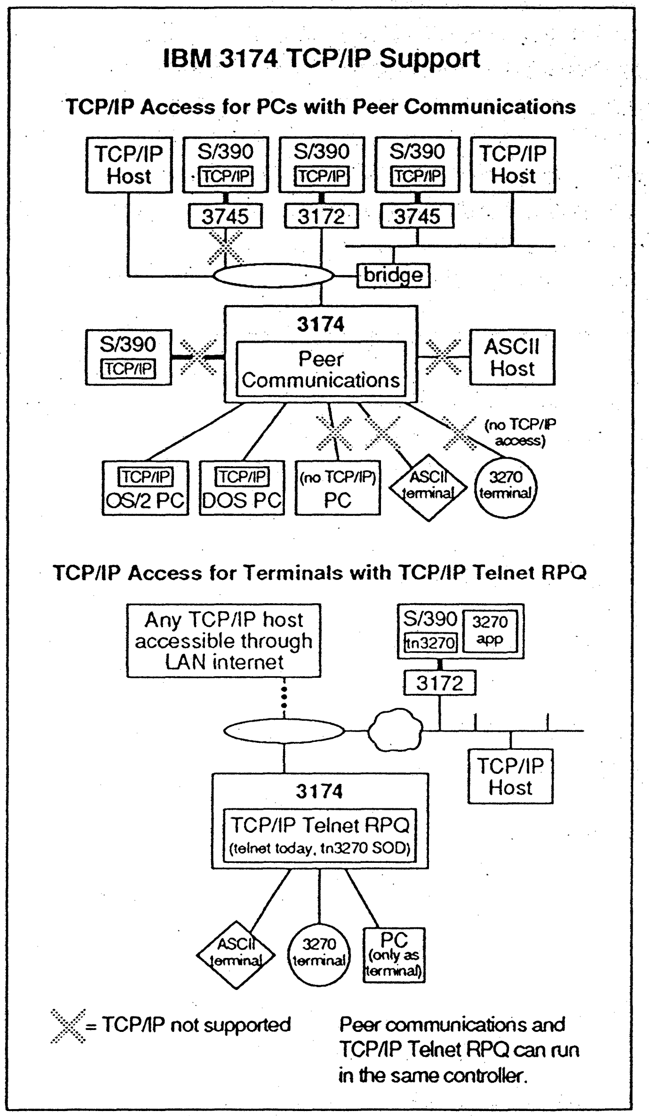

Figure 9

SNA Perspective expected this move in response to the popularity of Ethernet support by other vendors and because of the popularity of TCP/IP for the 3174. We expect the 3174 Ethernet adapter to ship by the end of 1993. IBM indicates that it will provide most of the support that the token ring adapter offers today. As with the other vendors' LAN support, the 3174 will support one LAN adapter; it will be able to attach to either token ring or Ethernet. Some provision will probably need to be made for Simple Network Management Protocol (SNMP) network management support.

SDLC Passthrough and Conversion

Older 3270 controllers such as the IBM 3274 do not have provision for direct LAN attachment. In the past few years, LAN support for these controllers has been provided by companies such as Netlink of Raleigh, North Carolina, Sync Research of Irving, Texas, and Ring Access of San Mateo, California. These companies' products allow existing 3270 controllers or other PU 2 devices to connect via SDLC and have their traffic sent across a LAN through reliable logical link control (LLC2).

The process is called SDLC passthrough if the SDLC traffic passes through to the other side of the link, as shown in Figure 11. It is termed SDLC conversion if the SDLC traffic is converted into LLC2 and presented to the 3745 on the LAN.

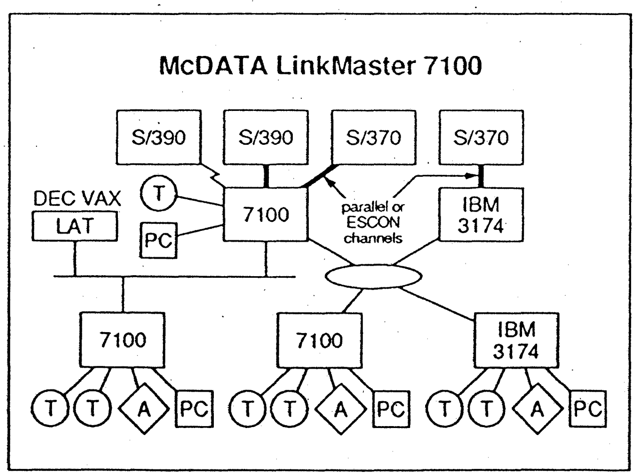

Figure 10

In addition to SDLC, Sync Research offers a product for QLLC-LLC2 conversion for access from X.25 networks. All these products can allow users to collapse multiple SDLC and LAN networks to a single network. Several issues still remain with these new products regarding performance impacts on both the SNA traffic and internetwork overhead.

Cisco Systems of Menlo Park, California, recently added SDLC conversion into a router. IBM offers a feature called data link switching on its 6611 multiprotocol router which supports SNA traffic over a router network through a process similar to SDLC conversion. These topics are further discussed in two recent *SNA Perspective* articles: "Data Link Switching on the IBM 6611" in August and "Optimizing SNA Traffic over Internetworks" in September. IBM encourages users to replace 3274s with 3174s because, even though SDLC passthrough or conversion allows these controllers access through the LAN, it does not allow them to upgrade their controller with any additional features.

APPN

IBM offers APPN network node support on the 3174 as a no-charge option with configuration support C. Upgrading the 3174 licensed internal code from configuration support B to C is $129 and APPN requires peer communications, another no-charge option. However, configuration support C also requires a second disk drive ($823), an additional 2 MB of memory (for a total of 4 MB) ($4,370), and a token ring adapter ($4,045). This makes adding APPN more expensive than it appears at first glance. A function called a hybrid link allows a 3174 to send both PU 2 and node type 2.1 traffic over the same parallel channel or SDLC link.

As discussed in the November 1992 *SNA Perspective* article, "Old Apps, New Nets: Supporting Dependent LUs across APPN," VTAM 4.1 provides APPN support for all existing 3270 controllers and other PU 2 devices. As long as they remain logically adjacent to the VTAM host or 3745 communication controller—that is, adjacent to its boundary function in all ways that are supported today with subarea SNA—VTAM 4.1 can deliver their dependent LU traffic to applications across APPN without any changes or upgrades and without APPN on the controller or other PU 2 device. (Even if the 3174 has APPN installed, it is ignored for dependent LU traffic with VTAM 4.1.)

In March 1992, IBM made a statement of direction that a future release of VTAM and of 3174 configuration support will contain a new capability, shown in Figure 12, called dependent LU server (dLS) and requester (dLR), respectively. The "Old Apps, New Nets" article also addresses dLS/R in depth.

The difference between the VTAM 4.1 support and dLS/R support is that, with dLS/R, there is no requirement that the 3270 controller remain logically adjacent to its boundary function. The dLR code on the 3174 APPN network node will act as a boundary function for its dependent LU. The SSCP control sessions will be encapsulated in an APPC session between the dLR and the dLS, but the LU-LU session data will run natively across APPN.

Several other 3270 controller vendors are considering APPN but have not made any product announcements. Since VTAM 4.1 can support existing controllers without APPN in them, *SNA Perspective* expects the other controller and gateway vendors may hold back on solidifying their APPN plans until IBM plans for licensing APPN and, perhaps, dLR become clear.

(continued on page 20)

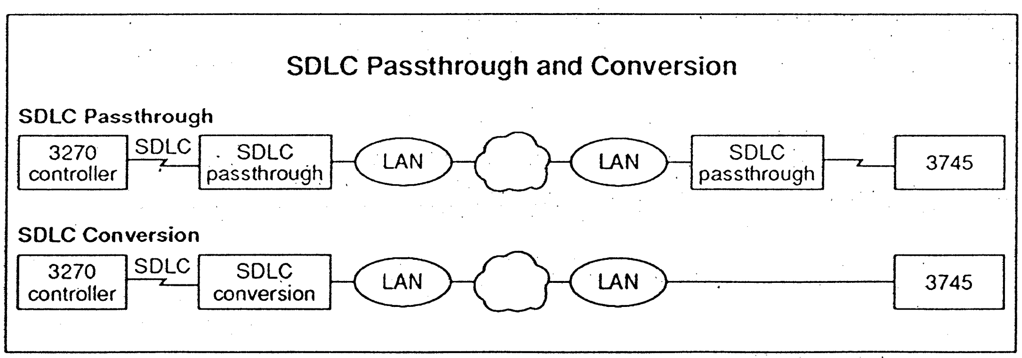

Figure 11

Architect's Corner

Out of Synch

by Dr. John R. Pickens

Conventional wisdom says that SNA cannot run over asynchronous links. Too slow. Too much overhead. Too unreliable. Synchronous modems and synchronous framing procedures are required—SDLC.

Yet IBM has recently proposed an architecture for asynchronous SNA—SNA-A—and is now supporting it in four products, NS/DOS, PC/Support (DOS), OS/2, and AS/400.

Why? What benefits can asynchronous SNA possibly offer? What has changed in the desktop/modem environment to enable it?

Also, is this just a flash in the pan? Or a significant architecture whose evolution should be watched closely? What about standards?

Note: the analysis that follows is based upon a document circulated at the APPN developers conference in August, 1992—"SNA-A: A Technical Description," dated August 5, 1992, by James J. Martin.

Reevaluating the Requirement

Certainly the traditional solution for isolated (single remote node) end systems has been to install a 2400/4800/9600-baud modem and run synchronous SDLC over the dial-up (or leased) line. The typical end point for the connection would be another synchronous modem front-ending a 3745-class device.

Isn't this traditional solution still adequate today? No.

This can be better understood by digging deeper into the characteristics of today's environment—particularly the desktop, modem, new media, and laptop.

First, the desktop itself. Almost every desktop system installed today has at least one asynchronous port (COM1) and most have two (such as a port for the mouse). The technology for such ports supports 9600-baud asynchronous hands down. Recent extensions are beginning to be shipped by vendors that enable 38.4-Kbps and 64-Kbps speeds also.

Next, the modem technology. A few years back 2400 baud seemed fast for asynchronous standards. But today 9600 baud is becoming commonplace; 19.2 Kbps is also common, and 38.4 Kbps is a short step away from becoming the de facto standard.

Next, new media technologies. Much ado is being made about wireless technologies (RF, infrared). Two interfaces are being offered, LAN and WAN. The LAN interfaces are based on variants of IEEE 802 standards (802.11, for example). The WAN interfaces are based upon the abstraction of the COM port—asynchronous again.

Finally, the laptop. Laptops are proliferating. Laptops also have COM ports. The dominant modem for laptops is asynchronous. Synchronous links require extra hardware and synchronous modems (which are more expensive because there aren't as many of them).

So, synchronous may have been the conventional way to support SNA, but asynchronous has become very appealing for two reasons—ubiquity and port cost.

The Asynchronous Requirement

If asynchronous is so pervasive, why not just do it? What are the technical challenges and requirements?

1. Reliable link—SNA requires a reliable connection-oriented service in the data link layer.
2. Efficient byte encoding—Since asynchronous communication uses less efficient framing (10 bits per byte), the requirement for efficiency is amplified.

3. Byte transparency—It must be able to operate in environments with 8-bit framing, XON/XOFF, control characters, and 7-bit framing (efficiently, I might note).
4. Negotiable—It must be able to negotiate and dynamically discover framing properties of communicating partners.
5. Layer transparency—The upper layers must see the asynchronous link in the same way as SDLC.
6. Standard—It must be a standard.

SNA-A Architecture

SNA-A, also called SDLC over asynchronous, is an extension to the data link layer. The method used is reminiscent of the IEEE 802 approach (minus SAP addressing) in which the data link layer is subdivided into two sublayers, the upper sublayer defining the SDLC elements of procedure, the lower sublayer defining the physical media interface (in this case asynchronous).

The lower sublayer handles such issues as transparency framing, negotiation at link initialization time, and checksum calculation. Some of the interesting properties of SNA-A include:

- The use of null-XID polling to negotiate common characteristics according to strict rules of precedence
- An 8-bit transparency scheme (e.g., public data networks) that efficiently packs seven characters into an eight-character stream (the last character contains the most significant bits of the previous seven)
- A character transparency scheme which (like BSC) uses a control escape octet to assure that information can be transmitted even if it conflicts with other control information

SNA-A is based on ISO 3309 (asynchronous HDLC), although its 8-bit transparency scheme is incompatible (the ISO standard changed after SNA-A was published) and extensions have been defined by IBM in order to support SNA link negotiation.

Open Issues and Evolution

How well does SNA-A meet the requirements? Pretty well, in my opinion. The architecture is transparent to upper layers, is flexible with respect to transparent framing requirements, and supports link negotiation. But there are a few open issues about the architecture and product support.

1. Standards alignment—As much as possible, SNA-A should be revised toward greater alignment with ISO 3309. Also, the link-initialization extensions should be generalized and submitted to OSI.
2. Specification ambiguity—During link negotiation there are many possible combinations between communicating systems—8-bit, 7-bit, ISO 3309 mode, framing and character transparency. Detail should be added to the current specification to cover all cases. This will improve interoperability.
3. Reliable links—Some modems already contain a reliable link sublayer—Microcomputer Network Protocol (MNP). In such cases, the SNA-A architecture works but is inefficient. Checksums are not required, for example, and SDLC retransmission logic is not required. A variant of SNA-A that supports MNP modems would be useful.
4. Product support—The current product support matrix is too limited. SNA-A support is needed in the 3745 and AIX. SNA-A implementations are required for 3270 emulators. Finally, a real opportunity exists to add SNA-A to SDLC conversion products (SNA-A to LAN).

SNA-A offers a chance to extend SNA services by an order of magnitude to asynchronous-based end systems—laptops, palm tops, remote systems. Other vendors have supported asynchronous links for SNA using gateways and customized host software. SNA-A can become the (de facto) standard way. If IBM expands its product support matrix, and if other system vendors incorporate SNA-A into their products, and if internetworking vendors provide SNA-A support in their products, conventional wisdom regarding the synchronous-only link requirement will be debunked for good. ■

(continued from page 17)

Other 3270 Support

ESCON Channel

Support for the enterprise systems connection (ESCON) channel attachment was announced for the 3174 when ESCON was unveiled in September 1990. McDATA is the only other controller vendor offering ESCON. The McDATA 7100 can support up to two channel connections from the 7100, either ESCON or parallel channel. In addition to increased speed, the ESCON channel supports attachment to multiple hosts through an ESCON director and allows attachment up to 43 kilometers away from the host.

Multiple Host Support

With the addition of the concurrent communication adapter, each 3174 can have up to three host interfaces; only one can be through a channel interface. Across X.25, a 3174 controller could access up to sixteen SNA hosts. From a downstream configuration with one LAN interface, a 3174 can connect to up to eight hosts across the LAN.

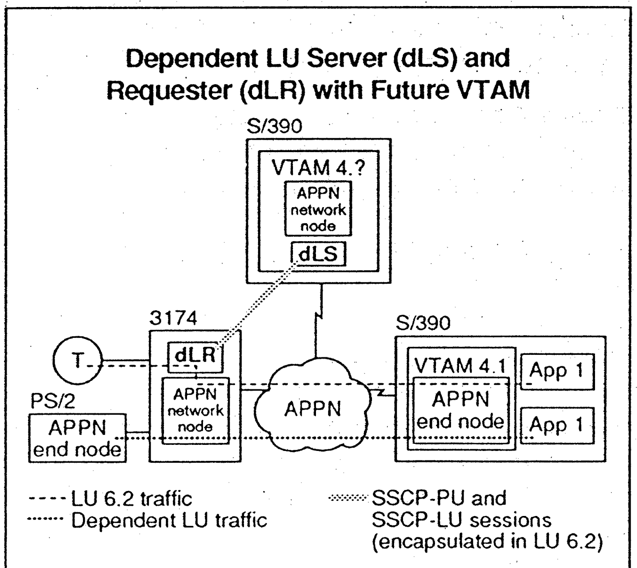

Figure 12

DDDLU

VTAM 3.4 includes support for dynamic definition of dependent LUs (DDDLU). With this support on the host, a 3270 controller or any PU 2 device can register its LUs at any time. Upon powering on and at any time a new LU device attaches to the controller, its vital product data can be sent to the host through a network management control vector (NMVT). With DDDLU, the LUs do not have to be statically defined in the host.

ISDN and Frame Relay

IBM offers an Integrated Systems Digital Network (ISDN) adapter for its 3174. Frame relay support has not been announced for any vendor's 3270 controller yet. But since frame relay is so important to IBM's networking strategy, *SNA Perspective* believes IBM might add this to the 3174. The 3174 can access frame relay today across a LAN through a frame relay gateway/router such as the 6611 or the RouteXpander/2.

Conclusions

Although shipments are dropping as the product line ages and competitive alternatives increase, keeping and enhancing 3270 controllers still makes sense in many environments. Companies should examine carefully the full cost of removing an existing installation, rewiring a building, and retraining MIS staff and end users. Users should be wary of solutions that require them to discard the old to embrace the new.

Although the 3174 was introduced in 1986, *SNA Perspective* does not expect a follow-on box that provides its functions in a similar way. Instead, within the next two years, we expect a new and very different platform that will provide some backward compatibility. ■

• 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, Basil Treppa • Marketing/Development: Alisson Walsh • Circulation Coordinator: Cheryl Roberts • Contributors: 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.