SNA Perspective

Volume 12, Number 3
March 1991
ISSN 0270-7284

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

Internetworking and SNA: Extended LAN Technologies

Enterprise networks in the future will be internets—sets of interconnected subnetworks which are the foundation for enterprise-wide computing. The growth of internetworking is challenged by the existing installed base of backbones—a major investment in hardware, software, and training which does not easily migrate to the emerging internetworking technologies that offer increased performance.

New technologies for LAN interconnection are changing the roles and paradigms of internetworking and enterprise-wide connectivity solutions. This article, another installment in *SNA Perspective*'s ongoing internetworking debate, looks at the emerging WAN technologies, including frame relay, fast packet, and switched multimegabit data service (SMDS). It examines the differences between these WAN technologies and traditional backbone architectures and evaluates their impact on SNA and changes that will be required of SNA if an IBM backbone is to remain a viable player in the new environment.

(continued on page 2)

APPC: The Long March

APPC has become the indelible refrain in IBM's communications anthem. Advanced program-to-program communication (APPC) appeared primarily as a marketing buzzword coined by IBM in 1983. It is often used as a generic reference encompassing a new generation of technology, capabilities, and products emanating from the logical unit type 6.2 (LU 6.2) and Type 2.1 node architectural extensions of SNA.

APPC embodies the spirit as well as the methodology of the new look and feel for SNA as it moves toward the 21st century: a dynamic, transaction-processing-oriented, peer-to-peer networking scheme. APPC is a significant part of IBM's ambitious computing model which revolves around cooperative processing, distributed data, interconnected heterogeneous processors, and total systems management. This article looks at APPC's beginnings, the resurgence of IBM's emphasis on APPC, the role of CPI-C, and the value of APPC LAN gateways and applications, and discusses some concerns about APPC's limitations.

(continued on page 10)

In This Issue:

(Continued from page 1)

Converging Trends

Several trends are converging to create a critical threshold that will change the spread and use of internetworking technology in the future, particularly in relation to power on the desktop and client/server architecture.

Usable Desktop Power

The first important trend is the increasing power on the desktop:

- Microprocessor price/performance has increased
- The availability of graphical user interfaces (GUIs) has improved ease of use
- Integrated applications are supplying a consistent environment

Although microprocessors have shown a steady increase in price/performance in the last decade, other elements have also come together now to create *usable* desktop power. The availability of GUIs such as Motif, OS/2 EE Presentation Manager, Apple Macintosh, and DOS Windows has simplified the use of highly functional desktop processors even for unsophisticated users. The development of integrated applications that combine spreadsheets, word processing, graphics, and communication makes it easier for users to perform different kinds of tasks in a consistent environment. These factors make it economical to invest in substantial desktop computing power throughout an enterprise.

Client/Server Architecture

Client/server architecture is beginning to gain the attention of large-scale enterprises. Client/server computing depends on a robust, solid internet service that provides enterprise-wide connectivity. A local client or desktop system can work with multiple servers concurrently to obtain data, processing, and other special services. The client handles the local presentation of the information, tailoring it to make it most useful for the consumer. The client also handles all the local coordination of information from various sources and the integration of that information into complex documents of text, graphics, and images.

Client/server architecture is shifting the role of midrange and mainframe systems toward the role of servers. These servers will manage and distribute information as well as provide specialized processing services for the clients. Most users will not directly access midrange and mainframe systems in the future.

The Extended LAN

Extended LAN is a new metaphor to describe internetworking at the enterprise level. The extended LAN has a desktop-centered viewpoint, providing access to all the needed resources across an enterprise. The development of new interconnection technologies extends the LAN across the entire enterprise, making any server available to the user as well as providing connectivity to public resources and other enterprises. Users at the desktop can access information in their local server on their workgroup LAN or from remote servers with little difference in user interface or performance.

The requirements for building an extended LAN as the foundation for enterprise connectivity include:

- Transparent reach—Information should be accessed in the same way whether it is stored locally or in a remote location across the enterprise.
- Minimal distance sensitivity—Behavior and performance should provide minimal sensitivity to remote references so that, although information is further away, the lag in response time is so minimal as to be almost unnoticed by the end user.
- Broadcast and multicast support—LANs are broadcast media (in which every system receives all the transmitted traffic and plucks its own traffic off the media). Extended LAN technologies should have similar support for broadcasting and multicasting (transmitting to a subset of all the possible destinations).

Extended LANs must match these local LAN characteristics in order to be a viable and effective part of the enterprise computing environment. The new LAN extension technologies will change the rules and concepts for building future enterprise internets. First, we'll examine the current status of WAN evolution and then discuss the new technologies that are changing the rules of the game.

Traditional Backbones

The initial backbones that interconnected LANs were the same as those that were originally designed to link two single systems across a backbone using packet- or circuit-switching technologies. Interconnected groups of systems have different characteristics from point-to-point connections, differences in:

- Computing power
- Data rates
- Variable behavior

Powerful groups of workstations on a LAN may have higher aggregate computing power than a midrange or mainframe system. This processing speed has been matched by much higher LAN data rates that may exceed 100 Mbps. A community of systems also shows more variable behavior than a single system supporting a small number of applications.

An ongoing dispute about wide area technologies, debated with almost religious fervor, pits those who advocate circuit-switching against those who believe that packet-switching is the most efficient means of interconnection.

Circuit-Switching

Circuit-switching is characterized by dedicated facilities that link two points across a backbone. Facilities may be dedicated as particular wires or trunks, or as a time-division multiplexed piece of a larger channel.

Circuit-switching is also characterized by relatively fast switching speeds. Any traffic arriving on a particular circuit is relayed on another circuit without processing delays since the route is already established. This performance has some drawbacks when particular session partners are not communicating—dedicated bandwidth is wasted when it could be used to satisfy the communication needs of other users.

Packet-Switching

Packet-switching, on the other hand, is characterized by dynamic bandwidth allocation—each packet receives the full bandwidth as it is switched. Many packets from various sources and intended for various destinations are serially interleaved on the same physical circuits. Each packet receives the full bandwidth for a small amount of time. The available bandwidth is distributed among those users who are actually communicating at any given time.

The packets that are multiplexed on the same physical trunk must be demultiplexed and individually routed by each packet switch. This additional processing results in a slower switching speed than circuit switching.

Both approaches have strengths as well as certain limitations. The new LAN extension technologies are being designed to take advantage of the strengths of each approach.

New Application Drivers

Emerging classes of applications have different characteristics from those traditionally supported by enterprise networks. For example, image processing is becoming increasingly important. Many organizations are scanning documents directly into their systems, capturing financial transactions, transferring circuit designs, processing purchase orders, and so forth. Image processing will require the transfer of multimegabit objects. The imminent introduction of handwriting recognition applications will make image processing even more important in the future.

For example, medical imaging is becoming increasingly important as a way to access medical expertise from areas that lack specialists. An X-ray can be digitized and transmitted to a distant medical center for diagnosis and consultation. A digitized X-ray would comprise approximately 2000 by 2000 pixels with an additional 12 bits per pixel for gray scale. Thus, a single X-ray would consist of approximately 48 megabits of data to be moved. If a traditional 9.6 kbps link were used, it would take approximately 1.7 hours to transfer a single picture—certainly not satisfactory in emergency situations. Using a T-1 link would reduce the transmission time to slightly more than half a minute, which would be much more appropriate for critical situations. Transferring other images may not be a matter of life or death, but customer satisfaction and business requirements often also dictate shorter transmission times.

Multimedia applications—combining voice, video, text, and graphics within a single integrated application—are also being introduced. Real-time visualization of computing processes is moving from the scientific community,where it was used for supercomputing, into the business community. Visualization helps give users an intuitive understanding of different business processes or simulations. Real-time visualization would require at least 16 megabits per frame, at a frame rate of 30 frames per second to sustain the appearance of motion. This is almost half a gigabit per second. Video compression techniques can substantially reduce this volume, but even with the best video compression techniques, visualization will require from one to two and a half times Ethernet capacity.

These kinds of applications can be characterized by the high volumes of data that must be transferred with minimum delay. Some applications, such as moving a single image, may be intermittent; others may require sustained rates when visualization or video are being used. The combination of low response time and high volume will require bandwidth on demand in order to effectively use expensive, high-capacity resources. This is markedly different from the traditional terminal-to-host applications where a user may enter a line or two of text and receive in return a screen at a time. These characteristics will require new communication technologies if the advantages of the LAN are to be extended throughout the enterprise.

Removing the Bottleneck

Traditional backbones, although they did provide connectivity, present a bottleneck for LAN interconnectivity. When compared to the transmission speed of a 10 Mbps Ethernet, a 56 kbps link is approximately 200 times slower. Moving large volumes of information with such a speed disparity severely impacts performance. Current T-1 technology, at 1.544 Mbps, reduces the disadvantage substantially but it is still a significant bottleneck for the emerging applications.

The advent of new WAN technologies is changing the basic rules of internet architectures. For the first time, these new backbone transmission facilities exceed the transmission capacities of most LANs. FDDI, at 100 Mbps, provides a powerful local backbone for interconnecting local workgroup LANs. T-3, at 45 Mbps, provides a high-speed wide area backbone which will soon be supplemented with SONET, at 150 Mbps. Compared to the 56 kbps circuits, SONET will be a speed improvement of almost 3000, three orders of magnitude in less than a decade! These new high-speed technologies can be used to build the first true extended LAN environments. Examples of the new WAN technologies include:

- Frame relay
- Fast packet
- SMDS

Frame Relay

Frame relay has been touted as the eventual successor to X.25, providing high-speed access to wide area networking facilities. Introduced by Stratacom and endorsed by many other vendors, including cisco Systems, DEC, and NTI, the frame relay standard is considered to be reasonably mature although some discrepancies exist between the American National Standards Institute (ANSI) and the International Telegraph and Telephone Consultative Committee (CCITT) versions of the proposed standard.

Advantages

Frame relay's T-1 speeds offer a substantial advantage over the 56 kbps limit of X.25. Frame relay is based on LAP-D frames to interface to a WAN switch. LAP-D is a subset of the standard high-level data link control (HDLC) protocol and is an important part of the integrated services digital network (ISDN) protocol technologies. All data is placed in a LAP-D frame and the addressing structure is used to specify a data link connection identifier (DLCI). In frame relay, the DLCI defines a target destination, or exit point, from the wide area switched network.

Current frame relay products are based on the permanent virtual circuit service where routes between pairs of entry and exit points are predefined. In addition, DLCIs can be assigned to multiple end points to supply a multicasting capability that is analogous to a LAN. Switched services could be installed by using a separate call signalling protocol between attached systems and the WAN. The Q.931 ISDN call set-up protocol, for instance, could be used. However, at this point, the standard only mandates permanent virtual circuit service.

Frame relay bypasses one of the layers used by X.25. Frame relay is a two-layer standard versus the three layers used in X.25 technology (see Figure 1). Within X.25, the LAP-B data link frames carry addresses that are basically fictitious—they are always the same for each attached system. LAP-B provides a reliable transfer service for encapsulated layer three packets. In contrast, the DLCI specifies a target address—eliminating the extra step of encapsulation and de-encapsulation. This provides a higher throughput between systems and wide area network services. Frame relay vendors claim a six-to-one cost/performance advantage over X.25.

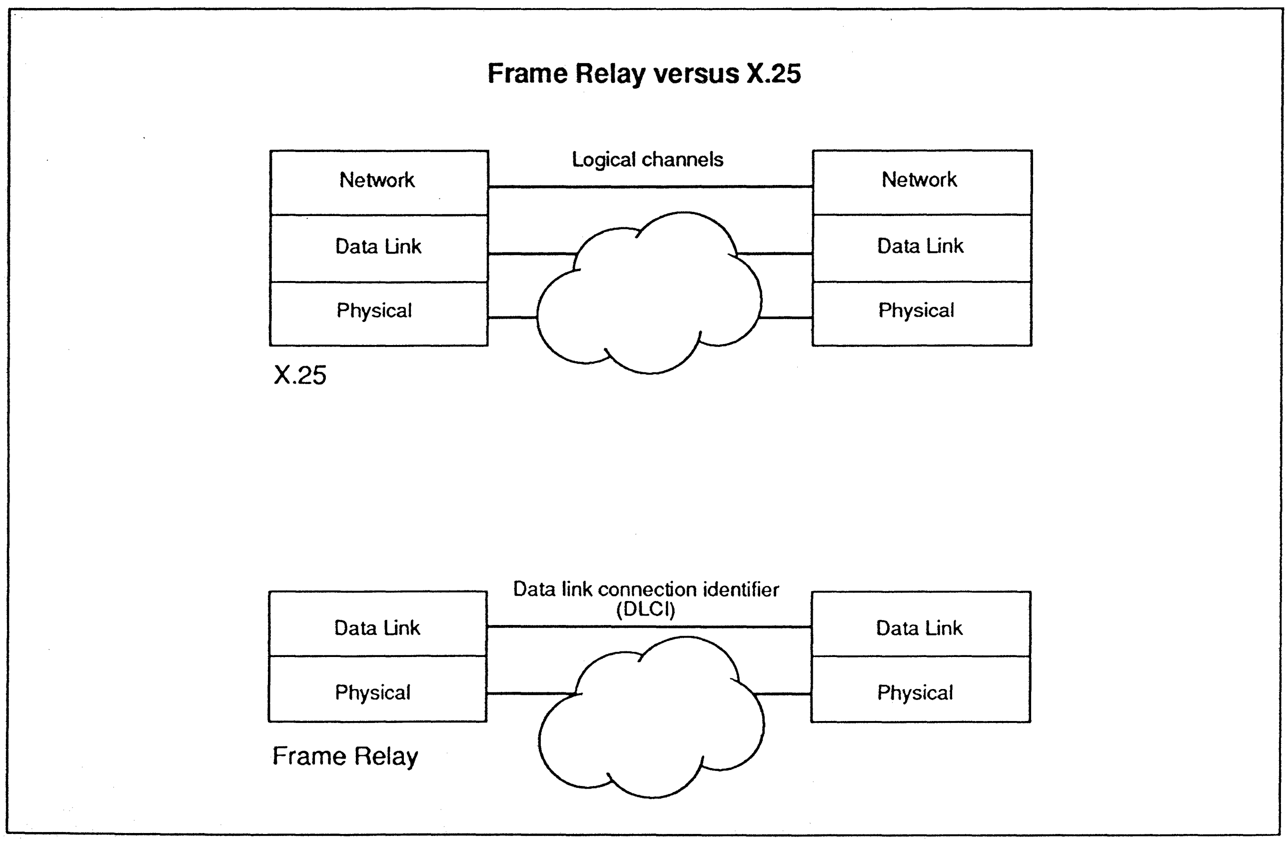

Figure 1

Limitations

There are also some limitations to frame relay. For example, the small DLCI field provides a limited set of addresses. Approximately 45 nodes can be supported if permanent virtual circuits are established between every pair. The small addressing capability has been considered best suited for private enterprise backbones. Some public carriers such as U.S. Sprint are offering public frame relay services in the first quarter of 1991.

Some discrepancies in the standards have made some users hesitant to use frame relay technology because they are concerned about interoperability problems. Vendors who participated early in this market began releasing products before the standards were fully stabilized, which leads to concern that some of these first products lack the management capabilities that are defined within the emerging frame relay standards.

Another concern is that frame relay may be too slow and too late. Other higher-speed technologies are moving to market more quickly than anticipated. This may restrict frame relay to the role of a niche player or cause frame relay to be relegated to the ranks of those technologies that experienced an initial boom but died out before they could become firmly established.

Fast Packet

A complementary technology has arisen with frame relay. Known as fast packet or fast packet-switching, it decreases switching delays across a switched backbone, thereby providing better throughput.

Traditional X.25 packet-switching networks use procedures based on the characteristics of copper circuits which have relatively high error rates. Because of this, each switch inside in the packet net exchanges acknowledgment messages, retransmits garbled packets, and resequences packets at each hop for every network connection. X.25 provides a highly reliable backbone service with a low undetected error rate.

However, signaling between the packet-switches consumes some network bandwidth, which cannot then be used for data packets. The additional processing, checking, buffering, and resequencing also slows the switching process. Holding packets that have arrived out of sequence until the proper sequence is restored also delays the transit of packets across the network.

In contrast, fiber-based backbones have an extremely low error rate so congestion rather than transmission errors will cause most packet loss. Fast packet-switching takes advantage of the reliability of fiber and uses a different algorithm for forwarding packets. Each switch simply moves the packet along to its destination as quickly as possible. Once the packet is sent, it is discarded; there is no holding-awaiting-retransmission or resequencing.

Any errors due to missequencing, transmission problems, and so forth, are handled by the end points. This, in a sense, is analogous to the classic roles played by the layer four transport protocol of TCP/IP's transmission control protocol (TCP) or by transport protocol class 4 in the OSI architecture. These protocols are able to detect sequence errors, retransmit undelivered packets, and provide a reliable end-to-end flow at layer four, relieving the lower layers of this responsibility. These capabilities must still be provided at some layer since traffic leaving a reliable fast packet backbone may still have to contend with less reliable networks on the way to its final destination.

The Future: Cell Relay

The ultimate successor to fast packet will be cell relay, or asynchronous transfer mode, which will be the fundamental packet-switching technology supporting broadband ISDN. Cell relay uses small, fixed-size cells to further increase the network switching speed. Since each packet is the same length, header processing can be optimized. Small packets also ensure that high-priority traffic is not backed up behind large, low-priority packets. Cell relay is an important technology which is expected to appear some time after 1992. Pilot projects are already under way. *SNA Perspective* expects that high-speed fiber, high-speed switches, and cell relay techniques will increase the capacities of current packet-switches from approximately 10-20,000 packets per second up to 100,000 packets per second and beyond within the next several years.

SMDS

Switched multimegabit data service (SMDS) is an exciting technology that is close on the heels of frame relay. SMDS offers some significant advantages, although it will be somewhat later to market. SMDS is currently a service offered by many of the regional Bell operating companies. Metropolitan area access of up to 50 km will be provided.

SMDS is based on the IEEE 802.6 metropolitan area network (MAN) access protocol, which is now an adopted standard. The 802.6 media access control (MAC) procedures, known as distributed queue dual bus (DQDB), is based on technology developed in Australia. DQDB is a technique for accessing high-speed MAN facilities usually through internet routers. Because it is part of the IEEE family of LAN standards, 802.6 uses the full IEEE addressing structure which provides compatibility for LANs and MANs as well as broadcast and multicast capability. An additional DQDB option will allow the full ISDN addressing scheme. SMDS can be part of the access structure for broadband ISDN in the future.

Each bus transmits data frames in a single direction (see Figure 2). The ends of each bus have special systems that generate the frames and handle timing and synchronization. The DQDB topology can be a bus or can be built into a ring by having a single system act as the head and tail for both buses.

DQDB is based on a slotted frame reservation technique, the basic principal of which is that a system can use any free frame for transmission This approach avoids contention for available capacity. If the number of free slots is relatively high, all systems will get adequate access without wasting any time and bandwidth resolving contention.

However, a drawback of this scheme is that it could give systems closer to the head of the bus a higher number of free slots than systems toward the tail. This is resolved by using a portion of the slots moving in the reverse direction for reservations. In this way, upstream systems have an idea of how many systems downstream require access and they must allow a certain number of free slots to pass without taking them.

Distributed Queue Dual Bus

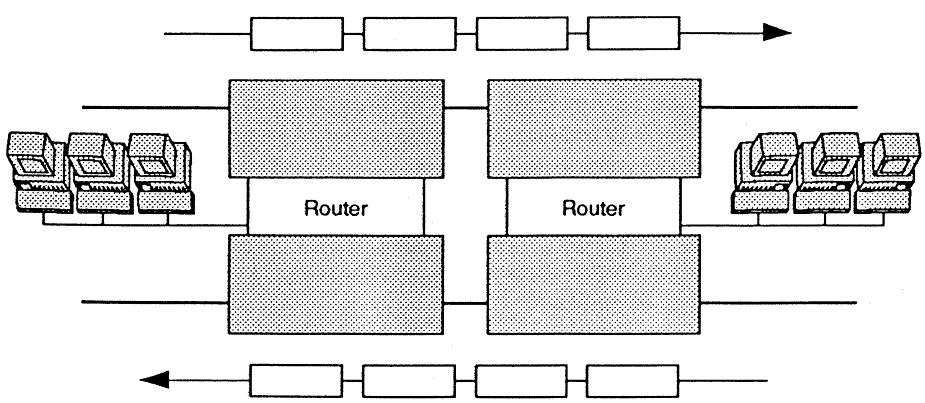

Figure 2

In addition, dedicated portions of each slot can be used for isochronous access to support real-time requirements such as video or voice. DQDB provides a relatively efficient access mechanism, although early reports indicate it may have a slight problem with very large frames.

SMDS will have an expanded role to play beyond MAN interconnection. Routers attached to the SMDS service can also interface to wide area backbones—particularly broadband ISDN in the future. This capability extends the scope of SMDS into a national or international means of interconnection.

Further, SMDS has a rich addressing structure that can support virtually unlimited numbers of attached routers and/or end systems. Thus, SMDS is targeted for environments in which many organizations can use a shared public backbone for internal interconnection and/or use the SMDS facilities as a means of interconnecting different enterprises.

Both frame relay and SMDS exhibit a new WAN adaptation of technology that was explicitly designed for LAN interconnection. Both provide high speed and the bandwidth-on-demand that are necessary to support the next generation of internetworking architectures. SMDS may be a quick successor to frame relay because of its much higher speed, its ability to access international standard facilities, and (the possible) tariff which would be based on the number of packets rather than fixed time or connect charges. If so, SMDS may become the major LAN interconnection technology for the second half of the 1990s.

Impact

The introduction of these kinds of technologies will have significant impact on enterprise information architectures. Removing the traditional penalty associated with using a backbone opens up new possibilities for the location and management of enterprise data. For example, LAN interconnect technologies can be successfully melded with the new generation of multiprocessor LAN superservers. A single superserver may be directly attached to a high-speed backbone and serve a community of local LANs rather than use a single server per LAN. This can leverage the cost/performance advantages of the superserver with the high-speed backbone. Some of the problems associated with distributing multiple copies of databases and keeping them consistent are also simplified by reducing the number of servers.

Other data distribution strategies may require rethinking. The trend in the 1980s was to distribute information from centralized locations. This strategy places information closer to desktop users and uses high LAN speeds. The new LAN interconnect technologies may reverse this trend since centralized data may now be easily accessed with lower time penalties.

Centralization simplifies management and the maintenance of data integrity. However, like most other strategies, we also may be changing the potential problem from the long delays associated with low-speed backbones to delays caused by congestion at the data site. In addition, centralizing sources of data exposes enterprises to a single point of failure. If that data becomes inaccessible for any reason, the impact will ripple throughout the entire enterprise. Multiple copies of data throughout the organization provide redundancy that protects against such failures.

In any case, users must keep in mind the fact that propagation delays remain constant even while transmission speeds are increasing. It still takes a finite amount of time for electrical or optical signals to travel over a distance. Additional time is also added by even minimal switching overhead. When small amounts of data are sent long distances, a substantial increase in bandwidth may not be directly reflected in better response times. The LAN interconnect technologies will leverage centralized approaches for extremely large amounts of data sent in large blocks over long distances. Users must carefully design their information architecture, message sizes, data distribution, and so forth in order to leverage these technologies in the most effective way. Where centralization is not feasible or desirable, information will still be distributed among local workgroup servers.

SNA's New Role

The rapid shifting of the enterprise computing paradigms brings up interesting questions about what role SNA will play in the evolution of enterprise computing. IBM has been slow to shift from its traditional mainframe-centered and hierarchical networking approaches. Certainly, the inertia imposed by a large installed base makes rapid changes of direction difficult for both vendor and customers. At the same time, IBM may have been trapped to some extent by its own success and a parochial view of the networking world. But even IBM has recognized that the subarea networking approach is not the solution for the next generation of internets.

Subarea networking depends to a large extent on preallocation of resources based on pacing windows for individual LU-LU sessions. This type of preallocation is much more analogous to circuit-switching behavior with dedicated capacity and limited dynamic capability, although IBM refers to its backbones as packet-switched. IBM is expected to make the subarea backbone more dynamic by using dynamic pacing windows in future releases.

At the same time, IBM has maintained a closed backbone that is controlled by proprietary protocols. Encapsulation techniques have been applied when the SNA backbone has been used to move other traffic. Products such as IBM's XI interconnect for public data networks and its recently announced LAN interconnection product are examples. The LAN interconnection product, for example, interconnects token rings across an SNA backbone. However, communication is controlled by an LU 6.2 session which imposes distributed transaction processing behavior on a backbone when the goal is simply to move data between two points. This approach imposes high overhead which impacts performance. The LU session constraints also limit the ability to deliver bandwidth on demand.

In contrast, several vendors are opening their backbone by using multiprotocol routing technology that allows a wide range of protocols to share the same physical facilities. Vendors such as NCR, Hewlett-Packard, Wellfleet Communications, cisco Systems, and Digital Equipment now offer multiprotocol routers. Multiprotocol routers also incorporate several backbone technologies so that frame relay and SMDS can be incorporated easily into existing internets. A backbone that can support several protocols and technologies can deliver data at the network's best speeds rather than being limited by a single higher layer protocol's constraints such as those imposed by LU 6.2.

IBM must make changes in its basic networking philosophy and products if SNA is to have a role in the emerging backbones. Certainly there will always be SNA backbones in some portion of IBM's market simply because some customers are conservative, depend heavily on IBM, or don't have the requirements that the new technologies can fulfill.

However, many other organizations may find it feasible to migrate from the proprietary IBM backbone and replace it with technology that is still capable of delivering SNA protocols. As a case in point, several multiprotocol router vendors, including cisco Systems, Wellfleet Communications, Vitalink Communications, and Proteon, have announced routers that can encapsulate SNA packets and *tunnel* them across different kinds of backbones. These products may start making significant inroads into IBM's traditional enclosed backbone niches.

Recently (including at the ComNet show in January 1991—see *IBM Announcements* in this issue)—IBM has begun to discuss its next generation architecture which it claims will provide superior multiprotocol routing at gigabit speeds. However, IBM has indicated it will be three years or more before these products will appear in the marketplace. Since viable alternatives are already available and more are on the way, it is not clear how significant IBM's announcement will be, particularly to leading edge users. Further, it is not clear that IBM's technology will be competitive when it arrives.

If IBM decides to build these products, they will be first generation products and will suffer the usual problems of being back on the learning curve—it takes time to learn and then refine products. However, IBM could cross license multiprotocol routers for LAN interconnect technologies from experienced vendors who have more mature products.

SNA Perspective believes the future of the SNA network's role within the new enterprise will require an inside-out perspective. Traditionally, IBM has maintained an SNA-centric viewpoint of the enterprise network whereby all traffic is centered around a proprietary, enclosed backbone. IBM must turn that perspective inside out and view SNA as only one of the kinds of traffic that will be carried on backbones designed to interconnect LANs. How quickly and how thoroughly IBM can move to the forefront of the evolving architectures is still an open question.

It is also possible that IBM is not deeply interested in the internetworking developments because of business priorities. There will be pressure for commodity-based pricing of multiprotocol routers as the internetworking market continues to mature. Perhaps IBM will stay above the internetworking fray and concentrate on enhancing its mainframes, strengthening NetView, and focusing on strategic applications. (See *Architect's Corner* in this issue.)

Users should begin to educate themselves about the new LAN interconnection technologies, the new architectural alternatives, and the organizational opportunities that they offer. Investigating viable alternatives to SNA backbones will be useful. IBM will have to prove its ability to adjust and adapt through the timely delivery of high-quality, competitive alternatives to the products that are already on the market. ■

(Continued from page 1)

A Faltering Start

SNA Perspective has never doubted the strategic significance, the technical merit, or the potential market impact of APPC. Nonetheless, APPC was far from an instant success. Facilities or applications that exploited APPC were few and far between. For nearly five years, until around 1988, APPC appeared to be just a futuristic architectural concept. IBM and third-party developers only seemed to be paying it lip service (see *SNA Perspective* September 1990 "LU 6.2 Growing More Slowly Than Anticipated"). IBM's implementational support for it, particularly in the crucial S/370 mainframe arena, was sparse and erratic.

APPC services offered during the mid-1980s on the S/36, the S/38, and the PC were each modeled on a product-specific subset of the architecture. Mutually incompatible, idiosyncratic programming interfaces were the result. True LU 6.2 protocol support on mainframes was restricted to CICS/VS. Document distribution utilities based on SNADS and DIA and embedded within IBM office automation systems such as DISOSS and Personal Services/3x were the only real users of LU 6.2.

Resurgence Within IBM

There has, however, been a upturn in IBM support of APPC over the last two years, culminating in the landmark APPC/MVS announcement in September 1990. A standard, built-in, high-level language, APPC interface will be available on IBM's flagship MVS/ESA operating system as of March, 1991. Ironically, *SNA Perspective* believes the much-maligned systems application architecture (SAA) has been responsible in part for this renaissance.

SAA's mission has been reclassified as the formalization and rationalization of client/server cooperative processing scenarios. Applications replete with intuitive graphical user interfaces, executing on powerful programmable workstations and freely and efficiently interacting with server applications on hosts across internetworked LANs and WANs to access or update large pan-enterprise databases are the future of information systems. A prerequisite foundation for the deployment of such distributed, functionally cleaved applications is a solid, peer-oriented, interprogram communications capability with powerful data flow control, error recovery, and inherent transaction check-pointing facilities. LU 6.2, in concert with Type 2.1 peer-to-peer nodes, was the obvious choice for IBM.

CPI-C—The Universal Interface

To enforce cross-system consistency—a primary goal of SAA—and to revitalize the flagging credibility of APPC, IBM unveiled, in October 1987, a new application programming interface (API) to APPC—the common programming interface for communications (CPI-C). It is the only communications-related API currently defined within SAA. It is now promoted and positioned as the single, universal programming access point for IBM networking as a whole, not just for SNA but for OSI as well.

The scope of CPI-C was extended in mid-1988 to embrace OSI distributed transaction processing and OSI networks. Now, with CPI-C, the same communications applications can be both SNA and OSI compatible. Software developers, especially in Europe, no longer need to agonize whether their future applications should be written to work with SNA or OSI (see Figure 3). CPI-C, in theory, provides a compelling coexistence and migration path. CPI-C, initially with just APPC support, will be available by the middle of 1992 on MVS/ESA, VM/SP, ASA/400, CICS/ESA, and IMS/ESA Transaction Monitor.

APPC LAN Gateways

With CPI-C bolstered by native LU 6.2 and node Type 2.1 integration support in ACF/VTAM Version 3 Release 2, there is now comprehensive coverage for APPC on IBM hosts as well as on midrange systems, particularly the AS/400. To ensure the effective exploitation of these host APPC facilities, an impressive array of APPC services and interfaces is now provided on PCs, PS/2, UNIX workstations, and LAN gateways.

The APPC-based LAN gateways are of special interest as they herald the next wave of LAN-to-mainframe interoperability technology. These APPC gateways can free workstations users from the limited functionality and configurability options of 3270 LAN gateways. They will now be able to fully utilize the processing, graphical presentation, and disk-storage capabilities of their workstations by being able to interact with host systems on a peer basis rather than having to continually masquerade as terminal users. The advantages of such distributed—rather than single-point—processing can be seen in remote database access, transaction processing, mathematical modeling, or file distribution applications.

APPC LAN gateways for MS-DOS and OS/2 servers are now available from a variety of well-known vendors, including:

- IBM—OS/2 SNA LAN gateway (OS/2 only)
- Digital Communications Associates Inc. (DCA)—Select Communications Workstation (OS/2 only)
- Eichon Technology Corp.—Access/TIC (DOS and OS/2
- Network Software Associates Inc.—AdaptSNA LAN gateway (DOS only)

An Overview of CPI-C

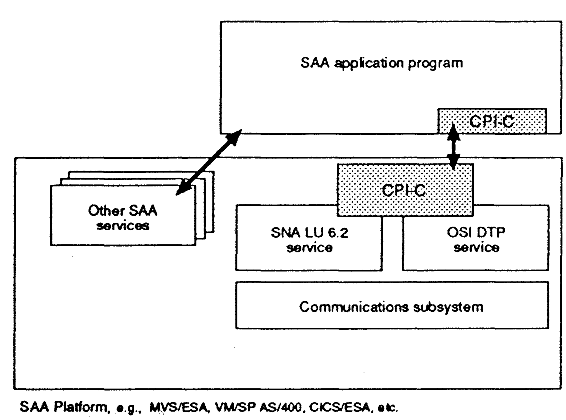

Figure 3

The Anatomy of an LU Type 6.2

LU 6.2 is the latest manifestation of SNA's architected services and protocols for realizing efficient end-to-end communications between application programs. It currently complements but will eventually replace the other SNA functional subsets, each of which is geared for a specific application-to-terminal, application-to-printer, application-to-I/O device, and terminal-to-terminal configuration such as LU-LU session types 1, 2, 3, and 4. Thanks to the ubiquitous microprocessor, most intersystem interactions in the future will be program-to-program as I/O devices and terminals become front-ended by control programs.

Another key feature of LU 6.2 is its protocol boundary. LU 6.2 includes the first and only formally specified architectural interface to SNA. Prior to LU 6.2, the LU interface through which end users gained access to an SNA environment, so that they could communicate with other end users, was implementation-specific and varied widely in structure, functionality, and scope.

As shown in Figure 4, the LU 6.2 architectural model consists of:

- A set of LU 6.2 services
- The LU 6.2 protocol boundary
- A set of LU 6.2–based service transaction programs ■

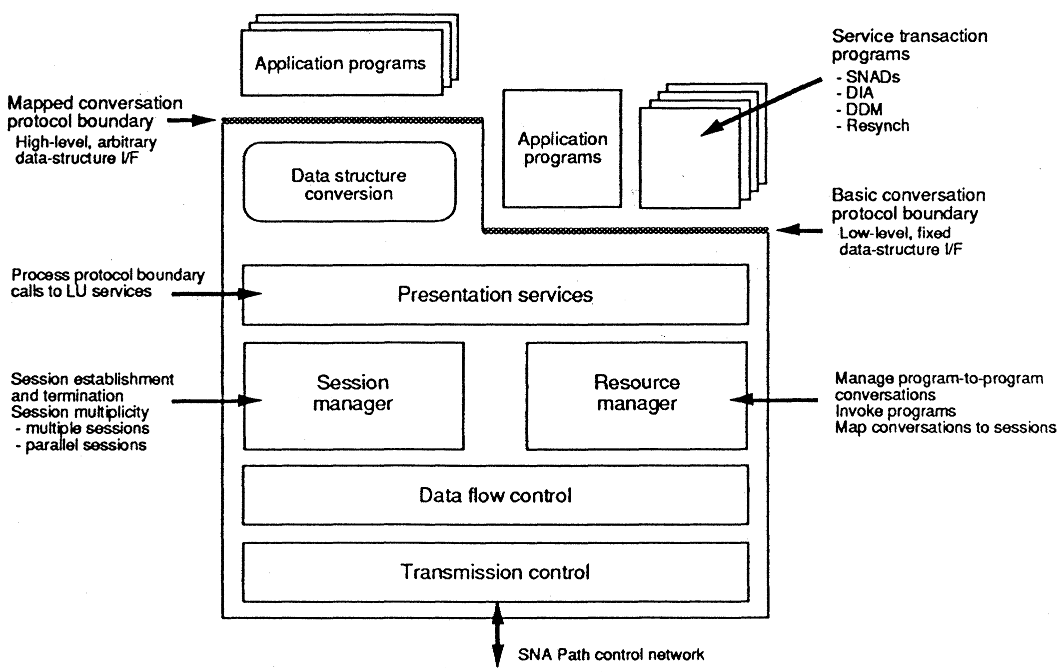

Figure 4

- Novell Inc.—NetWare SNA Gateway (DOS only)
- 3Com Corp.—3+Open Maxess SNA Gateway (DOS and OS)

An APPC LAN gateway should support a repertoire of LU 6.2 functions including both mapped and basic LU 6.2 protocol boundary verbs, parallel sessions, and as many multiple LUs as possible. These functions may be complemented by a Type 2.1 node implementation that permits peer-to-peer operation with mainframes, AS/400s, and S/3xs. (S/370 hosts with ACF/VTAM Version 3 Release 2 now offer some support for Type 2.1 node connections, with Type 2.1 node-resident PLUs, parallel sessions, and establishment of sessions that are network address independent. Such support was previously only available on AS/400, S/38, and S/36.) Ease of use, reliability, resilience, and serviceability are other important qualities to consider in an APPC gateway, as well as diagnostic, monitoring, configuration, network management features, and NetView compatibility. The use of IBM's standard API will improve the likelihood that IBM and third-party APPC/PC applications can be used without modification.

The APPC client applications interacting with host counterparts through a gateway would typically execute on a LAN-attached PC or PS/2 running OS/2 or DOS (see Figure 5). The APPC applications interact with the gateway via an APPC-compatible API. If the gateway does not require a dedicated server, APPC applications could also be run on the same PC or PS/2 in which the gateway is installed. The client PCs may be attached to the gateways through a variety of LANs, depending on the gateway. The gateway can then be connected to a host system via either an SDLC link or a LAN.

Some gateways offer a split protocol stack, with shared APPC and 3270 gateway capability to support migration and coexistence. With this feature, users can enjoy a nondisruptive and cost-effective migration path from the current 3270-based applications to the new wave of APPC applications without having to replace or upgrade their existing gateway. It also ensures that, for users who wish to continue using 3270-based sessions particularly for low processor usage, highly interactive applications can peacefully coexist and share the same remote access resources with APPC users.

Typical APPC LAN Gateway Configuration

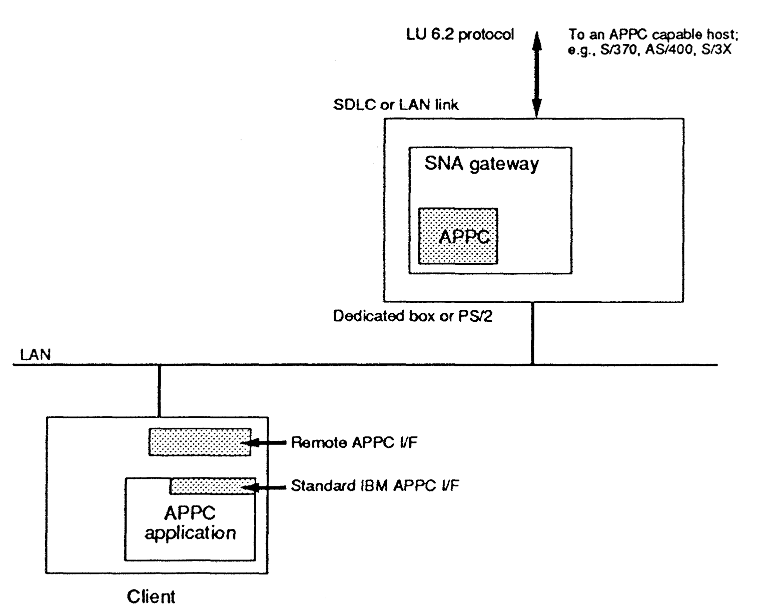

Figure 5

Selected APPC-Related Definitions

ACF/VTAM: Advanced Communication Function/Virtual Telecommunications Access Method, the S/370 communications subsystem required to realize SNA in S/370 host-centered networks.

APPN: Advanced Peer-to-Peer Networking; currently an AS/400 and S/36 product-specific extension to the Type 2.1 node architecture which epitomizes dynamic configuration and route selection in peer-to-peer SNA networks of the future.

Conversation: A short-duration, end-to-end, application program-to-application program logical pipe across which LU 6.2 interactions are performed. Conversations are serially multiplexed into the underlying SNA sessions between the LUs serving the various applications.

DDM: Distributed Data Management, an API-independent architecture for transparent access or update of a remote file or database which is implemented in the SNA environment as an LU 6.2 service transaction program (STP).

DIA: Document Interchange Architecture, a connection-oriented synchronous (i.e., real time) document distribution and document library management architecture which is implemented as an LU 6.2 STP.

LU: Logical Unit, a fundamental SNA construct. End users—application programs, terminal operators, or I/O devices—gain access to and interact with each other through services and protocols provided by LUs.

LU Type 6.1: The nonarchitected program-to-program predecessor of LU 6.2, pioneered by the CICS/VS Intersystem Communication (ISC) facility in 1980. LU 6.1, which is still used by the current releases of IMS/VS, was the prototype of most of the features now associated with LU 6.2.

Parallel Sessions: Multiple, concurrently active SNA sessions between the same pair of LUs, which are typically used to increase the capacity of simultaneous interactions possible between the LUs.

SNADS: SNA/Distribution Services, a connectionless, store-and-forward–oriented, asynchronous (i.e., non-real-time) message or file/document distribution architecture which is implemented as an LU 6.2 STP.

Service Transaction Program (STP): A standard systems utility of general use to other APPC applications, typically supplied by IBM. Examples include DIA, SNADS, and DDM. ■

APPC Applications

Buoyed by IBM's surge of support for APPC on both host and workstation clients, third-party software vendors and IBM itself have been joining the APPC applications bandwagon. Several APPC applications have already hit the market, and *SNA Perspective* expects many more to follow over the next eighteen months.

Consistent with market needs, the first wave of these commercial APPC products are addressing distributed database management and data distribution requirements. Not surprisingly, these are two areas ideally suited to make full use of the peer-to-peer, shared-processing functions of APPC. SQL-based PC LAN database servers that interoperate with IBM's host resident DB2 via APPC are already available in the form of Oracle Corporation's Oracle Server and Gupta Technologies' SQLBase Server. IBM has also intimated that its OS/2 EE database manager will be enhanced to provide a similar capability in line with its recently finalized SAA distributed data architecture.

On the data distribution front, Spectrum Concepts has led the way for over two years with its family of XCOM 6.2 products. XCOM 6.2, which is available on a wide range of diverse platforms including MVS, VM, PCs, S/3xs, DEC VAXes, and Wang VSs, offers a generalized file interchange capability. In the area of office automation, Soft-Switch Inc., a long-time leader in multivendor document interchange technology, is now offering a range of APPC- and SNADS-based, LAN-to-host interoperability products with its SNADS Gateway family which includes SNADS Gateway/3+Mail, SNADS Gateway/Banyan Mail, and SNADS Gateway/MHS (Novell). IBM also entered this host-to-workstation file/document interchange with the recently announced SAA delivery manager products.

These function-specific APPC products are complemented by a range of APPC-based application enabling technologies designed to facilitate and expedite the development of APPC client-server cooperative applications. The key applications among these are IBM's OfficeVision family with its application integration facility, IBM's OS/2 EE EASEL product, and the recently announced IBM SAA application connection service.

Limitations

The future for APPC has been enhanced by the new APPC products, SAA, CPI-C, OfficeVision, and APPC/MVS. It is, however, worth noting that APPC's LU 6.2 architecture still has a few drawbacks. LU 6.2 was initially developed expressly for program-to-program, batch-mode transactions for which response time is not critical. Thus it is still an inherently half-duplex protocol that favors an asynchronous (i.e., noninstantaneous) data-batching technique for end-to-end data exchange. The onus for torquing LU 6.2 to enable it to be effectively used in response-time critical, highly interactive applications is currently left to the expertise and resourcefulness of individual software developers. This is not optimum. There are some indications that IBM will introduce certain architectural and implementational extensions in the midterm future to facilitate interactive-mode APPC.

The complexity of the LU 6.2 protocol boundary (i.e., API) with its asynchronous nature has also been another area of concern. There had been much speculation and eager anticipation that IBM would rectify this with an LU 6.2 based remote procedure call (RPC) scheme along the lines of the RPC schemes so prevalent in the UNIX world. While this could still happen, the likelihood has diminished considerably with the announcement of APPC/MVS and CPI-C on CICS, IMS, and AS/400. Given the aggressive promotion of these SAA interfaces, IBM would obviously not want to muddy the waters by indicating the possible availability of another interface any time soon. ■

Architect's Corner

APPN '91: Tempest in a Teapot

by Dr. John R. Pickens

With all the recent talk about APPN, I am reminded of an experience which bears repeating. About four years ago one of my sleuthful consultant friends found an APPC bug in RT/PC AIX operating system. After hammering on IBM development, he finally succeeded in obtaining a special bug-fix release. During installation of the bug fix, he was surprised by a configuration screen which popped up, asking for APPN configuration parameters. "Aha," said he, "I deduce that APPN for AIX is in the offing." This was 1986. Little did he know...

Four years later, the official APPN communications mill finally begins to roll. IBM schedules events at ComNet '91 and elsewhere to promote the coming APPN story. Hints are made of OS/2-based APPN and of licensing arrangements with vendors. Yet major holes remain. APPN for the subarea backbone is in the distant future. APPN for mainframes is nowhere in sight. Even AIX, the source of my consultant friend's APPN anecdote, still has no announced date for support.

So what's all the hoopla? Has "new SNA" arrived?

No. Just a sneak preview.

Let me explain. Certainly aspects of the promised APPN enhancements are technically interesting. Roll-out on the OS/2 platform (whose development recently moved from Austin to Raleigh) is one important factor. Offering CPI-C as the upper layer interface to APPN (APPC) is another. And refining APPN to offer more system defaults for improved auto-configurability is yet another.

But much more is needed and expected:

- LU 0 and LU 2 encapsulation
- Subarea and mainframe support
- Virtual LAN support
- OSI convergence
- Dynamic LAN-layer name resolution protocols

LU 0 and LU 2 encapsulation—This allows PU 2.0 devices (e.g., banking systems and 3270 terminals) to attach to APPN network nodes, be routed across an APPN backbone, and subsequently be routed to subarea (3745) connections. Implementation of this feature requires a session encapsulation technique (sometimes called "tunneling" in the internetworking vernacular).

Subarea and mainframe support—Current subarea nodes and mainframes can only participate as PU 2.1 LEN nodes, with static preconfiguration only. What is needed is support of APPN directory and routing and topology update protocols within the subarea network and mainframe environment.

Virtual LAN support—Support of technologies such as frame relay and ATM/SMDS as valid SNA/APPN transports is needed. These converged virtual LAN technologies eliminate the need for much of the APPN routing function, and allow configuration (in theory) of large APPN networks.

OSI convergence—Not yet announced for APPN is convergence with OSI, first at the CPI-C layer and subsequently at the OSI/SNA routing layers.

Dynamic LAN-layer name resolution protocols—Current APPN (i.e., on the AS/400) requires static configuration of adjacent nodes. For example, in a large LAN with a hundred adjacent APPN network nodes, each node must be preconfigured with the name and address of the other 99. With the emergence of large LAN-layer backbones, such as FDDI, this technique does not scale up very well.

This final requirement can also be illustrated with an anecdote. Two years ago in a conversation with an

IBM SNA architect, I noted the above-mentioned restriction—the requirement for static configuration of adjacent nodes. He gave me a puzzled look, paused in deep thought, and then commented, "But I thought we gave the APPN developers an architected, broadcast-based, dynamic configuration function." So, once again, a gap exists between architecture and implementation.

How many of the above features will be delivered in the upcoming OS/2 releases? Time will tell. Possibly the LU 0/LU 2 encapsulation and dynamic LAN-layer name resolution protocols will occur in the OS/2 release. The others are probably further off. The real APPN storm, I believe, has yet to hit. ■

IBM Announcements

IBM at ComNet

At the Communications Networks '91 trade show and conference in Washington D.C., IBM gave a briefing on its directions in SNA, particularly with regard to advanced peer-to-peer networking (APPN). Little of the information was new; in fact, most of the slides were from an IBM presentation to the GUIDE users group in November 1990. IBM did add that it would announce and ship support for APPN on OS/2 within twelve months.

SNA Perspective believes it is useful to our readers to reexamine IBM's view of peer networking. The growth of workstations and nonmainframe nodes, which IBM refers to as small systems, in the SNA environment is a driving force in the evolution of SNA itself. APPN has developed within IBM as a means to support small systems, as first described in an IEEE paper dated May 1985 entitled "SNA Networks of Small Systems." IBM listed the small systems network requirements:

- Easy to install and use
- Decentralized peer control
- Arbitrary topology support
- Connection flexibility
- Interworking with subarea SNA
- Design simplicity
- Continuous operation

An additional attribute IBM cannot ignore is protection of the customer's investment by allowing peer networking to function on existing hardware, to build from existing SNA skills, and to provide seamless peer and subarea integration. In fact, this need to provide investment protection is the most difficult challenge to IBM in providing peer networking (see Figure 6).

To meet these requirements, IBM has developed node Type 2.1 (T2.1), also sometimes referred to as physical unit type 2.1. The T2.1 node allows two nodes to communicate without the need for VTAM intervention. (See *SNA Perspective* January 1990 for an in-depth look at how T2.1 works within subarea networks.) Such network communication (for example, between two PS/2 computers) is currently provided under IBM's low-entry networking (LEN), announced in 1986. However, LEN is at present point-to-point within a given subarea, and still requires static preconfiguration. With APPN, the communication is end-to-end. Currently, LEN implementations are available for AS/400, OS/2 EE, VTAM/NCP, S/36, S/38, S/1, S/88, RT PC, 3820, and APPC/PC. VTAM has had LEN support since 1987.

IBM intends to move SNA further toward APPN. APPN is built upon the T2.1 node and provides an extension of this architecture. IBM refers to LEN as an APPN-enabling function, as are recent VTAM enhancements of dynamic LU definitions and the multitail backup LU-LU sessions.

APPC and CPI-C

Although LU 6.2 provides the facilities for advanced program-to-program communications (APPC), it is a protocol set for the accessory network. Different LU 6.2 interfaces for different systems were developed independently. IBM soon saw that it needed to provide a common interface for all LU 6.2 devices, which became the common programming interface for communications (CPI-C) (see article on APPC in this issue). IBM plans CPI-C availability for OS/2 during 1991.

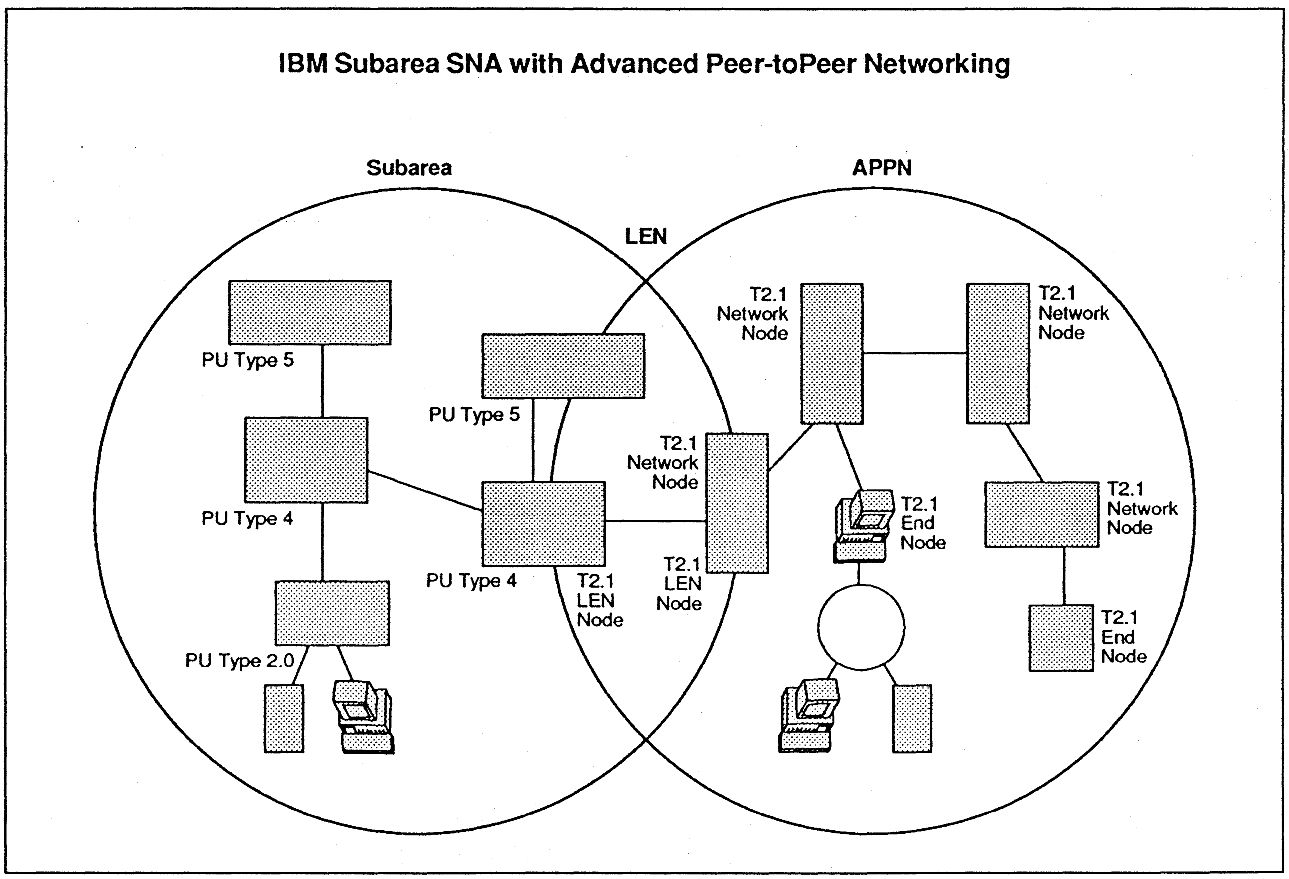

Figure 6

APPC and APPN

IBM used the analogy of a car to explain its view of the difference between APPC and APPN. APPC can be seen as the instrument panel—the end user interface for applications—and CPI-C provides a standard layout for the instrument panel in cars.

T2.1 nodes represent the engine, providing the network support for APPC and CPI-C. All T2.1 nodes will be divided into LEN nodes, end nodes (EN) (a LEN node with extra functionality), and network nodes (NN) (an end node with additional functionality). In the car analogy, LEN is a basic engine, EN represents a bigger engine, and a NN can be seen as a turbocharger.

The primary difference in functionality between ENs and NNs is that, while a T2.1 EN can *register* resources with the network, a T2.1 NN can *request* a directory of network services. ENs must have T2.1 and LU 6.2 support.

According to the IBM presentation, the three key functions of APPN are:

- Dynamic topology
- Directory functions
- Routing through the network

With APPN, each T2.1 NN maintains a copy of the network topology database. The T2.1 NNs will exchange topology update messages to keep all NNs informed about the current network topology. This means that users dynamically reconfigure networks without regenerating the network. Once set up, each T2.1 node informs the network of its existence and location; an EN only informs its adjacent NN.

The APPN directory allows resources to be defined only once. Resource locations are discovered via *locate* messages to all NNs and stored in the directory. Resources that are moved can then be noted on the network with minimal redefinition.

NNs can get information on the location of any given application, all available routes to that application, and the characteristics of each route. An EN would request a session with an application via a NN by sending a BIND request (see Figure 7). The NN would use the topology database, and consider the class of service required with the request to determine the route to the application. If the location of the application were not known, the NN would first send a locate message to update the directory.

APPN Routing

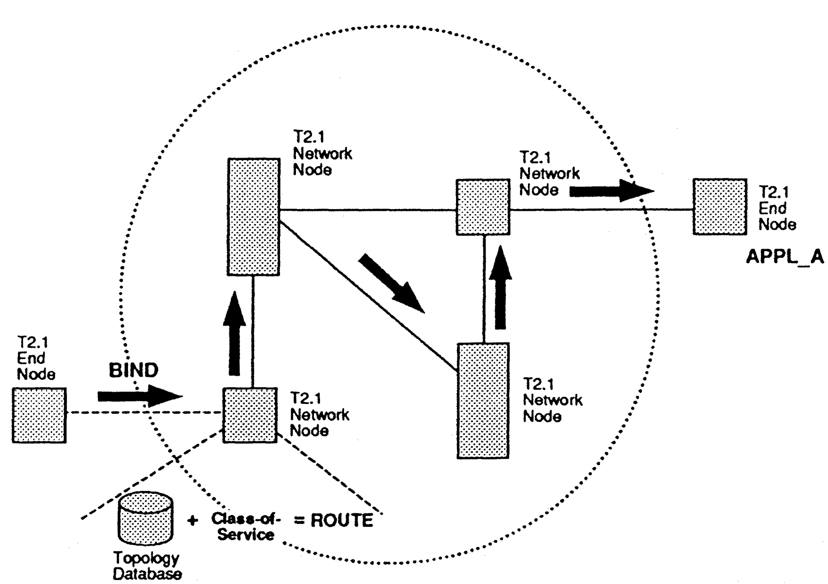

Figure 7

Routes are then dynamically determined with the following information:

- Address information from the directory
- Topology information from the database
- The class of service specified

The Future of APPN

At its presentation, IBM made several statements regarding the future of APPN.

- DOS PCs can be attached today as LENs and in the future can be T2.1 ENs. However, their EN functionality will be limited—they will not have the ability to dynamically register with the network. Specialized software available from IBM will be required to register DOS PCs with the network.
- T2.1 ENs will not do searching—no traffic flow-through via ENs will be provided.
- IBM intends to eventually support the following network links between NNs, in any combination: switched or leased SDLC, X.25 virtual circuits, X.21, Token-Ring, and Ethernet.
- OS/2 currently has LEN support. APPN for OS/2 will be announced and shipped in 1991.
- OS/2 Extended Edition will get both T2.1 functionality and a CPI-C interface.

The value of APPN to the end user is in the area of distributed, cooperative, and client/server computing:

- CPI-C provides a standard application interface
- LU 6.2 provides a peer-oriented protocol
- APPN provides a distributed peer network protocol

See *Architect's Corner* in this issue for commentary on IBM's APPN strategy. ■

Subscription: U.S. - one year $350, two years $520. International - one year $385, two years $590 • SNA Perspective is published monthly by CSI • 5400 Bayfront Plaza • Santa Clara, CA 95052 • Telephone (408) 764-5927 • Fax (408) 764-5001 • Managing Editor: Louise Herndon Wells • Associate Editors: Dr. John R. Pickens, Marianne Cohn, Suzanne D. Dowling • Contributors: Anura Gurugé, John McConnell, Roy Wells •