SNA Perspective

Volume 13, Number 1
January, 1992
ISSN 0270-7284

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

In This Issue:

Enterprise Network Management Part I: The SNA Subarea

Networked computing resources are perhaps the most valuable and yet least well managed resources within the enterprise today. End users operate in increasingly complex environments and require timely and uniform access to distributed applications. The competitiveness of the enterprise today depends on the coherent operation and management of its networking infrastructure.

IBM, aware that effective and efficient network management is key, is now positioning SystemView and NetView as the basis for integrated system and network management in complex interconnected environments. This first part of our two-part analysis examines the role of SystemView and NetView in providing for multivendor, multiprotocol management primarily within the host environment. Current and anticipated NetView capabilities are discussed. This article concludes with an introduction to issues of distributed network management. The second article will focus on system and network management in distributed networked environments.

(continued on page 2)

APPN Strategy Today

The success of IBM networking is closely tied to the future of SNA and the future of SNA itself depends on the success of APPN. IBM must meet several challenges in order for APPN to be the choice of its customers who need to move from subarea SNA to a peer oriented architecture. Although IBM will continue to own and direct development of SNA (including APPN), the company realizes that a significant degree of openness, access, and interoperability is essential for APPN to be successful in the market.

This article describes many internal and external forces affecting IBM's APPN development and marketing efforts. It explains how IBM is structuring its thinking regarding the major elements of networking systems and multiple architectures, such as decoupling applications from transport systems. IBM's moves toward APPN openness are addressed, including efforts to have elements of APPN adopted as ISO standards as well as the company's rumored plans to publish, license, and/or patent APPN.

(continued on page 12)

(continued from page 1)

Network Management Issues

Networked computing resources have assumed the roles of the central and peripheral nervous systems of the enterprise. These critical resources provide the infrastructure to deliver, format, present, and interchange data and information on a highly available, timely, and cost-effective basis. End users of networked applications include executive, middle, and operations managers, technical and administrative professionals, and clerical staff. End users throughout the enterprise increasingly require:

- Timely access to target application data, regardless of geographic location.
- Uniform access to applications, regardless of the underlying complexities of heterogeneous computing and networking platforms and infrastructures. IBM refers to this system-independent view as a single-system image.
- Object-oriented, window-based application views and interactions.
- Continuous availability of critical application data.
- Mean time between failures (MTBF) and mean time between outages (MTBO) approaching infinity.
- Mean time to repair (MTTR) approaching zero.
- Built-in fault tolerance in all computing and networking resources.

Critical Yet Least Managed

Paradoxically, networked computing resources are perhaps the most critical and yet the least well-managed resources within an enterprise. This paradox exists for several reasons:

- Network implementations not reflecting end user requirements
- Interdepartmental and enterprise-wide incompatibilities
- Undervalued networked computing resources
- Fundamentally incompatible multivendor, multiprotocol networks
- Isolated pursuit of solutions in multinational enterprises with multiple data centers and cost centers

Product and procedural implementations of the operational network generally do not reflect end-user application or operational requirements. From an MIS perspective, this is due to a tendency of design staff to specify, select, and implement configurations without first considering application or operational requirements and attempting to optimize deliverables against budget accordingly.

Improvements in cost-performance ratios in the underlying technology give rise to uncoordinated development. There is a marked tendency in local users and their direct management to acquire and implement local solutions in the complete absence of interaction with other departments or with MIS staff. This quickly generates interdepartmental and enterprise-wide incompatibilities.

Networked computing resources are still, to a great extent, considered "soft" assets by financial and accounting evaluators. That is, information and its underlying distribution infrastructure still remain difficult to quantify when compared to "hard" assets such as plant, property, and equipment. The result is that networked computing resources are, more often than not, undervalued generally at historic, tangible hardware/software purchase and implementation costs rather than from a net present value.

The overwhelming value of networked resources is timely, reliable, and cost-effective access to, and delivery of, appropriate data to the distributed decision maker, yet this critical feature of networks is generally not considered within a business evaluation framework. The absence of proper valuation of the networked asset leads to a tendency to improperly predict all tangible benefits, intangible benefits, tangible costs, and intangible costs associated with each of the phases of the network development life cycle.

Multivendor, multiprotocol networks are fundamentally incompatible and are becoming increasingly complex. End users find that they must somehow navigate through an ever-intensifying array of incompatible hardware, software, and interface offerings. These concerns are exacerbated by the proliferation of incompatible niche products seeking to address individual problems.

Many end users operate in corporate or governmental organizations that are multinational in scope. This geographic distribution often gives rise to multiple data centers and multiple management cost centers, both of which frequently result in isolated pursuit of solutions.

SAA Network Management Architecture

System and network management components of IBM's Systems Application Architecture (SAA) specify the management services necessary to enable planning, organization, and control functions in both SNA and multiprotocol networks. These architected components, properly implemented, can provide a reasonable and cost-effective basis for coherence in the enterprise networked computer resource.

SAA and AIX—Coordinated Network Management

IBM introduced SAA in March 1987 to begin elimination of a tradition of unique and incompatible application and networking solutions in its large-scale enterprise processor, midrange departmental processor, and programmable workstation platforms. The Advanced Interactive Executive (AIX) family was standardized in March 1988 to lend application and networking coherence to IBM's UNIX-based enterprise, departmental, and programmable workstation platforms.

IBM is currently engaged in a major effort to provide for interoperability among the various SAA and AIX networked application platforms. This effort involves network management in addition to other cross-platform application integration requirements.

SystemView

In September 1990, IBM introduced SystemView as its strategy for overall management of the enterprise information system, of which network management is a component.

Efficient network management has emerged as a key factor in the enterprise which affects its overall performance in domestic and international markets. IBM positions SystemView as the basis for integrated system and network management in increasingly complex connected environments. The SystemView framework has two major objectives:

- To guide development of IBM products to provide integrated application solutions
- To define and support an open development platform for integrated management of SNA, OSI, and TCP/IP applications and their associated network infrastructures

This second objective seeks to provide end users and system/network operators with a consistent look and feel across the managed system environment. Because SystemView is intrinsic to SAA, a significant value-add is that application developers should be able to implement one SystemView application on multiple platforms, thereby lowering the overall development effort and associated life cycle costs.

Coming Soon: A Common Management Interface

The SystemView open structure is designed to support multiple system and network management services and protocols through a common management interface (CMI) under development at IBM. The SystemView CMI is specifically intended to support the following management protocols:

- SNA Management Services (SNA/MS)
- OSI Common Management Information Services/Protocol (CMIS/CMIP)
- TCP/IP Simple Network Management Protocol (SNMP)
- LAN management protocols (e.g., CMIP over LLC (CMOL))

A significant provision of the SystemView CMI will be its capability to support these multiple system/ network services and protocols while keeping them transparent to application developers and end users alike.

SystemView Environment

The SystemView environment has three dimentions:

- End use
- Application
- Data

End-Use Dimension. The end-use dimension provides the user at a workstation with a consistent application view through the SAA Common User Access (CUA). End-user interfaces include graphic display, textual dialogs, and/or command language. SystemView end-use dimension tools include OS/2 Presentation Manager, OS/2 Dialog Manager, EASEL, or GraphicsView/2. Products which conform to the SystemView end-use dimension include:

- SAA Asset Manager
- SAA Delivery Manager
- Enterprise Systems Connection (ESCON) Manager
- ESCON Analyzer
- Information/Management and Information/System
- NetView Graphic Monitor Facility
- NetView Resource Object Data Manager
- Operations Planning and Control/Enterprise Systems Architecture (OPC/ESA)
- Problem Management Productivity Services (PMPS)
- Service Level Reporter
- Workstation Data Save Facility/VM (WDSF/VM)
- OS/400, AS/400 Systems Management Utilities
- SystemView System Manager/400

In June 1991, IBM announced development agreements to develop SystemView interfaces and services in conjunction with Goal Systems International, Inc., Candle Corporation, PLATINUM Technology, Inc., and Bachman Information Systems. All of these software vendors intend to integrate their products into the SystemView structure to include use of the NetView automation platform for System/390/370.

Application Dimension. The application dimension defines guidelines for implementation and integration of systems management applications, and groups system management tasks into the following disciplines:

- Business management
- Change management
- Configuration management
- Operations management
- Performance management
- Problem management

The SystemView application dimension is intended to support ISO 9595 (CMIS) and ISO 9596 (CMIP).

Data Dimension. The SystemView data dimension provides for standardized system management data definitions and access. Data definitions are specified through the SystemView data model, which is intended to be consistent with ISO 10165-4, Guidelines for the Definition of Managed Objects, and ISO 10165-1, Management Information Model. The data model incorporates descriptions of resource characteristics and their interrelationships. SystemView systems and network management data will be stored in an enterprise information base using the SAA Structured Query Language (SQL) database interface. Products that conform to the SystemView data dimension include:

- SAA Asset Manager
- SAA Delivery Manager
- Information/Management and Information/System

SystemView runtime environments include standalone LANs, multiple LANs, and interconnected LANs and WANs. SystemView support will include all SAA and AIX platforms across SNA, OSI, and TCP/IP networked computing environments.

SAA Network Management Elements

SAA Network Management Architecture (NMA) specifies the management services required to enable planning, organization, and control functions within SNA, OSI, and TCP/IP networked computing environments. The major elements of NMA include problem management, performance and accounting management, configuration management, and change management. These elements and their supporting subdisciplines are shown in Figure 1:

- Problem management is the process of managing network problems (unwanted changes) from their detection through to final resolution.

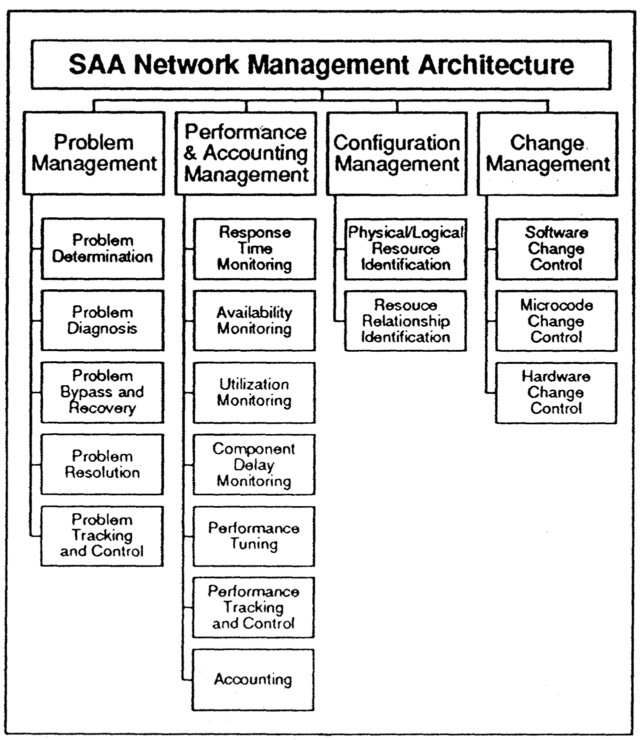

Figure 1

- Performance and accounting management is the process of quantifying, measuring, reporting, and controlling the responsiveness, availability, utilization, and usage charges of network components.
- Configuration management is the process of controlling information that is necessary to identify physical and logical networked resources and their interrelationships.
- Change management is the process of planning and controlling the additions, deletions, and modifications of networked hardware, software, and microcode resources.

These management components are executed by a network operator (either human or automated) which is associated with an SNA node (host processor, communication controller, cluster controller, or workstation at the end-point of a link or two or more links), OSI node, or TCP/IP node that contains appropriate network management facilities.

Focal Point, Entry Point, and Service Point

NMA distinguishes focal point, entry point, and service point components.

Focal Point. Focal points reside within System/390/370 hosts and make consolidated network management data available to centralized network management applications. Focal point product implementations include:

- NetView
- NetView Distribution Manager
- NetView Performance Monitor
- NetView File Transfer Program
- Automated Operations Control/MVS (AOC/MVS)
- CICS Automation Option/MVS
- IMS Automation Option/MVS
- Target System Control Facility (TSCF)
- Automated Network Operations/MVS (ANO/MVS)
- Host Command Facility

Entry Point. Entry points are distributed locations that provide network management services for themselves and for attached SNA resources and devices. Entry point product implementations include:

- AS/400
- System/36
- System/38
- Series/1
- 3174/3274 cluster controllers
- 37xx communication controllers

Service Point. Service points provide distributed points for management services to support non-IBM SNA and non-SNA access into SNA. In this sense, service points function as network management servers that collect network management data from non-SNA environments, architect the data into SNA management services data, and forward the MS data to a focal point. Service point communications with non-SNA resources do not necessarily use SNA management services or protocols. These may include OEM internal management, OSI management (CMIS/CMIP), or TCP/IP management (SNMP) services and protocols. Service point implementations include:

- NetView/PC (OS/2 and DOS)
- AIX NetView Service Point
- OEM SNA and non-SNA devices
- Rolm CBX II, 9750 BCS, and 8750 BCS

NetView: The Flagship

NetView is clearly IBM's flagship system and network management product set. As a focal point product set, NetView provides host-based, centralized management functions.

Advantages of Hierarchy

The host-based, hierarchical management nature of NetView reflects IBM's traditional host-centric philosophy with SNA. That is, subarea SNA networks are host-mediated such that all nonhost device and program access and control are orchestrated from a central point. In this role, all non-NetView and nonhost participating products are regarded as subordinates within a hierarchical scheme.

This host-centric nature of NetView with its hierarchical network implications provides a reasonable level of problem, change, configuration, and performance management for users whose networking requirements are predominantly program-to-device in nature. A significant portion of IBM's customer base relies heavily on subarea SNA infrastructures. The major advantages to hierarchical network management include:

- A single point of application access, control, and security
- Avoidance of hardware, software, personnel, and procedural redundancies
- Consolidation of network event, alert, and statistical intelligence
- A single point of control for network resource activation and deactivation
- A single point of recovery from unanticipated physical and logical problems

The Reality of Distributed Solutions

However, several SNA users are increasingly deploying distributed processing solutions. These solutions are quite often implemented at the department level independent of coordination with other departments or with MIS. Departmental distributed processing solutions are generally LAN-based, with a multitude of client/server workgroups located throughout the enterprise.

These distributed workgroup computing environments are often based on a series of client/requester programs running on workstations that interact with server programs. These server programs may be based in workstations, departmental processors, or System/390/370 hosts, that are either on the same LAN, a different LAN, and/or across a WAN or multiple WANs. NetView Version 2 begins to directly recognize and manage a variety of LAN networked environments. However, several issues of distributed control remain. These latter issues will be discussed in the second article in this series.

NetView Environment

Figure 2 provides an overview of the NetView elements. NetView runs as a collection of application programs within MVS, VM, and VSE operating environments. Operating system interfaces are MVS Subsystem Interface (MVS SSI), VM Programmable Operator (VM PROP), and VSE Operator Communication Control Facility (VSE OCCF).

Command Facility. NetView consolidated several previously disparate host-based management program products. The base NetView component is Command Facility, which is an enhancement of Network Communications Control Facility Version 2 Release 2 (NCCF V2R2). NCCF was enhanced for NetView to interoperate with other internal NetView components. NetView Command Facility simplifies and automates several tasks associated with operating system, subsystem, and network management through the use of command lists (CLISTs).

Hardware Monitor. NetView Hardware Monitor is an enhancement of Network Problem Determination Application (NPDA) V3R2. It runs under Command Facility and provides a problem management alert facility to monitor alert messages and automatically notify network operators (human or automated) of error conditions encountered that exceed predefined thresholds of acceptability.

Session Monitor. NetView Session Monitor is an enhancement of Network Logical Data Manager (NLDM) V1R3 and runs under Command Facility. NetView Session Monitor functions include:

- Detection of extended data stream terminals for full-color support
- Support for Extended Recovery Facility (XRF) to solicit, correlate, and present SNA Boundary Function trace data, Response Time Monitor (RTM) data, and session awareness data

The relationships between NetView Command Facility, Hardware Monitor, and Session Monitor are illustrated in Figure 3 (on page 8). Hardware Monitor is shown collecting event data, alert data, and statistical data from a wide range of SNA and OEM networked environments. Session Monitor collects SNA LU session and underlying connection data. Detected physical (Hardware Monitor) or logical (Session Monitor) problems are presented either to a network operator console or to an automation program (such as AOC/MVS, CICS Automation Option/MVS, IMS Automation Option/MVS, TSCF, or ANO/MVS). The human or automated operator then executes CLISTs through NetView Command Facility which are invoked to correct the detected problem(s).

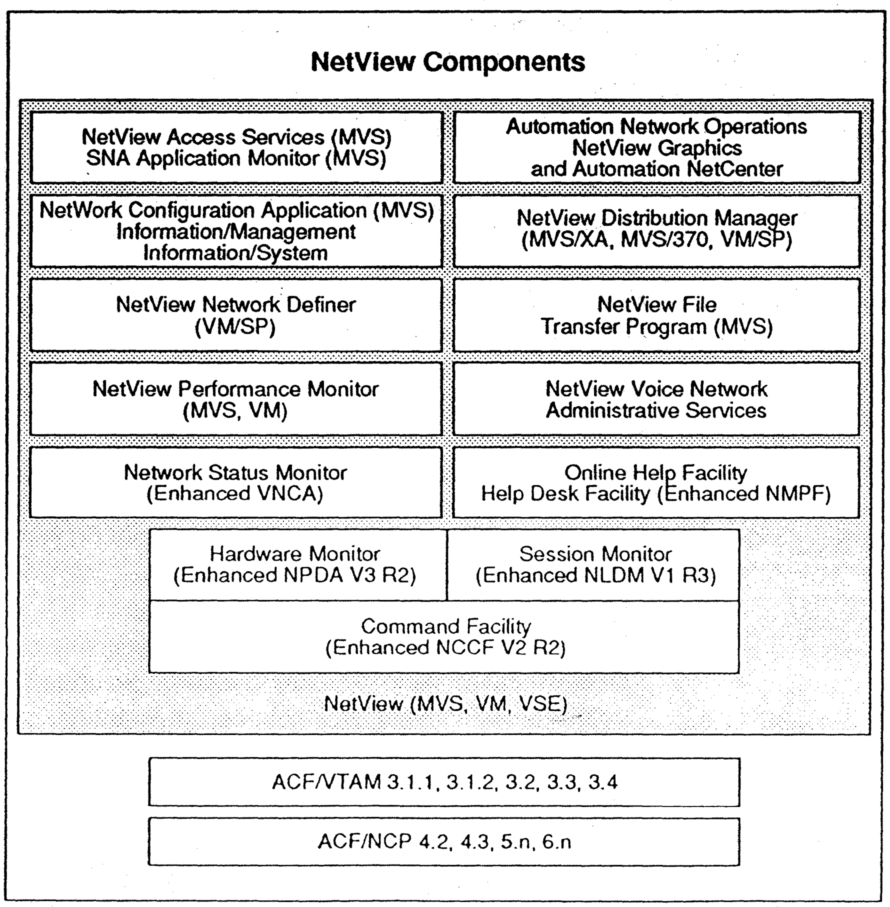

Figure 2

Consistency and Simplification Needed

SNA Perspective believes that IBM should further enhance NetView internal interfaces (NCCF, NPDA, and NLDM) to provide greater internal consistency between these pre-NetView platforms. In keeping with stated SAA and AIX directions, consistent APIs should be designed to support end-user and system/ network operator interfaces in a consistent way.

It is imperative to simplify NetView to the point where an MVS, VM, or VSE implementation can be installed and running within fewer than 40 staff hours, rather than the approximately 300 staff hours it has required in several sites. NET/MASTER from SCI, for example, requires an average setup time of only 30-40 staff hours and does not appear to suffer from the internal, pre-NetView interface complexities of NetView. NET/MASTER also supports the same CNM APIs as NetView, thereby ensuring provision for the entire range of NetView functionality as it emerges. *SNA Perspective* is not suggesting that NET/MASTER is necessarily a superior product to NetView. However, it is clearly more easily configured and internally more highly integrated.

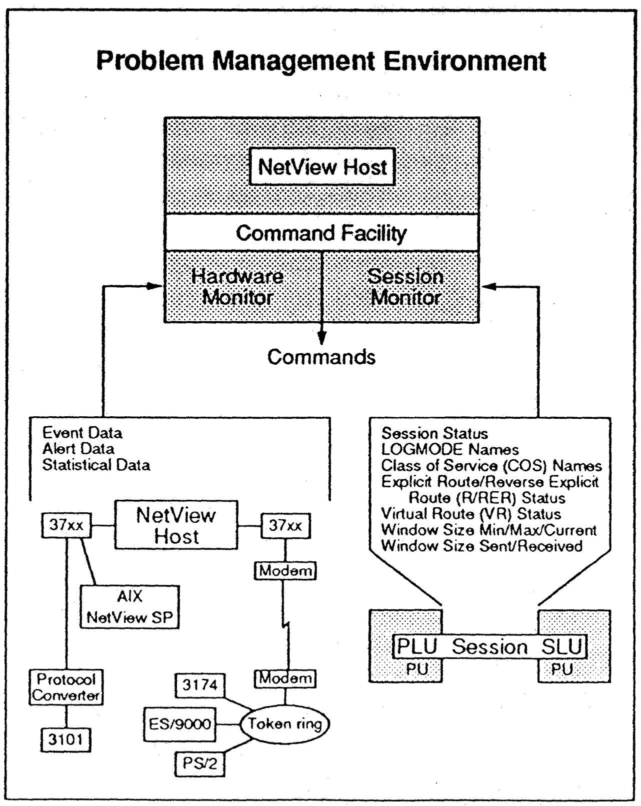

Figure 3

NetView integration of SNA, OSI, and TCP/IP system and network management will continue to increase internal complexities. The significant challenge facing IBM throughout this process is to increase the simplicity of the end-user and system/network professional view.

NetView Version 2

IBM introduced NetView Version 2 in September 1990. Table 1(on page 9) is a summary of new features in NetView Version 2 Releases 1, and 2. Version 2 provides several significant enhancements over Version 1:

- Central, distributed, and standalone NetView
- Graphic Monitor Facility (GMF)
- LU 6.2 interfaces

A central NetView site provides the full suite of NetView capabilities for a multihost network. Distributed NetView sites cooperate subordinate to a central NetView site within a multihost network. Standalone NetView provides a complete set of NetView functions within a single-host network.

GMF, shown in Figure 4 (on page 9), introduces an object-oriented network management interface to NetView environments. GMF runs as an OS/2 cooperative processing application. GMF provides a window view set of displays into SNA networks, from high level down to the physical unit (PU), and complies with SAA CUA. It functions as an integrated client/server application platform and supports LU 2/3270 emulation to a host operator task as well as LU 6.2 transport to a host graphics task.

The introduction of GMF into the NetView environment is significant in that graphical representations of system/network management data greatly enhance users' abilities to access, assimilate, correlate, and manage large volumes of complex management data. The GMF LU 2/3270 interface integrates network operator access to 3270 system/network applications such as NetView Performance Monitor (NPM), TSCF, AOC/MVS, and ANO/MVS.

LU 6.2 interfaces are provided in NetView V2R2 through SNA Management Services (MS) transport and high-performance transport. MS transport provides for short duration conversations and high performance transport supports bulk system management data flows.

New Features Of Release 2

Significant new features in NetView V2R2 include system/network automation and multiprotocol network management. NetView V2R2 provides the basis for managing multiple SNA and non-SNA environments which include:

- SNA networks
- OSI/CS Command Processor
- TCP/IP SNMP
- LAN Network Manager
- AIX NetView Service Point
- NetView/PC and applications (e.g., HubView/PC)
- Transmission Network Manager for Integrated Data Network Exchange (IDNX)

NetView-supported automation facilities, as stated earlier, include AOC/MVS, CICS Automation Option/MVS, IMS Automation Option/MVS, TSCF, and ANO/MVS. In general, these program environments provide expert system solutions within increasingly complex management environments. In so doing, they begin to replace human console-based interfaces with heuristic logic, which tends to reduce system failures through improved availability and reliability.

As a result of automated operations, network operator productivity tends to increase. This is due to the fact that, as the enterprise expands, the underlying network infrastructure becomes correspondingly complex, requiring a proportional increase in operators and consoles. Without

NetView Version 2

Release 1
- Graphic Monitor Facility (GMF)
- NetCenter Feature
- INFO/MGT Bridge Support
- STATMON Enhancement
- NPM R4 (MVS)
- FTP V2 (MVS), V1 (VM, VSE)
- NetView Access Services V1R3 (MVS/VM)
- Central, Distributed, and Standalone NetView

Release 2
- LU 6.2 Transport
- Automation Enhancements
- VTAM Persistent Session Support
- LAN Network Manager Command Enhancement
- GMF Support in VM/ESA
- Automation (AOC/MVS, CICS Automation Option/MVS, IMS Automation Option/MVS, TSCF, ANO/MVS)

Central NetView Functions
- Full Function
- NetView-NetView
  - Session Manager
  - Hardware Manager
  - Cross-Domain Support
  - Alert Forwarding

Distributed NetView Functions
- Cross-Domain Nodes
- Central NetView-Controlled
- Limited Operator Interface
- Local Automation
- Local Problem Management Recording
- Problem Management Alerts To Central
- No NLDM/NPDA Command

Standalone NetView Functions
- Single NetView Host
- No NetView-NetView
- Operator Interface
  - Session Monitor
  - Hardware Monitor

Table I

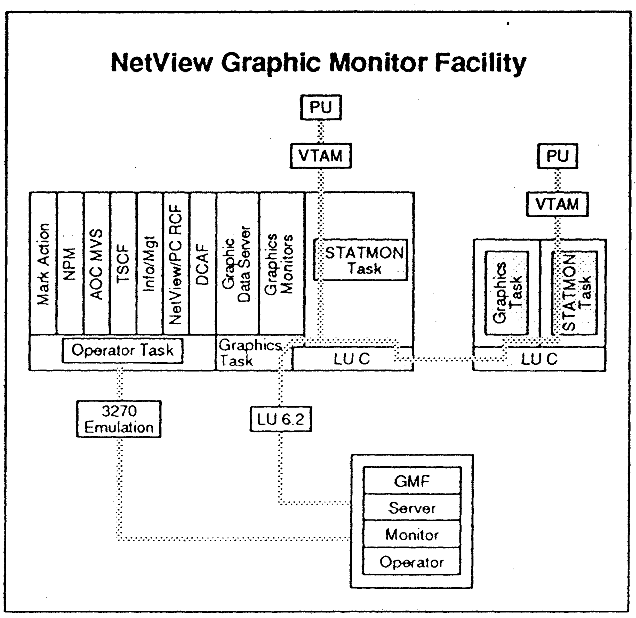

Figure 4

automation facilities, one direct result is increased coordination complexity, leading to longer flow times to reset problems, assess configurations, and implement changes.

NetView Can't Track APPN Changes

Unfortunately, while NetView V2R2 does provide for LU 6.2 transport (thereby enabling network management data to flow over LU 6.2 sessions rather than SSCP-PU/LU 0 sessions), there is as yet no NetView mechanism to account for the dynamic realities of a downstream APPN environment.

This problem becomes especially significant when we consider the dynamic and uncoordinated nature of implementing distributed processing solutions at the department level. It is significant that APPN network node capability is provided on the PS/2, while NetView cannot even recognize this highly dynamic downstream environment. At present, NetView can only maintain awareness of downstream environments through arbitrary updates or through orchestrated configuration solicitations of PU 4/PU 2 nodes with Network Asset Manager.

SNA Perspective believes that it is critical for IBM to incorporate APPN network node awareness in NetView. Further, *SNA Perspective* believes that, although AS/400 network management functions are improving, APPN network node management functions need to be provided there as well. Users would be well served to have APPN network node management capabilities incorporated into NetView and AS/400 environments concurrent with release of the SystemView CMI and Resource Object Data Manager API.

Failure to provide NetView with configuration and session awareness of APPN networked environments contradicts IBM's commitment to have NetView manage the entire enterprise.

Inconsistency Resolution Needed

It is imperative that IBM address the current inconsistency between SystemView and the SAA Repository Interface. Provision of a consistent look and feel for system and network management application developers and operators is just as vital as it is for system end users. Resolution of SAA consistency will go a long way toward convincing the system and network management community that NetView is, in fact, a viable and integrated member of the SystemView team.

NetView Directions

SNA Perspective believes that NetView V2R3 and beyond will provide several more improvements:

- Resource Object Data Manager. The Resource Object Data Manager API will provide integration of existing and emerging management applications. End users will perceive a consistent presentation of management data from a wide variety of software collection tools.
- SystemView CMI. The SystemView CMI, described earlier, will be available in future releases of NetView, further enabling consistent multiprotocol system and network management.
- Enhanced Graphics. In keeping with SystemView, NetView will increasingly provide for an integrated, graphics-oriented facility for monitoring and managing SNA as well as non-SNA (OSI, TCP/IP) physical and logical resources. The non-SNA support will continue to be based upon mapping alerts to represent resource status.

Prelude to Part II

Part II of this two-part analysis series will extend discussion of system and network management away from the System/390/370 host and consider:

- Client-server computing network management impacts
- APPN network management issues in depth
- Open network management issues, directions, and implications ■

Abbreviation Glossary

AIX	Advanced Interactive Executive
ANO/MVS	Automated Network Operations/MVS
AOC/MVS	Automated Operations Control/MVS
API	application programming interface
APPN	Advanced Peer-to-Peer Networking
AS/400	Application System/400
CICS	Customer Information Control System
CLISTs	command lists
CMI	common canagement interface
CMIP	Common Management Information Protocol
CMIS	Common Management Information Services
COS	Class of Service
CUA	common user access
DCAF	Distributed Console Access Facility
ESA	Enterprise Systems Architecture
ESCON	Enterprise Systems Connection
FTP	file transfer protocol
GMF	Graphic Monitor Facility
IMS	Information Management System
INDX	Integrated Data Network Exchange
Info/Mgt	Information/Management
LAN	local area network
LU 6.2	logical unit 6.2
MS	Management Services
MTBF	mean time between failures
MTBO	mean time between outages
MTTR	mean time to repair
MVS	Multiple Virtual Storage
MVS SSI	MVS Subsystem Interface
NCCF	Network Communications Control Facility
NLDM	Nework Logical Data Manager
NMA	Network Management Architecture
NPDA	Network Problem Determination Application
NPM	NetView Performance Monitor
OPC/ESA	Operations Planning and Control/ESA
OSI	Open Systems Interconnection
OSI/CS	OSI/Communication Subsystem
PLU	Primary Logical Unit
PMPS	Problem Management Productivity Services
PS/2	Personal System/2
PU	physical unit
RCF	Remote Console Facility
R/RER	Explicit Route/Reverse Explicit Route
RTM	Response Time Monitor
SAA	Systems Application Architecture
SLU	Secondary Logical Unit
SNA/MS	SNA Management Services
SNMP	Simple Network Management Protocol
SSCP	System Services Control Point
SQL	Structured Query Language
TCP/IP	Transmission Control Protocol/Internet Protocol
TSCF	Target System Control Facility
VM	Virtual Machine
VM PROP	VM Programmable Operator
VR	Virtual Route
VSE	Virtual Storage Extended
VSE OCCF	VSE Operator Communication Control Facility
WDSF/VM	Workstation Data Save Facility/VM
XA	Extended Architecture
XRF	Extended Recovery Facility ■

(continued from page 1)

IBM's Quandary

Corporate and Economic Environment

IBM is facing many fundamental challenges, evidenced by its poor financial results and falling market share over the past several years. IBM management has publicly acknowledged these issues for some time and the company has stated that 1992 will be a year of major reorganization. These challenges are compounded by the slump over the last several years in the computer and data communications industry. The Networking Systems group, though, has been faring better than many IBM lines of business in the face of these environmental challenges.

Balancing Acts

IBM needs to balance several competing concerns which affect APPN and its potential for success. These balancing challenges include:

- Protecting its customer investments in existing SNA networking without strangling its ability provide customers with new technology competitive with offerings from other vendors
- Making a financial return on its considerable investment in APPN and yet encouraging multiple vendors to support, and users to purchase, APPN in order to achieve market success
- Maintaining its control over APPN in order to provide its customers with the benefits of a proprietary technology (manageability, flexibility, functionality, security, reliability) while supporting user demands for standards, connectivity, and interoperability
- Providing connectivity between APPN and subarea SNA, both for APPN traffic across a subarea as well as subarea traffic (e.g., 3270) across APPN
- Supporting coexistence and interoperability between APPN, OSI, and TCP/IP

How Open?

IBM considers SNA (including APPN) to be its proprietary architecture, "proprietary" in that IBM owns it and directs its development. Since its introduction, however, the company has made many of SNA's specifications available. Three of the most notable exceptions are at the heart of SNA control and routing: PU 5, PU 4, and APPN network node.

SNA Perspective expects that IBM will continue to see APPN as a proprietary architecture. The first question is how open IBM will allow APPN specifications to be—specifically, will it open network node? The second question is whether IBM's answer to the first question will be sufficient, together with IBM's other efforts (e.g., marketing, incentives, alliances, migration support, pricing), to motivate the market to accept APPN.

Market Success Yet To Come

Though introduced over five years ago, APPN has not yet found widespread success in the market for several reasons:

- IBM's host and communication controller APPN offerings are still limited to LEN node, though the company says that it is likely to announce full APPN for both during 1992.
- Because of delays in major IBM products over the last few years, such as OfficeVision and the Repository, users are wary of depending on IBM for major product introductions.
- Because of the lack of a single voice between IBM divisions, culminating in the approaching major IBM reorganization which will significantly decouple major IBM business units, users are concerned that architectures such as APPN which are developed by the Networking Systems unit may not necessarily be fully supported by the systems units.
- Users have voted with their standards dollars not for OSI, which is where IBM has invested most of its development for open networking, but rather in TCP/IP. (See *Architect's Corner* in this issue for an examination of this phenomenon.) Therefore, IBM's plans for coordinated support between SNA/APPN and OSI is not the selling point the company had hoped for.
- LU 6.2 has not caught as quickly as many expected since its debut in the early 1980s. A larger installed base of LU 6.2 by now would have created a stronger market demand for APPN.

This last point deserves a bit more discussion. Many studies in the mid-1980s, including a major user survey in 1987 by Communications Solutions, Inc., the publisher of *SNA Perspective*, revealed user plans to have a majority of their SNA development transitioned to LU 6.2 before 1990. However, that kind of growth did not materialize (see "LU 6.2 Growing More Slowly Than Anticipated," in *SNA Perspective*, September 1990). Estimates of actual development are closer to twenty percent today, though the growth rate is ramping up. Some reasons for this dampened early market enthusiasm included lack of off-the-shelf software, lack of LU 6.2 across mainframe environments and on other platforms, inconsistency of LU 6.2 implementations on different platforms, implementation problems (e.g., APPC/PC was too big and inefficient), and lack of IBM developer support.

However, LU 6.2 support was significantly enhanced across IBM's product line in 1990 and 1991 and IBM is now devoting resources to APPC market development. *SNA Perspective* has called 1991 "the year of LU 6.2" and we expect LU 6.2 growth to accelerate.

APPN and Subarea SNA

The sidebar entitled "SNA Routing: Subarea SNA and APPN" provides a brief overview of APPN and its relationship to subarea SNA routing.

Rick McGee, manager of communication systems architecture in IBM's Networking Systems line of business, whose team has responsibility for APPN architecture development, says the company has a long-term APPN strategy but has not set the details more than two years out. IBM has been openly discussing three generations of APPN architecture. These have been referred to as APPN, enhanced or version 2 APPN, and fast-packet APPN.

Each of these will offer increased functionality and flexibility. APPN network node enhanced LEN node capabilities with dynamic directory, dynamic routing, and intermediate node routing. *SNA Perspective* expects that enhanced APPN will include LU 6.2 transport for LU 2 and other dependent LU sessions, more network management capabilities, and improved intermediate node routing. Fast-packet APPN will be adapted even further to effectively support the emerging very high transmission speeds characteristic of SONET, ATM, and multigigabit LANs.

Architecture versus Implementation

Users need to be aware, however, of the difference between architecture and implementation. IBM's architectural staff can influence but not direct product development. Therefore, although architectural solutions may have already been developed that address a particular user's concerns, the product teams retain the choice of whether and when to implement those elements of the architecture. For example, APPN has long been architected to support dynamic configuration but LEN node still required static configuration of adjacent nodes. Further, the full addressing capability available in the architecture is still significantly limited in implementation. We will address these and related issues in more detail in an upcoming issue of *SNA Perspective*.

APPN and Subarea SNA Integration

Currently, with only LEN node on VTAM and NCP, an entire subarea acts as one composite T2.1 node. IBM is strongly suggesting that APPN with both end node and network node for VTAM and NCP will be announced in 1992 and available hopefully in late 1992 or at least in 1993. Enhanced APPN is likely to be announced in 1993.

IBM has also been hinting that NCP will constitute a smaller portion of the APPN picture. APPN VTAM seems central when IBM publicly discusses APPN strategy, but not so for NCP. Further, IBM has said that its multiprotocol router will not support subarea SNA routing (PU 4) but, rather, will use some form of encapsulation of SDLC to deliver subarea traffic. Finally, IBM has been talking about the need for SNA, if it is to continue as an enterprise backbone, to be able to support other protocols more effectively. This won't be effective unless APPN rather than PU 4 is the dominant architecture in the major SNA routing nodes, currently the 3745. *SNA Perspective* expects that the 3745 will be replaced with a new hardware platform in 1992, based on the same technology as the RS/6000 and the upcoming multiprotocol router. The new platform's NCP must certainly support PU 4 because subarea SNA will stay installed well into the next century. However, both the platform and the software will be optimized for current and future generations of APPN.

However, even with full APPN on the host, subarea SNA SSCP control flows will still be needed from the "owning" host for LU 2 sessions, at least to assist in establishing the session. This is because control flows required to establish a 3270 session flow over an SSCP-PU sessions and these SSCP-PU sessions are not allowed in APPN. Also, the PU 2 supporting the secondary LU for a 3270 session must be logically adjacent to an NCP or the host. To support a migration and maintain customer investment in 3270 applications, IBM needs to extend these control flows over an arbitrary APPN topology.

Two Approaches to 3270 on APPN

There are two ways that LU 2 (and other non-6.2 dependent LU types) could be routed across APPN backbones. The first approach, favored by the IBM architectural team, would encapsulate the 3270 datastream, carrying LU 2 sessions unmodified within APPC sessions. It makes sense that they favor it, because it is architecturally elegant and compliant and is SAA conformant, which makes it a "native" implementation.

We call the second approach LU 0123 tunneling. In this approach, APPN nodes present a PU 4 interface downstream to cluster controllers or a gateway and a PU 2 interface upstream to the host side. This is a more pragmatic solution and has similarities to successful approaches by several multiprotocol router vendors tunneling 3270 sessions across TCP/IP networks.

SNA Perspective believes that IBM is focusing on the first approach, first with adjacent nodes and then across multiple APPN hops. *SNA Perspective* understands that IBM has already architected a solution to this problem, which we believe involves the APPC general data stream (GDS).

Depending how much documentation or code IBM makes available on APPN network node as well as the market success of APPN, other vendors are likely to offer the second approach. Further, *SNA Perspective* thinks that even IBM products may support the second approach, at least initially, for marketing reasons.

Support of 3270 is one of several APPN challenges for IBM to handle. Another challenge is continuing moves toward interoperability of SNA (APPC/APPN) and OSI, which is happening with CPI-C at the API layer and seems to be in IBM's long-range APPN plans at the routing layer, as discussed below.

Direction: Decoupling Applications from Transport

Not only with 3270 but in general, IBM is moving to decouple application services from transport, and from data exchange facilities such as LU 6.2, OSI transaction processing (OSI/TP), and Remote Procedure Call (RPC). The ability to have mixed stacks is not a new concept. This would be analogous to what is already happening at the data link layer where SNA, for example, can run over SDLC, X.25, and LANs. It can be accomplished through the use of multiarchitecture boundaries at the application, transport, and data link layers. What would be the results of this decoupling?

3270 can be carried today over TCP/IP (see "tn3270: Another Interoperability Option" in *SNA Perspective,* June 1991); conversely, TCP/IP application services such as FTP could be carried over APPN. IBM says it has a model of this running in the lab.

Perhaps, in addition to 3270 in LU 6.2 pipes as discussed above, it could also be carried in OSI/TP pipes. Further, perhaps OSI/TP could run over APPN like LU 6.2 does. IBM says it has an architectural approach to deal with this.

But users are again cautioned about gaps between architectural capability and product availability. One gap is time: architectural prototypes in the lab precede product development by several years. Another gap is market considerations: vendors do not develop products on the basis of architectural elegance but to provide a set of solutions they expect users to buy. For example, IBM is unlikely to

SNA Routing: Subarea SNA and APPN

Advanced Peer to Peer Networking (APPN) is an IBM architecture for peer-oriented routing. It was initially developed in the early 1980s to support peer communication between IBM midrange computers (System/36 and then AS/400) so they would not be limited by the static and hierarchical nature of subarea SNA. A prototype APPN was first introduced in 1986 and was also referred to as small systems networking.

With the evolution in the computer field toward client/server computing, which requires peer-oriented networking, APPN's scope has been increased to include providing the migration path for users of subarea SNA.

APPN and subarea SNA routing are both part of the lower layers of networking architecture models. In the SNA seven-layer architecture, they both exist at layer three and a portion of layer four. Although not directly comparable, they compete with products based on OSI and TCP/IP-related routing protocols as well as other proprietary protocols such as NetWare IPX with RIP, at layers three and four of the OSI Reference Model.

Subarea SNA and PUs

Subarea SNA routing is structured using a hierarchy of networking software components, called physical units (PUs). PUs have been implemented by IBM in different software products.

PU 5, the highest level, is the managing component and resides in the host; it is implemented in the System Services Control Point (SSCP) which is part of the Virtual Telecommunication Access Method (VTAM) software product which resides in a mainframe computer.

PU 4 is the routing component; it is implemented in the Network Control Program (NCP) software product which resides in communication controllers such as the 3745. It is equivalent in internetworking terminology to an intermediate system.

PU 2 acts as an access point to an SNA network. Originally, it supported terminal access to the mainframe through being implemented in microcode on terminal or cluster controllers like the 3174. PU 2 is also emulated on 3270 protocol converters, PC-to-mainframe packages, and LAN-SNA gateways.

IBM has never used PU 3. PU 1 is essentially obsolete.

A PU 2 must be logically adjacent to a PU 4 or a PU 5. Even if the PU 2 is implemented on an intelligent device such as a midrange or personal computer and wishes to communicate with a similarly intelligent device, subarea SNA requires that its communications be managed by a PU 5. This was the impetus for developing APPN

APPN—A Single Node Type

APPN was developed as an extension to SNA using a single enhanced node type called node type 2.1 (NT2.1). This node type is implemented in Control Point (CP) software. It was based on a peer orientatation rather than the hierarchical structure of PU 5/4/2.

The initial implementation of this node type was referred to as a low entry networking (LEN) node. The LEN node was limited in that it could only communicate with logically adjacent nodes.

APPN has since been enhanced to include two node types: end node (EN) and network node (NN), which are closely analogous to the ISO end system (ES) and intermediate system (IS), respectively. An APPN end node can communicate directly with any other logically adjacent end node. To communicate with an end node that is not logically adjacent, it must go through an APPN network node, which contains network routing tables.

The LEN node and end node implementations of APPN do not actually contain a PU nor therefore a CP, which is why they can only communicate with logically adjacent nodes.

Physical units (PUs) were given this designation because they used to be more closely aligned with physical networking devices. IBM has been moving toward calling them nodes, mostly for marketing purposes. Hence PU 4 is now often referred to by IBM as node type 4. This terminology is not catching on for the older PU types. However, for APPN, the term type 2.1 node is widely used. ■

develop support for OSI/TP over APPN in the near future because it has so many other product needs with higher priorities to create a basic SNA/APPN product set.

However, the company has been increasingly supportive over the past several years in assisting other vendors who may want to develop such implementations. In addition, IBM is also active in the development of standards, another way it can promote use of its architectural developments beyond its own product plans.

APPN: An OSI Standard?

IBM has always actively participated in standards bodies both in the United States and internationally, though its demeanor has become more that of a team player as its dominance in the industry has declined. Recently, IBM noted that 1,700 of its staff members were involved directly in standards development activities. IBM has several times been successful in having its proprietary (i.e., company developed and owned) architectures adopted as standards, after being adapted to some extent by standards bodies.

For example, high-level data link control (HDLC), which was published by ISO in 1975, is based on IBM's synchronous data link control (SDLC), publicly announced in 1974. From the beginning, SDLC was architected with the ability to support all three classes of procedure supported by HDLC but, to date, IBM has chosen to implement only one (unbalanced normal mode).

Another example is token ring, which IBM fought strenuously and successfully in the early 1980s to have adopted by the IEEE 802 committee, which would have preferred to select Ethernet as the single LAN data link standard.

A more recent example is the APPC/LU 6.2 specifications, which were proposed to the ISO transaction processing (TP) working group and which, again with modifications, have been accepted as the OSI/TP standard. OSI/TP is expected to be formally accepted into IS status in mid-1992.

Similarly, IBM is working with APPN in the ISO and ANSI committees dealing with the transport and network layers. It characterizes this involvement as being willing to assist the bodies as requested with parts of the company's APPN work (both completed and underway). IBM is not submitting APPN as a whole for standardization. This is in part because it is unlikely that an architecture of such size and complexity would be considered and in part because that would be equivalent to publishing APPN, which IBM would much prefer not to do (as discussed below). Rather, IBM is specifically proposing selected elements of APPN (not current APPN but the forthcoming enhanced and fast-packet versions) to ISO and ANSI.

These APPN elements most likely relate to connection-oriented protocols. For example, they may concern some major reworking of X.25, which is the current OSI connection-oriented network protocol (CONP). (OSI also has a connectionless network protocol (CLNP) which was developed largely from TCP/IP.) The company probably also has ideas to offer regarding using the OSI IS-IS standard which, so far, is used only to support connectionless protocols (e.g., IP, OSI CLNP, DECnet) but, in principal, could support connection-oriented protocols.

IBM has stated that it is most involved, with regard to APPN in the standards bodies, in developing the support that is needed to exploit very high speed transmission, which carriers are expected to offer by the mid-1990s. For these gigabit and multigigabit networks, IBM believes that even OSI IS-IS and current APPN are inadequate. *SNA Perspective* expects that this work would involve development of a new transport/network layer standard for very high speed networks which can support both connectionless and connection-oriented protocols. IBM's proposals in this area are also likely part of its fast-packet APPN.

APPN: Publish or Perish?

IBM published the APPN end node in March 1991 at the same time it announced APPN for the 3174 and the PS/2. APPN end node lets end systems access the APPN network, but does not provide the technology for routing. This is roughly equivalent to OSI end system to intermediate system (ES-IS) routing but not IS-IS routing.

The company had decided at that time not to publish APPN network node in its entirety. However, in response to strong concerns voiced by users, vendors, and consultants alike about this decision, IBM announced three months later, during its June 1991 announcements which focused on the theme of openness, that it was "reconsidering its ability to publish APPN network node." At the time, IBM said that it would make this determination within a few months; however, the company has not yet announced its decision. Though expected by many, *SNA Perspective* believes that IBM will not make an announcement regarding this APPN network node publishing controversy at the announcement of its multiprotocol router this month.

Licensing APPN Network Node Probable

However, since October 1991, senior IBM staff have been discussing with users, other vendors, and consultants the possibility of licensing APPN network node. IBM is unlikely to be discussing the idea so openly if it were not very likely to move in that direction. The discussions are probably to gather feedback about packaging and pricing as well as to keep users interested in an APPN migration path in the interim. *SNA Perspective* does not expect an announcement from IBM on this subject during the first half of 1992.

Licensing would allow IBM to make a return on its significant investment in APPN architecture and product development while encouraging APPN market development by supporting other vendors' participation. *SNA Perspective* believes the licensing would be in the form of source code rather than documentation, both saving IBM the effort of documentation and saving vendors the effort of developing and updating the APPN code.

Patents for APPN

IBM has received many patents and has applied for others on several components of APPN, which the company has always done for technologies it believes have value. With an architecture as complex as APPN, it would be impossible to patent it as a whole. However, these patents serve to support the company in charging license fees and prosecuting vendors who use but do not license APPN.

Support from Other Vendors

Though licensing is a far better alternative than not allowing other vendors to support APPN network node at all, *SNA Perspective* believes that IBM is still taking a risk by not publishing APPN.

- **OSI Isn't Selling.** Vendors have invested in OSI for several years and have yet to see a return on that investment, even though it is an open, published standard.
- **TCP/IP Is Hot.** TCP/IP and the routing protocols that support it, such as OSPF, are gaining significant market share and many vendors are moving from their proprietary protocols to jump on the TCP/IP bandwagon.
- **Multiple Protocol Stacks Are Hard To Support.** The effort involved in supporting several protocol stacks is significant for small and large vendors alike.

These factors make vendors understandably cautious about taking on support of yet another architecture, especially one that is not completely open and free (as OSI is) *and* does not already have a large installed base (as TCP/IP does).

On the other hand, there are several positive signs:

- **Early Vendor Support.** Several significant vendors have already stood up in support of APPN, including Apple, Novell, Siemens, and Systems Strategies.
- **SNA Migration Potential.** The majority of enterprise backbones today are SNA. APPN is designed, in part, as a migration path for SNA networks. If a significant percentage of these users commit to a transition to APPN, then potential future revenues are enormous. Vendors are watching these users to see whether they turn to APPN, OSI, and/or TCP/IP.

(continued on page 20)

Architect's Corner

Industrial Strength Standards

or, Why De Facto Standards Are A Good Thing

Prediction: 1992 is going to be viewed as a year of standards for IBM, both de jure and de facto (or as someone appropriately suggested, de ibmo).

Contrary to historical popular perception, IBM does take standards seriously. Not only international standards per se, but standards that succeed in the marketplace. IBM's de facto track record is reasonably successful (but not without a few notable failures—who remembers PC Network?). Its track record in international standards is improving, certainly if measured by "head count." But, in the competitive world of standards overall, how does one define success?

Defining Success in Standards

First, a model.

In a recent speech, Craig Burton, principal of the Clarke-Burton consulting firm and former high-ranking Novell executive, argued for a two-dimensional model for standards (see Figure 5). The first dimension is the degree of openness, ranging from closed to open; the second dimension is the communal design center, ranging from proprietary to public domain. Common perception would argue that the ideal standard is open and is designed in the public domain. Experience would argue otherwise.

An observation: The success of a standard is measured by its degree of implementation.

- Who implements standards? Product vendors.
- Why do they implement standards? To expand market share.
- If no vendor has implemented the standard, why? There is no market for it.

Implementation Before Standard

Conclusion? In order for a standard to be successful, it must be implemented before it becomes a standard. This sounds like a chicken-and-egg problem!

Note: Yes, there are exceptions. (TCP is notable in this regard.) Sometimes a standard does precede implementation (actually, in TCP's case the two are concurrent). But, on balance, there are many more proposed "standards" than there are successful product implementations of them.

So if degree of implementation is the definition of success, who are the real standards winners? Simple. Those vendors who, based upon achievement of market dominance with proprietary solutions, open their solutions to other vendors and international standards bodies (in that order).

Is this heresy?

Look at the experience of Novell. The world leader in file servers (based on installed base); began with a closed proprietary solution; is now publishing it and working successfully with other vendors to migrate it to all platforms.

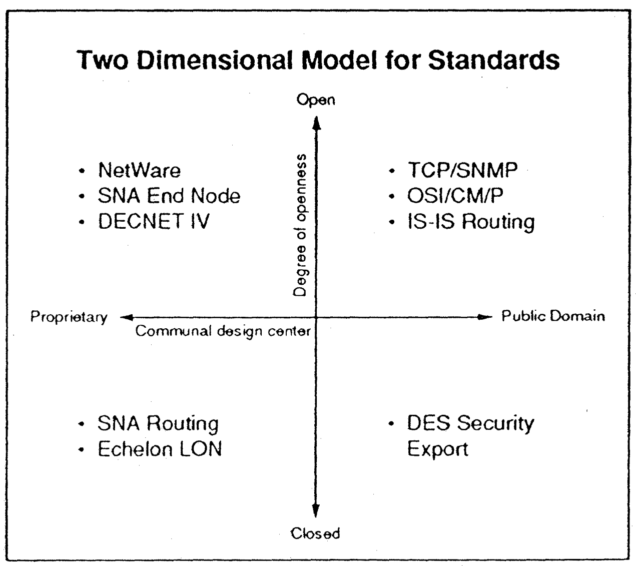

Figure 5

Look at the experience of IBM with SNA. The world leader in large networks; the company began with a proprietary architecture; subsequently published it; other vendors followed suit with implementations.

IBM's SNA and Novell's Netware have probably the world's largest installed bases of workgroup and enterprise networks, respectively.

One other observation (Figure 6). As a technology migrates from proprietary to public domain, it invariably loses function based upon the process of compromise—what remains is the lowest common denominator. Thus, when asked to choose between a proprietary and a "standard" solution, the user must often choose between greater and lesser functionality. (Network management is a real demonstration of the loss-inducing character of the standardization process.)

Based on the above analysis, the vendors that must drive the standards are those that have previously succeeded with proprietary solutions. Further, on the way toward public domain standardization, these proprietary designs must be opened by those vendors and made accessible to other vendors.

How is IBM doing with respect to standardization? Pretty good in my book. But there's still plenty of room for improvement.

Token Ring

Example, token ring. IBM was successful in influencing adoption of the basic token ring architecture by IEEE (802.5), but unsuccessful in adoption by IEEE of its extended network management features. Nevertheless, these extended management features are openly published, licensed to chip suppliers, and further, are a common market requirement in multivendor products in the token ring marketplace.

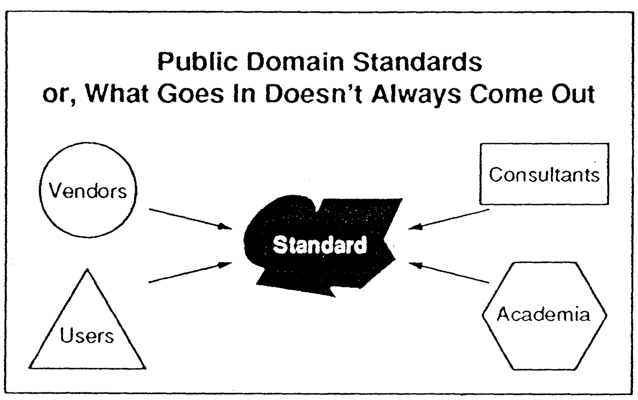

Figure 6

APPC

Example, APPC. APPC has been open from its inception. Novell, Apple, DEC, HP, etc., all are implementing it. IBM submitted APPC as a contribution to OSI, and, while it changed significantly in the process, the resultant Distributed Transaction Processing (DTP) protocol bears remarkable resemblance to APPC. And IBM has said it will upgrade APPC to become DTP compliant.

SNA Network Management

Example, SNA Network Management. Network management has always been considered IBM's competitive edge. Nevertheless, SNA Management Services (MS) was (belatedly) published. SNA/MS support exists both within vendors products and through "gateway" products such as NetView/PC. This story takes an interesting twist, however. Recognizing the benefits of public domain standards alignment, IBM has chosen CMIP as the future evolution for SNA management; has been very active in the CMIP forums—e.g., ISO, OSI Network Management Forum, IEEE 802.1B. Oddly, CMIP, which has not been a marketplace success to date, may turn out to succeed precisely because of IBM's decision to deploy it on a wide scale. And the common CMIP MIBs will be, you guessed it, those proffered by IBM (SystemView)—further demonstrating that standards are made or broken by the installed base.

By now, I'm sure I've seriously challenged those who promote open public domain standards. But don't misunderstand my position. I, too, am in favor of such standards. But not in the absence of market implementations. IBM not only has a duty but also an opportunity to influence standards. First, by opening its architectures to vendors and then by cross-fertilization of technology with public domain standards forums—ISO, IETF, ANSI, ECMA, IEEE, etc.

IBM's big test, in this regard, will be how it handles multivendor access to its SNA routing "standard." In which quadrant will SNA routing emerge? 1992 should be a telling year. ■

(continued from page 17)

Conclusions

Peer networking support is important to SNA users. IBM is working to convince them that it would best come from APPN. However, although leading users are quickly moving from PU 4, subarea SNA will not disappear. IBM must offer sufficient support for subarea SNA in APPN to make it more attractive to customers than switching to other architectures entirely. It must also enhance SNA to act as a backbone for other protocols, including APPN, OSI, and TCP/IP.

IBM is trying to stave off SNA defections to other architectures, in the face of APPN's long development cycle, by being more open about its long range strategy than it usually is.

IBM is also moving through standards bodies to have parts of APPN adopted as standards and to make sure that other standards will support APPN as well as other architectures. *SNA Perspective* applauds IBM's cooperation in the standards arena.

IBM is not altruistic; it is a business established to make money. Therefore, although it could publish APPN network node, *SNA Perspective* believes IBM will more likely license it, defending the license rights with patents for as much as possible of APPN.

We believe that license fees may hamper APPN's market acceptance, unless the fees are minimal and are coupled with early, quality releases, vendor and user development support programs, and other incentives. Therefore, *SNA Perspective* expects IBM will eventually be forced by market pressures to publish APPN network node, probably during 1994. ■