SNA Perspective

Volume 12, Number 4
April 1991
ISSN 0270-7284

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

Peer SNA: Getting Closer?

IBM has made several announcements recently about peer networking in SNA's future. In particular, the company has been brushing the dust of neglect off advanced peer-to-peer networking (APPN) with a flurry of recent trade articles, speeches, advertising, and announcements of APPN for the PS/2 and the 3174 and further integration of APPN with mainstream SNA.

Peer networking is catching on in a big way in the information systems market, fueled by the trend toward open systems. IBM's installed base of SNA customers is beginning to drift; some large corporations are turning to TCP/IP for multivendor, peer networking solutions for which they feel SNA is unsuited. In addition, while OSI may currently be little more than a distraction, it is a sufficient distraction that many all-Blue or very-Blue shops are seriously reexamining their choice of networking technology. In both cases, SNA's market share is being eroded and will continue to be eroded until IBM arrives at a satisfactory answer to the growing demand for peer networking in SNA. This article explores IBM's options and intentions for peer networking. It particularly notes various requirements that should be met by peer SNA for automated configuration management, self-healing networks, and adaptive routing protocols.

(continued on page 2)

Internetworking Vendors Target SNA

In early 1991, several internetworking vendors made announcements regarding support for IBM communications, particularly SNA. These companies had been adding increasing numbers of protocols to their multiprotocol bridge/routers; support for IBM protocols is a logical next step.

This article, the first in a two-part series, looks at the products and announced plans of selected internetworking vendors to support IBM connectivity. The second part, in the next issue of *SNA Perspective*, will address the technical issues involved in SNA routing in a multivendor environment.

(continued on page 10)

In This Issue:

(Continued from page 1)

IBM's Objectives for Peer Networking

What does IBM want from a peer SNA? *SNA Perspective* sees these as four principal objectives of IBM's architects of peer networking.

- Automated configuration management
- Self-healing networks
- Adaptive routing protocols
- Support for follow-on hardware sales

For each of these four objectives, this article will examine IBM's goals and user needs for peer SNA and then, for the first three, note the APPN issues and trends with regard to that objective.

First, IBM will want any peer architecture to facilitate the further automation of configuration management, especially keying the data-defining network resources. Second, it must address user requirements for the self-healing network: high-availability transport services with automatic recovery and real-time reconfiguration after component failure (e.g., when a link goes down). A corollary requirement is for adaptive routing protocols. Current SNA allows some recovery by switching traffic to backup routes, but all the routes must still be defined when each network control program (NCP) and virtual telecommunications access method (VTAM) are generated (sysgen). With adaptive routing protocols, routes through the path control network will be recalculated as the network or its traffic patterns change. Finally, a new, peer SNA will help sell more hardware—by using up more cycles and storage than master/slave equivalents, peer software will drive the sales of communication processors.

APPN Design Objectives

The APPN design objectives follow from positions taken on peer networking as long ago as 1985 by IBM's Low Entry Networking (LEN) group at Yorktown Heights Research. IBM identified the most likely area for peer networks—namely, the small system environment—and performed extensive market research of customer requirements. This resulted in the following design objectives:

- Customer installed and maintained: IBM engineers need not be involved in installation or maintenance.
- Decentralized control: no focal point(s) running network operations; requires the equivalent of distributed VTAM.
- Support for frequent changes: dynamic environment—plug 'n' play, the antithesis of NCP sysgens.
- Limited technical support: few if any systems programmers, network operators, etc., are required.
- Lower traffic volumes: original design was for small networks of small systems (PC, PS/2, S/1, S/36, S/38, AS/400), not SNA backbone networks.

The LEN group's efforts developed into APPN which was shipped for the S/36 in 1986. These small systems design objectives determined APPN's protocols. The two most important are decentralized control and lower traffic volumes, which define the peer orientation as well as the assumed operating environment. By assuming low traffic rates, certain options were available for designing the peer control protocols that are, unfortunately, difficult to scale up to the sizes and utilizations of today's large SNA backbone networks.

Automation of Configuration Management

The reason that automated configuration management heads the list of design goals for peer SNA has less to do with peer networking requirements and more to do with market pressure. Customers complain to IBM about all the data they must enter into countless configuration screens when adding or deleting devices in SNA networks. What IBM is hearing at GUIDE and SHARE meetings is that manual configuration management is no longer acceptable. At customer briefings on the S/390's communications facilities, the item that stirred the most interest was the new ability for MVS/ESA to automatically enroll configuration data when a machine powers on. Although now limited to channel-attached devices, the future of SNA configuration management can be seen in the new capabilities of MVS/ESA. Further, the new APPN for OS/2, announced in March 1991, has reduced the number of configuration screens from more than fifty to three.

How does this relate to IBM's drive for peer networking? It complicates it immensely. Recall that, in SNA's hierarchy of control, there is one focal point where configuration information is ultimately stored, namely the system services control point (SSCP) running its domain. Devising a database structure for storing and retrieving information that will be kept in one place is much easier than than partitioning and/or replicating it among all the peers. This is part of the argument for basing any future SNA on today's implementation, even if it means staying hierarchical and not bowing to "peer pressure."

It is possible and, in fact, easier to automate configuration management in a hierarchical than in a peer network. On the other hand, cumulative storage needed for configuration data is greater with peer than with centralized networks. So in designing automated configuration management, IBM has a choice between ease of implementation versus driving more hardware sales (its fourth objective for peer networking).

APPN Configuration Management

In a peer network, configuration data at each node must be accessible to other nodes. This implies a sophisticated query and search mechanism. APPN, for example, has each network node maintain a directory of the network related resources on itself as well as the network related resources on any attached end nodes. APPN's architects chose to handle this with a three-tier directory structure, partitioning configuration data between end nodes (internal directory only) and network nodes (directory of internal and attached end-node resources). End nodes automatically enroll their network resources with the attached network node(see Figure 1). Each network node participates in directed or broadcast search requests for configuration data not found on a local network node, again automatically. Note that not all of APPN's configuration management is automated; plain Type 2.1 nodes cannot automatically enroll their resources in the attached network node's directory—this must be done manually .

Self-Healing Networks

If centralized architectures have the advantage in configuration management, peer networks have the advantage in fault tolerance since a central point of control also means a single point of failure. Indeed, the idea of peer networking started as a means of achieving more robust networks. The peer approach to networking employed by TCP/IP and certain other architectures originally derives from the early U.S. Department of Defense ARPAnet development, when the protocols were designed for survivable battlefield communications. In the original TCP design framework for the ARPAnet, nodes were assumed to be mobile, intermittently faulty, and frequently subject to hostile fire. Datagram routing was employed because a virtual circuit might not stay up long enough to complete transmission and the return of acknowledgements. At the same time, no central point of control means no single point of failure, hence the advantage of peer over nonpeer when it comes to self-healing networks.

By way of contrast, the relative fragility of centralized control is exemplified par excellence by none other than SNA. Early SNA was the antithesis of self-healing: if a link went down it could (and did) take down the upper layer connections that were using it. Peer networks are more robust than hierarchical networks because centralized control implies a single point of failure. Again, in SNA's earlier versions, if the host running the SSCP went down, the entire domain it controlled also went down. This is not to say that IBM cannot make SNA more flexible while still keeping the hierarchical control. With current SNA, it would be possible

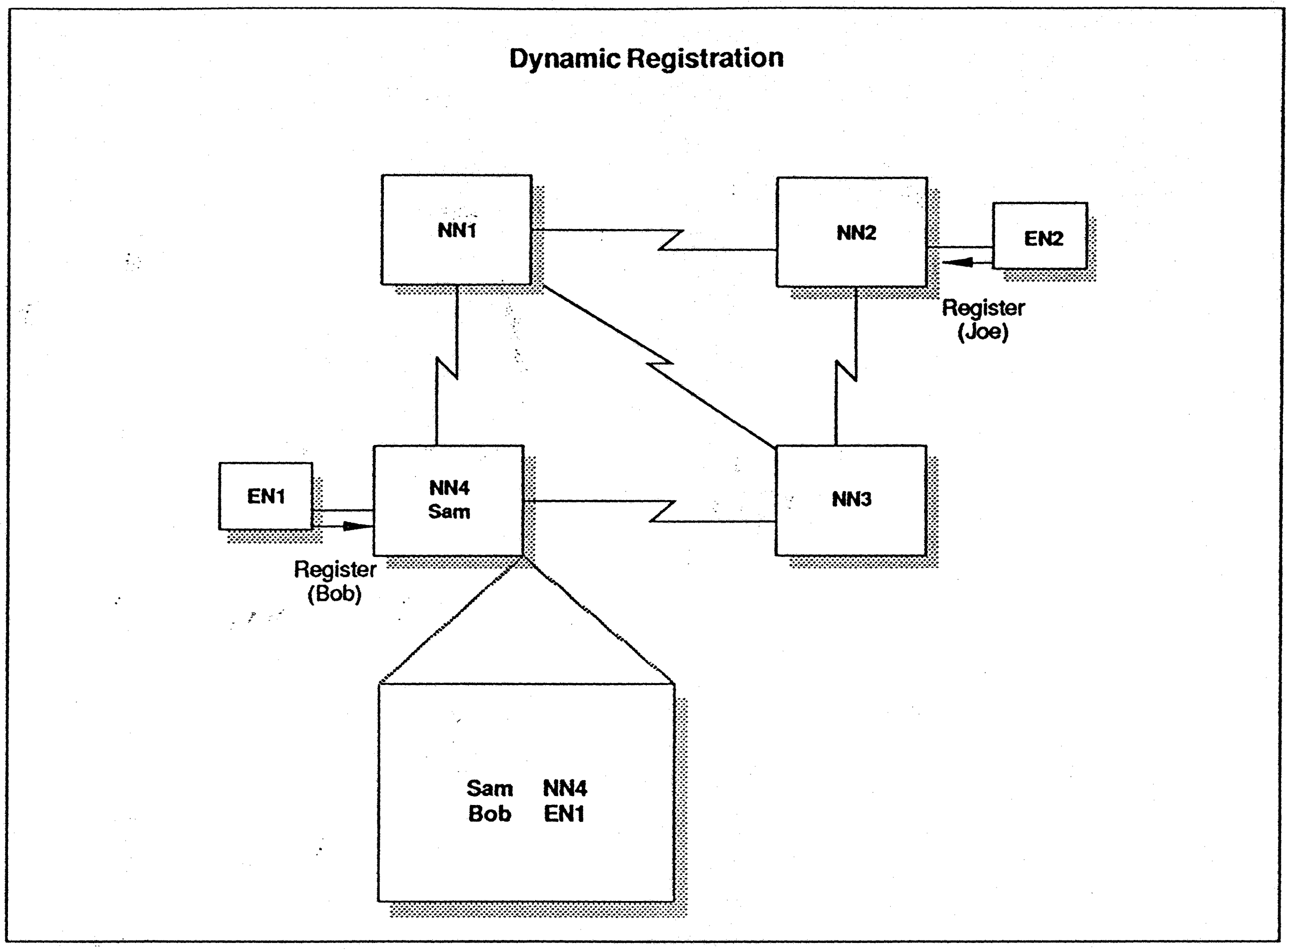

Figure 1

to automate sysgens in real-time (an adaptive reconfiguration) as the network changes with nodes coming up and leaving as they choose. However, IBM has chosen to focus on other priorities rather than incorporate these capabilities.

How APPN Does Fault Tolerance

IBM's customers today are asking for fault tolerant networks that are resilient in the face of failures in network components. Peer networking protocols are the most direct way to get there. At least part of IBM agrees: APPN, which was designed from the outset to be robust, is a peer architecture. And APPN can be accurately termed self-healing by virtue of its adaptive routing protocols.

APPN uses an architected mechanism for exchanging information concerning the network topology, a database of which is kept by each network node (see Figure 2). This replicated information is kept synchronized by the exchange of topology database updates (TDUs). When a change occurs—for example, if a link fails or comes back up—its locally attached nodes will broadcast a TDU throughout the network informing all the distant network nodes to modify their topology model of the network's interconnections. This is important because APPN network nodes contain a component called route selection services that calculates routes dynamically whenever a session is requested. Route selection services uses the current information in the topology database to avoid links that are down or even just congested.

Exchange of Topology Data Updates (TDUs)

Network Node 4 Joins an APPN Network

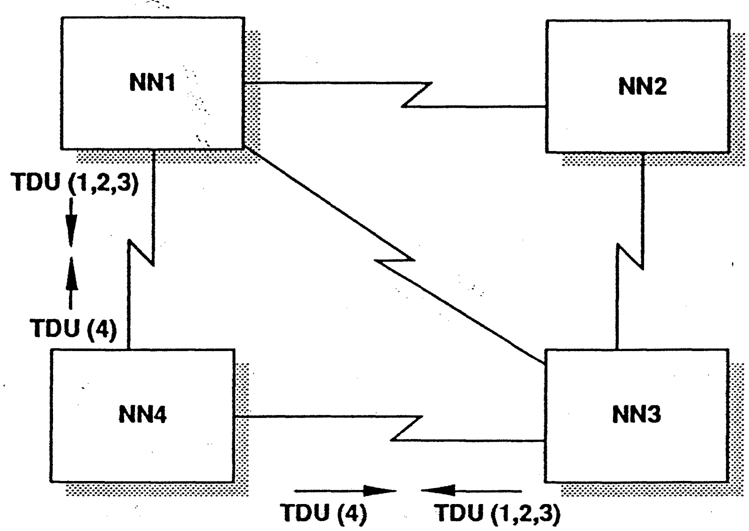

Figure 2

Adaptive Routing Protocols

An essential element of self-healing networks is provided by adaptive routing protocols, protocols that reconfigure the end-to-end connections in a network. Adaptive routing means that the way traffic is routed through the network changes as the network changes—for example, when a key link is congested or goes down. Changes in the network can be passed on to someone or something to calculate the new best end-to-end routes for carrying user traffic. Note that this, too, can be implemented adequately in a nondistributed, nondecentralized way. As with configuration management and fault tolerance, adaptive routing can be executed with a centralized control node such as SNA's SSCP.

Global versus Local State Information

The challenge in peer networking is the a time lag between when an event occurs, such as a fault that causes a data link to fail, and when information about the event has propagated to all the routing nodes in the network. What each node is aware of immediately is the state of its local resources. This is called local state information. For example, when a link goes down, this fact will be known to those nodes attached to it. But, to make decisions, more than local information is needed—each node must have some knowledge of the global state of the entire network. In effect, the global state is equal to the union of all the local states of the individual nodes. By sharing local state information, each node can reconstruct an estimate of the network's global state.

The problem with decentralized or peer execution of adaptive routing is that each node must construct for itself an estimate of the network's global state. With a centralized point of control, all the local state information can be forwarded to a single destination, i.e., the routing node. In contrast, with decentralized control there is no single destination for local state information; it must be shared among all the nodes that calculate routes for the network.

There can be a substantial downside to decentralization—the exchange of local state information can consume much of the network's carrying capacity—which is why decentralized adaptive protocols do not perform as well in heavily loaded networks. Also, as a network grows in size, communications bandwidth is increasingly devoted to such overhead rather than to moving user data.

APPN's Overhead in Adaptive Routing

This overhead problem is depicted for subarea SNA versus APPN in Figure 3. The two curves plot the relative overhead of adaptive versus fixed routing algorithms. As can be seen, fixed routes entail a certain amount of irreducible overhead but this tends to grow slowly as more users are added. In contrast, APPN's overhead starts out low but, as users (and nodes) are added, it climbs almost exponentially. For large SNA accounts with hundreds of mainframe computers and tens of thousands of users, the cost of this approach is prohibitive. IBM has privately disclosed to major accounts a future in which the two routing schemes would coexist in a network with distinct operating regimes—APPN when lightly loaded, subarea SNA when heavily loaded, and a "phase transition" between where the two are mixed in some proportion.

Hardware Requirements for Peer SNA

Anyone who sells hardware may spend a certain amount of time trying to devise ways to increase the loading of customers' installed equipment. This might be called the "baby needs a new pair of shoes" design philosophy. IBM is probably aware of the hardware ramifications of its choices for future networking protocols because it is always trying to sell hardware. Hence IBM does not likely consider increased overhead to be necessarily an impediment to peer SNA's development.

Peer networking protocols are more complex than their master/slave equivalents. All other things being equal, nodes in a peer network must execute more cycles and store more bytes because additional code must be executed in each of the network nodes and because additional information is necessary for distributed/peer algorithms to work. When each node contains essentially a mini-VTAM, it is going to have more overhead running the networking software than if all the hard work is off-loaded to a single SSCP.

Similarly, a packet in a datagram network has greater overhead than a packet in a virtual circuit

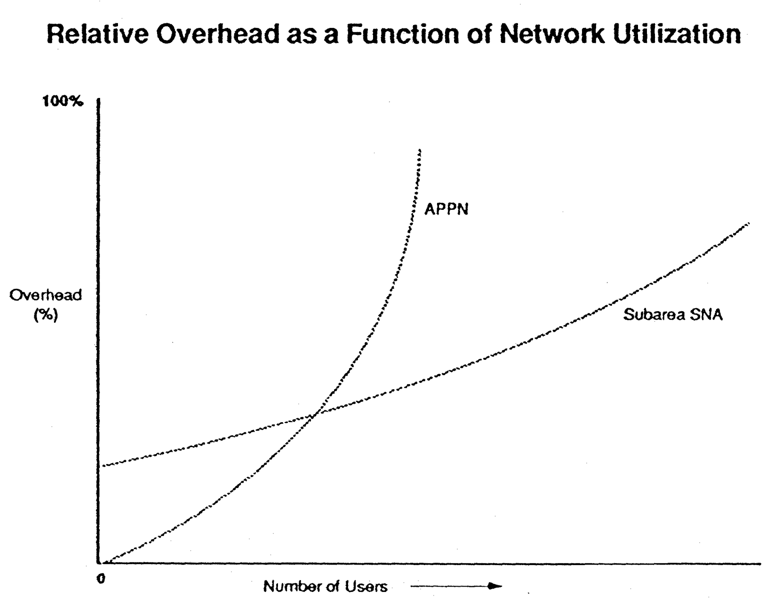

Figure 3

network (see Figure 4). A packet in a datagram-oriented network must carry full network identifiers for both the origin and destination. In contrast, a packet in a virtual circuit network need only carry its VC identifier, which is mapped by all intermediate nodes to incoming and outgoing data links. Finally, an added load is placed on the links of the network when decentralized routing protocols are employed. We have seen how the overhead created by APPN's TDUs, as well as broadcasts and searches, rises faster as the traffic loading increases.

SNA and APPN: Coexistence, Convergence, or Overlap?

What are IBM's options?

- Continue to evolve SNA, increasing its automatic and fault-tolerant features.
- Grow APPN into a fully developed replacement for SNA.
- Merge the two architectures.
- Wait until a third architecture is available, notably the OSI routing standards that are still under development.

Putting aside the OSI option for the moment, the choice is essentially between fleshing out APPN with NCP's Type IV subarea routing and flow control versus reengineering SNA's most important internal protocols (those controlling the transmission groups, explicit routes, and virtual routes that constitute SNA's path control network) so that they are more flexible and adaptive in real time. That is, whether to add SNA to APPN or to add APPN to SNA.

How IBM resolves this decision is heavily dependent on the resolution of protracted internal political conflicts.

IBM's Internal Communications War

IBM employs a strategy of "management by contention" in which two or more groups within the company build products for an identical purpose and the fittest makes it to market. Traditionally, IBM has been divided by a competition over SNA and APPN between, respectively, the Communications Products Division (CPD) group in Raleigh and a coalition of forces including the Research and Application Business Systems division.

The SNA partisans assert that IBM can solve what are often thought to be peer requirements by means

Relative Overhead of Datagram versus Virtual Circuit Packets

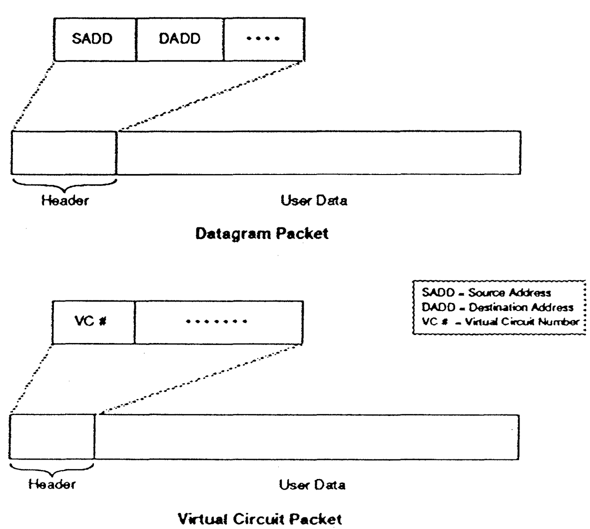

Figure 4

of a nonpeer, centralized architecture to be developed from today's SNA. APPN advocates respond that APPN already has many of the peer features necessary.

The argument between the APPN revolutionaries and the SNA evolutionaries has been further muddied by an uncertainty over what people mean when they say peer networks (see Figure 5). Along with "client/server" and "object-oriented programming," "peer networking" is an industry term that many endorse but few agree about what it means. As with beauty, peer is in the eye of the beholder, and one person's peer architecture may be another's master/slave. To IBM's APPN advocates, peer is synonymous with decentralized execution and distributed control; it is incompatible with the centralized control architecture on which SNA is built. Traditional SNA supporters within IBM argue in reply that users may say "peer" but they want is flexibility in running their networks and automation of many of the tedious set-up and maintenance tasks. They argue for accommodation of these small systems within the hierarchical SNA of today. Whichever prevails in the internal architecture, *SNA Perspective* believes the selling points of the new SNA will be user needs; users buy for functionality, not for architecture.

What About OSI?

Could the new peer SNA be built using OSI protocols as they become available? It could, and this would be consistent with the five-year horizon IBM has been discussing for peer SNA's arrival (see Table 1). It may be, in fact, that one of the reasons why the APPN network node has not been published is that internal debates are still ongoing regarding the extent to which OSI protocols would be incorporated.

There are problems, though, with staking such a major revision on an unproven technology. For one thing, TCP/IP's momentum may continue to grow—in five years it could be established as the de facto peer networking standard and OSI could be left at the gate. Will SNA's current users wait five years for what may be an OSI solution? Then, too, there could be uncertainty within IBM itself about

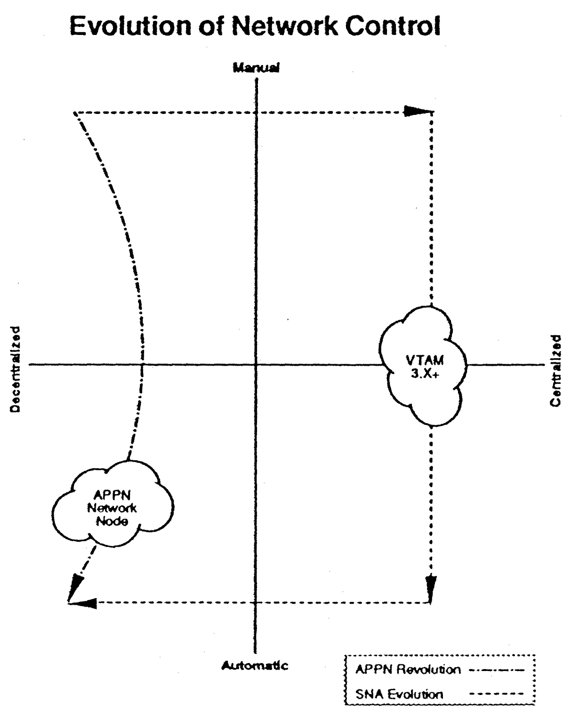

Figure 5

the future viability of OSI, at least in North America. IBM's OSI work is being transferred to Rome, Italy, far from Raleigh. Is this emancipation or banishment? But that's a topic for a future article .

Conclusions

Peer SNA is heating up. The question is, what's cooking? Has IBM at last reached a consensus on the APPN route, the SNA route, or something in between? IBM has been giving increased prominence to APPN recently, but may still be splitting its bets between a future emerging from subarea SNA and a future developed from peer APPN. This split could divide critical resources, especially scarce programming talent, so that neither is successful.

The arguments are well rehearsed. The evolutionists want to graft peer features onto SNA gradually. The revolutionists are the champions of APPN, which want SNA capability added on to APPN, primarily for customer investment protection and migration purposes. *SNA Perspective* believes that IBM has not resolved the internal debate between its networking evolutionists and revolutionists.

SNA Perspective believes that IBM's talk about five years before full emergence of new SNA will be unacceptable to customers, which would certainly challenge the potential for an OSI solution.

Automated Configuration Management

From what we know of IBM's interpretation of what its customers want when they say "peer," it differs from today's SNA in its automation of configuration management (including enrolling network addressable units, applications, new devices, and so forth). Manual configuration management—entering all the data pertaining to a network's physical and logical resources, their locations, identifiers, and such administrative impedimentia as serial numbers—is tedious work. Self-identifying hardware is already here. What have been missing are the corresponding software "hooks."

Summary of Respective Architectures

	Architecture	Openness	Automatic Configuration	Self-Healing	Adaptive Routing
Subarea SNA	Hierarchical	Proprietary/ Open	Advantage	Disadvantage	Uncertain
APPN	Peer	Proprietary	Disadvantage	Advantage	Uncertain
OSI	Peer	Open	Unknown	Unknown	Unknown

Table 1

Self-Healing Networks

The new SNA will move more toward greater availability by further automating problem management and recovery, creating self-healing networks that automatically reconfigure themselves when a fault occurs. This, in turn, requires that the new SNA abandon its fixed, predefined route definition and use instead an adaptive routing protocol or protocols. Peer networking implies distributed control—each network node makes routing and flow control decisions for itself—but adaptive control does not necessarily imply peer network protocols.

Adaptive Routing Protocols

Decentralized control is more robust than centralized control in the face of failure because centralized control allows a single point of failure. This accounts for most of the bias toward peer networking to be found in the open-system, multivendor community. Nevertheless, users are demanding that IBM deliver automated, automatic networking products. This objective does not mandate scuttling SNA's investment in protocols and redefining a new, peer path control layer along the lines of APPN or some future APPN extensions. Further, APPN's distributed routing protocol consumes an increasing amount of the total network bandwidth as size and traffic increase. The consequence of this is that APPN's metamorphosis into a full-fledged large networking alternative to SNA may also require an extensive reworking of its present routing scheme.

Peer is in the Eye of the Beholder

There are as many definitions of peer networks as there are people who talk about them, but the market is clearly sending IBM the message that SNA must go peer. Part of IBM believes that the solution is to change SNA to a peer architecture, while part seems to feel it would be enough that it is sufficient for SNA to provide the features and flexibility users would want from a peer network.

How long will IBM's strategy of "management by contention" lead it to continue to split its bets between SNA and APPN? IBM will probably leave us guessing until the last minute—what else would you expect from the world's most consummate marketing company? ■

(Continued from page 1)

The challenge of bridging SNA networks is not as simple as it may seem to those coming from a multivendor background. IBM's synchronous data link control (SDLC) is a serial synchronous protocol like X.25's HDLC. However, this barrier was broken early in the 1980s when multiplexer vendors added SDLC support. The delay in adding of SNA support by the internetworking vendors until now was not due to SDLC issues as such as much as limited processing capabilities of the earlier bridge/routers and vendor inexperience in multiprotocol routing for even non-proprietary protocols.

Routing SNA traffic, that is performing some or all of the functions above layer two, the data link layer, is an even more formidable task. IBM does not provide an open interface to SNA except at the data link layer and then again at the session layer with logical unit interfaces. More on this will be discussed in the next issue.

Wellfleet Communications

Wellfleet of Bedford, Massachussetts, is one of the leading suppliers of bridge/routers. Like Cisco Systems, Wellfleet is a young company that was founded in the mid-1980s to exploit the opportunties in internetworking and has vaulted to the top ranks of router vendors.

In January 1991, Wellfleet announced a new sync pass-thru feature to support SDLC and high-level data link control (HDLC). Transparent Sync Pass-Thru allows users to combine both HDLC (on which X.25 is based) and SDLC (on which SNA wide area communications is based) traffic, along with the other protocols supported on its bridge/routers. This allow a customer to use a single wide area backbone to support local area network bridging and SNA traffic as well as X.25 communications. Wellfleet has not announced an alliance with any experience SNA supplier, and SNA Perspective believes this would be a reasonable next item on the company's SNA agenda.

Vitalink

Founded in 1980 and based in Fremont, California, Vitalink Communications was a pioneer in remote bridges. A significant portion of its revenue for some time has come from Digital Equipment Corporation, which OEMs its remote bridges. While Vitalink still leads in the remote bridge market, the company has financially stagnated in the last few years as it lagged behind other internetworking suppliers in providing multiprotocol routers and newer features in its bridges. Vitalink has recently introduced a new high-bandwidth platform, the Enterprise Network Switch (ENS), which the company intends as its major entry into the high end of the multiprotocol routing market.

Vitalink does not have announced SNA products nor a well-developed SNA marketing story, but it has development efforts underway and has announced an strategic partnership with an OEM SNA connectivity supplier—Netlink, Inc. At the Comnet '91 show in January, the company announced an agreement with Netlink of Raleigh, North Carolina, to jointly develop SNA support on the ENS. Vitalink is also investing in Netlink.

In 1986, Netlink introduced SNA_Gate, an SNA gateway product. The SNA_Hub product line, introduced in 1987, is centered around a SNA network concentrator/router which connects SNA and non-SNA lines into SNA networks.

3Com Corporation

3Com pushed into internetworking with its acquisition of Bridge Communications in 1987. However, though formerly the only company dominant in both local and remote bridges, 3Com has seen significant erosion in its formerly dominant position in the bridge market due to its corporate focus on the workgroup client/server market. For similar reasons, 3Com has also lost ground to start-ups Wellfleet and Cisco in the fast-growing multiprotocol router market.

As part of a corporate restructuring in early 1991, 3Com has spun off Communications Solutions Inc. (CSI), which it had acquired in 1988 for CSI's SNA expertise which led to development of 3Com's Maxess gateway product. Though this change takes 3Com out of the PC workgroup-to-mainframe SNA market, 3Com maintains an SNA development team to focus on internetworking.

Likely because of this major recent shake-up in its IBM connectivity group, 3Com is the only internetworking company of those discussed in this article that has not officially announced its SNA strategy, nor would it say when such a statement would be forthcoming. 3Com is still shipping its CS-1/SNA, a communications server based on the Series/1, which is essentially a 3270 gateway from devices on TCP/IP Ethernet LANs. *SNA Perspective* expects that 3Com will follow the market trend toward supporting IBM (Token Ring) bridging and, eventually, SNA routing, either on its CS-1 platform or, more likely, its newer Linkbuilder platform.

Proteon

Based in Westborough, Massachussetts, Proteon began as the first token ring network supplier in 1981, four years before IBM's Token-Ring introduction. It added routers in 1985 and is a leading vendor in both markets. In January 1991, Proteon introduced its new CNX RISC-based multiprotocol router family. The second release of the product, in late 1991, is scheduled to include SNA support.

Cisco Systems

The most well developed marketing presentation as well as the most ambitious plan of these vendors comes from Cisco Systems.

Growing quickly since the mid-1980s and based in Menlo Park, California, Cisco Systems first product was a local LAN bridge. Cisco now finds itself among the leaders in the LAN internetworking market.

Support for IBM's SDLC protocol was added to Cisco Systems' internetwork bridge/routers in

January 1991. This was announced as part of Cisco's five phase IBM connectivity strategy (see Table 2).

Phases I, II, and III represent current or announced products from Cisco; Phases IV and V are for future developments. Phase I consists of Cisco's current support for token ring and IBM's source route bridging (SRB), added in 1990. Phase II adds bridging support for IBM's SDLC layer-two protocol, which Cisco provides by encapsulating SDLC within IP packets. Phase III provides for support of the IEEE 802.1d source route transparent (SRT) bridging standard, which will allow a standard means to bridge Ethernet and Token Ring LANs. Phase III also includes interface to the IBM LAN Network Manager, IBM's product for managing LANs and interfacing between LANs and NetView, for management of SRB and SRT bridges.

Phase IV adds SNA routing, with the product to be jointly developed with Brixton, and will include a more direct interface to NetView. Finally, in Phase V, Cisco will support physical unit (PU) type 2.1 on its multiprotocol routers in order to provide APPC support.

In January 1991, stating, "SNA routing is the cornerstone of Cisco's IBM strategy," Cisco unveiled this strategy for extending its internetworking to IBM-based environments. At the same time, it announced a joint development agreement with Brixton Systems, a small operation in Cambridge, Massachusetts. Brixton has a product which allows Sun workstations to run a scaled-down version of the physical unit (PU) type 4, which is implemented by the network control program (NCP) on IBM's 37x5 communications controllers. At this time, Brixton's does not promote its product as supporting intermediate node routing nor transmission groups; however, it says that the code

Cisco Systems Five Phase IBM Connectivity Strategy

Phase	Local Area Network (Token-Ring/Ethernet)	Wide Area Network	Network Management
Phase I	4 Mbps & 16/4 Mbps Token-Ring; Routing and Source Routing bridging (SRB)	Routing and Remote SRB	SNMP
Phase II		Synchronous Data Link Control (SDLC) transport	
Phase III	Source Route Transparent (SRT) bridgind		LAN Network Manager support
Phase IV	Systems Network Architecture (SNA) routing		NetView support
Phase V	SNA Advanced Peer-to-Peer Communications (APPC)		

Source: Cisco Systems

Table 2

licensed by Cisco does include these capabilities. Because of its focus on the Unix market which is primarily Ethernet, Brixton does not have Token Ring support which Cisco would need to add or develop jointly with Brixton. *SNA Perspective* understands that part of the Brixton code will be ported to Cisco routers, but that a Sun workstation with Brixton code on an Ethernet network will be required for each Cisco bridge/router providing SNA routing.

IBM Corporation

From public and private statements by IBM representatives, including Ellen Hancock, IBM seems willing to tolerate multiprotocol routers supporting SNA as long as they don't try to move in on the SNA network itself. That is, they can exist on the periphery but not integrally. As another indication of this preference, even with the March 12 APPN announcement, IBM released the end node specifications, but not for the network node, at least not yet.

SNA Perspective has seen indications that IBM is close to selecting a third party to supply a router to support the RS/6000, which exists for the most part outside SNA networks. At press time, our best guess for the likely supplier is Cisco Systems, which already has an agreement with IBM regarding token ring technology.

Conclusions

Most internetworking vendors have announced an SNA networking strategy, if not products, and *SNA Perspective* believes that those who have been silent thus far will make their intentions known shortly. All these vendors are likely to discover, as other "plug-compatible" suppliers have experienced over time, that though the potential rewards are large, it is risky to eat at IBM's table. Also, the difference in the selling process for the IBM arena is common stumbling block for new entrants.

SNA Perspective readers who are considering SNA-capable bridges and multiprotocol routers would be wise to proceed with caution, as they come from companies unfamiliar with the SNA market. At this early phase, however, larger customers could influence these companies to provide more valuable products by becoming involved in the product development cycle.

Cisco Systems has the most ambitious SNA routing strategy announced, but *SNA Perspective* has its concerns whether even this internetworking star might have bitten off more than it can chew. It appears extremely unlikely that Cisco could ship a PU 4 emulator with even minimal functionality, within eighteen months to two years, given the nacent stage of its internal SNA development team, even if a large part of the effort is "only" porting Brixton's code to its platform.

There are certainly many benefits to including SNA protocols on multiprotocol routers, particularly a more fluid routing architecture and topology. The second part of this two-part series, in next month's issue, will examine the technical issues involved in bridging and routing SNA protocols in a multiprotocol environment. ■

Architect's Corner

Routes But No Tunnels

by Dr. John R. Pickens

Reference: IBM Advanced Peer-to-Peer Networking Architecture and Product Family Overview, Programming Announcement 291-079, March 5, 1991.

Finally. APPN is elevated to the status of "architecture."

Predictions True—So Far

Looking back on my predictions in the January 1989 inaugural Architect's Corner column I might lay claim to prophetic insight—but for one minor detail. I predicted then that IBM would publish end node APPN protocols. True enough—APPN end node is published. I also predicted that IBM would not publish network node protocols. Also true—network node protocols did not get published. However, hints are leaking that IBM does plan to publish the network node protocols also—some day. So I may eat my words on this one.

[Editor's comment: In response to our gentle editorial persuasion, Dr. Pickens has agreed to literally eat his words—a seasoned archival copy of the issue containing the inaugural column—should APPN network node be published. Please stay tuned for this event. Dining "protocol" and condiments to be supplied by CSI.]

So, the dynamic peer protocol cat is out of the bag. The first open element of modern SNA dynamic routing services—SAA APPN end node—is published.

Big News

In the recent APPN announcements, I noted a few especially interesting items:

- OS/2 versions of both APPN end node and APPN network node (Programming Announcement 291-080)
- APPN network node in the 3174 (191-018)
- New development infrastructure—a new version of the Teleprocessing Network Simulator (TPNS) to simulate APPN network node functionality for testing APPN end node systems (291-082)
- The statements of four vendors planning to implement APPN end node (one an OEM supplier) (291-079).

All in all, this announcement is big news to the SNA technical community. I would venture that, at the current announced level of APPC/APPN functionality, SNA is maintaining pace with—and in some areas, such as security, even advancing ahead of—TCP/IP and OSI peer networking architectures. This APPN capability lays the foundation for future work toward convergence of OSI and SNA (and possibly TCP) at both the applications layer and the transport/path control layer.

Still Missing

Just as notable is the continued lack of APPN support in the 3745. The 3174 becomes the first implementation of APPN in a true communications processor platform, beating the 3745 out of the gate.

But, not to dampen the enthusiastic reception of this event, some pieces are still missing—and they are missing from the core of APPN network node (NN) function (perhaps a factor in its not-yet-published status). Close study of the 3174's APPN feature in the announcement best highlights one missing piece, which is simply stated thus:

> "It is not possible to route 3270 (and other non-LU 6.2 datastream) across APPN backbones."

One of the 3174's strengths, but highlighting the overall APPN architectural weakness, is its ability to run the so-called "SNA gateway" feature and the APPN feature concurrently on subarea boundary links. APPC traffic can of course be routed by the

APPN feature (APPN NN protocols across token ring; LEN protocols on SDLC link). But 3270 traffic cannot be passed through APPC/APPN on any link. Instead it continues to operate on its own SDLC poll address. Downstream physical units (PUs) continue to be mapped to separate multidrop poll addresses. Thus the 3270 datastream continues to be constrained to the routing capabilities of subarea networking. Any 3174 through which 3270 datastreams pass must be directly boundary-connected to a subarea node.

3270 Support Needed

So what is the architectural solution to this problem? There are two ways that 3270 (and other non-6.2 logical unit (LU) types) could be routed across APPN backbones:

- 3270/LU 6.2 encapsulation. I've described this alternative in the past—the 3270 datastream is carried unmodified within APPC sessions. This technique is analogous to the 5250 PC Support Program which encapsulates 5250 datastream within APPC (see Figure 1).
- LU 0123 tunneling. Present a PU 4 interface downstream to cluster controllers and PU 2 end nodes, and a PU 2 or subarea interface upstream to communications controllers. Encapsulate the non-LU 6.2s in APPC sessions. This alternative is probably new to most readers (see Figure 2).

The first option is the preferred SAA solution—with 3270 as an SAA datastream and LU 2 not an SAA datastream, an LU 6.2 version is mandatory. Why hasn't IBM delivered 3270/LU 6.2 to date? Probably because of the diversity of platforms on which it must be implemented for completeness—multiple mainframes and mainframe subsystem environments, cluster controllers, AS/400, PS/2, RS/6000, etc. I doubt it's because of any architectural reason.

The second option, tunneling, is the most pragmatic transition solution. APPN routing services can be used to advantage for existing devices—cluster controllers, banking terminals, retail store controllers, etc.

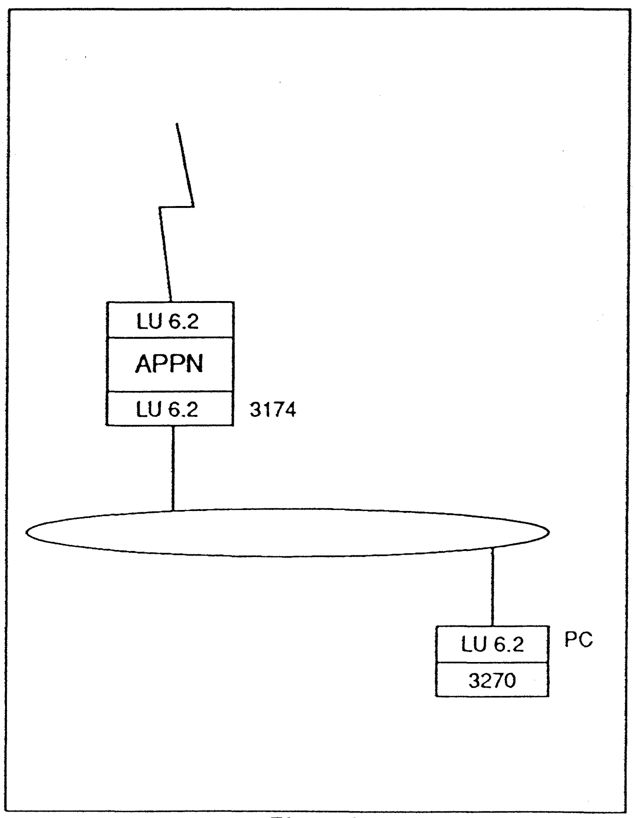

Figure 1

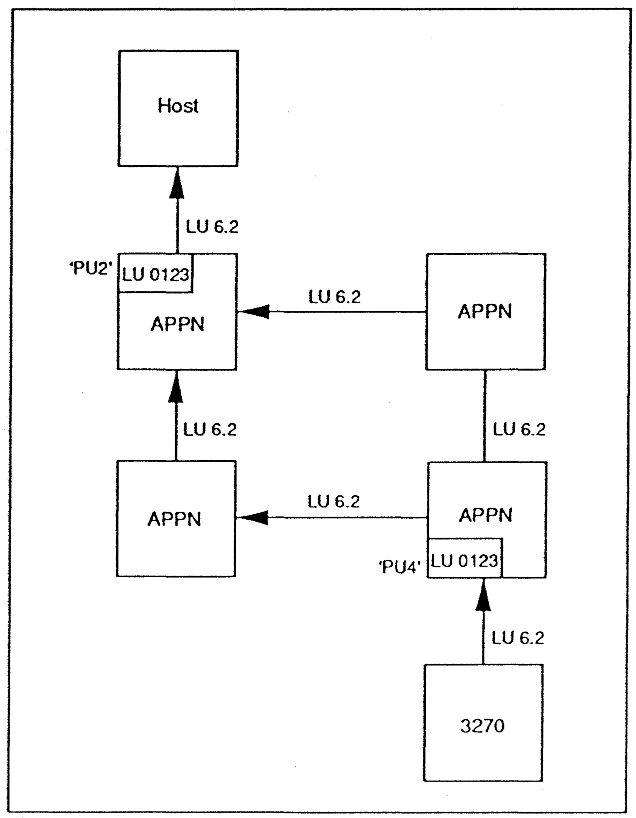

Figure 2

Which technique will appear first? And when?

Over a year and a half ago, at an IBM sponsored analysts briefing, I was told in a lunch conversation of a planned enhancement to APPN—LU 0123 tunneling. APPN network nodes would be enhanced to provide routing services for existing non-LU 6.2 devices. I was expecting this feature in the recent announcements. But, anticipating the prospect of having to eat my words, I won't predict this one. ■

• Subscription: U.S. - one year $350, two years $520. International - one year $385, two years $590 • SNA Perspective is published monthly by CSI • 5201 Great America Parkway • Santa Clara, CA 95054 • Telephone • (408) 562-6031 • Fax (408) 562-5745 • Managing Editor: Louise Herndon Wells • Associate Editors: Dr. John R. Pickens, Vince Busam • Contributors: Philippe Byrnes, Anura Gurugé • The information and opinions within are based on the best information available, but completeness and accuracy cannot be guaranteed.