SNA Perspective

Volume 12, Number 7
July 1991
ISSN 0270-7284

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

Internetworking and SNA: How Much Type 4 Do You Need?

As the information systems market has grown over the past two decades, the number and diversity of systems have exploded, as have the networks supporting them. Of concern to many are the issues involved in bridging and routing SNA from, to, and through other networks, and *vice versa.*

This article focuses on internetworking issues related to SNA, particularly sending SNA across non-SNA networks. It identifies user requirements and discusses different strategies for addressing these requirements.

(continued on page 2)

IBM Network Management: New Platforms and Directions

Watching IBM is like observing whales—a fluke emerges here, a breach there. These are the only clues that help us predict where Big Blue is going. It's easier when IBM makes an explicit announcement, but even then what *isn't* announced is often as revealing as what is. In the case of major changes in direction, however, we usually don't have to wait for an announcement to see that a change is in the works—we can see Big Blue's bow wave.

SNA Perspective believes that we are witnessing such a bow wave now in IBM's approach to network management. It hints of Big Blue moving sharply toward a network management strategy based on open systems and open protocols. *SNA Perspective* has believed for some time, based on analyzing a myriad of clues, that IBM has decided to gradually abandon its proprietary SNA management architecture and join the movement toward open network management. This article explores what such a move means to IBM and its customers and how we expect it will be carried out.

(continued on page 11)

In This Issue:

(Continued from page 1)

The Need for Internetworking

Increasingly, throughout the 1980s, major communications suppliers in non-SNA environments began supporting communication into SNA. At the same time, companies primarily serving the SNA market, including IBM itself, had been adding access to systems and networks outside the SNA network.

These products were often designed for communication between two specific architectures such as SNA and X.25 or SNA and DECnet. They were provided primarily for two applications (as shown in Figure 1):

- To allow devices on one type of network to access resources on the another type of network
- To allow one protocol to "tunnel" through a network supporting another protocol

The first option requires some degree of translation from one protocol to another. The second may use translation or the first protocol may be encapsulated and delivered unchanged at the other end.

Emergence of Multiprotocol Routers

The late 1980s saw the emergence of products, particularly multiprotocol routers, that were specifically designed to assist with internetworking the variety of protocols in the primarily non-SNA LAN world. These products were developed to support routing of several protocols including OSI, TCP/IP, IPX, AppleTalk, DECnet, X.25, and XNS.

There are some analogies between these new internetworking products and X.25 networks. Both take information from several different protocols and route it across a network in a single protocol. However, while X.25 packet assemblers/disassemblers (PADs) take information at a bridging level (i.e., layer 2 or the data link layer) and route it at layer 3 within the network, multiprotocol routers actually interpret incoming protocols up to layer 3 and translate them into their own routing protocol. These routing protocols have been proprietary and will gradually be replaced by emerging standards, including the routing information protocol (RIP), open shortest path first (OSPF), and OSI's intermediate system-to-intermediate system (IS-IS) protocol.

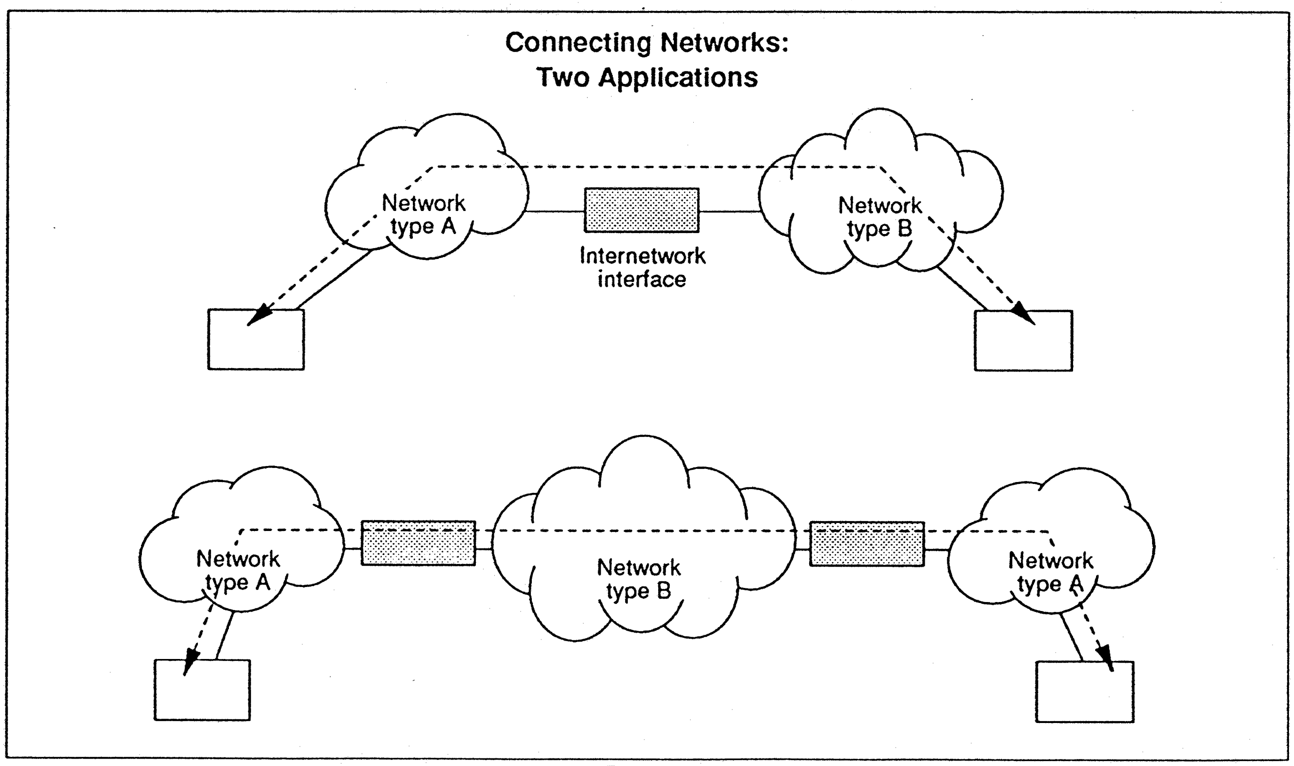

Figure 1

Benefits of Internetworking

Internetworking can provide several benefits for SNA users since few companies have homogeneous SNA networks today. Existing communications facilities, SNA or non-SNA, can be leveraged to increase their usefulness without increasing costs (see Figure 2). (The third option shown is only supported today with IBM's APPN.) Some redundant facilities can be eliminated or anticipated new investments can be avoided. For example, one company uses internetworking technology to take greater advantage of an existing transatlantic SNA link to send TCP/IP traffic between its United States and European offices.

Issues and Concerns

Lowest common denominator—Different networks have different strengths and weaknesses and combining them can lead to problems of the lowest common denominator—the resulting network may support only the functions that can map to both networks.

Lack of standards—Standards rarely exist for implementing internetworking, even for connecting two standard protocols, so proprietary solutions are often the only available choice. This leads to problems of interoperability between solutions developed by different vendors. With increasing frequency, usually at the strong urging of users, vendors have worked together to develop a standard to solve a particular problem. For example, LAN vendors gathered under the auspices of the IEEE to develop a standard approach to dealing with the incompatibility between IBM's source route bridging used between Token-Ring LANs and the transparent bridging used between most Ethernet LANs. The outcome was source route transparent (SRT) bridging.

Network management—This emerges as a more significant concern as networks become larger and more mission-critical for users. Depending on implementation, increased integration of networks could also improve the integration of network management. However, in most cases, each portion of the internetwork is managed separately. For example, the use of synchronous pass-through on

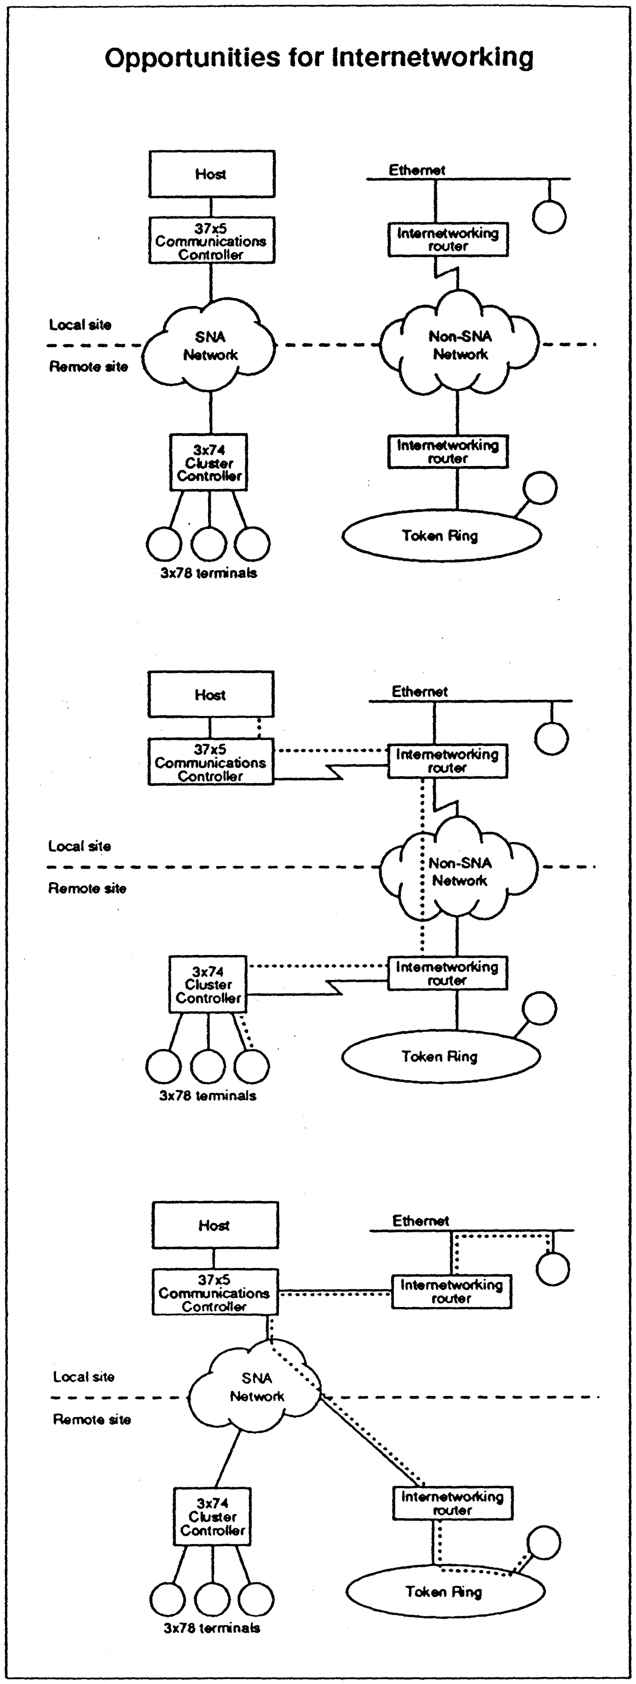

Figure 2

multiprotocol routers for SDLC traffic between a 3174 cluster controller and a 3745 communications controller would create a NetView network management "hole" for the portion of the non-SNA network between the two routers. This issue can be addressed if vendors provide mechanisms for supporting network management communications in their internetworking products. However, most multiprotocol routers are SNMP-based and, though IBM has announced its intention to have NetView more fully manage SNMP devices, today's solution is far from complete.

Experienced staff shortage—Another significant concern for users is staff expertise. Staff with experience in managing both SNA and non-SNA networks are rare, so implementation and maintenance of internetworking solutions can be difficult.

SNA-specific issues—In addition to the generic difficulties of internetworking, there are some specific issues involved in internetworking with SNA. First, traditional subarea SNA is hierarchical, so implementations need to take into account reporting themselves and the devices they support to SNA control programs (which is discussed below). Second, IBM has not published Type 4 node specifications (i.e., FAP) for ten years, so implementations will be proprietary and incompatible between the various vendors.

User Requirements

There are five primary requirements for supporting IBM/SNA communications on non-SNA networks:

- Support all SNA-capable link types:
  - SDLC
  - Token Ring
  - T-1
  - X.25
  - Ethernet
  - Frame Relay
- Enable "link-layer" routing of end systems (especially synch pass-through)
- Interoperate with "standard" IBM intermediate systems
  - Subarea (Node Type 4)
  - APPN (Node Type 2.1)
- Support multiple LU types (LU 0, 1, 2, 3, 6.2)

Strategies

There are four strategies for addressing these user requirements for support of IBM/SNA communications on non-SNA networks:

- LAN bridging
- Synchronous pass-through
- Partial Type 4 support
- Full SNA routing—Type 4 or APPN

SNA Perspective believes the first two have equal priority for users. The fourth strategy is quite advanced—few companies in the SNA arena, Amdahl, NCR Comten, and Unisys, fully emulate IBM's Type 4 node and these emulations lag behind IBM's ACF/NCP releases.

The third option implements some subset of the functions of the Type 4 node, or "spoofs" SNA into thinking these functions are being handled. When tackling either this option or the fourth one, the developer's primary question is, "How little Type 4 can we get away with?"

LAN Bridging

Bridging provides interconnection between two networks, usually LANs, at layer 2, the data link layer, of the International Standards Organization (ISO) Open Systems Interconnection (OSI) model. Since this happens at the lower sublayer of the data link layer, called media access control (MAC), it is also often called MAC layer bridging.

Bridging support is increasingly required for LAN-attached SNA end systems on Token-Ring and Ethernet because increasing numbers of SNA end systems—be they workstations, midrange systems, or mainframes—will be on LANs.

There are two primary protocols for LAN bridging: transparent bridging (also known as spanning tree bridging) and source route bridging. Transparent bridging is the primary technique used for Ethernet networks and source route bridging is IBM's standard for Token-Ring bridging. They are mutually incompatible primarily because they use different methods to learn about other bridges on the network. As discussed above, an IEEE committee worked on this issue. The committee agreed on a resolution in 1990, which has been announced but not shipped by any vendors yet. This resolution, called SRT bridging, essentially states that:

- Any bridge connecting only Token-Ring networks will use source route bridging
- Any bridge connecting at least one Ethernet will use transparent bridging.

The resolution does not, however, specify how to translate transparent bridging headers into source routing headers.

Synchronous Pass-Through

Supporting SDLC on a multiprotocol bridge or router is more difficult than supporting asynchronous protocols. This capability is called synchronous pass-through, SDLC encapsulation, or SDLC tunnelling (see Figure 3). SDLC is encapsulated, for example, within IP packets and given an IP address for the multiprotocol router at the other end of the link. These IP end points must know *a priori* that they are carrying SDLC frames. These products have to perform what is called SDLC "spoofing," that is, tricking each side of the SDLC link into believing that it is directly connected.

Several vendors, including 3Com, Wellfleet, Proteon, and Cisco Systems, provide or have announced a synchronous pass-through capability for their internetworking products (see *SNA Perspective* April 1991 "Internetworking Vendors Target SNA"). It is widely believed that IBM and Wellfleet are jointly developing a multiprotocol router, although this has not yet been announced. *SNA Perspective* believes it will be based on the IBM RS/6000 and expects it will be introduced soon, but certainly during 1991. A multiprotocol router from IBM will doubtless include this capability.

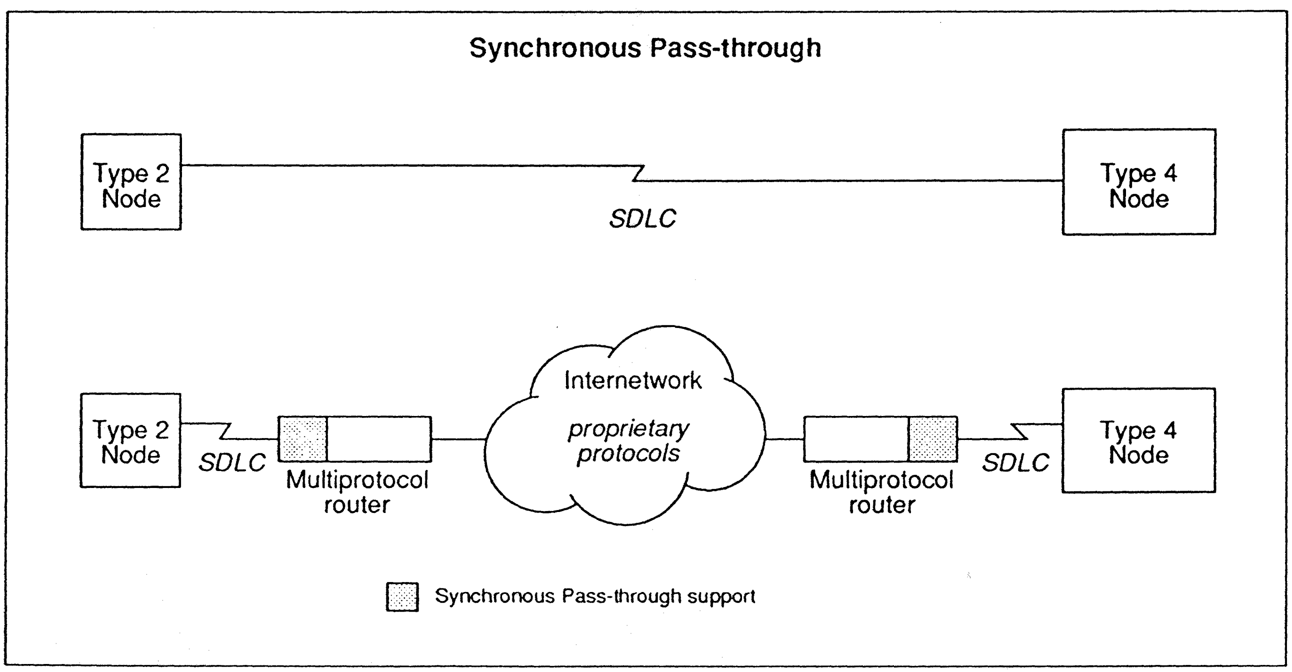

Figure 3

The synchronous pass-through capability is useful for a link extension for a point-to-point connection between two SNA devices such as a Type 2 node and a Type 4 node, as shown in Figure 3, to run SNA applications across an existing non-SNA network.

This is often referred to as an extension cord approach. There is no flexibility for the Type 2 node to dynamically change its destination Type 4 node. The internetworking nodes (multiprotocol routers in the example shown in Figure 4) are configured so that communication from one Type 2 node always is sent to the same Type 4 node.

This "one link in to one link out" constraint may not be a problem for many configurations, for instance where a remote cluster controller is configured in the SNA subarea network to always link to a particular communications controller. However, intelligent workstations—such as PS/2s running the OS/2 EE Communications Manager—have the capability to change their partners, so this would be a concern.

Another possibility presented by SDLC support is connecting SNA nodes with mixed link technologies such as Token-Ring or X.25. For example, as shown in Figure 4, the multiprotocol router supports an SNA node with an SDLC interface at one end to an SNA node connected to a Token Ring at the other end. *SNA Perspective* believes that SDLC to SDLC is the most important need today, followed by SDLC to Token-Ring. Following in a distant third place is SDLC to X.25 support.

We see X.25 as such a low priority because it is already a backbone technology, as are SMDS and frame relay. Multiprotocol routers can *either* route over X.25 networks *or* displace them. In the long run, as multiprotocol routers become more capable and pervasive, they may begin to displace X.25. However, users are reluctant to displace what they have unless there is a significant benefit and small enough risk. Most business sites already have X.25 access now. Furthermore, IBM already supplies the capability to communicate through X.25 networks, including NCP Packet Switch Interface (NPSI) and qualified logical link control (QLLC). This support

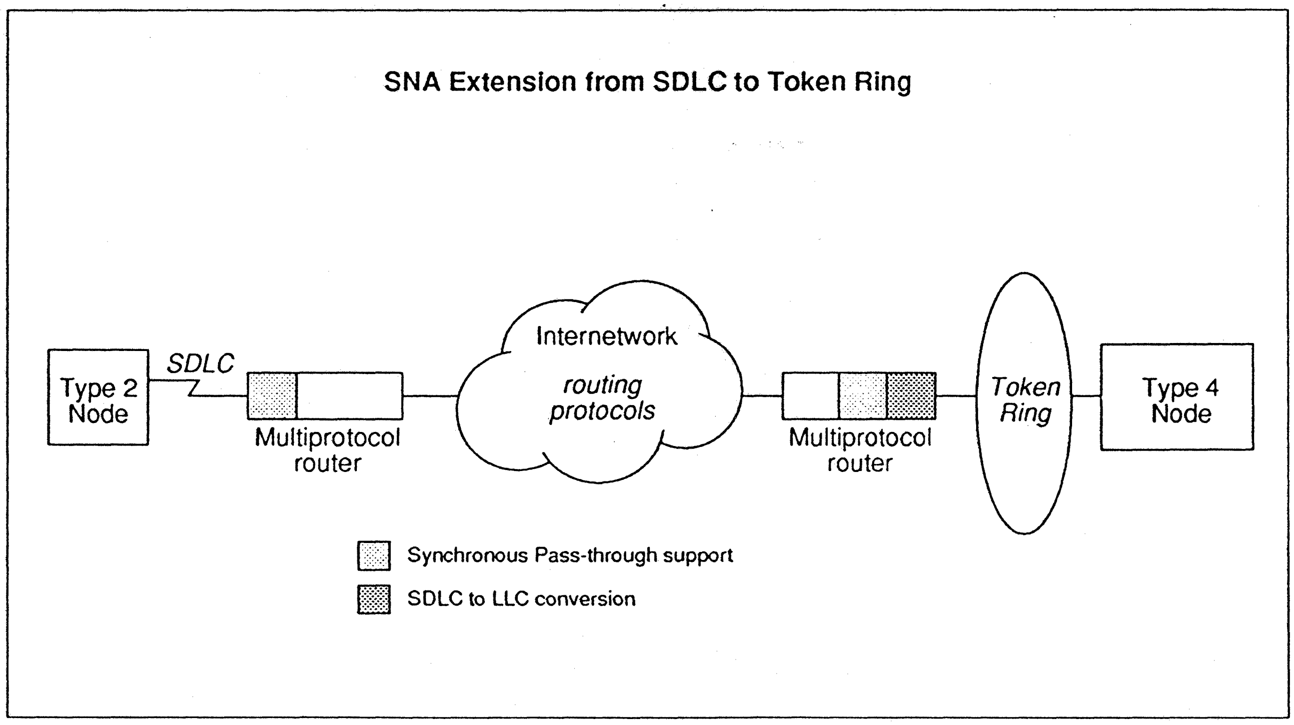

Figure 4

has been available for years and has allowed many of the problems with SNA through X.25 to be worked out in the real world. Therefore, users would be more likely to use IBM's X.25 support to connect SNA nodes directly to the X.25 network, as shown in the upper part of Figure 5, rather than through a multiprotocol router, as shown in the lower part of the figure.

Partial Type 4 Routing

Vendors provide several intermediate steps between SNA link-layer support and full SNA path control routing on multiprotocol routers. Type 4 nodes perform three kinds of functions:

- Type 4 nodes which *only* communicate with other Type 4 nodes perform path control network (PCN) routing functions.
- Type 4 nodes which communicate with hosts (Type 5 nodes) perform front-end functions in addition to routing.
- Type 4 nodes which support peripheral nodes (i.e., Type 2 or 2.1 nodes) perform boundary functions in addition to routing.

Partial Type 4 support would require the multiprotocol router to handle some or all of the boundary functions and routing functions. Most routers supporting synch pass-through already "spoof" SNA so that the internetwork does not get flooded with SNA polling messages between the routing node and its peripheral nodes.

In a multiprotocol router with certain Type 4 node functions, each LU could be configured to go different routes and to different destinations across the multiprotocol router's network. They might even be able to establish several separate sessions with different hosts if the router is capable of convincing the hosts that the LU "belongs" to them. However, the multiprotocol router on the workstation side would have to explicitly state all the LU names it accepts. In addition, the network would need to provide at least a primitive SNA directory scheme somewhere in that network (at the host or workstation side). This is needed because a multiprotocol router network with only partial Type 4 support would not have the LU name-to-route mapping inherent in a full Type 4 node in an SNA network.

Further, since there is no standard for how all this is done, any of these solutions would be vendor specific, which would add another level of incompatibility. Today's multiprotocol routers are already incompatible with those of other vendors in their intermediate system-to-intermediate system protocols, although standards are emerging.

Full Type 4 SNA Routing

Advantages—The advantages of providing full Type 4 routing on a multiprotocol network include the following:

- Compatibility as native node routers in the SNA network
- The router could be placed anywhere on a subarea network
- Each multiprotocol router with SNA routing would be independent of all other routers—there would not be "one link in to one link out" across the internetwork
- Full Type 4 addressing, with full name-to-address mapping, would be supported
- Full staging (session routing or session switching) would be provided
- Support of transmission groups would be provided, which concentrates several lower speed links into a wider virtual link
- The router could participate in subarea routing, handling explicit routes, virtual routes, and transmission control groups
- Network management visibility would be provided for the whole network
- A Type 4 node in an SNA backbone participates in flow control of virtual routes and of explicit routes

Disadvantages—Some disadvantages of using multiprotocol routers with full Type 4 SNA routing include:

- The router probably provides only a subset of full Type 4 functionality. For example, it could be missing:
    - Dynamic loading
    - Adaptive routing
    - Node Type 2.1 support
    - SNA-X.25 support
    - Trace facilities
    - Group poll support for 3174 Token-Ring gateways
- The router may face challenges when presented with:
    - Pacing
    - Multiplexing
    - Flooding
    - Congestion
- The router may not have a way to know whether the above problems are occurring.
- It is a challenge to emulate SNA, since IBM is not consistently open in providing specifications. For example, IBM has provided new releases of its NCP code less than two years apart for the past several years, but it has not published its Type 4 node specification in ten years. Some of the more recent features, such as Dynamic Path Update, are totally unexplained in IBM publications.
- NCP code is extremely complex, consisting of some three million lines of code. Users need to take note of which features are supported or which version or release level the multiprotocol router supports. If the code is based on a version or release earlier than ACF/NCP Version 5 Release 2 (for 3745s) or Version 4 Release 3 (for 3725s), for example, it would not contain support for LU 6.2 and node Type 2.1 for peer communications.
- Because no standards have been developed for SNA routing on multiprotocol routers, each implementation will be proprietary and incompatible.

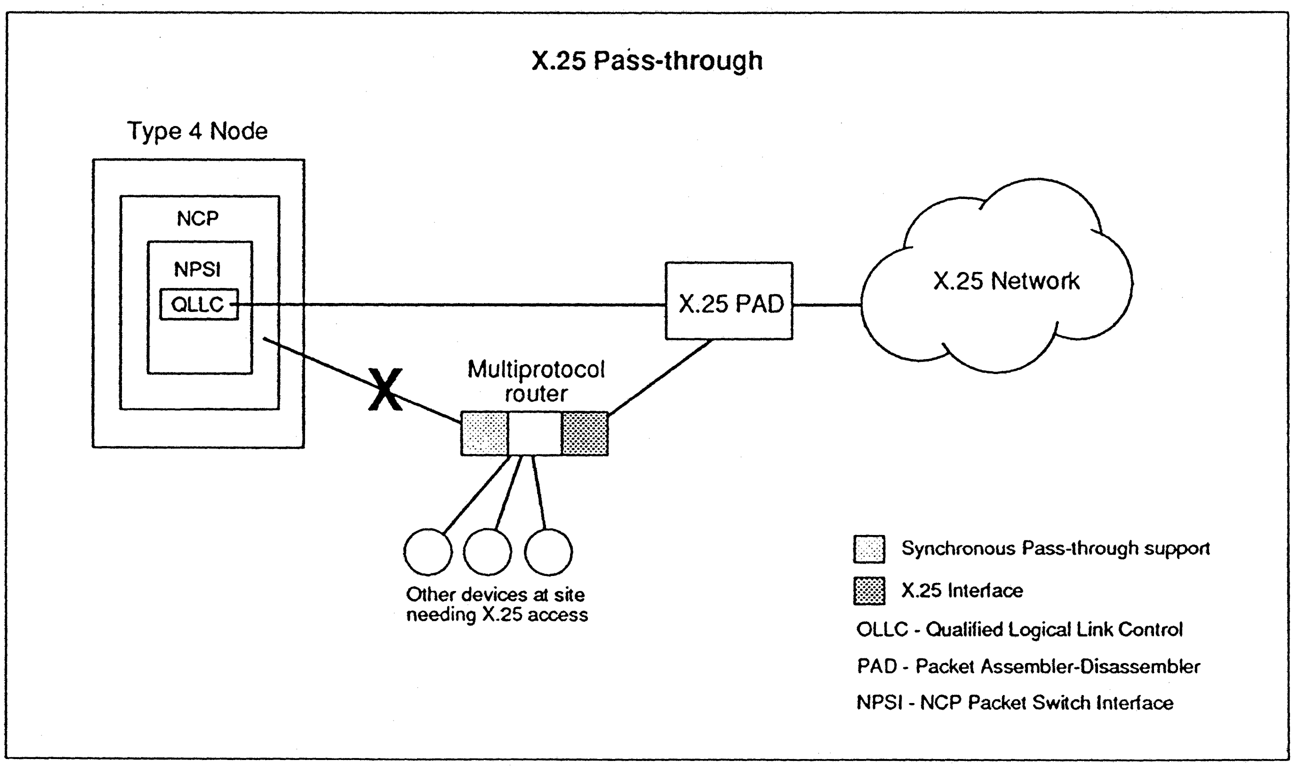

Figure 5

- To operate as Type 4 nodes in an SNA subarea network, these devices would need to be configured on the host. (This configuration would not be needed for synchronous pass-through.) Further, these host tables are now dynamically downloaded, implying that multiprotocol routers must understand NCP gens.

An Example

It is beyond the scope of this article to provide detailed, complete listings of all the vendors that supply or have announced plans to develop products that provide some or all of the SNA Type 2 or Type 4 internetworking functions described above. The briefest list would include such varied vendors as Amdahl, Cisco Systems, Digital Equipment Corporation, Interlink, McDATA, OpenConnect Systems (Mitek), NCR Comten, Systems Strategies, Wellfleet, Vitalink, and IBM itself. We believe it would serve our readers better to consider one company that has taken on the Type 4 routing challenge, Brixton Systems, which has contracted to supply this capability to Cisco Systems. This company, its focus, and its SNA products will be discussed in a separate mini-article in next month's issue.

What's Still Missing

Standardized SNA Interface—There is no standard yet for synchronous pass-through, so any solution will be vendor-specific and will need to be paired with products from the same vendor. IBM's product with Wellfleet may provide a basis for the emergence of a standard.

Peer SNA Networking—None of the current multiprotocol routers' SNA support includes LU 6.2, APPN, Node Type 2.1, or LEN node, though several companies have stated their intention to do so. A few companies provide APPC on their standalone or LAN-based SNA gateway products.

Network Management—Standardized network management in non-SNA environments is narrowing to OSI's CMIS/CMIP and TCP/IP's SNMP. IBM is also working (see the Network Management article in this issue) to develop closer ties between NetView and these standards.

Conclusions

The Decision Process

Even given the known benefits of internetworking, SNA users are understandably reticent to embrace multiprotocol routers for SNA. Hierarchical SNA was designed to be a closed system for IBM equipment and that assumption is apparent in the networking elements, which are strongly intertwined. Even though IBM may support greater openness in its newer developments which address internetworking issues on the drawing board, implementations in real-life installations have caused significant, unanticipated problems. It took IBM years and several iterations to support SNA robustly on X.25 networks.

Several approaches are becoming available to transport SNA in multiprotocol environments. The user should take into account several issues before deciding which solution is appropriate:

- What degree of SNA and non-SNA network integration is required?
- Is the need for occasional traffic flow or constant availability?
- Is the need for reduction in cost by eliminating some redundant lines or is it for more integrated network management of both environments?
- Is the need for access for all users to all resources in both environments, or is the need primarily for one-way flow from users in one environment to resources in another?
- Is the need for connection between SNA and only one other protocol, such as TCP/IP, X.25, or DECnet, or for integration of several protocols? For the former, a protocol-specific product from, for example, Interlink or Mitek, might provide more capabilities than a more generic product.
- Is the need for a long-term strategy or a short-term solution to a specific set of issues? With regard to SNA over multiprotocol routers, these products are perhaps too new to be the basis for a long-term design—the quality and extent of upper layer support in the real world remains to be seen. ■

(Continued from page 1)

Open Platforms, Open Protocols, and Open Partnerships

IBM's new network management strategy has three principal components—open platforms, open protocols, and open partnerships (see Figure 6).

Open Platforms

As a cornerstone of this move to openness, *SNA Perspective* believes IBM will make the "open" RS/6000 with AIX its principal platform for future network management development. The reasons for this move are partly technical and partly logistical, as we will see below. Consistent with IBM's overall strategy of positioning the mainframe as a centralized database/data server in a network of workstations, a mainframe running NetView will retain a role in IBM's network management strategy—but it will be recast. IBM is likely to create a distributed client/server NetView functioning as a management information base (MIB) for a constellation of satellite RS/6000s running the next generation of network management applications such as performance management and security management.

Open Protocols

Complementing open platform and distributed NetView, IBM shows signs of having decided to migrate to OSI's Common Management Information Services/Common Management Information Protocol (CMIS/CMIP) as its internal management protocol. Even SNA products operating in SNA networks will use CMIS/CMIP to exchange data and commands. The management information to be exchanged between the RS/6000 client and the NetView server would be in CMIS/CMIP format, and this change would happen faster than the internal NetView transition to CMIS/CMIP.

Open Partnerships

SNA Perspective sees this trend to embracing open protocols making possible the third element in IBM's new strategy—the recent network management alliances it has concluded with traditional

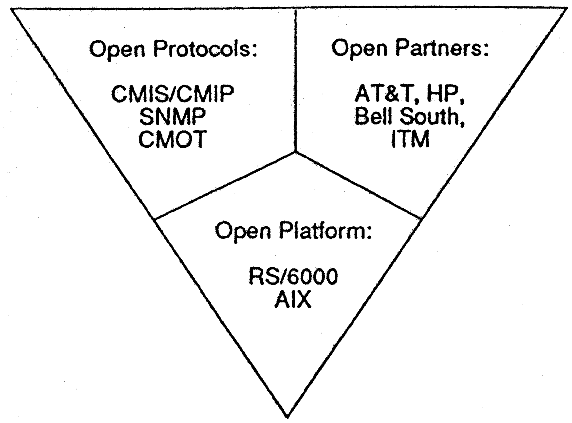

Figure 6

rivals including AT&T and Hewlett-Packard as well as other third parties. As IBM has discovered over the years, it's hard to get the other kids to play with you if you insist on using your ball and your rules. IBM's recent forays into partnerships and cooperative development is only part of this larger change.

The Networking Market

For IBM, which is at heart a marketing company that happens to be in the information systems industry, the hand that guides its actions is the logic of the marketplace, not the elegance of the technology. Today the networking marketplace is undergoing a significant shake-up, giving SNA its greatest challenge since its introduction in 1974. Open, nonproprietary protocols like TCP/IP (current) and OSI (future) have emerged, and present the possibility of replacing most proprietary protocols. The critical mass afforded by this consolidation has changed the market realities confronting IBM. Whereas SNA's market share was largely immune to inroads from other proprietary protocols, it has not been so fortunate with regard to open protocols.

Few information systems installations are all-Blue today. Companies want to manage their non-IBM equipment and facilities in the same way they manage their IBM products. IBM does not want to have its network management products, particularly its flagship NetView, shut out of the growing open systems segment of the computer marketplace. Most other computer and communications suppliers are converging on OSI's CMIS/CMIP and/or TCP/IP's Simple Network Management Protocol (SNMP). This places IBM in a predicament. It has two choices: continue with proprietary management protocols or embrace CMIS/CMIP and SNMP in a big way. Continuing down the proprietary management road would preserve its advantage in the SNA marketplace but shut it out of the growing open market.

Open Platforms: Switching Horses

If we assume IBM has decided that open network management is the way to go, one part of the move is to an open platform, which *SNA Perspective* believes will be the RS/6000. Two advantages come from this move—one is technical, the other logistical.

Technical Benefit

The technical benefit is in matching the task with the appropriate tool. The RS/6000 is a very fast machine that is more capable of executing high-computation algorithms than the mainframe. The mainframe is designed for high numbers of small-volume users rather than a single high-powered application (see Table 1).

Network Management Platform Strengths: RS/6000 versus S/3x0

	RS/6000	S/3x0
Portable Application	+	—
CPU Power/ Programs	+	—
Database/ MIB	—	+

Table 1

This is a big change for IBM. Since it began working on network management over fifteen years ago, the platform has always been the mainframe. Certainly, at that time, it was the only appropriate platform. Leaving aside some tentative IBM moves with the AS/400, until now NetView on the mainframe has not had serious competition. The RS/6000 has changed that.

Logistical Benefit

Another principal benefit of the open platform approach is access to a wider third-party development community. Much third-party network management software already exists for open platforms and third-party software houses can be contracted to write more.

There are two realities of third-party software today. The first is that it is advantageous to write for the open system market. Not only is this market segment growing, it is easier to write portable code in an open (read UNIX) environment. It stands to reason that a program that can run on any open box will sell more copies than one written for just one specific operating system on one specific machine (unless that specific machine has a market dominance that IBM no longer enjoys). As a result, third-party developers are shying away from doing projects on proprietary operating systems.

The second reality of third-party software is that the mainframe is less frequently used today in third-party development environments. Most development is done on workstations and mainframe support, when needed, is provided off-site by service bureaus. Familiarity with the mainframe sufficient to develop systems software is not attained over a switched line—programmers end up knowing more about the modem than they do the host. But since many network management applications on an IBM operating system such as MVS or VM are privileged, their programmers must have a deep understanding of memory management, process scheduling, and file access and manipulation in order to avoid conflicts with other systems applications running concurrently. Also, NetView and the host applications before it have all been designed as VTAM applications, with the attendant challenges, such as using S/370 assembler and using the 3270 data stream in writing user interfaces. This has eased somewhat with the emergence of C, PL1, REXX, and Knowledge Tool in the last few years, but programmers for the mainframe are still hard to find.

What all this means is that IBM cannot hope to attract the sort of programming interest it is counting on if the target environment is the mainframe. (This would be so even if NetView is eventually adapted internally to conform with the OSI CMIS/CMIP standards, as we also believe will happen.) That is why moving as much as possible to the RS/6000 as quickly as feasible is a smart move, and IBM is moving down that road. In the last six months it has assembled, without much attention, a considerable array of third-party relationships to develop products for the RS/6000 (see Open Partnerships below) that clearly foreshadows the product's role in IBM's strategic plans. IBM seems to have convinced these developers some time ago that RS/6000 is a, and perhaps *the*, target platform for IBM's future network management applications. *SNA Perspective* has heard, for example, that IBM is using the workstation for a new voice network management system under development jointly with Siemens. Increasingly, new IBM network management applications such as performance management, accounting management, and automated operations will likely run on the RS/6000.

Creating a Client/Server NetView

So what would IBM do with NetView if the company commits to UNIX workstations for its network management future? Two hurdles stand in the way—technological inertia and competitive exposure. In a world where software doesn't rust, somebody is always going to cling to old software. It is difficult to make the transition from one family of applications to an incompatible successor, especially since one of IBM's primary goals is to maintain its customers' investment. In addition, whenever a company tries to effect a major transition, its customer base, previously locked into its product line, is exposed to the predations of its competitors. This levels a playing field in which the company used to be king of the hill.

Putting NetView's proprietary nature aside for a moment and comparing it with competitive network management systems, NetView is superior in most areas and competitive in the rest. NetView has been the workhorse of IBM's network management for six years. There are no major fundamental flaws with NetView itself. But it is locked-in in some ways that limit its future. Let us assume for the moment that IBM is focused on an open strategy based on the RS/6000. To contemplate porting NetView to UNIX is lunacy. So what will IBM do?

We anticipate that IBM will start a careful stripping of functions from NetView, retiring the old code some time after developing the new implementation for the RS/6000. In other words, stabilize NetView, migrate its functions from the mainframe to new applications for the RS/6000, and use the mainframe as a database server—a network management information base (MIB).

This is a reasonable plan. Despite all the boxes on IBM's transparencies, NetView as it is presently implemented has only four basic components—a VTAM operator's console, a hardware problem management package, a session/logical resource problem management package, and an underlying configuration database. (The Command Facility is merely an interface to these four components.) Migrate the problem management applications to RS/6000s (new applications, not merely ported NetView) and redirect the VTAM console to an attached RS/6000. This would leave on the mainframe a largely irrelevant Command Facility but, more importantly, that part of NetView that arguably *should* be on an IBM mainframe—the database of configuration and administrative records. This new NetView may take on more the appearance of a DB/2 application than a VTAM application.

It should be noted that the idea of breaking up NetView is nothing new. Partitioning NetView into a distributed application has always been IBM's intention, but this has long been frustrated by two problems—a lack of requisite peer (APPC) implementations and a shortage of programmers to carry out the plan. (More later about the shortage of programmers.) The original plan, however, was to have the partitioned NetView execute in a homogeneous environment of S/370 mainframes, depicted in Figure 7.

The new plan with the RS/6000 strategy is a more extensive change although, ironically, it may be easier (see Figure 8). NetView will continue to function as a central storage for administrative and configuration data. In this capacity, IBM has said that NetView will be one of the first users of the Repository Manager.

"Open Protocols, Openly Arrived At..."

We already discussed two reasons for IBM to move to open network management. The first was demand side—customers are opting for open systems, and these are growing faster than the proprietary market segment. IBM wants to be in fast growing markets, not fast diminishing ones. The second reason has to do with supply side—IBM is software "bound," unable to finish what it has already committed to do, let alone produce the menu of network management applications customers will need in the 1990s.

IBM increasingly sees the value it adds for its customers is as a systems integrator like EDS. If it doesn't have access to the software tools, internally or externally, to do the job then it can't win the business. This dictates enlisting third parties to help develop the tools. Such open partnerships require that the platform and the protocol both be open. As with the RS/6000 move discussed above, *SNA Perspective* believes that IBM has decided (internally but not publicly) to move from its proprietary SNA management protocols because of market pressures.

There is some evidence of this move, in addition to the OSI/CS product. IBM is having new communications hardware implement CMIS/CMIP. *SNA Perspective* understands, for example, that CMIS/CMIP has been implemented in 3174 microcode as the first of many such implementations (although it operates over Token Ring and not LU 6.2 yet). Of course, IBM could opt to implement and maintain both SNA and OSI management protocols across the product line. But, as we have discussed, it already has too few programmers to do the work it has committed to, let alone add the load of coding, testing, and supporting a second management protocol. This shortage is, in fact, another impetus for IBM's openness drive, as we will see in the next section.

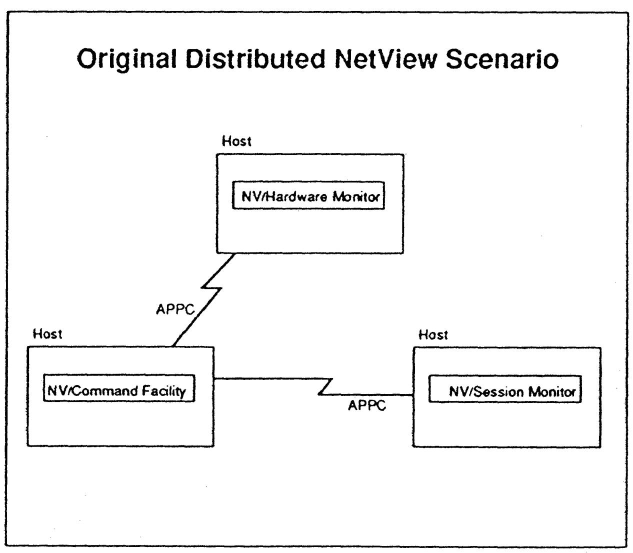

Figure 7

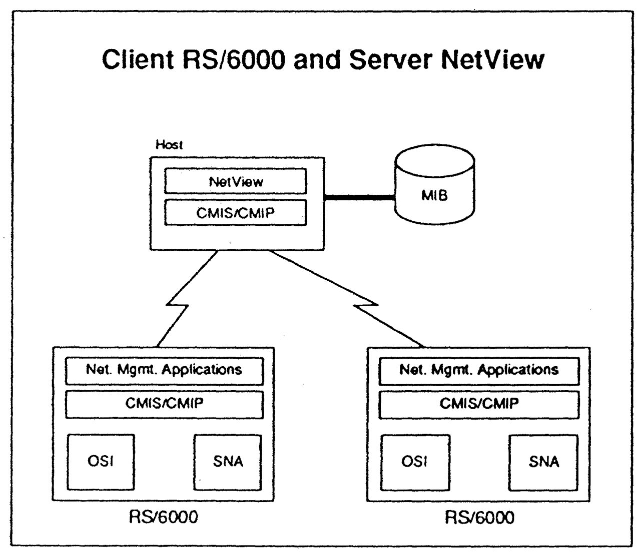

Figure 8

Open Partnerships for Management Applications

To put things plainly, software development at IBM is a bottleneck. Four years after its announcement, the company's Systems Applications Architecture (SAA) has been only partly realized and the flagship SAA products—OfficeVision and Repository Manager—are behind schedule. To make matters worse, IBM's top management has embarked upon a reorientation of the company from one that makes its money on hardware to one that makes its money on software. Confronted by its inability to produce the volume and quality of software it needs for this undertaking, IBM has been driven to do something it has been loathe to do before—recruit outside companies and their technologies to help implement IBM's vision. IBM has already invited third-party collaborations for its SystemView initiative and, as evidenced by deals in the last six months, it is taking a similar tack with its new network management strategy. Some of the more intriguing ones are discussed below.

Bell Atlantic Software Systems

Bell Atlantic Software Systems, a subsidiary of Bell Atlantic, the regional Bell operating company of the same name, has announced a product for the RS/6000 called IntelliManager. IntelliManager is an expert system that acts as an "intelligent mailbox" for the NetView operator by capturing the network management messages directed to the mainframe, and acts as an intelligent interface between the operator and the rest of NetView.

This use of artificial intelligence to build an "operator's assistant" keeps IBM very busy. At a network management conference in November 1989, IBM disclosed work on an expert system for diagnosing performance problems in SNA networks. Similar projects include YES/MVS, the Yorktown Expert System implementation of an MVS operator. For its part, Bell Atlantic has discussed adding modules to IntelliManager that would allow it to access on-line databases of the telephone company's management data for its switched, leased, and data networks. IntelliManager is an APPC application; IntelliManager and NetView communicate using APPC over an SDLC link by means of the NetView Program-to-Program Interface (NPPI) (see Figure 9).

International TeleManagement Corporation

International TeleManagement (ITM) and IBM plan to jointly market ITM's Multivendor Automated eXpert Management (MAXM), another product developed to work under AIX on the RS/6000. As the name indicates, MAXM is an open (multivendor) management tool. It functions as an SNA service point, much like NetView/PC, by converting non-SNA management data into SNA formats and management commands. Downstream connectivity to the non-SNA devices is via asynchronous links to the managed devices.

As a nice touch in distributed computing, ITM developed MAXM as two modules, the MAXM Server and the MAXM Concentrator, that can execute in the same or different RS/6000(s). Unfortunately, the MAXM Server uses the SSCP-PU session to communicate with NetView, indicating either that development of the product was started before the NPPI was available or that there was something about the NPPI implementation that could not support MAXM's needs. Since MAXM can be connected to an OS/2 EE via independent LU to independent LU APPC sessions (as a service point to gather network management information),

IntelliManager to NetView

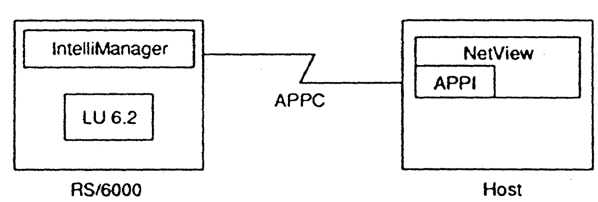

Figure 9

the absence of such peer connectivity to NetView is noteworthy. The second module, the MAXM Concentrator, contains a rule-based expert system that the customer can modify to execute local procedures for alarm and alert filtering, as well as script-based automation.

Wellfleet Communications

Wellfleet makes a line of bridges and multiprotocol routers. It is strongly rumored that IBM and Wellfleet are negotiating an agreement for the latter to port its routing software to the RS/6000 to act as a standalone multiprotocol router. This may be the most interesting deal of all. Additional support pointing to the RS/6000 as the platform for IBM's standalone multiprotocol router is the NSFnet. Funded by the National Science Foundation, NSFnet is a wide area TCP/IP network connecting university and government research centers. Since 1988, NSFnet has been operational and all the network's TCP/IP-based routers are IBM's, first based on the RT/PC and then the RS/6000.

If the RS/6000 is configured as a multiprotocol router with subarea SNA as one of the routing protocols implemented, then the Networking Systems product line may be in for some rough competition, especially the 3745. Even compared with the June announcement that supports an increase in processing speed of 50% to 150% for the communications controllers, the RS/6000 is still a blazingly fast box.

One question would be the proprietary nature of the subarea routing protocols: would Wellfleet be given access to the protocols or would IBM license the software and do the SNA itself? We expect that latter is the more likely scenario.

Imagine an RS/6000 with the intermediate routing functions of the APPN network node and with FDDI and frame relay interfaces. SNA networking would take on a vastly different appearance than it has today. This could be the sleeper announcement of the year.

IBM's Strategic Alliances

We must address one last area to which IBM's new direction in network management is taking it—the strategic alliances it has been concluding with such computer rivals as AT&T and Hewlett-Packard. IBM has traditionally scorned such efforts as almost collaborating with the enemy and has, in any case, been reluctant to concede that it might benefit from cooperation with its competitors. Indeed, one of the jokes about IBM's notorious insularity is that it even has its own term for "not invented here." Given this mentality, IBM has traditionally not sought out partners for purposes of technology acquisition, let alone product development. All this is changing. Although IBM and alliances have been rarely uttered in the same sentence, it is likely to become increasingly common. IBM has become prominent in most of the computer standards work such as the Network Management Forum and the OSI Implementers Workshop, as well as becoming increasingly active in many committees—ANSI, ISO, and even CCITT. Of the network management collaborations announced recently, two stand out.

AT&T Accumaster

Traditionally rivals, AT&T and IBM are not natural candidates for a cooperative development agreement. Yet in March of this year they announced just that. The pact covers two areas: joint development of software allowing NetView and AT&T's Accumaster to exchange configuration data, alerts, and management commands; and a mutual commitment to CMIS/CMIP for future versions of their respective products. Not surprisingly, Systems Center's Net/Master announced it will have equivalent capabilities. The first fruits of this deal will be realized using LU 2, which is unfortunate. But the plan is to move to LU 6.2 and SNA/DS. Depending on the timetables involved, it may be that OSI's DTP and X.400 support will displace these APPC parts of the project. Each side will implement gateways to the other's protocols—IBM will convert SNA formatted alerts into AT&T's Network Management Protocol, and AT&T will convert Accumaster commands and formats into NetView/SNA formats.

Hewlett-Packard OpenView

In April this year, IBM also announced an RS/6000 license of portions of Hewlett-Packard's OpenView network management software. The portions licensed include the OpenView Network Node Manager and the OpenView Network Manager Server. Both use the SNMP to monitor and control nodes in a TCP/IP network. The server also supports CMIS/CMIP over TCP/IP (CMOT), which is significant; IBM and Hewlett-Packard both have declared their intention to implement CMIS/CMIP, and CMOT is a natural management bridge to the TCP/IP world. IBM plans to migrate functions from its own Network Management/6000 product to the OpenView product. This announcement is a useful reminder that the open systems community has not repudiated TCP/IP.

Conclusions

Some big changes are in the works at IBM, changes that result in a new direction for IBM's whole approach to network management. These changes will see NetView recast in a client/server mold and function as a centralized database server for satellite clients operating primarily in RS/6000 UNIX workstations. New network management applications are under development for these RS/6000s and we expect we will eventually see all the non-Repository tasks now executed by NetView migrate to these RISC workstations. IBM has chosen to base its network management strategy on the RS/6000 because of its high performance CPU, which will allow more computationally intensive network management software than can be written for NetView executing in the host. At the same time there is a logistical dimension to IBM's move to the open RS/6000. By switching to the RS/6000, IBM will ease the task of third-party software developers who want to write network management applications that are portable, at least among the workstations running OSF UNIX. Simply put, developers are becoming remarkably resistant to writing applications that tie them in to a given manufacturer's proprietary offerings.

IBM has been very successful in getting some interesting partners to write network management applications for the RS/6000. All this is consistent with an IBM that no longer believes it can do it all alone. It has been working with third-party software and systems houses to develop pieces of SAA and SystemView and now is obviously seeking similar collaborations in network management.

The key here is open systems. In recent years, developers' reluctance to tie their products to a proprietary operating system have sought refuge in the portability of UNIX. Developers want to write for UNIX systems. As it happens, IBM now has a very hot UNIX workstation, the RS/6000. Seems like a good fit, doesn't it?

But a funny thing happens when a company opens up closed systems. Its installed base becomes open to competition. IBM's competition may be another group better able to compete with the RS/6000 than the S/3x0 mainframes. On the other hand, if IBM can be successful in the open systems market then it will probably keep much of its SNA installed base even when the management is open. This is definitely a market-driven change.

What All This Will Mean to You

Current and prospective users of IBM's network management will be affected differently by these changes. If you already have an extensive investment in IBM's network management and are happy with it, then sit tight. IBM is going to make the transition as transparent and as painless, if not as inexpensive or as quick, as possible. In any case, NetView users, fear not: the mainframe will not be entirely superseded.

On the other hand, if you are unhappy with the procrustean bed aspect of IBM's current network management, be assured it is going to change. If your goal is open network management there is no sense in cutting your requirements to fit IBM's cloth when IBM is abandoning those constraints anyway. An investment in open network management systems that conform to international standards (ISO or RFC) should be a safe investment. ■

Architect's Corner

APPN Treasure Hunt: LU to APPN Name Resolution

by Dr. John R. Pickens

I'm dying to know. How many of you have tried to establish a (documented) connection between IBM's peer logical unit specifications (LU 6.2) and APPN end-node specification? Based, that is, upon published specifications? I have, and came up empty-handed.

Hmmm. "But I thought it was all published," you say, "APPN end node is open, isn't it?" But is it really? Is something missing here?

The Background

Real APPN products allow applications to issue allocates using Partner_LU_Names that are not defined a priori in the local LU database. Transparent to the application, the APPC environment somehow resolves undefined Partner_LU_Names using the services of a serving APPN network node. This capability is an important requirement. So, from an architectural specification perspective, how is it done? How does an allocate() issued through a TPRM-compliant interface get routed to the APPN Directory Services function?

Check the Specs

Well, the logical strategy is to first check the current published version of LU 6.2 psuedo-code: Peer Protocols (SC31-6808-1). Checking the handling of the allocate verb we find (simplified):

```
Procedure ALLOCATE_PROC
      If ALLOCATE parameters valid ...
   .... set ALLOCATE_RCB.LU_NAME ...
   Send ALLOCATE_RCB request to RM

Procedure RM
  When ALLOCATE_RCB
    Call ALLOCATE_RCB_PROC (ALLOCATE_RCB)

Procedure ALLOCATE_RCB_PROC
  ...
  initialize everything in RCB
```

No (psuedo-code) checks for Partner_LU_Name validity in this section.

Later, when a verb is issued which requires an actual conversation (such as RECEIVE_AND_WAIT), GET_SESSION_PROC is called within the Resources Manager:

```
Procedure GET_SESSION_PROC
  Find the PARTNER_LU identified by the
RCB.LU_NAME
  If the mode is closed then
    Send a SESSION_ALLOCATED record with
    return code of UNSUCCESSFUL_NO_RETRY to PS
```

Nowhere is there any logic to detect an undefined Partner_LU_Name, other than creative interpretation of the logical comparison involving "mode is closed".

Obviously, this psuedo-code provides no hints. (Nor would it be expected to. This version of peer protocols was published prior to the APPN end node specification.)

So what is our next tactic for determining how and where undefined Partner_LU_Names get resolved?

Now, a comment. IBM developers have suggested to me that, in certain product implementations, BINDs are sent from the end node to the network node, and the network node always does all the work. Given the lack of a check for undefined Partner_LU_Names, how might that be done? And why isn't it in the pseudo-code?

This approach is a red herring, not the answer. Implementations operating at the LEN node level require the user to predefine all Partner_LU_Names as mapping to the LU name of the APPN network node directory server. Some implementations have tricks, like allowing LU names ending with wild card characters to be sent to the network node, which resolves the final destination or generates a negative response.

Good pragmatic solution, but not the answer.

And, no, the answer is not magic. The answer?

The Answer

No (that's right, no!) published pseudo-code exists. The mapping from the LU to APPN name resolution is an exercise for the reader.

Perhaps the real issue here is where do the requirements of open published architecture end and closed implementation begin? This seems to be one of the first cases where the full operation of an SNA architecture is not described in pseudo-code. True, the protocol is specified. True, the implementation details would not be terribly difficult. But the specification is nonetheless lacking.

Will IBM publish updated pseudo-code? Perhaps. Informal memos must exist within IBM describing how the linkage between the LU and APPN might be provided. But no immediate plans apparently exist to publish updated LU documentation. Certainly internal IBM product developers have figured it out. It is therefore only a matter of time before non-IBM vendors do the same.

I think, on this issue, that a minor architectural hole has been allowed to open—perhaps inadvertently. This feature does have impact all the way to the transaction program API specification (a parameter on allocate and mc_allocate allowing the application to disable the directory server query, for example). An update to the LU and TPRM specifications would seem to be required. ■

IBM Announcements

IBM Enhances 3745s, Adds Multiprotocol Router to T-1

On June 19, 1991, in a worldwide news conference based in London, IBM made a variety of communications announcements. IBM's Communication Systems line of business, now headquartered in Southampton in the U.K., was renamed to Networking Systems. All of the announcements came from Networking Systems.

Details on selected announcements are provided below. The announcements, in summary, are:

- 3745 communications controller—two new models
- Ethernet LAN adapter card for the 3745, to ship in the third quarter of 1992
- Statement of direction to add Ethernet support to NCP
- LAN/WAN Exchange—bridge/router enhancement to IBM's T-1 multiplexers
- Statement of direction to add Frame Relay to 3745s
- IBM joined the Frame Relay Forum
- X.25 Interconnect (XI)—primarily a currency release to support the CCITT X.25 1988 standard, ACF/NCP Version 5 Release 4, and X.25 NPSI Version 3 Release 4
- Non-SNA Interconnection—new (and final) release of NSI to support enhancements in ACF/NCP Version 5 Release 4
- X25Net Switch—PS/2 software to build a private X.25 network of IBM and non-IBM terminals and hosts

- NetView Performance Monitor—two new releases with enhanced support for X.25 and LANs
- AIX NetView Service Point, which does for AIX and UNIX what NetView/PC does for OS/2 and DOS
- NetView Graphics and Automation Offering—a combination of software and services to support custom network management solutions
- NetView Call Accounting—enhancements including a centralized U.S. and Canadian tariff database

Hot Topics

TCP/IP—Ellen Hancock, IBM Vice President and General Manager of Networking Systems, stated that IBM was increasing its investment in TCP/IP because it has been told by many customers that they do not consider TCP/IP an interim step. When asked if TCP/IP would be added to SAA, she said TCP/IP is already on all SAA platforms and IBM may reassess including TCP/IP in SAA.

APPN network node—In response to a question regarding the APPN network node during the question and answer session, Ms. Hancock said that IBM was reevaluating its decision not to publish (she actually said "our *ability* to publish") the network node specifications and "within a few months" will announce the results of its reassessment.

CICS—Ms. Hancock stated that transaction processing was one of five primary emphases of IBM Networking Systems and that CICS was the base for all transaction processing within the company. During the question and answer period, it was stated that IBM was considering porting the CICS application programming interface (API) to other systems. When stating the importance to users of stable APIs, IBM listed three APIs as examples: CPI-C, CICS, and CallPath.

Standalone router—When asked about a future standalone bridge/router with multiprotocol support, Ms. Hancock said that the bulk of development would be IBM based, but that several companies have expressed a willingness to help in its development. *SNA Perspective* believes this is the project that Wellfleet is assisting with, even though Cisco's software is used in the LAN/WAN Exchange multiprotocol router for IBM's Transmission Resource Manager (NET's T-1 multiplexer).

More Details

3745—IBM announced two new models of its 3745 communications controller family: 310 and 610. Performance is enhanced by the new Central Control Unit (CCU), which replaces the TCM installed in Models 210 and 410. IBM tests indicate that the performance improvement is fifty percent to one hundred and fifty percent faster; the higher results are found if the customer also moves to the Buffer Chaining Channel Adapter from the Channel Adapter. The new controllers are priced at $189,550 (310) and $294,550 (610). Converting a model 210 to 310 costs $42,000, and from a model 410 to 610 costs $72,600.

Ethernet—The IBM Ethernet LAN adapter card for the 3745 will not be available until the third quarter of 1992. It will support all seven models of the 3745 and will connect to Ethernet Version 2 or IEEE 802.3 LANs. IBM also made a statement of direction to extend NCP to support the new Ethernet LAN adapter and to route Internet Protocol (IP) traffic either across an SNA backbone or to IBM hosts running TCP applications. The Ethernet LAN adapters are priced at $21,420 (for 3745 models 210, 310, 410, and 610) or $8,030 (for 3745 models 130, 150, and 170).

LAN/WAN Exchange—This is a bridge/router enhancement to IBM's Transmission Resource Managers, the IDNX T-1 multiplexer line that IBM OEMs from Network Equipment Technologies (NET) of Redwood City, California. It integrates packet-switched LAN traffic with circuit-mode communications such as voice or video. The LAN/WAN Exchange is an NET product which is resold and serviced by IBM. IBM has also licensed its Token Ring Network Bridge Program technology to NET for future product developments.

The LAN/WAN Exchange's multiprotocol bridging and routing technology and protocols are from Cisco Systems and have been ported into NET's own code base while being optimized for WAN traffic. The network protocols supported are: TCP/IP, DECnet, Novell IPX, ISO/CLNS, XNS, AppleTalk, Apollo Domain, X.25 (1992), Chaosnet, and PUP. The routing protocols supported are RIP, EGP, BGP, HELLO, IGRP, and OSPF (1992). Both transparent bridging and source route bridging are supported, as well as concurrent bridging for both Ethernet and Token-Ring environments. LAN interfaces are provided for both Ethernet and Token-Ring; the Token-Ring card for the LAN/WAN Exchange is the only element provided by IBM. Available in December 1991, the LAN/WAN Exchange is priced from $8,320 to $13,420, in addition to the cost of a configured IDNX, which is between $16,000 (model 20-S) and $100,000 (model 90).

Frame Relay—IBM made a statement of direction to add a Frame Relay data termination equipment (DTE) function" to NCP. By having DTE function and not DCE function, IBM's communications processors will be able to connect to frame relay nodes, but they will not act as frame relay nodes themselves. IBM and NET joined the Frame Relay Forum as charter members.

X.25—The X25Net Switch is a program for the PS/2 with Micro Channel with Realtime Interface Coprocessor cards. Developed for IBM by AMNET of Framingham, Massachusetts, the X25Net Switch lets customers build a private X.25 network of IBM and non-IBM terminals and hosts based on PS/2s. It includes an integrated X.25 packet assembler/disassembler (PAD). A companion product, X25Net Manager, can manage these networks by itself or in conjunction with NetView. X25Net uses adaptive routing techniques and changes routes nondisruptively. In a statement of direction, IBM said these products are the initial members of a family of X.25 networking products and that the company intends to complement the current announcement by adding a multiprotocol concentrator (on the low end) and to support, in the first half of 1992, V.35 and X.21 high-speed line adapter cards, and later add X.75 and X.25 1988 (the current product complies with X.25 1984), as well as perhaps Frame Relay. X25Net is priced based on the aggregate bit rate of all ports, from $6,000 to $30,000. Also on the X.25 front, IBM announced a new release of X.25 Interconnection (XI), which is primarily a currency release to support ACF/NCP Version 5 Release 4, the new 3745s, and CCITT X.25 1988.

NetView Performance Monitor—Two new releases were announced of NPM, which allows customers to collect, monitor, analyze, and display performance data on SNA networks. NPM Version 1, Release 4.1 allows management of X.25 networks, and is available now. NPM Version 1, Release 5 offers acceptance of commands from NetView, dynamic configuration support, and the ability to collect statistics about data traffic sent over bridges between LANs. It will be available in February 1992.

AIX NetView Service Point—Designed to do for AIX and UNIX what NetView/PC does for OS/2 and DOS, the AIX NetView Service Point is an AIX or UNIX system supporting non-SNA equipment to exchange information with NetView. It will run on the IBM RS/6000 or on Sun SPARCstations. Priced at $3,150, it is available in July 1991. For more implications of this announcement, see the article on network management in this issue.

NetView Graphics and Automation Offering—This offering is a combination of automation software, graphics software, installation services, on-site training, and support to enhance customers' network management systems. For $49,800, the customer gets an upgrade to the current version of NetView; license charges, installation, and customization of the IBM Automation Network Operations (ANO/MVS) and NETCENTER graphics administrator workstation software; on-site training; and one year of toll-free telephone support. ■

 Subscription rates: U.S. - one year $350, two years $520. International - one year $385, two years $590 • SNA Perspective is published monthly by CSI • 2071 Hamilton Avenue • San Jose, CA 95125 • Telephone (408) 371-5790 • Fax (408) 371-5779 • Managing Editor: Louise Herndon Wells • Associate Editors: Vincent Busam, Basil Treppa, Alan Kucharski • Circulation Coordinator: Cheryl Roberts • Contributors: Dr. John R. Pickens, Philippe Byrnes • The information and opinions within are based on the best information available, but completeness and accuracy cannot be guaranteed.