SNA Perspective

Volume 12, Number 8
August, 1991
ISSN 0270-7284

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

SNA End Systems Services: 3270 and APPC

The advent of the APPC architecture in the early 1980s coincided almost perfectly with the introduction of the IBM personal computer. At the time, it appeared to be not merely a coincidence; it seemed that the distributed processing capability within SNA that was heralded by APPC could now reach all the way to the desktop.

But the revolution has not happened. What happened instead must be described as an evolution. This article examines the reasons for continued dominance of 3270 and the relative advantages of 3270 and APPC for end-system services.

(continued on page 2)

Get The Picture? Managing SNA With Graphics

As SNA users have learned over the years, management of an SNA network is not easy. The sometimes cryptic and always tedious textual messages required by VTAM or NetView are hard to learn. Further, with the naming conventions used most commonly, it is not intuitively obvious where failing components are located in the network.

Today, users demand graphics or pictures to manage their networks. Pictures make a complex network easier to comprehend. With current "point-and-shoot" technology, operators can spot the problem device, see its context in the network, point to the component with a mouse, and enter commands without having to know the device name or address.

This article explores two graphics interfaces to IBM's NetView, in addition to a text mode interface for operating SNA networks and attached non-SNA components.

(continued on page 6)

In This Issue:

(continued from page 1)

APPC and 3270 both provide end-system services within SNA. By end systems, we mean the software infrastructure that provides services to the end user in the form of access to the SNA network. In this context, intermediate systems SNA takes the form of the SNA subarea and APPN. IBM has made considerable progress over the past three years in upgrading intermediate systems SNA without adversely affecting end users. However, different organizations are at different points along the evolutionary path of end-systems SNA. Despite the considerable energy expended over the past ten years discussing APPC and preparing for its eventual dominance in the world of SNA, well over 90 percent of all SNA traffic today is still 3270.

Reasons for Continued Dominance of 3270

The major cause of 3270's extended life is that there has been no compelling reason to change. Users have been able to take advantage of advancing technology such as PCs and LANs while retaining their 3270 host access. Likewise, most of the host services that users need are still accessible through 3270.

Coax Boards

Soon after the IBM PC arrived on the scene, users could place a relatively inexpensive coax emulator board in it and replace a 3270 terminal with a machine that could do much more than just terminal emulation. Many companies did quite well manufacturing these coax emulator boards, including IBM itself. Placing a coax board in a PC is often cited as the first step in the migration away from traditional 3270 *terminal* connectivity, yet it also allowed 3270 host applications to continue to flourish.

HLLAPI Extends 3270 Life

IBM extended the capabilities of the 3270 end systems by defining a programmatic interface to 3270 called high-level language application program interface (HLLAPI). This interface was originally intended to serve as a rudimentary method for DOS applications to interact with the 3270 presentation space being maintained in one of the SNA sessions on a 3270/PC. However, this interface has further extended the life of 3270 by permitting the user to access host applications with the tools available in graphical user interfaces such as icons, radio buttons, and dialog boxes. No longer is the 3270 user constrained by the rigid, field-oriented structure of the 3270 data stream.

Host Access Via LAN

After coax adapters, the next step along the evolutionary path of SNA end systems is connecting PCs to LANs and using a LAN-based SNA gateway to access the host. In addition to replacing the 3270 terminal, the 3x74 cluster controller and the requisite coaxial cabling can also be replaced with either Ethernet or Token Ring LANs.

Many vendors who offer 3270 terminal emulation via a LAN-attached PC also offer logical unit (LU) 6.2 services. This is the first configuration along the path that provides users with the choice of different methods to access host services, and users must decide which architecture—3270 or APPC—best fits their computing needs. Even though APPC has been a strong selling point, the majority of the users of LAN-attached PCs still use 3270 terminal emulation and not the APPC services.

APPC—The Latest Step

The latest step in the evolution of applications for SNA end systems builds upon APPC services. Even though APPC is a distributed transaction processing architecture, it does fit other computing paradigms including:

- the simple local/remote interprocess communication model
- the client/server model
- the distributed (transaction) processing model

While the APPC architecture and its manifestation in SNA, LU 6.2, provide adequate services to support all of these computing models, there is a dearth of applications that use APPC in a form other than either the interprocess communication (IPC) model or a relatively simple form of client/server communications.

Table 1 lists some of the APPC applications that are available on the market today from both IBM as well as third-party vendors. Many of today's applications that use APPC services use it as a simple IPC protocol. Client/server products using APPC are just now showing up and completely distributed transaction processing products do not yet exist.

Computing Models and SNA's Logical Units

Master/slave computing is the time-honored form of computing where the knowledge and power are on a central host computer and the remote device is nothing more than a user-interface tool. 3270 terminal emulation exemplifies master/slave computing. The underlying LU services for 3270 display emulation are LU 2 and, as we shall see below, the services provided by LU 2 readily accommodate this master/slave relationship.

As with all SNA LU types, a session is actually composed of two half sessions—one on each end of the dialogue. The side that creates the session is known as the primary logical unit (PLU) and the side that receives the session creation is called the secondary logical unit (SLU). Further, for LU 2, the two half sessions are not equal in capability—the PLU is the master and the SLU is the slave. An example of a session characteristic that manifests the power of the PLU over the SLU is the ability to control the direction of the flow of data. The PLU can send data directly to the SLU while the SLU must ask the PLU for permission to send data. As one might expect, the PLU for 3270 communications is on the host and the SLU is on either the cluster controller or its LAN gateway counterpart.

APPC Applications and their Computing Model

Product	Company	Computing Model
XCOM 6.2	Spectrum Concepts	IPC
SQL*Net	Oracle	client/server
SQL Network	Gupta Technologies	client/server
PC/SQL-Link	Micro Decisionware	IPC
Remote Data Services	IBM	client/server

Table 1

IND$FILE

This master/slave relationship manifests itself in some unwanted ways with certain applications such as file transfer. On workstations that contain a file system in addition to 3270 terminal emulation (such as a PC), the capability exists for transferring files between the host and the workstation. The traditional file transfer program used with LU 2 is known as the IND$FILE program. It essentially looks to the host like an operator hitting the Enter key for a screen's worth of file data to be transferred between the two systems. This is because the file transfer was built on top of the existing LU 2 services which were designed specifically for communications between a host application and a 3270 display station.

With IND$FILE, when data is to be sent from the PC to the host, the PC's file transfer program places a screen's worth of data along with some control information into the screen buffer and simulates an operator hitting a host Attention key. This causes the host to issue a 3270 read command to read the contents of the screen buffer, thus transferring approximately a screen's worth of data from the PC to the host. Whenever the host sends data to the PC, the host fills up the screen's buffer and then unlocks the 3270 keyboard.

In both directions of data transfer, there is an overall unidirectional flow of data but a bidirectional flow of control information at the 3270 application level. The effect of this control overhead has been minimized in recent years with the introduction of 32 kbyte buffers for data transfer, but the end-to-end lock-step operation continues. When APPC is used for file transfer operations, however, there is no need for this extra bidirectional, application-level handshaking. APPC file transfer programs rely on the underlying LU 6.2 session capabilities if end-to-end synchronization is required, thus increasing the efficiency of the operation.

Strengths of 3270

LU 6.2's strengths do not spell the end for 3270. There are many applications for which the unbalanced operation of LU 2 is a natural fit. For instance, stock quotes are essentially unidirectional messages from a central site (e.g., a host computer) to an output-only terminal (e.g., a 3270 without a keyboard). The PLU's ability to blindly send data to the SLU without first asking for permission means that the host can continually update the 3270 presentation space in rapid succession.

APPC Symmetry

Even though LU 2 and LU 6.2 sessions have PLUs and SLUs, the capabilities of each LU 6.2 half session are identical. Even though one of the LU 6.2 half sessions might be on a mainframe and the other half session on a PC, the power of each half session is the same when viewed from the SNA perspective. This symmetry is necessary for the distributed processing requirements found in both client/server computing as well as distributed transaction processing. Figure 1 illustrates the differences in symmetry between LU 2 and LU 6.2.

APPC can be used very effectively for interprocess (program-to-program) communication. While it is true that there are several other communication protocols that would also suffice for simple remote interprocess communication, APPC is more extensive than others and reaches into more environments. Named pipes, NetBIOS sessions, and TCP connections all offer reliable, bidirectional data transfer, but APPC also scales up to reach mainframe systems. Because most computer vendors need to connect to IBM systems in order to sell into the business marketplace, they must support APPC on their systems. Thus APPC is one of the most widely available protocols for interprocess communication.

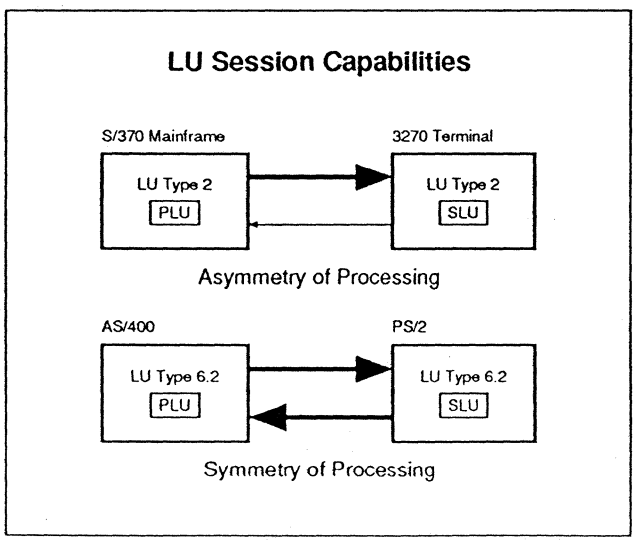

Figure 1

APPC also specifies a mechanism to generate a process at a remote system. This capability of creating a process at a remote site is a powerful feature of APPC and, when used in concert with APPC's security mechanisms, allows for secure, unattended operation of server-oriented systems. The error reporting and process synchronization facilities built into APPC are beyond the services offered by any other remote interprocess communication mechanism.

Client/Server

Client/server is a loaded term in today's potpourri of industry buzzwords. Industry experts can't always agree on what is required to make an environment conform to client/server computing and what exactly distinguishes client/server computing from the more general term cooperative computing. One thing the experts do agree on, however, is that APPC plays a very important role in client/server computing.

Client/server computing in some ways looks very much like master/slave computing. In its degenerate form, the client is nothing more than a user-interface tool and the server is a host computer. However, client/server computing has the potential to be much more. One of the best examples of client/server computing is database record retrieval. The client in this case is a program that interacts with the user and formulates a database query without any specific knowledge of where the data is stored. The server could very possibly be a LAN-based database management system (DBMS) running on a file server. If the database is large and contains diverse information, it is possible that the LAN-based DBMS contains only a portion of the database. If so, the DBMS on the file server must forward the request on to yet another DBMS somewhere else in the network. More often than not, the ultimate repository of enterprise database information is on a mainframe.

When the LAN-based DBMS forwards the request to the mainframe DBMS, the LAN DBMS performs the client role in this phase of the transaction. Because of the dual role of the LAN-based file server (i.e., it is both server and client), the underlying communication protocol must be symmetrical. For instance, if LU 2 were to be used as the underlying transport for client/server computing, any node that takes on the dual role would have to implement two variants of LU 2 (i.e., both the PLU and the SLU) rather than operating with a single LU 6.2 implementation.

The effect of asymmetrical versus symmetrical communications protocols on distributed client/server computing is shown in Figures 2 and 3. Figure 2 illustrates the LU 2 implementation requirements if that protocol were to be used in support of client/server computing. In Figure 2, the intermediate node, labelled "LAN-based DBMS," would have to implement two vastly different forms of LU 2: its *slave* form in the SLU and its *master* form in the PLU. On the contrary, LU 6.2 provides symmetrical capabilities and the intermediate node needs to have only one LU implementation (see Figure 3).

The dynamics of the computer industry will favor client/server configurations for quite some time. IBM is motivated to provide mainframe-based solutions and has made advances in repositioning the mainframe as a high-availability data repository (see *SNA Perspective*, October, 1990). However, with the improved price/performance of networked PCs, other vendors are left to whittle away at IBM by providing services on downsized platforms.

With the recent explosion in graphical user interface (GUI) technology, there is a strong desire to move the front end of mainframe-based applications onto platforms with user-friendly interfaces such as PCs. The resulting heterogeneous environments further fuel the need for client/server computing.

While LAN-based solutions, such as the DBMS example cited above, might serve a division or a workgroup quite well, they cannot rival the storage and power of a mainframe-based system. But exercising complete control over an entire database at a central site yields it own set of problems with access time and ownership of data. Customers will probably want the best of both these worlds and client/server computing holds great promise for meeting their needs.

Another phenomenon holds hope for client/server computing and thus for making greater use of APPC services. There are a variety of PC-based applications that provide the user with greater flexibility and control than 3270-oriented applications—witness the success of spreadsheet

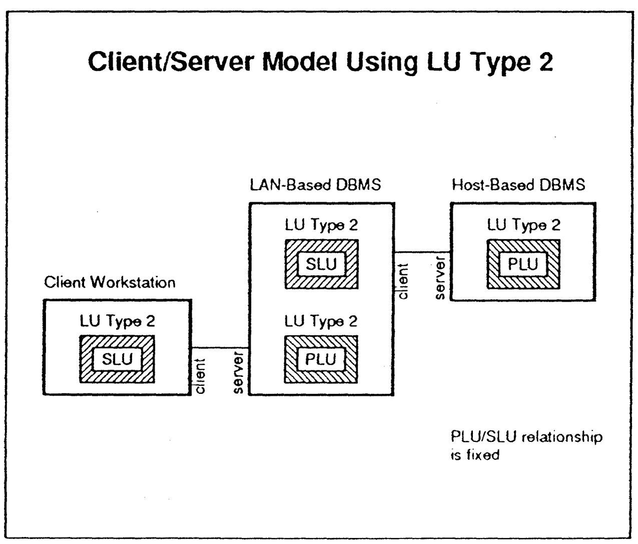

Figure 2

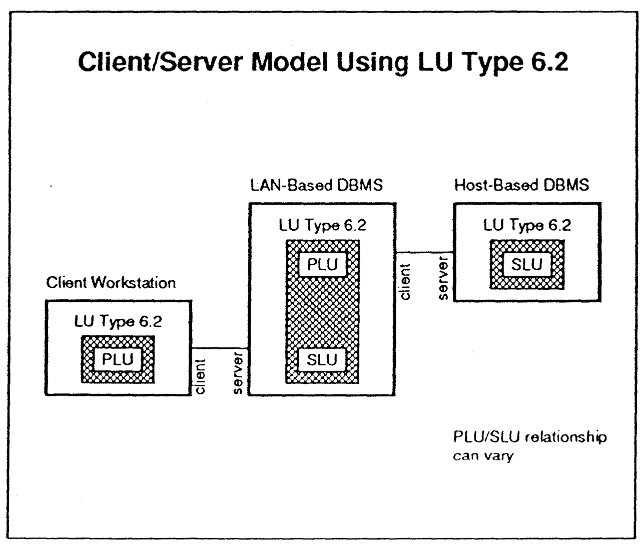

Figure 3

programs such as Lotus 1-2-3 and Microsoft Excel. In a corporate environment, the data needed as input to these programs is often contained in a mainframe database somewhere within the company network. What these PC-based applications need is a simple, programmatic interface to get at this corporate data. This is one of the strongest features of APPC—its interface is at least regular across different client platforms even if it is not standard. And CPI-C makes the interface standard across client platforms.

Distributed Transaction Processing

The need for balanced communications protocols becomes even more acute with distributed transaction processing. While client/server computing implies a rigid hierarchy of computing elements, distributed transaction processing implies equality of all cooperating nodes. There is always some degree of network transparency associated with distributed processing—a feature that is lacking in client/server computing. (The client always knows the location of its server on the network.)

Distributed computing not only requires protocols with symmetrical capabilities but also requires services beyond those found in the older LUs. Some of the network capabilities necessary in a totally distributed environment are security, synchronization, sophisticated error recovery, and support for heterogeneous nodes. All of these features are found in LU 6.2.

Perhaps one of the most eagerly anticipated applications in the computer industry is a reliable, totally distributed DBMS. This product would permit the user to perceive a database as a single monolithic entity much like today's mainframe-based database management systems, but would allow different pieces of the database to reside on geographically dispersed computers. Support for a product like this is the ultimate test of a distributed processing technology. Only today's APPC and tomorrow's OSI/TP can deliver this functionality. ■

(continued from page 1)

Graphics Network Management Tools From Other Vendors

Graphics interfaces for non-SNA network management products have existed for some time. Most vendors of communications products have built very sophisticated graphics packages, primarily on UNIX workstations. These workstations typically manage networks of T-1 multiplexers, modems, DSUs, common carrier services, and other network communications components.

Most of these tools, however, manage only parts of the network, not the global network that NetView attempts to manage. While most of these vendors attempt to manage as many component types as they can, their systems lack the knowledge of SNA components, SNA naming conventions, and SNA alerts flowing on SNA sessions.

IBM SNA Graphic Products

IBM, in keeping with its tradition of lagging behind the market in network products, finally has some products in this area. In fact, IBM currently offers two graphics workstation products that work in conjunction with NetView:

- NetCenter
- GMF

These two products are quite different from each other and choosing one over the other depends largely on the makeup of the network and future network plans.

NetCenter Graphics Network Monitor

NetCenter was built by US West several years ago. US West marketed this product on its own for a while, then sold exclusive marketing and product rights to IBM.

NetCenter is a host-based distributed application. It uses the host to store network definitions and views and automatically sends new and changed views to operators' workstations whenever those operators log on to NetView. NetCenter incorporates the concept of an administrator station and a user station. Administrators can build new views and forward them to the host; users can share the views that are stored on the host and downloaded to them.

Advantages

NetCenter is a thorough implementation of a graphics network monitor (especially considering the time frame of its development). Its graphics are excellent. It offers good icons for the various network components, components that change colors based upon status, and pull-down menus for more in-depth information and commands. Pictures can be built either with a wide variety of maps that are provided or via pictures scanned in to the GEM environment.

NetCenter has one major advantage. It manages both SNA *and* non-SNA components, and both of these can be shown on the same screen using the same point-and-shoot technique. This means that the SNA network, including applications, hosts, terminals, controllers, and links, can be shown on the same view as non-SNA components. Operators don't usually need to know the difference.

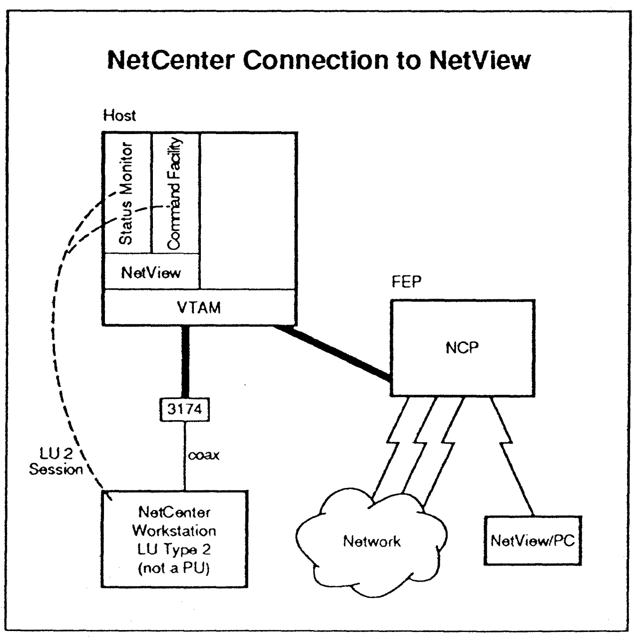

Figure 4

Another advantage of NetCenter is that point-and-shoot commands can be accessed from pull-down menus above the graphic status display. These commands can be used to display more information about components as well as to enter commonly used commands without typing the component name—a real productivity enhancement.

Disadvantages

NetCenter does not take advantage of newer operating systems. It runs on a DOS 3.3 platform using GEM graphics and Attachmate terminal emulation software. It uses an LU 2 coax connection through a 327X control unit (see Figure 4), which is not as sophisticated or efficient a method for program-to-program communications these days as APPC.

NetCenter requires quite a lot of customization, especially for non-SNA components. SNA component configuration data can be obtained from the NCP gen; drawings (called *views*) can then be built with the information. Non-SNA components must be manually configured from start to finish.

Before NetView Version 2 Release 1, NetCenter was marketed as a separate product from NetView. As such, NetCenter could be installed on top of NetView

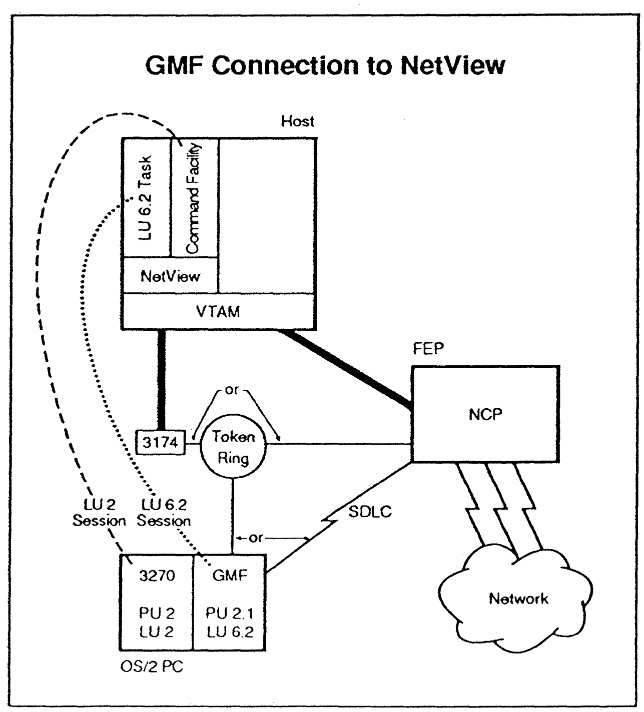

Figure 5

Version 1 Release 3 by purchasing both the host product and the workstation components (administrator and user). When IBM announced NetView Version 2, the NetCenter host portion of the code was bundled with NetView. However, there is still a $3,000 per workstation charge for the workstation portion.

NetView Graphics Monitor Facility

The Graphics Monitor Facility (GMF), the second offering from IBM, comes from a completely different background. It was built by IBM on an OS/2 platform using Presentation Manager graphics. GMF, in the current release, comes packaged with NetView Version 2 Release 1.

GMF is IBM's strategic product for graphically monitoring networks. That's the good news. The bad news is that GMF is not yet as complete as NetCenter; hence users must decide which to implement. More on that later.

Advantages

Some of the things that GMF does well are:

- graphics handling
- SAA conformance
- LU 6.2 connection to NetView
- easy-to-build views

Graphics handling—Because of the OS/2 Presentation Manger platform, GMF has an excellent graphics handling capability. This means that windows can be individually sized, overlapped, etc. Since each window is an OS/2 task, many windows can be open at once.

SAA conformance—GMF also conforms to the SAA presentation services scheme. This could be important in the future, as more applications conform to this structure, so that users don't need to learn their way around using new screen, mouse, and keyboard rules.

LU 6.2 Connection to NetView—GMF uses LU 6.2, a program-to-program communication vehicle to talk to NetView (see Figure 5 on page 7). LU 6.2, or APPC, is a more efficient protocol to use when programs talk to each other. The LU 2 (3270) protocols used by products like NetCenter make use of a protocol designed to be used by humans (interactively) rather than programs.

Easy-to-Build Views—GMF views are about as easy to build as is possible. The SNA configurations are extracted from the user's VTAM definitions and drawn for you (see Figure 6). Users can select prebuilt maps, or portions of maps, and then superimpose the components of a network on the map. Then with the mouse they can point and drag the component to the desired location on the view. This makes building views easy; it even "feels" right when it's done this way.

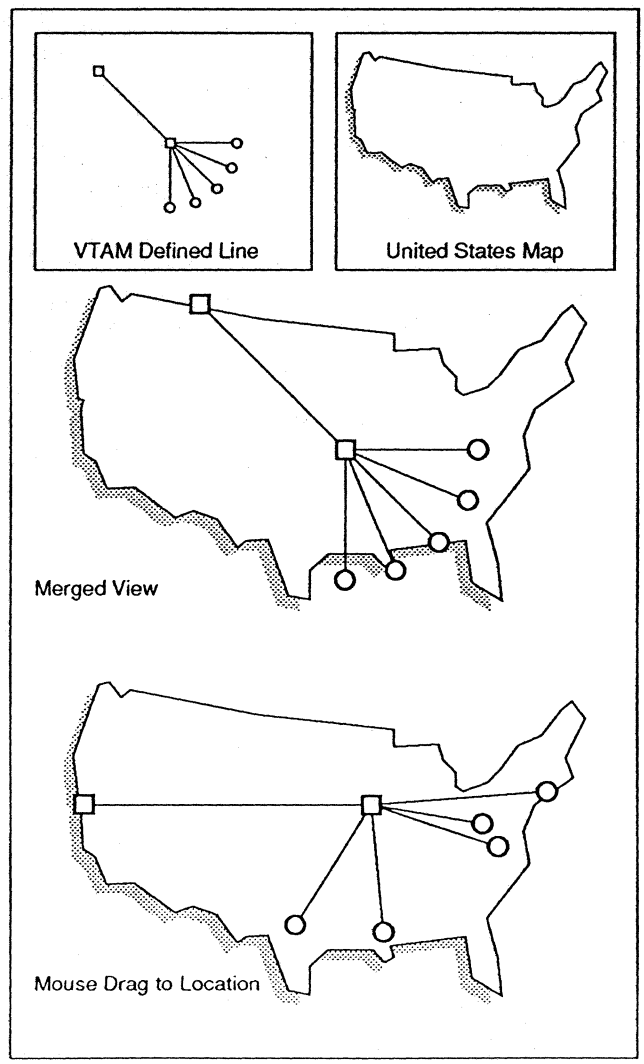

Figure 6

Disadvantages

On the other hand, GMF has some serious drawbacks at the moment, although IBM has announced that it will fix them. Existing problems are that GMF:

- is SNA only
- has no command interface
- has limited icon flexibility

SNA only—Unlike NetCenter, GMF does not provide a way to configure and monitor non-SNA components of the network.

No command interface—GMF provides status display only. Again, unlike NetCenter, there is no command interface to NetView. Users can recognize failures in the network by color changes on the views, but they must go to a NetView text screen to enter any commands, such as VARY ACTIVE or VARY INACTIVE. This is much less efficient than command entry with a mouse, since the names must be typed in.

Limited icon flexibility—GMF icons, like NetCenter icons, are fixed in shape. Users can see small pictures of the component type if they zoom in on a view, but the general icons are fixed geometric shapes like five-pointed stars, triangles, squares, and circles. The colors are good but the shapes are quite basic.

The Future of NetCenter and GMF

To help users decide which graphics monitor to install, IBM has stated that GMF is its strategic platform. IBM also says that the functions of NetCenter that are not now included in GMF will be merged into GMF on the OS/2 platform. IBM has not said exactly when this will occur, but *SNA Perspective* believes it won't be until 1992 (after being announced some time in 1991).

What this means is that non-SNA product management will be added to the GMF platform, as well as a pull-down menu command interface. Because of the greater power of OS/2 and the LU 6.2 interface, this is a great idea but, as usual, it's taking a long time.

Where Does the IBM RS/6000 Fit In?

Based on what we have said about NetCenter and GMF, users can probably make an appropriate decision about which one to implement today, primarily based upon how much non-SNA management is required. But there is another factor that muddies the waters—the IBM RS/6000.

The RS/6000 is a much more powerful platform than a PS/2 (although an 80486 system is no slouch) and has powerful graphics capabilities. Also, the RS/6000 has AIX (IBM's UNIX) and inherent connections to TCP/IP. AIX has an SNMP network management package called XGMON, and the RS/6000 has a NetView service point application similar to NetView/PC.

As we stated in last month's article on network management and the RS/6000 (IBM Network Management: New Platforms and Directions, *SNA Perspective,* July 1991), the RS/6000 appears to be a strong candidate as a network management workstation. How does this affect the NetCenter-GMF decision? In 1991 and even 1992, using the RS/6000 as a full IBM NetView-integrated network management workstation is probably premature. But IBM has been working with other vendors in procuring UNIX-based applications and expertise. (IBM recently signed agreements with Hewlett-Packard for OpenView code and with 3Com for HLM co-development.) If more applications are added to the RS/6000, it may be a strong choice in the future.

TCP/IP and SNMP

IBM already provides a host-based adapter for conversion of SNMP frames to and from SNA NMVT frames. This allows NetView operators to manage TCP/IP networks from the NetView console. With NetCenter, graphics can be added to the SNMP management provided by MVS-TCP and VM-TCP products. But the host-based TCP products are expensive, especially if users only want the management function.

On the AIX platforms (RS/6000, PS/2, and RT/PC), IBM's XGMON does SNMP network management.

Both of these platforms provide both SNMP manager and agent functions.

OSI CMIP

In a very similar manner to MVS-TCP, IBM markets an OSI product called OSI/CS (OSI Communications Services). OSI/CS running under MVS provides connectivity to other OSI systems for data transfer. More importantly, OSI/CS provides network management manager and agent functions as well. OSI/CS, however, is also an expensive solution, since it requires IBM's NPSI product in the FEP.

Handling TCP/IP and OSI in Graphics

For both TCP/IP and OSI icons on a graphics workstation, NetCenter is currently the only IBM choice because, since it supports non-SNA devices, it supports TCP/IP SNMP and OSI CMIP from the mainframe environment. XGMON at this point seems to be an isolated product. GMF doesn't support any of these features at the moment, but should in its next release.

Summary

Future Product Capabilities

For both SNA and non-SNA network management, the future workstation is likely to assume much more responsibility in the area of user interfaces and problem analysis. Expert systems technology is likely to be used increasingly in these workstations, since the expert system will, in most cases, be doing tasks on behalf of a single user.

However, for global tasks like gathering network-wide problem, change, and configuration data, the mainframe will still reign supreme. This means that heavy cooperative processing will be the wave of the future. The workstation MIPs will be used for individual operators and the mainframe MIPs will be used for global tasks.

NetView Implications

NetView will probably remain functionally like it is today, but with better connections to sophisticated workstations and other network management systems. An example is the LU 6.2 connection that has already been announced.

One of the tasks that NetView does well, for both systems and networks, is automated operations. This is likely to either remain in the mainframe or be split between the mainframe and the workstation. Many experts feel that automated operations should be done as close to the problem source as possible, so moving automated operations further away completely into a workstation does not seem to be the right choice.

Some of the individual subtasks of NetView will slowly fade away. The following functions, in particular, seem to have short lives:

- StatMon
- Hardware monitor (NPDA) Alerts function

StatMon—StatMon, NetView's Status Monitor, is probably the most likely task to become obsolete. Since StatMon shows network-wide status in a text mode, graphical views will easily prevail.

Hardware Monitor (NPDA) Alerts Function—Not likely to go so quickly, but also suspect, is the hardware monitor function, particularly the operator screen interface. The Hardware Monitor, however, plays a big part in gathering of Alerts and provides Probable Cause and Recommended Actions for them. This is an important function and a major portion of this activity is likely to remain in the mainframe. It would make sense for the workstation to be able to ask for system-wide probable cause and recommended action text for specific alert codes only, as required by operators.

Since IBM wants to manage it all, NetView will continue to have good interfaces to non-SNA data, especially SNMP and CMIP. It will continue to manage the large amounts of data gathering on the mainframe. The trend for the foreseeable future will be to move graphic displays, expert systems analysis, and other front-end applications to smarter and smarter workstations. ■

Brixton Systems and Node Type 4 Routing Emulation

Brixton Systems, Inc., a privately held company based in Cambridge, Massachusetts and founded in early 1989, calls itself a UNIX networking software company. Brixton builds products that integrate UNIX applications and users with a variety of networking and computing systems such as IBM's SNA, internetworks, PC LANs, and X.25 networks. The company was founded primarily by people from Bolt, Beranek, and Neumann (BBN), a company known for its considerable strength in UNIX and TCP/IP but not for SNA expertise. The company currently employs about 20 people.

Products

The Brixton SNA internetworking products include cluster controller emulation, 3270 emulation, LU 0 program-to-program facilities, Token Ring LAN support, network management interface, SunNet Manager connectivity to NetView, and communications controller (Node Type 4) emulation for TCP/IP over SNA.

All Brixton products run on Sun Microsystems workstations. Since the code is written in C, it can be ported by OEMs to other platforms, as Cisco is planning for its multiprotocol router platform.

The Brixton SNA products and their relationship to each other are shown in Figure 7. Other Brixton SNA products use the BrxPU4 SNA Server to:

- Route TCP/IP traffic directly across SNA networks (BrxSNA/IP Router)
- Emulate SNA 3270 devices (Brx3270) and 3770/RJE devices (Brx3770)
- Manage locally attached SNA peripheral devices (BrxSPD)

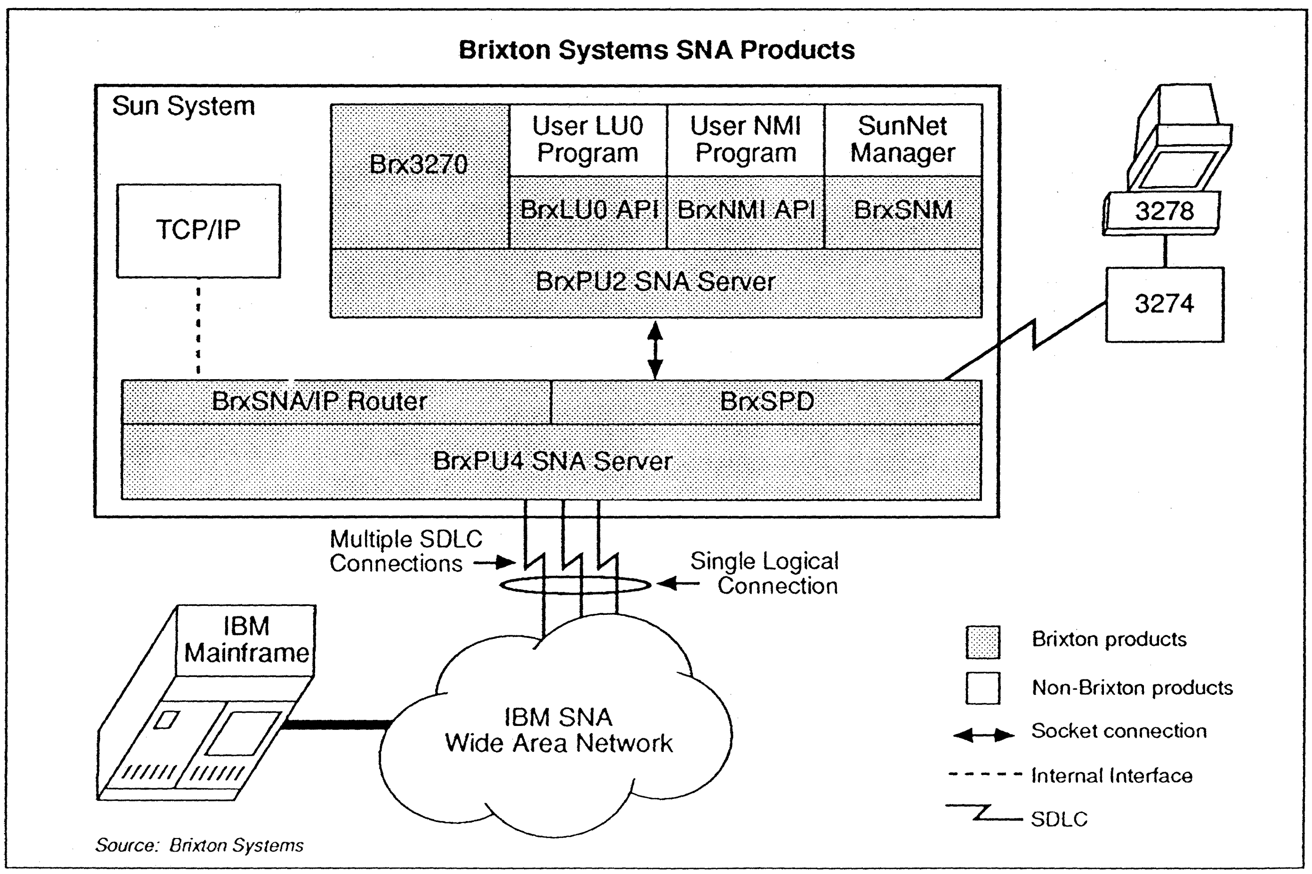

Figure 7

- Coordinate UNIX network management tools and IBM's NetView (BrxSNM for SunNet Manager and BrxNMI for other network management systems)
- Facilitate development of user SNA programs (BrxLU0)

BrxPU2 SNA Server—Emulates peripheral node common services for LUs. Connects directly to SNA network via SDLC lines (not shown in figure) or attaches to BrxPU4 SNA Server (via BrxSPD) to access mainframe.

BrxSPD—The SNA Peripheral Devices (SPD) interface works in conjunction with the BrxPU4 SNA Server. Where BrxPU4 SNA Server concerns itself with the connection to the SNA backbone, BrxSPD attends to the devices needing access to the backbone. It allows attachment of IBM or emulated SNA peripheral devices, such as 3174s or PCs running 3270 software. They can attach to the Sun system in two ways: directly through SDLC lines or through a TCP/IP LAN.

BrxPU4 SNA Server—Allows a Sun system to emulate an SNA routing node. It shares routing data with other SNA routing nodes, establishes routes through the SNA network, and maintains multilink transmission groups into the SNA network.

BrxSNA/IP Router—Designed primarily to interconnect TCP/IP-based LANs across an SNA network, it also allows connection of devices on TCP/IP LANs to resources on SNA hosts. It is used in conjunction with BrxPU4 SNA Server. BrxSNA/IP Router appears as an IP router to the TCP/IP network, while its partner BrxPU4 SNA Server appears as a PU 4 node to the SNA network. No communication with a host is required for tunneling IP traffic across SNA. The TCP/IP traffic can be assigned to different priority levels in the SNA network. The BrxSNA/IP Router does not require any specific release of VTAM or NCP in the SNA network.

Brx3270—A basic 3270 emulator. It currently emulates only 3274 C series cluster controllers and 3278 Model 2 terminals, and Brixton reports that it is expanding the range of emulations. Brx3270 cannot stand alone; it operates in conjunction with BrxPU2 SNA Server, which performs all data and session set-up.

Brx3770—Allows users to perform Remote Job Entry (RJE) batch file transfer over SNA networks by emulating the full functionality of an IBM 3777 workstation device. Works in conjunction with BrxPU2 SNA Server.

BrxLU0 API—A series of library routines that allow the user to write programs to interact with IBM mainframe applications. Works in conjunction with BrxPU2 SNA Server.

BrxNMI API—A series of library routines that allow development of applications to interact with network management programs on the mainframe. It facilitates communication with the BrxPU2 SNA Server and its control sessions (SSCP-PU).

BrxSNM—Captures alarms generated by Sun's SunNet Manager and reports these alarms to IBM's NetView in a Network Management Vector Table (NMVT) format as SNA Alerts. Can also send a NetView operator's commands to SunNet Manager. It operates in conjunction with BrxPU2 SNA Server.

More on Type 4 Router

In combination with either BrxSPD or BRXSNA/IP Router, BrxPU4 SNA Server provides the following SNA support: 37xx communications controller, PU Type 4, boundary function support, SSCP-PU sessions, inbound and outbound pacing, one-stage and two-stage pacing, virtual route traffic priorities, multilink transmission groups. The following SNA request/response units (RUs) and boundary functions are currently supported:

BrxPU4 SNA RUs

ACTPU, DACTPU, ACTLU, DACTLU, BIND, UNBIND, ACTLINK, DACTLINK, CONTACT, DISCONTACT, CONTACTED, ACTTRACE, DACTTRACE, VR_INOP, NC_ACTVR, NC_DACTVR, NC_ER_ACT, NC-ER-ACT_REPLY, SETCV, REQMS, RECFMS, ROUTE_TEST, SESSEND, XID, IPR

SNA Boundary Function

SNA commands: ACTPU, DACTPU, ACTLU, DACTLU, BIND, UNBIND, SDT, CLEAR, RQR, STSN, CANCEL, CHASE, LUSTAT, SHUTD, SHUTC, BID, BIS, SIGNAL, QC, QEC, RELQ, RSHUTD, RTR, SBI, FM data, REQMS, RECFMS, NMVT, NOTIFY

SSCP-PU, SSCP, LU-LU sessions

Segmentation and Assembly

The BrxPU4 SNA Server is configured with VTAM like a generic NCP, with a seven character NCP name. The object code resulting from a system generation (sysgen) is not downloaded to the Sun workstations as it is to other IBM communications controllers; it creates its own object under UNIX.

The communications boards currently used by Brixton maintain up to 8 SDLC lines each running at 64 kbps, which can be combined into transmission groups; multiple boards can be configured per Sun server. Brixton is working to integrate higher speed boards. BrxPU4s can connect to other BrxPU4s or with IBM 3745s and can pass information between 3745s.

Communication with NetView is provided across an SSCP-PU2 session (in a manner comparable to NetView/PC) or the user can write an application which can utilize the SSCP-PU4 session. To control the resources maintained by the BrxPU4 SNA Server, NetView uses the SSCP-PU4 session as a matter of course.

More than one BrxPU4 SNA Servers can be interconnected. A BrxPU4 SNA Server can connect with more than one upstream FEP. The product can route information around a failed link in the SNA backbone.

What Was Left Out

Brixton explains that the BrxPU4 SNA Server is designed to support a certain subset of the basic routing functions of NCP. It does not, therefore, support a particular release level of NCP.

The BrxPU4 SNA Server does not support extended partitions, NPSI, NTRI, XI, SNI, or NTO (BSC or asynchronous communications).

Neither SNA peer communications (APPN or Node Type 2.1) nor Token Ring LANs are currently supported, though Brixton says these are under development.

Several of the dynamic features available in recent releases of NCP, such as dynamic path table changes, are not included in BrxPU4 SNA Server. The product would have to be generated to incorporate changes, which may interrupt SNA paths passing through this PU 4.

It supports point-to-point links on PU4 but not switched subarea links. "Round robin" polling is used to communicate with peripheral nodes—the link scheduler does not implement a special service to allow priority for any PU on multipoint lines.

Current Brixton pricing is as follows:

BrxPU4 SNA Server	$3,000.00
BrxSNA/IP Router	3,000.00 *
BrxSPD	1,500.00 *
BrxPU2 SNA Server	500.00
Brx3270	2,000.00 †
Brx3770	2,000.00 †
BrxLU0 API	2,000.00 †
BrxNMI API	3,000.00 †
BrxSNM	3,000.00 †
SBus board 8 lines	1,900.00

(Sun workstation pricing depends on configuration.)

* requires BrxPU4 SNA Server
† requires BrxPU2 SNA Server

Market

Before the announcement of its agreement with Cisco Systems heightened market awareness of its SNA internetworking software, Brixton was more known for its network management interface products. The company works with Cabletron and Bridgeway, among others, to provide its support for

SNMP managers to communicate with NetView. *SNA Perspective* understands that, according to its agreement with Cisco, Brixton has certain constraints in providing its Node Type 4 code to other companies in the multiprotocol router market.

Brixton does not claim that its Type 4 node emulation can replace 3745 controllers. Instead, the company sees it as a useful product where a site needs only the subset of Type 4 node functions that its products provide *and* there is also a need for UNIX/TCP/IP to SNA communications. Brixton believes that, in such environments where it participates as a peer in both SNA and internetwork domains, a BrxPU4 SNA Server can maintain the integrity of each network while promoting the merging of the two.

SNA Perspective believes that Brixton appears to have a good understanding of SNA and presents its products' capabilities and limitations responsibly. It has wisely chosen to implement the features its target market seemed to need rather than trying to be all things to all people. Users considering Brixton products will need to inventory their current network's feature usage carefully to see whether these products provide appropriate support for their environment.

As discussed in the April 1991 *SNA Perspective*, Cisco received considerable press for its SNA strategy, including its agreement with Brixton, which appeared to us more than a little ambitious. Since then, Cisco seems to be revising customer expectations more in line with Brixton capabilities and realistic time lines to transition Brixton products to Cisco platforms. We expect that performance may be a concern in the Cisco port, since the Brixton product will probably run as an application on top of the Cisco router operating system, which would result in two routing processes in each node. ■

Architect's Corner

Hitchhiker's Guide to the SNA/OSI Galaxy

by Dr. John R. Pickens

The recent IBM 3745–based announcements go a long way toward confirming IBM's moves in the direction of support of multivendor-based and standards-based architectures. These announcements include:

- Ethernet
- Frame relay
- IP routing (but hardly a credible full-service router)

But with the first being over a year away (late 1992) and the other two being but statements of direction, we are hardly "blown away" by IBM's fast pace.

Even SNA itself is advancing slowly:

- no subarea/APPN Network Node convergence
- no open APPN Network Node
- no 3270 over LU 6.2
- no 3270 over OSI-TP4

The biggest news seems to be partnerships—HP OpenView network manager, NET licensing of split-bridge technology, AT&T network management compatibility, rumors of third-party router partnerships, etc. I guess these might be termed "architecture by liaison." So, while we wait, where is there room for conjecture?

Contemplating this question, I imagined a guide—a guide to the future—a hitchhiker's guide to the path toward SNA/OSI integration. Here might be some of such a guide's projected contents.

Foreword

Over time, many little convergences (did I almost say concessions?) toward OSI will be seen. Significant ones. Workable ones. Some as soon as 1992. Most later.

Individual chapters detail these moves. Here are three.

Chapter 1 - OSI/SNA Network Management

Network management OSI/SNA convergence is the strongest bet. The first implementation will be at the data link layer: CMIP over LLC type 1.

Note: It appears that the IEEE 802 LAN/MAN management protocol, based upon this model and strongly supported by IBM (see *SNA Perspective* January 1991 Architect's Corner), is on track toward standardization by late 1991 or early 1992. During the standardization process, this joint IBM/3Com submission received a lightweight convergence layer, used to preserve CMIP's requirement for logical associations between nodes.

With products based on CMIP/LLC probably to appear on the heels of the emerging IEEE standard, support of both CMIP/OSI and CMIP/SNA cannot be far behind. Thus we see IBM's convergence on OSI CMIP for network management.

Chapter 2 - OSI/SNA Directory Services

There are two levels of naming resolution services—at the application level and the network/subnetwork level. (And there's possibly a third, operating at the transport (OSI) or LU (SNA) layer.)

At the application level (seen through CPI-C), the need is for converged OSI/SNA use of X.500 directory services. I would expect an ability for both symbolic_destination_names (SAA CPI-C) and local_alias names (OSI) to be resolved through a common X.500-based directory service with local Directory User Agent (DUA) and remote Directory Service Agent (DSA). The X.500 DSAP protocol would require both SNA (APPC) and OSI communications profiles. Directory service convergences between OSI and SNA are several years away at a minimum.

At the network/subnetwork level, the need is for converged name-to-address (and reverse address-to-name) resolution services. What form this might take is wide open at this point. ES-IS connection/connectionless alternatives exist; an 802-level discovery protocol could also be utilized for LANs. Whatever the scheme(s) adopted, it/they should operate over LANs (802.3, 4, 5; FDDI), MANs (802.6), and WANs (SDLC, HDLC, frame relay, ISDN). I won't predict yet the final outcome(s) at this layer.

Chapter 3 - OSI/SNA Integrated Node (ES/IS)

Both trivial and nontrivial variants of integrated SNA/OSI nodes will exist, for both end systems and intermediate systems.

The trivial cases exist already, such as fully parallel, nonintegrated stacks (dual stacks) within end systems: ES-9000, OS/2, AIX, OS/400. Convergence exists at the level of the network interface adapter, no further. X.25 NPSI is the first foreshadowing of future integration within intermediate systems (X.25 carried within SNA).

The nontrivial cases will feature three levels of enhancements. All these nontrivial OSI/SNA integrations will also be several years away.

First to appear will be a converged link/subnetwork access layer. This layer converges OSI layer 2 (e.g., IEEE 802 LLC) and layer 3a (e.g., X.25) for use of multiple upper-layer protocol stacks—both OSI and SNA (and possibly TCP, NetBIOS, and others). A first view of this layer was shown in the heterogeneous LAN management (HLM) documents released in 1990. Ultimately, this layer will include access to newer services such as FDDI, frame relay, and even voice/data bandwidth management.

Second will be a sharing of common services, including node operations services such as directory services, network topology management, route selection, and configuration. System management services both of the network layers and other system functions.

Third will be the integration at the CPI-C interface—SNA/OSI access to symbolic name resolution, conversation services, asynchronous messaging services, and remote procedure call.

Epilogue

Who knows when such a galactic guide will really be required. Certainly it will be many years in the making. But it will be required. And more chapters will be needed:

- SNA/OSI addressing
- SNA/OSI CPI-C
- SNA/OSI security
- SNA/OSI and OSI/SNA tunneling
- SNA/OSI file transfer
- SNA/OSI multimedia messaging

Meanwhile, I believe SNA will continue to be the beneficiary of IBM's hottest technical protocol and architecture innovations—but OSI will soon follow.

Aside from SNA and OSI timing, the primary open question in my mind will remain: "What will be the role of TCP?" Support of TCP is growing, even within IBM products—albeit at the level of "tactical market demand." But this is a question for future treatment. I wonder, in the context of a hitchhiker's guide series, will TCP require simply a footnote, or rather a full supplement? ■