# SNA Perspective

Volume 13, Number 10
October, 1992
ISSN 0270-7284

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


# The Rites of Autumn: IBM's September Announcements

As surely as the leaves change color in Armonk, IBM makes new product announcements and statements of direction each autumn. The September announcements were quite wide-ranging, from 3270 "credit card" adapters to a new expansion chassis for the 3745. Furthermore, IBM has made a concerted effort to define these new products in terms of their place within its networking blueprint.

Seven major networking areas were addressed in these announcements: enhancements to the 3745 and its Network Control Program (NCP), a low-end PS/2-based bridge/router, VTAM 4.1 enhancements, several network management extensions, new frame relay support, a multiprotocol LAN hub, and many LAN adapters—including a statement of direction for Ethernet on the 3174. This article examines the aspects of these announcements that are most relevant to SNA users and discusses the benefits and challenges they present.

*(continued on page 2)*

# Users Plan for New Technologies: APPN and SNA over Routers

SNA users face the challenge of incorporating new technologies and different protocols. These include decisions regarding integrating their SNA networks with LANs, TCP/IP, OSI, NetBIOS, and frame relay, to name just a few. Two areas of significant interest today are Advanced Peer-to-Peer Networking (APPN) and SNA support over multiprotocol networks through routers. Now that IBM has announced APPN for VTAM, more users are seriously considering whether, when, and how to migrate to APPN. Also, since multiprotocol router vendors started in 1991 adding SNA support in a variety of ways , many users are wondering whether to trust their SNA traffic to routers and which approach to use.

In New York City, in conjunction with the recent New SNA Conference sponsored by InterLAB of Sea Girt, New Jersey, *SNA Perspective* held a focus group with executives from five large organizations—a major shipping company, a major pharmaceutical company, a major metropolitan area management authority, a major lodging and food services company, and a major public utility—to discuss user concerns about APPN and about transporting SNA across bridges and routers.

*(continued on page 14)*

**In This Issue:**

(continued from page 1)

Although the mid-September networking announcements were extensive, we sense a change in the air in 1992. We believe the company is spreading its fall lineup over several months—expect to see more networking announcements in October and November, especially on the 6611, LAN management, and software distribution. *SNA Perspective* is not in the habit of reporting on announcements, but we have decided to analyze those aspects of these announcements with the greatest impact on SNA users. Selected announcements are summarized in Table 1.

### Not Covered Here

This article touches only lightly on the many network management announcements; they will be examined in a future issue of *SNA Perspective* where we will analyze integration of SNA and TCP/IP management. We also do not cover the Customer Information Control System (CICS) announcements—IBM stated that CICS applications will be supported under OS/2; a future *SNA Perspective* article will address the many networking issues of transaction processing. Finally, we do not analyze the new PC token ring and Ethernet adapters nor the new credit card–size token ring, Ethernet, and 3270 adapters for portable PCs.

## Expanding the 3745

Of central interest to SNA users in these announcements were new capabilities for the 3745 Communication Controller. The main elements were a new expansion chassis, Enterprise System Connection (ESCON) support, enhanced token ring support, increased total memory of 16 megabytes (MB), and a new PS/2-based Maintenance and Operator Subsystem–Extended (MOSS-E) service processor. IBM also announced a new version of NCP, the software for the 3745, which is discussed separately below.

### 3746-900 Expansion Chassis

The key to the ESCON support and token ring enhancements lies with the new IBM 3746 model 900 expansion chassis, or expansion frame, for the 3745. The 3746-900 puts to rest rumors of a 3765 which would have replaced the 3745; instead, this new hardware frame enhances the 3745. IBM is still discussing, however, a gigabit APPN switch which has its roots in the PARIS switch technology that *SNA Perspective* expects several years from now.

The machine type 3746 designates a 3745 expansion chassis which is designed to provide additional scanners for the 3745. Five current 3746 models are available—A11, A12, L13, L14, and L15. Expansion chassis are designed only for the larger 3745 models—210, 310, 410, and 610.

The 3746-900 differs significantly from these other expansion chassis. It is based on a different hardware architecture and supports Intel 80486–based adapters, while the 3745 adapters are based on the Motorola 68000. These new adapters are faster, attach to the 3745 direct memory access (DMA) bus, and have a level of intelligence that enables them to offload data link control processing from NCP.

The 3746-900 can only connect to the larger 3745 models—210, 310, 410, and 610, although IBM stated that it intends to support these new token ring and ESCON adapters on the 3745-170. The MOSS-E service processor and NCP 6.2, both discussed below, are prerequisites for the 3746-900.

Two adapter types are currently provided for the 3746-900—ESCON and token ring—and are discussed in separate sections below. A 3745 or 3746 adapter is defined as one processor and one or more couplers. Each 3746-900 ESCON processor can support one ESCON coupler and each token ring processor can handle two type 3 token ring interface couplers (TIC3). The TIC3 is a different coupler from the TIC2 that runs on the 3745 itself. The 3746-900 adapters connect to both the 3745 DMA bus and input/output controller (IOC) bus, as shown in Figure 1 on page 4.

The maximum number of adapters depends on the 3745 model they are attached to. The difference between the number of ports supported on the different models is as complex to describe as most of the tradeoffs on the 3745. The base 3746-900 comes with four processor slots and a controller bus and service processor, which has one coupler

# Selected IBM September Announcements

• **3746-900 Expansion Chassis for the 3745**
A new frame supports ESCON and higher performance token ring adapters. Long the source of rumors of a "3765" to replace the 3745, the 3746-900 instead enhances existing 3745s. Best for sites with large networks, data transfer applications, both token ring and Ethernet host access, or memory constraints because of X.25 NPSI and/or token ring support.

• **ESCON for the 3745**
Five times faster than standard parallel channel adapters for data transfer, but about the same for interactive traffic. The 3745 with ESCON can be up to 27 miles (43 km) from a host and can connect to up to sixteen hosts with one ESCON connection. ESCON adapters run only on the 3746-900.

• **Enhanced token ring for the 3745**
New token ring adapters offload DLC processing from NCP and connect to the DMA bus. Two to three times faster than current adapters for file transfer. New adapters run only on the 3746-900. Nine token rings through the 3746-900 are supported in addition to or instead of eight token rings on the 3745 base unit.

• **Doubled maximum memory for the 3745**
16 MB total memory per CCU. 12 MB maximum NCP load module.

• **MOSS-E service processor for the 3745**
New PS/2-based, token ring–attached service processor and operator console for the 3745. Can be operated locally or remotely. One MOSS-E can support up to four 3745s. MOSS-E is a prerequisite for the 3746-900, ESCON, new token ring adapters, and increased memory.

• **Network Control Program Version 6 Release 2**
Enables new 3745 features such as the 3746-900. Supports frame relay frame handler. Supports mixed-media multilink transmission groups. Supports APPN composite network node (with VTAM 4.1).

• **VTAM 4.1 Enhancements**
Additional support for the release announced in March and scheduled to ship in the first half of 1993. Better APPN support for end nodes on LANs through connection networks. Simpler methods for limited dynamic VTAM/NCP reconfiguration, VTAM default changes, and problem determination. Command Tree/2 uses a mouse and prompts to help users construct VTAM commands. Several other miscellaneous additions and enhancements.

• **8250 Intelligent Multiprotocol Hub**
An OEM LAN hub from Chipcom for token ring, FDDI, and Ethernet. A related hub management program runs on the RS/6000.

• **LAN Adapters**
A statement of direction for Ethernet on the 3174. Ethernet adapters and better token ring adapters for PCs. Credit card–sized adapters for portables—Ethernet, token ring, and 3270. A wide area network interface for RouteXpander/2.

• **RouteXpander/2**
Low-end OS/2-based device driver software that supports frame relay, source route bridging, and multiprotocol routing. Supports multiple protocols. Appears to applications as a token ring LAN, appears to a 6611 as another 6611. Supports up to 200 frame relay logical links over a single physical link.

• **Network Management**
A new release of AIX NetView Service Point adds LU 6.2 to NetView. A new release of NetView Performance Monitor includes frame relay, Ethernet, and LAN segment support. A new release of NetView/6000 offers the X Windows/Motif GUI, manages OSI and SNMP, and supports four APIs. New Systems Monitor/6000 works with NetView/6000 for distributed management. New features of LAN Network Manager 1.1 include graphics for the 8230 hub, a new event filter, and HLM support. ■

*Table 1*

connected to the 3745 buses and one token ring coupler. If the 3746-900 is attached to a 3745-210 or 310, it can support a combination of four ESCON and/or token ring processors in the four processor slots, which means up to four ESCON channels or up to eight token ring ports (plus one token ring port on the 900 base). With the 410 or 610, however, which have a dual central control unit (CCU), one of the four processor slots is required to contain a token ring processor with one coupler in that processor dedicated to communicating with the second CCU.

### *Pricing—3746-900 like a Second 3745*

The pricing of the 3746-900 channel and token ring adapters is comparable to the price for the parallel channel and TIC2 adapters in the 3745 base unit, as shown in Table 2 on page 5. However, the total cost, including the 3746-900 chassis and NCP support for it, is much higher. Also, as is increasingly IBM's custom, the purchase of one product requires the purchase or upgrade of several other products. The 3746-900 requires VTAM 3.4 or an enhanced VTAM 3.3. Also required are NCP 6.2 and the MOSS-E service processor.

The minimum cost to enhance a 3745 model 210 to add a 3746-900 with one ESCON channel and two token ring ports is therefore $144,720, plus the cost of upgrades to NCP 6.2 and VTAM 3.4. This is in addition to the 3745 base price of $147,550 for model 210. At this price, it makes sense only as an alternative to an additional 3745.

### *New 3745 Designations*

IBM has elected to add just a little bit of confusion to all of this. When used in conjunction with the

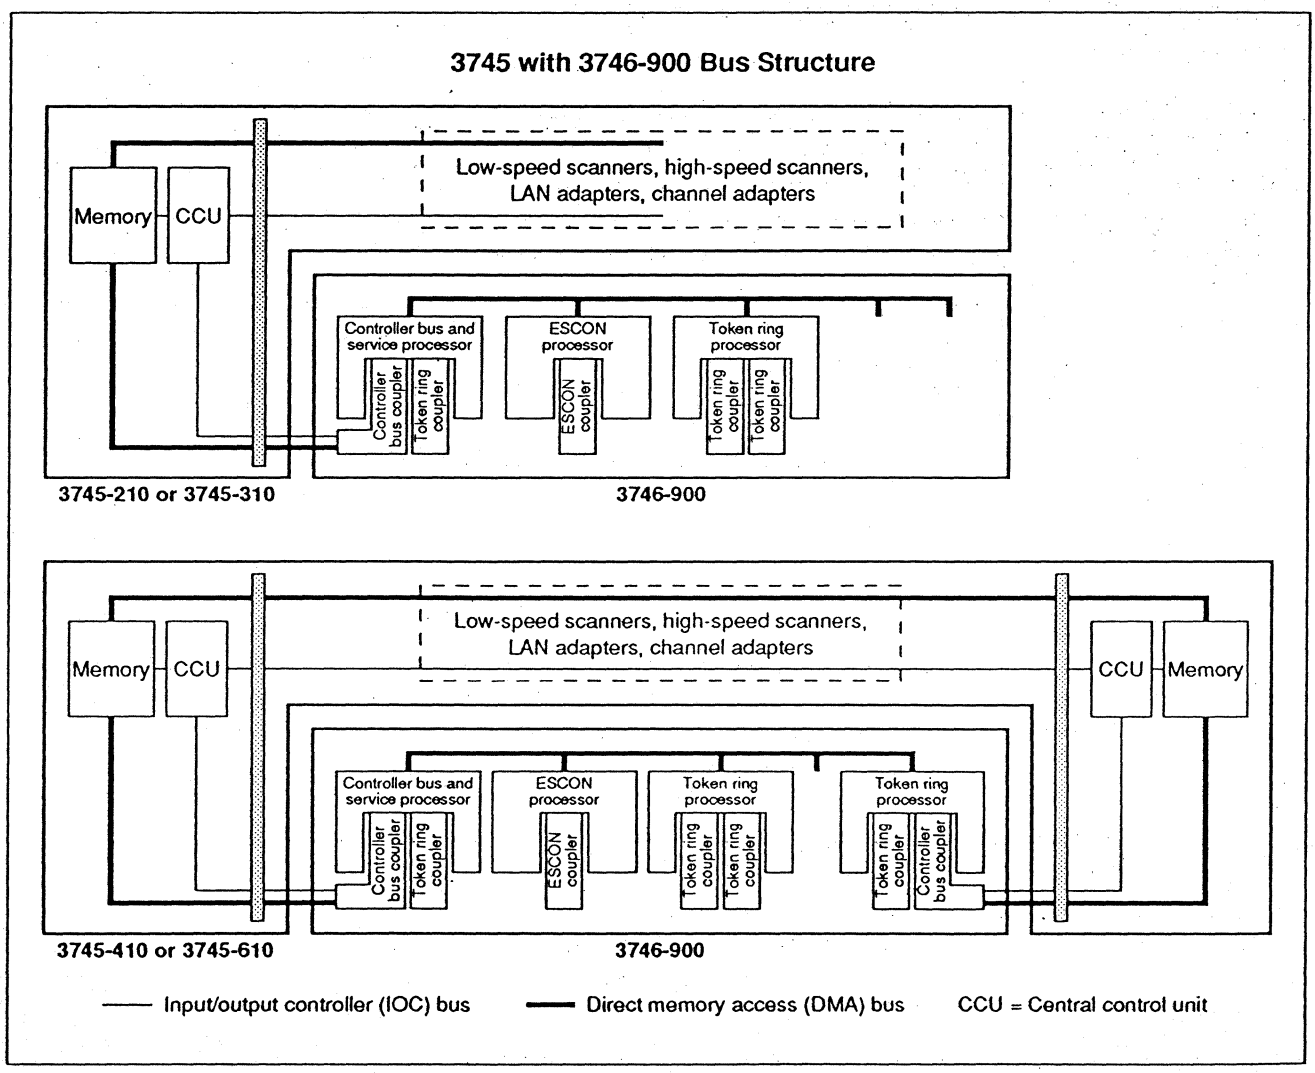


*Figure 1*

new MOSS-E service processor, the 3745 models 210, 310, 410 and 610 are now designated the 21A, 31A, 41A, and 61A. The model redesignation is because the 3745 base hardware is changed to support a token ring interface to MOSS-E rather than the current RS-232 connection to the existing service processor. Since MOSS-E is a prerequisite for the 3746-900 and for the increased memory, any 3745 model with those features will also be redesignated to the administrative (A) series model numbers.

### *ESCON at Last*

The 3746-900 enables the addition of ESCON to the 3745. ESCON is a high-speed, 17 MB per second (≈140 megabit per second), fiber-optic channel for CPU communications over distances far greater than those that can be provided via bus-and-tag cables, where distance limitations are generally less than 400 feet (≈125 meters) unless a channel extender is used. As shown in Figure 2, channel communications can be extended up to approximately 27 miles (43 kilometers) with ESCON, ESCON directors, and the extended distance feature. As indicated in the figure, the 3746-900 does not support the extended distance feature and so must be within three kilometers of a host or an ESCON director.

IBM benchmarks indicate that, in high-volume (i.e., file transfer) environments throughput on the 3745 with ESCON is about five times that of standard parallel channel adapters and about two and a half times that of a buffer chaining channel adapter. However, users with a mainly interactive or mixed communication environment will find that ESCON provides little or no improvement over parallel channel communications for the 3745.

Another ESCON benefit is that, if connected to an ESCON director, one fiber can connect the 3745 to up to sixteen hosts. Current parallel channel support would require a separate cable for each host. Furthermore, in dual CCU models, one fiber can be used to support both CCUs, which currently require two parallel channels. Finally, the addition of ESCON can be an evolutionary step: customers can add ESCON in addition to current 3745 parallel channels. This multihost support can take some of the sting out of the high price tag for the 3746-900.

The 3745 is the last of IBM's three controllers to support ESCON—both the 3174 and 3172 got ESCON at its unveiling in September 1990. *SNA Perspective* believes that the market growth for ESCON has been limited because of its absence from the 3745.

### *Token Ring: More! Better! Faster!*

Another feature for the 3745 with the 3746-900 expansion chassis is an enhanced, high-performance token ring adapter. In addition to higher token ring performance, transferring token ring from the 3745 proper also frees up other 3745 resources.

**3746-900 Pricing**

| | |
|---|---|
| 3476-900 basic expansion chassis | $37,400 |
| NCP 3746-900 feature | 47,520[1] |
| ESCON processor | 13,200[2] |
| ESCON coupler (one per processor) | 6,600 |
| Token ring processor | 13,200[3] |
| Each TIC3 (up to 2 per processor)[4] | 3,300 |
| MOSS-E service processor | 12,500 |
| 3745 upgrade to support MOSS-E | 11,000 |
| | $144,720 |

[1] One time charge or $990 monthly fee.
[2] Compare to $11,800-$19,650 per parallel channel adapter in the 3745.
[3] Compare to $21,420 for a type 2 adapter with two ports in the 3745.
[4] 900 base includes one TIC3 port.

*Table 2*

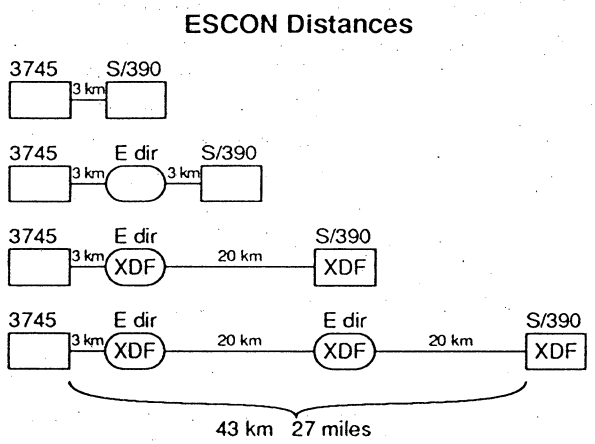


*Figure 2*

IBM's existing 3745 token ring interface coupler (TIC2) connects into the slower speed IOC bus, an interface system more commonly associated with low-speed serial connections than with the high-speed connections required by a LAN. Ethernet adapters and T1 scanners, in contrast, connect to the DMA bus which bypasses the CCU. The new token ring adapters connect to DMA in a similar way to existing Ethernet and T1.

A 3745 today is limited to eight token rings, while the new 3746-900 allows up to nine, or a combined total of seventeen token rings.

Alternatively, instead of combining token rings, users can transfer all token ring support to the 3746-900. This increases the available slots for T1 and Ethernet interfaces on the 3745 because the current token ring adapters compete for four of the eight slots that can be used by T1 and/or Ethernet. The use of even one token ring adapter allocates all four of these slots for token ring and away from T1 and Ethernet. Of course, this has not been a problem to date for Ethernet users since NCP 6.1, which is required for Ethernet, began shipping only in August 1992.

While an increased number of token ring ports is useful, there is an additional benefit to the new adapters on the 3746-900—they offload processing from NCP. Currently, NCP Token Ring Interface (NTRI) processes the logical link control and error handling for TIC2. IBM states that using TIC3 instead of TIC2 can remove as much as seventy percent of the load from the CCU in a token ring environment. This means that more CCU cycles are available for other communication functions such as X.25 processing by the NCP packet switching interface (NPSI), which is a heavy NCP user.

In addition, more traffic can be sent through a single token ring port with less impact on the 3745. According to IBM, the maximum LAN-to-host file transfer rate available for the 3745 with TIC2 and a parallel channel adapter can be increased by two to three times with the 3746-900 with both token ring and ESCON. But the CCU load is not increased in proportion to the additional traffic.

Of course, there is a hefty price tag for these advantages that far exceeds the cost of the adapters themselves, as discussed above under Pricing. Also, even the new token ring adapters still support only SNA traffic. Token ring adapters on the 3172 can handle SNA or internet protocol (IP) traffic, and *SNA Perspective* had hoped these new adapters and the new NCP would support this also.

### Storage Increase Welcome for NPSI, Token Ring, SNI, and APPN

Twice as much storage for the 3745 models 310 and 610 was announced by IBM. A total of 16 MB of memory per CCU will be available. In conjunction with this, IBM announced the capability for NCP load modules of up to 12 MB. This doubles the current maximum support of 8 MB memory and 6-MB NCP load modules. The upgrade makes it possible for these controllers to connect more users and to concentrate more traffic.

The 210 and 410 models cannot be upgraded to 16 MB of memory. *SNA Perspective* believes this is because their slower processors could not take advantage of the additional storage. Users would have to upgrade to models 310 or 610 to use the additional memory. Pricing for this 8 MB to 16 MB upgrade is $27,150 and IBM has scheduled availability for the new storage capabilities for June, 1993. This compares to current 4 MB storage increments on the 3745 cost $12,150.

*SNA Perspective* believes that the additional memory will most benefit users with NPSI or token ring TIC2, and especially those with both. NPSI and TIC2 use a significant amount of 3745 CCU cycles and NCP resources. Because of this, many users have found it essentially unfeasible for NPSI to coexist with token ring within a single 3745. It is possible on a large 3745 with the current maximum of 8 MB of memory and relatively few supported devices, but even then performance is limited. Some sites needing both NPSI and token ring have even been encouraged by IBM to buy two separate 3745s. This problem is addressed by the 3746-900's new token ring adapter, which both offloads NCP processing and uses the DMA bus which bypasses the CCU.

The additional memory will also benefit customers with several SNA networks supported with SNA Network Interconnect (SNI) and those moving to APPN and need the memory for flexible network definitions.

### MOSS-Extended

IBM announced a new PS/2-based, token ring–attached service processor/operator console called MOSS-E that enhances the 3745's MOSS and replaces the existing MOSS operator console.

The current MOSS operator console connects to the 3745 through an RS-232 interface and operates the hard disk, diskette, and control panel, loads NCP, and stores dumps of the 3745. It also provides system procedures for failure notification and tools for assisting in problem determination. MOSS-E provides these same functions and some new ones. Its primary benefits are easier maintenance and remote operation.

A single MOSS-E service processor can support up to four 3745s, only one of which can include a 3746-900; a separate MOSS-E is needed for each 3745 with the new expansion chassis. *SNA Perspective* expects that IBM will eventually remove this limitation.

Operation of the 3745 MOSS-E can be handled locally or remotely. A remote operator can establish sessions with the service processors and access the various control features and console functions. This can be handled in several ways—from a NetView Graphic Monitor Facility workstation running the Distributed Console Access Facility (DCAF) or from a PC on a LAN running DCAF—as well as the current access through switched lines.

MOSS-E enables the 3745 to report problems to the IBM service center, receive microcode changes, and inform the network operator of these—all automatically. The operator can then activate the microcode changes if desired.

### Laying the Groundwork: SODs

IBM made two statements of direction (SODs) about the future of the 3745. First, IBM intends to support more ESCON and high-performance token ring adapters on the 3746-900. Second, IBM also intends to provide both the ESCON and high-performance token ring adapters announced with the 3746-900 for the 3745-170 as well. All statements of direction discussed in this article are summarized in Table 3. *SNA Perspective* expects that IBM will also add additional processor types for the 3746-900 such as T1 and, later, Ethernet and perhaps even low-speed scanners, further offloading NCP and CCU cycles.

**IBM Statements of Direction in September Announcements**

**3745**
- Support more ESCON and high-performance token ring adapters on the 3746-900
- Provide both the ESCON and high-performance token ring adapters announced with the 3746-900 for the 3745-170

**NCP**
- Allow the 6611 bridge/router APPN network node to interoperate with the VTAM/NCP composite network node through a frame relay connection
- Allow a peripheral SNA device such as a 3174 to connect to its NCP boundary function through a frame relay connection

**VTAM**
- Support for APPN connection network virtual node definition
- Support for APPN communication over subarea routes

**3174**
- An Ethernet attachment and tn3270 support for the 3174

**Network management**
- Add an SNMP agent to LAN Network Manager in the future, and permit it to communicate with NetView/6000
- Provide NetView Performance Monitor with a SystemView-compliant, object-oriented graphical user interface that will run on an OS/2 platform under Presentation Manager.

*Table 3*

# NCP 6.2

Now that IBM has just begun shipments of Advanced Communication Function/Network Control Program Version 6 Release 1 (NCP 6.1), what could be more natural than a new release? Actually, the announcement of NCP 6.2 was expected and welcome. It supports some previously announced features—this is the "then-current NCP release" presaged in the March 1992 announcement of APPN composite network node with VTAM 4.1 (see *SNA Perspective*, April 1992). It also enables several products and features announced concurrently—3746-900, MOSS-E, and 16-MB memory. In addition, it supports capabilities within NCP, including frame relay and transmission group enhancements, which are discussed below.

## Frame Relay Enhancement

The first of these capabilities is enhanced frame relay support. Until this announcement, a 3745 could function in a frame relay network only as frame relay terminating equipment (TE), which is approximately analogous to data terminal equipment (DTE) in X.25. This meant that the 3745 could serve as the source or destination for data in such a network.

With NCP 6.2, this is changed. Now the 3745 can participate as a frame relay frame handler (FH), which is approximately analogous to data circuit-terminating equipment (DCE) in X.25 terminology. A comparison of frame relay and X.25 packet formats is shown in Figure 3. By adding this capability, IBM has enabled the 3745 to become a frame relay switch within a private frame relay network. Although we use terms analogous to X.25 as a start, the reader is cautioned against considering frame relay only in X.25 terms because this limits frame relay's more extensive capabilities. For example, one node can support both terminating equipment and frame handler functions, and it can support them both over the same line at the same time.

Frame relay has been one of the hottest topics of the early 1990s as it represents a new packet switching focus to provide higher speed technology for information systems that are increasingly time sensitive. With speeds presently up to T1 (1.544 Mbps) and soon to move into the T3 (45 Mbps) arena, frame relay represents a step that IBM must take, and a logical step in light of the past connection between IBM equipment and X.25, a technology that frame relay seeks to replace. A significant benefit is that, unlike X.25, frame relay can natively support any protocol, such as SNA, TCP/IP, and non-routable protocols such as NetBIOS or Novell's IPX encapsulated in IP.

Currently, IBM's support is limited to frame relay connections set up as permanent virtual circuits because the standard for switched virtual circuit for frame relay has not been finalized. *SNA Perspective* believes that the availability of switched virtual circuits will significantly boost frame relay's popularity.

## Mixed-Media Multilink Transmission Groups

This new NCP feature provides a new means for defining transmission groups. Previously, multiple

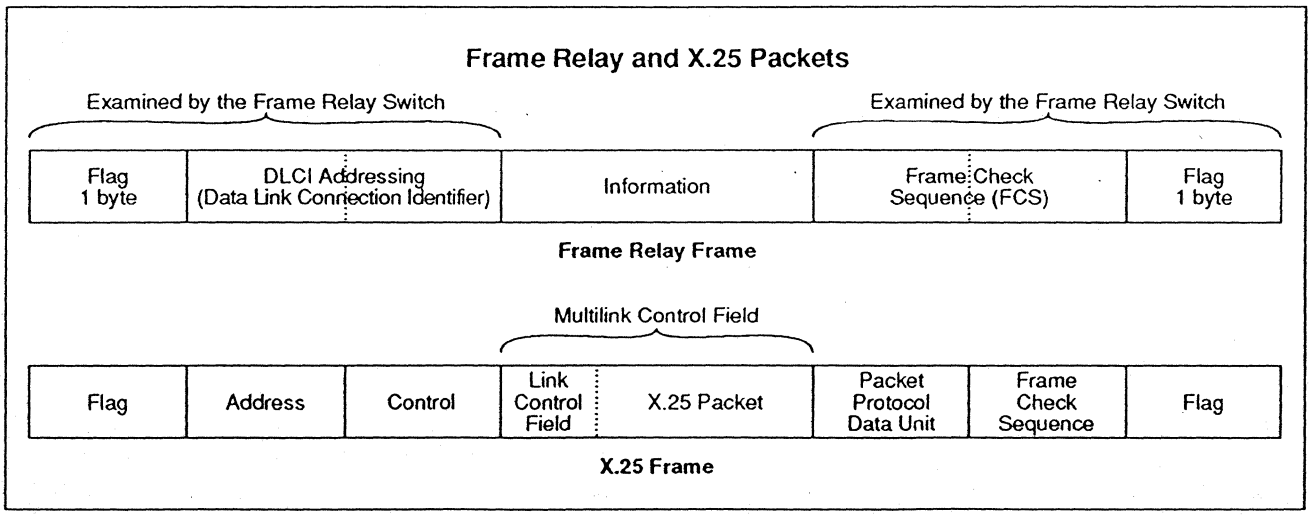


*Figure 3*

SDLC links between two 3745s could be associated as a single transmission group, providing a means of load balancing and permitting functional links to take on the load from failed links (see Figure 4). However, a token ring network had to be defined in a separate transmission group from SDLC links, and token rings were limited to one network per transmission group. Frame relay networks were similarly limited. X.25 networks could not and still cannot participate in transmission groups.

The first new element is the multilink transmission group. Now, not only can multiple SDLC lines be associated as a group, but multiple token ring or frame relay networks can be associated in a single transmission group, as shown in the second part of Figure 4.

The next new element is mixed media, as illustrated in the third part of Figure 4. With NCP 6.2, a transmission group no longer need be homogeneous. SDLC links, token ring networks, and frame relay networks can be associated in a single transmission group.

### *Statements of Direction*

IBM made two statements of direction about the future of NCP, both regarding the use of frame relay. First, NCP will allow the APPN network node on the 6611 bridge/router to interoperate with VTAM/NCP composite network node through a frame relay connection. Second, NCP will allow a peripheral SNA device, such as a 3174, to connect to its NCP boundary function through a frame relay connection.

These will permit consolidation of communication lines, expenses, and operations into a single system and allow compatible frame relay connections (such as 6611s and PS/2s with RouteXpander/2) to pass SNA traffic to the NCP peripheral boundary support.

Pricing of NCP 6.2 is highly dependent on the configuration and hardware, with one-time charges ranging from around $19,000 to just under $95,000 or a monthly license charge ranging from around $400 to nearly $2,000. Availability is scheduled for June 1993.

## *VTAM 4.1*

Although IBM announced ACF/VTAM Version 4 Release 1 (VTAM 4.1) for MVS in March of this year, several enhancements were added in September to the planned release that will be included at general availability in the first half of 1993. Selected enhancements are discussed below. Other additions involve dynamic reconfiguration, network default setting, trace enhancements, and multipath channel-to-channel support.

It should be noted that VTAM 4.1 is for MVS/ESA Version 3 Release 1.3 or later, with VTAM 4.1 for VM still under a statement of direction. We do not expect VTAM 4.1 for VM to be announced until the MVS version ships. IBM stated it has no plans for VTAM Version 4 for MVS/370 or MVS/XA.

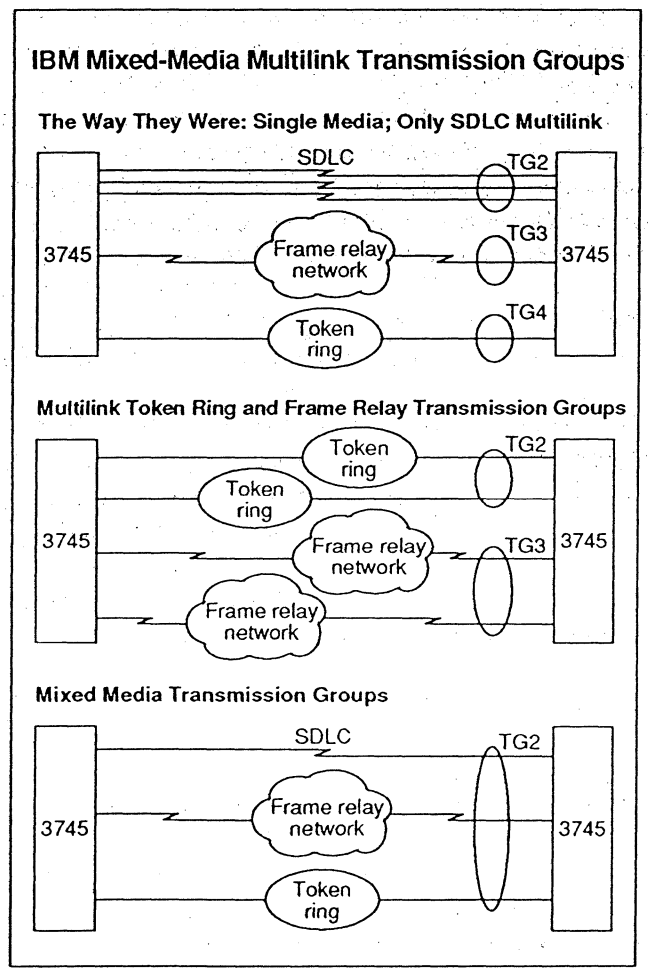


*Figure 4*

### Connection Networks: APPN for LANs as Virtual Nodes

An important capability is an addition to the APPN functionality for VTAM 4.1 announced in March. This new function relates to APPN *connection networks*. First, APPN connection network support allows an APPN end node to define its connection to a LAN as one *virtual routing node* rather than defining its connections to every other node on the LAN. Connection network support then allows a network node to calculate routes for two end nodes that have defined the same LAN as their virtual node. VTAM 4.1 will support connection networks as a network node. IBM made a statement of direction that a future release of VTAM will participate in connection networks as an end node.

### Command Tree/2

IBM has simplified entering VTAM commands through Command Tree/2, which runs on a PS/2. Command Tree/2 permits users to enter VTAM commands via prompts without having to remember command options or syntax. This function is already included in the NetView Version 2 Release 3 Graphics Monitor Facility, announced earlier this year, but it can now also be purchased separately. It should be noted that Command Tree/2 is designed to work with VTAM 4.1 and not prior releases.

### Third-Party Early Testing

To allow third-party software vendors to obtain early publications, access macro libraries, or remotely test applications against VTAM 4.1 before its general availability, IBM announced the Early Test Program, which will begin this month and run through June 1993.

### Statements of Direction

IBM made two statements of direction regarding VTAM. First, as discussed above under "Connection Networks: APPN for LANs as Virtual Nodes," a future release of VTAM will support APPN connection network virtual node definition. This means that a VTAM host which is operating as an APPN end node and is LAN-attached will not need to define its connection to every other APPN end node on that LAN but, rather, can identify the LAN as a virtual routing node when it registers with its APPN network node server.

Second, a future release of VTAM will support APPN communication over subarea routes. This will allow two APPN nodes to set up an APPN transmission group between them that runs over a subarea virtual route. This will support users who want the dynamics of APPN while keeping the session management options of subarea routes. With VTAM 4.1, APPN transmission groups can run over type 2.1 links or, alternatively, interchange nodes can send traffic from an APPN network through subarea transmission groups and routes, but APPN traffic is not supported directly over subarea routes. (An interchange node is a gateway between an APPN network and a subarea SNA network.)

## 3174 Ethernet Statement of Direction—The Other Shoe

IBM made a statement of its intention to make the 3174 capable of attachment to Ethernet networks, both as a downstream node and as an upstream gateway. Since McDATA of Broomfield, Colorado, has a 3174-compatible product already capable of such an attachment, this enhancement is overdue by IBM, but is nonetheless welcome. We began to anticipate Ethernet for the 3174 as we watched its increasing TCP/IP features, waiting for the other shoe to drop. 3174 Configuration Support C supports TCP/IP traffic from coax-attached PCs running TCP/IP and, with the addition of request-for-price-quotation code, supports Telnet sessions for attached 3720 or ASCII terminals. In a related SOD, IBM announced that it would support tn3720 to its TCP/IP support on the 3174. Where TCP/IP is growing, can Ethernet be far behind? *SNA Perspective* expects IBM to move this Ethernet development process quickly enough to ship in 1993.

## 8250 Intelligent Hub: A Chipcom Off the Old Block

IBM has announced a new entry into the multiprotocol intelligent hub market—the 8250. An intelligent hub can be thought of as a backbone LAN-in-a-box. The 8250 comes with either a 6- or 17-slot chassis,

each capable of supporting several modules for establishing a LAN.

Modules available for the 8250 include several multiport modules for Ethernet in various media types including optical fiber, token ring, or FDDI, as well as modules that provide LAN management and SNMP for the system. Other modules permit interconnection via Ethernet bridging. According to IBM, the new 8250 can be combined with existing 8230 and 8240 hubs to form a single token ring or FDDI segment in single or multiprotocol LAN environments.

It is important for users to remember that, while Ethernet, token ring, and FDDI networks can be supported on a single hub, these remain essentially separate networks. As yet, no modules provide translation within this hub from one network topology to the next. Customers who wish to interconnect them can do so only by connecting to external systems, such as 8209 Ethernet to token ring bridges, that provide this function.

The IBM 8250 is an OEM version of the ONLine Intelligent Hub from Chipcom Corporation of Southborough, Massachusetts. Both systems have the same number of slots and support the same basic options, down to the availability of redundancy in the hub controller and redundant, load-sharing power supplies as well as the ability to perform "hot swaps" of modules into the system without the need to power down or reboot. IBM had announced in July 1992 an agreement to work with Chipcom on future products. The 8250 represents the first fruit of cooperation between the two companies.

Noticeably missing from IBM's list of options for the 8250 is the module offered by Chipcom that provides a single wide area network port for routing. This module is supplied to Chipcom under an agreement with Cisco Systems of Menlo Park, California.

Why did IBM elect to work with Chipcom instead of the hub industry leader SynOptics or number two Cabletron? We suspect that SynOptics did not appear to be a desirable partner because of its relatively close ties to router manufacturer Cisco Systems. While Cabletron did not present any specific negatives, we believe the features of the Chipcom hub, specifically high availability and fault tolerance, were preferable to IBM.

### Hub Management

System management for the 8250 is provided through Hub Management Program/6000 which operates in conjunction with AIX NetView/6000, IBM's platform for SNMP management. This software, which IBM also acquired as an OEM product from Chipcom, will enable network operators to manage hubs, bridges, routers, and LANs from a single station.

## Network Management

Many announcements were in the domain of network management and are summarized in Table 4 on page 12. *SNA Perspective* will analyze these announcements in a future article as they relate to the integration of SNA and TCP/IP network management.

## RouteXpander/2

RouteXpander/2 is a low-end OS/2-based software product that supports bridging, multiprotocol routing, and frame relay. It is a challenging product to grasp because what it is depends on how it is being used—it can function in several ways. Furthermore, several of its benefits are obscure because they rely on the emerging and unfamiliar frame relay technology. The following are a few possible RouteXpander/2 configurations:

- It can act as a remote token ring source route bridge supporting NetBIOS, IPX, and other protocols that cannot be routed.
- It can support routing for IP or peripheral SNA traffic.

Users can take any PS/2 with OS/2 2.0, add a wide area adapter (such as the IBM Wide Area Connector —a synchronous serial communications adapter also

**Network Management Summary**

- A new release of AIX NetView Service Point Version 1 Release 2 supports LU 6.2 sessions with NetView.
- A new release of NetView Performance Monitor Version 1 Release 6 includes support for frame relay, Ethernet, and LAN segment performance data provided through the 3745.
- Statement of direction: NetView Performance Monitor will have an object-oriented graphical user interface (GUI) running on OS/2 under Presentation Manager.
- A new release, NetView/6000Version 2 Release 1, offers a new GUI based on X Windows and Motif, can manage OSI as well as SNMP, and supports three industry-standard application programming interfaces (APIs)—XMP (X/Open Management Protocol), the Carnegie-Mellon SNMP API, and the End User Interface API from Hewlett-Packard—as well as an IBM API.
- A new package, Systems Monitor/6000, supports distributed management in conjunction with NetView/6000 by collecting data, providing it selectively (through filters and thresholds) to NetView/6000, and responding locally and automatically to preestablished conditions.
- LAN Network Manager 1.1, long scheduled to ship in December, will have additional features including a graphical representation of IBM's 8230 token ring hub, a new event filter, and support for Heterogeneous LAN Management.
- Newly announced LAN Network Manager 1.2 offers additional graphics support and will also ship in December. Not really an enhanced release, LAN Network Manager 1.2 is an alternative to LAN Network Manager 1.1 for different graphics environments.
- Statement of direction: IBM will add an SNMP agent to a future release of LAN Network Manager, permitting it to communicate with NetView/6000 and any other SNMP monitor.

*Table 4*

announced in September—$795) and RouteXpander/2 software ($795). The result is a low-end bridge/router/frame relay interface. The RouteXpander/2 software itself is essentially two device drivers—a frame relay device driver and a token ring source route bridge device driver.

RouteXpander/2 does not contain any protocol stacks itself, but rather can be used by any communication software on the same PS/2 to connect to frame relay as if it were connecting to a token ring. It does this by inserting itself under the communication software and looking like a token ring but actually changing the token ring headers to frame relay format. Think of it as a voltage converter: if you visit another country, you may need a voltage converter to adapt between what your appliance expects and what the local power system provides.

Any communication product that can access token ring via the network driver interface specification (NDIS) interface can use the RouteXpander/2. Because of this, applications that use token ring today need not be changed to support frame relay. Several examples of communication products that can run in conjunction with RouteXpander/2 are listed below:

- To support APPN routing, the PS/2 with RouteXpander/2 must also have OS/2 Extended Services.
- To support TCP/IP routing, the PS/2 with RouteXpander/2 must also have IBM TCP/IP for OS/2.
- To support NetBIOS, the PS/2 with RouteXpander/2 must also have either OS/2 Extended Services, TCP/IP for OS/2, or LAN Server 2.0.

Bridging or routing through RouteXpander/2 can be done through its one physical connection to either a wide area point-to-point connection or a frame relay network. It can communicate through this connection or across the frame relay network with another RouteXpander/2, to a 6611, or to another vendor's router.

Though it only supports one physical connection, RouteXpander/2 can support up to 200 frame relay logical links, called data link connection identifiers (DLCIs), over the one physical link. Further, it can support multiple protocols over each of these logical links. (A frame relay DLCI is roughly equivalent to a LAN media access control (MAC) address.)

RouteXpander/2 probably provides poor routing performance compared with dedicated routers on the market. Its actual performance will depend in part on the power of the PS/2 used but the product should be examined in context. To its credit, it can use existing PS/2s which need not be dedicated to RouteXpander/2. It could be valuable in sites where inexpensive, flexible connectivity is important and minimal performance is acceptable. Sites needing better performance can choose a 6611 or another router.

We do not believe that RouteXpander/2 itself is a strategic product but is, rather, a low-cost entry point for routing and frame relay. *SNA Perspective* expects, however, that RouteXpander/2 capabilities will probably be incorporated in a future low-end member of the 6611 family and perhaps in other platforms.

The fact that the RouteXpander/2 can route peripheral SNA traffic is significant, because no other product today can do it. However, this new type of SNA routing is currently supported only between two RouteXpander/2s, so it is not of major current value to users. However, IBM made a statement of direction on the same day that NCP will be enhanced to support this SNA routing and will allow a peripheral SNA device to connect to its boundary support through a frame relay connection. Further, as discussed above, *SNA Perspective* believes that IBM will add RouteXpander/2-type capability to other platforms. *SNA Perspective* will discuss this issue further in a future article.

## Conclusions

Rather than repeat the product description and analysis, we will conclude with a broad look at the announcement. The most important themes were frame relay (NCP and RouteXpander/2), multiprotocol LANs (Ethernet adapters, new token ring adapters, 8250 hub, Ethernet SOD for 3174), and network management (distributed NetView/6000, frame relay and Ethernet monitoring on NetView Performance Monitor, LU 6.2 between AIX NetView Service Point and NetView, SNMP coming for the 8230 and LAN Network Manager).

The most disappointing part of the announcement to us was the price of the 3746-900, when all the prerequisites and software are tallied in. Nearly doubling the base price of a large 3745 seems a high price tag, even for the benefits of ESCON and enhanced token ring support.

We found the most interesting parts of this announcement were, unfortunately, hidden between the lines:

- The ability of frame relay connections between two 3745s to go through other 3745s is a significant enhancement. The NCP statements of direction continue this trend. Such additional flexibility for existing 3745s is much appreciated.
- APPN connection networks is an important addition in order for APPN to become accepted in the LAN environment, but the addition of connection networks to VTAM 4.1 was surprisingly delayed from the March announcement and buried in a series of trace enhancements.
- Support for SNA routing over frame relay was buried in IBM's effort to describe an admittedly intriguing product, the RouteXpander/2. ■

(continued from page 1)

*SNA Perspective* extensively analyzed VTAM APPN in its April issue. We have also considered several times the costs and benefits of running SNA over routers (see *SNA Perspective*, August and September 1992). For this article, we chose a focus group format so that our readers could benefit from the decision process and experience of several other large SNA users.

"The IBM network has 20,000 LUs defined...Our standard LAN is Ethernet...We'd like to integrate the two networks onto a single network."

## Network Environment

**SNA Perspective:** Let's start with a brief overview of your network and the quantity and types of systems (mainframes, data centers, PCs, Macintoshes, and terminals; whether you have only SNA or lots of SNA and lots of TCP/IP, and so on) so we know your interests and concerns.

**Utilities:** We're primarily an SNA shop. We have some TCP/IP and some Ethernet here and there but we're in the process of converting from the traditional SNA hierarchical network with dedicated lines to cluster controllers to a flat, bridged LAN-WAN environment. We use that for host access and for LAN activities. In addition, we're running two 3090-equivalent Amdahl processors in our data center and we service all host access out of the data center. We don't have any minicomputers, only mainframes or micros. We have one 9370 but it's on its way out.

We have traditional SNA links over SDLC. Our entire LAN environment is token ring, with IBM local and split bridges. We still have a significant number of traditional SDLC lines going out to cluster controllers—probably in the range of 250 lines and something like 300 control units—scattered over a two-state area.

In most places where we have the WAN, the token ring network is extended too. We moved our host connectivity onto the token ring wherever we had installed it. At the other locations, we still use the SDLC links. We realize that we had been in plug-and-play mode up to this point; now it isn't so much plug and play anymore.

**Shipping:** Our network consists of fifteen mainframes, everything from Hitachi, IBM, and Amdahl. We have six different data centers located across the United States, in addition to our headquarters.

For the most part, our SNA network is concentrated around those seven data centers where we have two different SNA networks—one that's built with IBM front-end processors and another that's built with NCR Comten front-end processors. The IBM network has approximately 20,000 LUs defined, somewhere around 1,000 lines.

We have not started migrating toward an internetwork because we need to look at that technology. In addition to the IBM network, we have approximately 300 LANs spread around the country that are currently interconnected using bridges that carry DECnet and IPX. Our standard LAN is Ethernet, although we do have probably 40-50 token ring LANs as well.

We'd like to integrate the two networks onto a single network. But I don't see us being able to get rid of our front-end processors because we also have about 200 Wangs spread around the country that are not on LANs. They will probably still have to be supported using Wang emulation on 3270. We also have probably 300 sites that use bisync and 3780 stuff that we'd like to dispose of.

**City management:** Our metropolitan area has a 50-subarea network—approximately 50 NCPs (running on 37xx communication controllers at remote sites) attached to two IBM hosts, one of which is running this whole network all the time. The other host is a backup. Off this we have numerous

Systems Network Interconnect links with numerous hosts that come into our central SNA network. We have hundreds of AS/400s on the network. We're now looking to expand this central network for our metropolitan area to become a packet switched data network (PSDN), in effect, and take on anything that is not SNA: the old LANs that are not attached at this point and all the non-SNA hosts—async hosts, DECs, and so on. We have approximately 1,800 circuits on SNA and plan to put in about 800 non-SNA circuits to make this a PSDN and provide interoperability from SNA to async hosts and to LANs (and vice versa) via gateways, etc.

The remote NCPs are dual linked so in case there's a failure the other link on the other system takes over. We have a lot of APPN low-entry network nodes. The AS/400 is a low-entry network node, at this point. APPN coming to the mainframe is going to change the way we do business, if we choose to use it.

We are more of a service bureau organization, so we don't have very much control over the remote user environment and there's a real mix out there. That's part of our problem. One of our goals is to establish certain criteria that people have to satisfy in order to be connected to the network.

We've just completed the design of the non-SNA strategy and we're going with routers across the front. As far as possible, we can't allow any end network to have any influence on the WAN—we'll have to analyze every proposed integration to ensure that our WAN is not affected. Basically, the SNA backbone is now 56 kb and we're bringing it up to T1.

**Pharmaceuticals:** Our company is divided into five business sectors, although the primary income is based on pharmaceuticals. It's an IBM SNA environment in the corporate offices are located both in the U.S. and the U.K. R&D is all DEC; the other business sectors are a mix of DEC and IBM, mostly IBM.

As a result of a merger, my operation, which was mainframe-based, was moved to an AS/400. The AS/400 is very interesting in that it's LU 6.2. The company also decided to go with LAN Server and OS/2 1.2. OS/2 seems to prevail right now—we're using an OS/2 1.3 platform called "Ducas Bear"

> "It seems that every new bridge you plug in at a different site has something running around in the background that you're unaware of."

which is based on the Philadelphia forecasting. It's a good client/server application in which data resides on the AS/400 and the PCs allow the graphical front-end processing.

The company's U.K. corporate office is primarily mainframes. They just put in an ES/9000 and they have some 3090s. There are four mainframes in the U.K. and four in the United States, with the new one coming on-line right now. The U.K. office has just managed to do something that IBM said was a first in the history of IBM—they've managed to have an AS/400 in a remote site connect to the U.S. data center, back to the U.K., then to Belgium, and establish a sign-on session (an end file transfer) from Belgium.

APPC and APPN may be a solution for more of this. For example, new hires expected to use the AS/400, which is about a $25,000 package, to get files from the mainframe. The alternative is to go to PC gateways, but gateways seem to be cumbersome and we don't really want to do that especially if there's an alternative. R&D is using new bridges and multiplexers because we have a lot of videoconferencing

**Lodging/food:** Our company is primarily in the hotel/lodging and food services industries. We currently have three 3090s and an ES/9000 at our data center to support our network and applications. We also have a midrange 93xx, several DECs, System/38s, System/36s, AS/400s, and anything IBM can sell us spread out across our network.

We have approximately 1,800 PUs on approximately 530 SNA lines that are really running through an X.25 packet switching network. The X.25 network is international, going as far as Tokyo and Australia and, in the other direction, to Saudi Arabia. We also have a satellite data center in New York with

# Brief Overview of Networks and Systems

**Utilities**
- Primarily an SNA shop.
- Two 3090-equivalent Amdahl processors in data center servicing all host access.
- No minicomputers except one 9370 on its way out.
- SDLC lines to 300 cluster controllersover two-state area.
- Entire LAN environment is token ring, with bridges. Some TCP/IP and Ethernet.
- Converting from hierarchical SNA network to flat, bridged LAN-WAN environment for host access & LAN.
- Token ring at most sites reached by WAN; sites without LAN still use SDLC.
- Moving host connectivity onto token ring.

**Shipping**
- Fifteen mainframes—Hitachi, IBM, and Amdahl.
- Six data centers across U.S.
- Two different SNA networks—one with IBM 37xx's and one with NCR Comten's.
- 20,000 LUs defined, 1,000 lines.
- 300 LANs across country interconnected using bridges that carry DECnet and IPX.
- Standard LAN is Ethernet, but have 40-50 token rings.
- Want to integrate the two networks: SNA and Ethernet with DECnet and IPX.
- Not yet migrating toward internetwork; need to look at technology.
- Cannot get rid of 37xx—need to support 3270 for 200 Wangs across U.S. not on LANs.
- 300 sites still use bisync and 3780—would like to dispose of it.

**City management**
- Metropolitan area with 50 remote NCPs attached to two IBM hosts.
- One host is running network, the other is a backup.
- Remote NCPs are dual-linked for backup.
- Numerous SNI links with hosts accessing central SNA network.
- Hundreds of AS/400s.
- Many APPN low-entry network nodes, including AS/400s.
- SNA backbone now 56 kb, moving to T1.
- Will expand central network to become packet switched data network (PSDN).
- PSDN will also support non-SNA: older LANs and non-SNA hosts—async, DEC, etc.
- 1,800 circuits on SNA now; plan to add 800 non-SNA circuits to create PSDN.
- PSDN–interoperate from SNA to async hosts and LANs and vice versa via gateways, etc.
- Just completed design of non-SNA strategy—going with routers.
- Act as service bureau; don't have control over remote user environment—a real mix.
- One goal: establish criteria in order to get network connection.

**Pharmaceuticals**
- IBM SNA environment in corporate offices—U.S. and U.K.
- U.K. corporate office has four mainframes—three 3090s and just added ES/9000.
- Four mainframes U.S.—including new one coming on-line.
- R&D is all DEC.
- Other sectors—mix of DEC and IBM, mostly IBM.
- One mainframe-based operation was moved to AS/400—LU 6.2.
- Company standardized on LAN Server and OS/2 1.2—OS/2 prevails now.
- APPC and APPN may be a solution for mainframe files access.

**Lodging/food**
- Three 3090s and an ES/9000 at data center.
- A midrange 93xx, several DECs, System/38s, System/36s, AS/400s, etc.
- 1,800 PUs on 530 SNA lines running through X.25.
- International X.25 network—Tokyo, Australia, Saudi Arabia.
- Satellite data center in New York with another ES/9000 and a 3725.
- Transaction processing facility shop—must stay up 24 hours.
- Primarily Ethernet LANs—less than five token rings (one connected to 3725).
- Ethernets go through SNA gateways and routers to access backbone LAN to 3725. ■

*Table 5*

another ES/9000 and a 3725. Because we're a transaction processing facility shop, we need to stay up pretty much 24 hours a day.

Our primary LANs are Ethernet with probably less than five token ring LANs, only one of which is connected directly to a 3725. The other Ethernets are backboned and bridged, going through SNA gateways and routers to connect to the one LAN that's connected to the 3725. The only APPC we use is really under the covers in the standard IBM applications for NetView file transfer and NetView DM. I'm here primarily to find out how we can better use APPN and APPC in our network.

**SNA Perspective:** All but one of the five companies here currently have parallel SNA and non-SNA networks, right? Do you have now, or plan to add, bridges and/or routers to support your SNA traffic?

**Shipping:** Yes, we have them and plan to add more.

**Pharmaceuticals:** Yes, we have them now and plan to add more.

**City management:** We're planning to add some. We don't have any at this point.

**Lodging/food:** We have them now and plan to add more.

**Utilities:** Same.

## Concerns about SNA over Routers

**SNA Perspective:** What are your issues and concerns with SNA over bridges and routers?

**Utilities:** From my company's standpoint, I'm concerned about bridging, since we're in a flat network environment. What I'm concerned about is the way traffic just washes back and forth from one end to the other. We don't have a particular interest in the complexities that routers bring because we like the simplicity of bridges—you just plug them in and they work. However, we're concerned because of the way traffic flows. We need a firewall and we'd just as soon not do it with routers if we could avoid it.

**Pharmaceuticals:** Our biggest concern right now is the management of the control ring and the group. It seems that every new bridge you plug in at a different site has something running around in the background that you're unaware of. Especially after the corporation goes through a merger like ours.

**City management:** We're very worried about preserving the integrity of the LAN. Therefore, we're looking at routers to actually firewall and preserve the bandwidth across the WAN.

| Do you have now, or plan to add, bridges and/or routers to support your SNA traffic? | |
|---|---|
| Shipping: | Have them and plan to add more. |
| Pharmaceuticals: | Have them and plan to add more. |
| City management: | Don't have any at this point, planning to add some. |
| Lodging/food: | Have them and plan to add more. |
| Utilities: | Have them and plan to add more. |

*Table 6*

| What are your issues and concerns with SNA over bridges and routers? | |
|---|---|
| Shipping: | Will be installing over 1,000 routers; concerned about network design and management and how to build the architecture. |
| Pharmaceuticals: | Biggest concern is management of the control ring and the group because of different protocols and applications at different sites. |
| City management: | Concerned about preserving the integrity of the LAN, so looking at routers for firewall protection. |
| Lodging/food: | On the VTAM side, so no comment. |
| Utilities: | Concerned about bridging because of the way traffic flows; no interest in the complexities of routers. |

*Table 7*

**Lodging/food:** I'm primarily on the VTAM side and I'm not enough involved with them to really know.

**Shipping:** We have a number of concerns. Network management, how we're going to build the architecture. I forgot to say earlier that we presently have about 55 IDNX T1 multiplexers. How are we going to integrate the drivers into that architecture, and how many levels down? We're looking at potentially installing over 1,000 routers. The other concern is how to design the network. It's going to be somewhat hierarchical, like the IBM network, because of the numbers. We can't see doing a flat, thousand-node network.

## SDLC Passthrough or Conversion

**SNA Perspective:** Are you going to do synchronous (SDLC) passthrough or conversion? Which makes most sense to you?

**Utilities:** You've been talking about SDLC transport. I'll tell you the way we're handling it and you tell me whether it's passthrough or conversion. Wherever possible, we're running OS/2 devices which, with Communication Manager, handle their own activity. Since we come right back to the data center and our 3745s are connected to our data center backbone, that's how we get our traffic in. In terms of control units, we bring the token ring interface couplers (TICs) on the control units.

We have all 3174 control units—no 3274s (except to support the consoles in-house). So, as we expand the LANs and the WAN out into various locations and a ring becomes available, we'll change the 3174 to token ring connectivity, drop the 9.6 line, and come across the WAN. But again, it's staying flat. So now with the 3174 being downstream of the token ring, its upstream link is to the token ring. I don't consider that to be an SDLC gateway.

**Shipping:** At this point, we would prefer conversion—having the 3174 SDLC converted to LLC2 at the remote location and then carried as LAN traffic up to the host. That seems to make more sense to us than passthrough. There are too many things I don't like about passthrough—I'm concerned about the time-outs; I'm not sure about the polling; and I'm not sure about the response time. IBM going that direction with its router says a great deal to us.

We have SDLC passthrough in a pilot now, though, and we're not actually experiencing really noticeable problems with response time. But especially after the APPN announcement in March, we really would like to just take the step to SDLC to LLC2.

**City management:** Basically, we're keeping SNA separate from non-SNA so we'll probably not use either SDLC passthrough or SDLC conversion. Maybe later on, for performance reasons, we might but at this point we're keeping it separate. When these products become more stable and more available, we'll probably look at it. Whatever works really well on the circuit-support end of it, whatever will provide the lower cost circuit and other connections, is probably the way we would go.

**Lodging/food:** Conversion just seems to add a couple of extra steps and complexity. Our network experience is that any time you have to deal with spoofing you never really know what's going on out in the network from a management standpoint. We also have a lot of 3274s still in our network. I think we'll probably work better in a passthrough environment than conversion.

**Pharmaceuticals:** We're currently using passthrough to allow AS/400s to access each other

**Are you going to be using synchronous (SDLC) passthrough or conversion?**

| | |
|---|---|
| Shipping: | Conversion. |
| Pharmaceuticals: | Doing passthrough now, would like to migrate to conversion. |
| City management: | Neither—keeping SNA separate from non-SNA. |
| Lodging/food: | Passthrough. |
| Utilities: | Neither—replacing SDLC with LANs. |

*Table 8*

on both sides of the ocean. In the context of peer-to-peer, I would like to see them go through conversion when we're getting down to the TIC part; I would like to see us migrate toward conversion.

## APPC and APPN

**SNA Perspective:** Let's shift the focus now to APPC and APPN. Do any of you have any implementation plans for LU 6.2, any applications that use LU 6.2 today, or pilot projects underway?

**City management:** We have PU 2.1 connections to an SNA network for AS/400s and their low-entry network nodes. We have tremendous concern about what will happen to our backbone once type 2.1 and LU 6.2 go from AS/400 to AS/400 in a big way. If we do introduce APPN, we're going to see much faster use of this AS/400-to-AS/400 communication. We may have to decide to somehow separate this type of network from the rest of the network. The SNA bandwidth is going to have to jump drastically. Also, we're basically going to be running two networks within VTAM. They're going to be bridged or meshed together somehow but, in effect, there has to be double the amount of code in order to support some of this.

**Do you have any implementation plans for LU 6.2, any applications right now that use LU 6.2, or pilot projects underway?**

| | |
|---|---|
| Shipping: | Have several LU 6.2 sessions—for e-mail, Wang computers, a CICS application, Tandem computers. |
| Pharmaceuticals: | Evaluating image processing application; leaning toward APPC. |
| City management: | Have PU 2.1 and LU 6.2 for AS/400. Concerned about bandwidth when proliferation of type 2.1 and LU 6.2 comes when we APPN. |
| Lodging/food: | Using NetView DM and a NetView file transfer program, which use an LU 6.2 platform. Have an LU 6.2 interface for e-mail. |
| Utilities: | Have an LU 6.2 prototype in place now, running an APPC application, testing another LU 6.2 application. |

*Table 9*

**Utilities:** We have a prototype in place right now that's using LU 6.2—IBM's ImagePlus application. We're running that APPC application over the WAN from one of our power plants back to our headquarters over a 256-kb link. And there's another LU 6.2 application we're still testing with the vendor that does backups to the mainframe for the main controllers and file servers.

**Lodging/food:** We use IBM's NetView DM, which functions somewhat on a 6.2 platform, as well as IBM's NetView file transfer program. We also have a package on the PC called Metaphor that interfaces, I believe, with DB2 or one of the IBM database packages for transferring inquiries and storing data. We also have an LU 6.2 interface with Soft*Switch and cc:mail for generating e-mail down through peripheral nodes.

**Pharmaceuticals:** We're leaning toward APPC right now. Image processing is one thing we're evaluating. Also, APPC and Networking Services/2 on OS/2 1.3 is now incorporated in OS/2 2.0 Extended Services. Since we're an IBM shop, I'm sure that we'll pursue it by the fall.

**Shipping:** We also have Soft*Switch as our e-mail integrator and use LU 6.2 sessions there. We also have 200 Wang computers running Wang Office. We then have a Customer Information Control System (CICS) application that allows the different Wang computers to talk to each other using LU 6.2 sessions. A Wang package that runs in CICS allows remote logons between the Wang minicomputers and electronic mail on the Wangs. We have several other 6.2 sessions. I think we have one with Tandem computers, and our IBM mainframe is on an international express delivery application.

**SNA Perspective:** Will IBM's 1992 APPN announcements or statements of direction affect your APPN plans?

**Pharmaceuticals:** It looks to me as though a lot of the products announced are mainframe-based and that IBM will move that way to create true peer-to-peer networking.

**SNA Perspective:** You've done a lot of research on APPN in your metropolitan area.

**City management:** We'll most probably wait and see what happens before we make any decisions. I believe our city is going to take a conservative approach, especially since we don't have that much bandwidth at this point. We know that a six-month lag (after announced shipment date) in IBM installation of products is a safe bet in general. We don't want to be the guinea pigs.

**SNA Perspective:** But even though it may be a long way off, do you see yourself definitely going in the APPN direction?

**City management:** Eventually we will, by default, have to and it will definitely be a plus. The dynamics of APPN should save us money. The people doing 50 NCP gens won't have to do that, so it will probably be construed as a budget savings in terms of personnel. But right now we're concentrating on bringing in the non-SNA, and that's a priority at this point.

**Will IBM's announcements or statements of direction affect your APPN plans?**

| | |
|---|---|
| Shipping: | APPN fits in well with network. Will move toward APPN eventually. |
| Pharmaceuticals: | IBM is using mainframe-based products to create peer-to-peer networking. Have to look at TCP/IP path now for our LANs. Will look for router vendor with APPN solution. |
| City management: | Will wait and see, but will eventually move in the APPN direction which will save money. |
| Lodging/food: | Some interest in APPN eventually but, because of delays, already planning to convert to TCP/IP. |
| Utilities: | Probably won't utilize APPN. Doesn't seem necessary or applicable for flat, bridged network through 3745s. |

*Table 10*

"Eventually we will have to (go in the APPN direction) and it will definitely be a plus. The dynamics of APPN should save us money."

**Utilities:** Because of the nature of our network, the part that I'm having the biggest problem with is trying to understand how to apply what we have. Having a flat, bridged network going right through 3745 wasn't part of APPN as I understood it. And since we don't use gateways or anything like that, I'm struggling to see where we're going to utilize it, especially from the host access standpoint.

**SNA Perspective:** So you don't see APPN applying to what you have at this point or in the future?

**Utilities:** Not right now.

**Shipping:** I can't speak for our software division, but I think APPN fits in fairly well with our network and that we would probably eventually move toward it. At least in terms of what we want to do with going from SDLC to LLC2, APPC would be the next logical step. But I don't know that we're going to have a choice; anyone who keeps a mainframe will eventually have to migrate in that direction. I don't think IBM is going to discontinue traditional SNA for many, many years to come. But if you want the new stuff, the bigger and better features IBM is bragging about, you're going to have to move forward. Eventually, though, IBM will stop supporting it. If you've got a bug they'll fix it but it's just like the 3725 to 3705—at some point, they'll draw the line.

**SNA Perspective:** No more enhancements, that type of thing.

**Shipping:** Unless we see the 3765, which is supposed to be announced soon, which may change the whole picture. But the sooner IBM supports the 6611, the better I'll feel. You know what scares me? IBM comes out with the software for the future and the hardware platform, but it seems that one hand isn't talking to the other in terms of software support. I spend $10,000, but some of you spend $100,000 on it.

**Lodging/food:** We would have some interest in APPN, but if IBM delays any longer in implementing what it's talked about, most of our network is likely to be converted to TCP/IP.

## TCP/IP or APPN

**SNA Perspective:** You said the "T" word, TCP/IP! That was my next question. If IBM takes much longer in implementing APPN, will you choose something else?

**Lodging/food:** TCP/IP development is under way for all our systems; it's running our properties; and it's based on an IBM platform. We need to talk to our mainframes and access our reservation system using a 3270 interface; we also need to talk to each other from property to property, from property to the reservation center, and from regional and district offices back to any group of properties. That sounds like a great APPN/APPC network, other than the fact that it's on a non-IBM platform and the software to do it is not readily available on the IBM machine.

**Will you migrate to APPN through TCP/IP?**

| | |
|---|---|
| Shipping: | Won't wait for APPN. May use TCP/IP and migrate to APPN if it has advantages. |
| Pharmaceuticals: | Need a multiprotocol environment right now, looking at TCP/IP. |
| City management: | Positioned to switch if need be. |
| Lodging/food: | TCP/IP very likely. APPN more doubtful. |
| Utilities: | Nothing driving us to APPN, nor to TCP/IP. |

*Table 11*

> "We have the software (to migrate to TCP/IP) already developed. We're that close, so IBM is already too late—and knows it."

**SNA Perspective:** The first version of VTAM with APPN is supposed to be available by the middle of 1993 and the next one probably not being available until at least a year after that, so it's the middle of 1994 before you can actually install it and use the dependent LU support IBM is talking about. Is that within a time frame that would work for you?

**Lodging/food:** No, we needed it last year.

**SNA Perspective:** So switching TCP/IP is a very real possibility?

**Lodging/food:** Very real.

**SNA Perspective:** So IBM's announcement letting you know the time frames actually makes it more likely that you might consider alternatives?

**Lodging/food:** Yes, because IBM is saying "We have nothing now and don't expect to any time in the near future."

**SNA Perspective:** You don't see a problem in transferring over because support is available today?

**Lodging/food:** It's always a problem.

**SNA Perspective:** Yes, but migrating to APPN would be a problem too. Do you see the transition to APPN, if it were available today, being less painful than transferring over to TCP/IP?

**Lodging/food:** I don't know. I believe we have most of the software already developed. It's a

*(continued on page 24)*

## Architect's Corner

# Independent Study Project

*by Dr. John R. Pickens*

> **Wanted:** Graduate or undergraduate student to undertake analysis, design, and implementation of 3270 datastream in APPC environments. Learn APPC, 3270 session utilization paradigm, analyze tradeoffs versus dependent LU requester (DLR) and tn3270, determine gateway versus mainframe (VTAM) methods of host attachment. Contribute protocol design to the IETF (travel to exotic places). Follow-on business opportunity potential.

An anecdote from Internet folklore bears telling. In the early days of the Internet, the U.S. Defense Advanced Research Projects Agency (DARPA) funded an information service for an Internet (then called ARPAnet). Sometime in the early 1980s, in an effort to streamline and save costs, this function was dropped. Recognizing an opportunity, an enterprising individual, Dan Lynch, founded Advanced Computing Environments (now Interop Inc.) and expanded the Internet information service into a profitable business venture with a tenfold, no, hundredfold increase in information flow.

The point I want to make? Often the original owner of a technology or service does not recognize its potential.

Now, a similar situation exists in the IBM internetworking milieu. In the interests of serving the greatest need—the installed base—IBM has (apparently) abandoned its product plans for 3270 over LU 6.2 in favor of dependent LU requester. Who is going to capitalize on the opportunity?

## Why 3270 Over LU 6.2 is Better

Current SNA 3270 is manual configuration–intensive, depends on a subarea node (e.g., a 3745) for flexible network connectivity, and has the wrong session establishment paradigm (the mainframe initiates the session). In modern internetworking for all protocols—peer SNA, TCP/IP, and OSI—the requirements are very different. No preconfiguration requirements (especially for terminal datastream access), flexible choice of internetworking device (e.g., routers), and session establishment controlled by the client's node (where the user lives).

A comparison of current SNA 3270 with tn3270, a protocol for 3270 over TCP/IP, demonstrates the differences.

In SNA 3270, the user configures a datalink to the subarea boundary function (e.g., in a 3745). All session establishment is under control of a single SSCP. A second session is required between the user's node and SSCP in order to control session establishment (e.g., log on to a different host). The user's 3270 node must have a fixed configuration, which has to be preconfigured in the host. (An aside: some relaxation of preconfiguration requirements has recently occurred—dynamic LU definition—but this moves the burden of configuration from the mainframe to the user's device.)

In tn3270, the user establishes a TCP/IP connection directly to the the host of interest. No control session. No preconfiguration. The client node controls everything. Internet routers perform all sessions and routing. No hierarchical SSCP session to the workstation is required. No separate control session is required.

As you can see, tn3270 is more consistent with the principles of modern internetworking technology.

The benefits of supporting 3270 over LU 6.2 in a manner similar to tn3270—let's call it appc3270—are identical to the benefits of tn3270.

## Historical Diversion

This is not the first time I have written on appc3270. This is also not the first time I have been proven wrong in my predictions. Several years ago I generated the following hypothesis—since 3270 is an SAA service and since LU 6.2 is the only SAA transport type (for SNA), it followed that 3270 over LU 6.2 was a foregone conclusion. I expected products to follow soon.

This hypothesis was destroyed by the convergence of two architectural events.

The first event was the new ability for APPN (not APPC) to carry LU 2 session traffic. To accomplish this, there was the minor matter of extending the LU 2 BIND to carry route selection control vectors. Further, there was the issue of allowing independent LUs to initiate sessions to dependent LUs.

The second event was the increased focus by IBM architecture on the design requirements for the installed base, which led to the DLR model. DLR basically encapsulates the SSCP-LU session in LU 6.2, but lets APPN carry the LU-LU session as a native session type.

So my conclusion was not foregone. IBM believes that APPN (not APPC) and DLR satisfy the requirement for migrating the installed base to APPN (not APPC).

I agree. But it only goes part way. 3270 over LU 6.2 is still required by many customers because it puts SNA clients in the same league as tn3270.

## Design Requirements

OK, so I've convinced you, an independent study student, to take on this project. What needs to be done?

First, get hold of two books, the 3174 Functional Description (IBM document GA23-0218) and RFC1041 on the tn3270 protocol. This will suggest one solution for the TCP/IP community that has a significant installed base today.

Second, get hold of the printer extensions to tn3270, developed originally by OpenConnect Systems and submitted as a technical contribution to the IETF.

For fun, participate in the IETF working group which is attempting to extend tn3270 into all aspects of 3270 operation and control. For example, this group is contemplating extracting certain BIND parameters and carrying them as Telnet negotiation variables. This will show how deep a hole can be dug if the requirements space is not properly constrained.

Finally, design a protocol that operates over a single conversation (well, perhaps multiple conversations could be used—an infrequent one for control traffic and a long-standing one for 3270 session traffic). Make sure #INTER, the interactive class of service, is used (response time guarantees). Encode the negotiation stuff in GDS variables. Carry the 3270 datastream intact—no translation or change. Implement in client system (3270 emulator) and mainframe (VTAM service).

## Conclusion

At the recent APPC/APPN developers' conference sponsored by IBM's APPC market enablement group, I asked two questions. First: Is IBM planning to implement 3270 over LU 6.2? Answer: No, not in plan. Second: Might IBM be interested in a university or other partner collaborating in its development? Answer: Yes.

So, here is a classic case of an opportunity awaiting an enterprising individual. Any sponsors? ■

*(continued from page 21)*

matter of completing our pilots and getting ready to roll it out. We're that close, so IBM is already too late—and knows it. We're that close, so IBM is already too late (with APPN)—and knows it.

**City management:** We've just completed design of the implementation of non-SNA into the network creating the "any to any" and we're doing it with routers and time-division multiplexers and manager multiplexers. We're definitely not waiting for APPN. We will definitely enhance our SNA part with APPN, eventually.

**SNA Perspective:** But other possibilities are available, so you'd switch over to whatever is available now?

**City management:** Right. Also, IBM's router is not full function at this point, which is another problem. I'm sure that what we'll do in the future, though, is constantly hone the system. If it looks as though something we've converted via the non-SNA support that we provide later on falls into a better scheme with APPN, we'll probably phase some of it over. Whatever works.

**SNA Perspective:** Will you migrate to APPN through TCP/IP?

**City management:** We're positioned to switch it, if need be. We're going to put in the TCP or X.25 backbone, but we are flexible in that realm. Since APPN will be supported on routers, anyway, we can always redirect our goals.

**Shipping:** We won't sit back and wait for APPN to become available. We'll go ahead, put the multiprotocol routers out in our network; we may use TCP/IP. When APPN becomes available, and if it seems to offer us some advantages, I'm sure at that point we'll take a look at what it would take to migrate to it. Then, if it's to our advantage, we'd do so. What we'd like to do is make sure that we get a multiprotocol vendor that will support APPN once it's published, if IBM publishes it.

**SNA Perspective:** So you're saying that, rather than waiting for APPN and not doing anything, you're going to add multiprotocol support now and get support from a multiprotocol vendor that has an APPN plan so you can move back to that.

**Shipping:** Right. That's probably our direction. We can't afford to wait two years for the product to come to life when there are products out there today that can fit the bill.

**Pharmaceuticals:** With our Finance and Marketing relying so heavily on our R&D, we're forced into doing something on a multiprotocol environment right now. My question is going to be what will Wellfleet and others be able to offer me in terms of an APPN solution? I have to look at the TCP/IP path right now. We use some Wang systems for imaging and I don't know how it's going to get shipped around. But right now we're looking at multiprotocol routers and the only solutions seem to be TCP/IP. I would like to think that a company like Wellfleet is far enough ahead in the industry to have some insight into what IBM's direction is. But we can only find out when I ask the questions.

**Utilities:** I guess we have the biggest luxury of all in that we're structured enough to dictate to the various locations how things will go in, rather than being driven from the opposite end through acquisitions or home-grown LANs, or whatever. So right now we have a small amount of TCP/IP traffic going across our token ring LAN to support a few scientific applications, and that's it. We don't see that expanding for a while, so there's nothing driving us to APPN at all.

**SNA Perspective:** Any other concerns about APPN, or APPC, or peer-to-peer SNA in general?

**City management:** Just an opinion, but I'm afraid that if IBM continues to delay people may just want to go off on their own. If you have a business application, you can't wait for someone who says they'll be there in a year and then have a six-month delay after that. ■