

Evolution of an open communications architecture

by R. J. Cypser

An overview of the current IBM communications paradigm for interconnecting computer networks is presented. Emphasis is on the incorporation of multiprotocol, multivendor facilities in an integrated architecture. This paper presents an overview of key elements of the evolving communications architecture.

The world steadily becomes more interconnected, more interdependent, more aware, and (hopefully) more cooperative on a global scale. In such an environment, almost every part of an organizational enterprise benefits from the assistance of computers. Parts of an enterprise become directly or indirectly linked to the other parts via networks of computers. We are witnessing a gradual networking of the thought resources of humankind. An ambitious vision of the future has been simply stated, as follows:

It is the vision of a world on-line, any to any, with instant responsiveness, not limited by bandwidth or proprietary considerations—a world in which the network conforms to our work habits rather than the other way around, with people manipulating images as easily as they manipulate words and data today—and all with continuous availability, high reliability, tight security, and automatic network management.¹

The architecture for such an evolution to near-universal networking must strive for commonality through international and industry standards.

It must also accommodate heterogeneity and multiple subarchitectures. There is simultaneous extension and convergence. The architecture, therefore, is a living entity, large in scope, but integrating all the vital components in meaningful ways.

This paper attempts to present a concise overview,² from a communications perspective, of key elements in the architecture for open enterprise distributed systems. The underlying thesis is that an evolution is proceeding rapidly in complementary sections. The result is not homogeneity, but a mixture of technical advances and convergence to standards. Although the process is technology-driven, it represents nevertheless an evolving architecture for effective cooperation of heterogeneous elements. The following key features of this evolving process are tools for achieving unity amid diversity:

- Common application services, such as distributed databases, office services, and facilities for multimedia information interchange, designed with common directory, recovery, time, and security facilities, and usable with multiple communication protocols

©Copyright 1992 by International Business Machines Corporation. Copying in printed form for private use is permitted without payment of royalty provided that (1) each reproduction is done without alteration and (2) the *Journal* reference and IBM copyright notice are included on the first page. The title and abstract, but no other portions, of this paper may be copied or distributed royalty free without further permission by computer-based and other information-service systems. Permission to *republish* any other portion of this paper must be obtained from the Editor.

Figure 1 Potential service switch points at the API, transport, and link services boundaries




Figure published with permission of Addison-Wesley Publishing Co. (see Reference 2).

- Multiprotocol system management of integrated systems, involving local area, metropolitan area, and wide area networks
- Application layer gateways for interoperation among applications (like messaging) that were designed for different communication protocols
- Common application programming interfaces,

providing programmer independence from communication protocols

- Common multiprotocol naming and addressing structures within international standards
- Transport service switching that allows applications to run on networks other than the ones for which they were designed

- Bridges and multiprotocol routers that allow interconnection of networks with similar and dissimilar transmission services
- Automated network services, like dynamic topology update, optimum route selection, and user registration, that reduce training and management costs
- Common transmission facilities like local area networks (LANs), X.25 packet switched data networks (PSDNs), Broadband Integrated Services Digital Networks (BISDNs), and frame relay, which serve all the higher-level protocols of Open Systems Interconnection (OSI), Systems Network Architecture (SNA), and Transmission Control Protocol/Internet Protocol (TCP/IP); and strong emphasis on the use of high-speed transmission facilities, occasioned by the widespread use of optical fiber
- Fast-packet³ integration of voice, text, image, video, and data

A multivendor coordination of this evolution to heterogeneous interoperability is being aided by organizations such as the Open Software Foundation (OSF), OSI/Network Management Forum, X/Open Company, and Corporation for Open Systems (COS).

Potential switch boundaries

Nondisruptive growth requires an architecture that facilitates the incorporation of new facilities and the interoperation of new and old facilities. Investments in application processes and supporting application services (or subsystems) must be protected. These must continue to operate on new systems with different architectures and new communications facilities. Different transport service providers must smoothly interoperate with one another. The three potential switch boundaries that offer the best opportunities for such future function sharing and interoperation are shown in Figure 1. These are:

- *Application programming interfaces (APIs)*, at the presentation services boundary, which serve the application processes and the common application services. The ultimate purpose of this boundary is to make the application process and supporting application services independent of the system architecture and independent of all the underlying communication protocols. This is today the Systems Application Architecture* (SAA*) Common Programming Interface (CPI).
- *Transport services (TS) boundary* in the neighborhood below the application-enabling layers (OSI layers 7, 6, and 5). The ultimate purpose of this boundary is to give the application, with its end-to-end data-exchange facilities,⁴ the opportunity to use alternative transport service providers. Those programs that are written to a common transport level API also potentially obtain independence of underlying transport service protocols. The X/Open** Transport Interface (XTI)⁵ is near this boundary.
- *Link/subnetwork-access (LSA) boundary* below the networking-oriented layers (layers 4 and 3c).⁶ The ultimate purpose of this boundary is to make all the lower-layer link and subnetwork services (such as LANs, X.25, BISDN, and frame relay) commonly available to the protocol stacks, layer 3c and above.

Within this architecture, an enterprise (or portions thereof) may use all OSI protocols, all TCP/IP protocols, all SNA protocols, all NetBIOS (Network Basic Input/Output System) protocols (or other LAN-oriented protocols), or any combination of these. The architecture, then, seeks commonality, particularly in management and control functions, common application services, shared use of common subnetworks, and common programming interfaces.


The use of these three switch boundaries has the potential of providing the simplification of and the preservation of investments in application processes and application services; varying amounts of transport protocol coexistence and interoperation; controlled stages of migration among the dominant transport service provider protocols; and the staged use of advanced technologies.

Network evolution

The evolution of networks in the present decade focuses on four main thrusts:

1. Interoperability of diverse workstations and mainframes, unhindered by specific transport protocols
2. Sharing of transmission facilities among multiple, independent, coexisting layer-3 protocols
3. Consolidation of protocol diversity so as to improve manageability and efficiency of operation

Figure 2 Evolution to an open systems network architecture



4. High-performance and multimedia (voice, text, image, video, and data) capabilities

The inclusion of the fourth thrust affects the direction taken by the first three.

A prerequisite to any evolution plan is the assurance of the preservation of the customer's return on investment. This applies particularly to the investments in applications and in terminals and/or workstations. The evolution, accordingly, must be able to proceed in orderly steps that preserve much of the existing systems. Moreover, the time schedule for staged deployments must be tailorable so as to maximize the value of customer information. Finally, the cost of application development and network operations must be reduced in the evolving networks.

An overview of the SNA evolution to a true open systems network architecture⁷ is sketched in Figure 2. The current phase is that of Advanced Peer-to-Peer Networking (APPN), along with TCP/IP, OSI, and LAN protocols like NetBIOS, Internetwork

Protocol Exchange/Sequential Packet Exchange (IPX/SPX), and AppleTalk^{**}. This stage is characterized by the integration of LANs with wide area networks (WANs), the emphasis on ease of use, network dynamics for topology changes and user location, and full peer-to-peer operation.

APPN, in turn, is a stage in the evolution to the next phase, that of high-performance fast-packet switching with full multimedia capabilities, integrating voice, text, image, video, and data. Fast-packet wide area networks, in concert with local area networks (LANs) and metropolitan area networks (MANs), will provide the high-performance, multiprotocol, multimedia networks that are needed. Multiple layer-3 protocols (TCP/IP, SNA, OSI, IPX, or AppleTalk) may continue to exist, but consolidations will occur.

The transition to this last phase may involve a series of concurrent or overlapping steps. These may include the following:

- Integration of both common application serv-

Figure 3 Evolution of the Common Programming Interface for Communications (CPI-C), common application services, and common link/subnetwork-access facilities





Figure published with permission of Addison-Wesley Publishing Co. (see Reference 2).

- ices and subnetworks so as to be common to multiple layer-3 protocols
- Coexistence of multiprotocols in both LANs and WANs, using multiprotocol routers
- Operation of applications across transports other than those for which they were designed (e.g., sockets applications over SNA)

- Interoperation of end systems, by using gateways between concatenated transport protocols (e.g., between NetBIOS and SNA transports)
- Upgrades of performance, by taking advantage of very large scale integration (VLSI) and high-speed transmission facilities
- Transmission with bandwidth management fa-

Figure 4 IBM 6611 Network Processor configurations




cilities using multimedia (voice, text, image, video, and data)

- Gradual consolidation of protocols in each segment of the high-performance, multiprotocol, multimedia network, so as to make one or more protocols dominant
- Management and training simplifications by the reduction of multiple configuration, routing, and recovery systems

Integration

In the IBM integration process, SAA systems (i.e., Operating System/2* [OS/2*], Operating System/400*, [OS/400*], Multiple Virtual Storage [MVS], and virtual machine [VM] systems) and AIX*/OSF (Advanced Interactive Executive*/Open Software Foundation) systems⁸ are primarily involved, but interoperation with other equip-



ment manufacturers (OEMs), is also greatly facilitated. To illustrate, an idealized picture of the growth of commonality at various levels, for the combination of SNA and OSI, is that shown in Figure 3. At three levels, there is a gradual expansion of the commonality with OSI and SNA. In general, we have first a selection of optional *application services*, such as distributed file services and distributed database services, a growing number of

which could be commonly used with TCP/IP, OSI, and SNA communications. Second, there is a *Common Programming Interface* for Communications (CPI-C) that can be independent of the underlying TCP/IP, OSI, or SNA facilities. In the TCP/IP area, the X/Open Transport Interface (XTI, at the transport services boundary) provides an API that is common to both TCP and OSI transport facilities. IBM has proposed to X/Open an architecture for the extension of XTI so as to make these programs usable on SNA and NetBIOS transport services as well.⁹ Third, there are *common transmission facilities* that can service either the OSI or SNA end-to-end data-exchange facilities. All of the OSI link/subnetwork-access facilities can be common. Subnetworks, like X.21 circuit switched networks, X.25 packet switched data networks (PSDNs), synchronous data link control (SDLC) wide area networks, multiple LANs, frame relay, and BISDN, would be available to all applications, whether they used OSI or SNA option sets in the upper layers.

In the LAN area, applications written to the NetBIOS end-user interface (NetBEUI) can run across NetBIOS or IPX transports.

Equal coexistence

Network coexistence is the concurrent use of multiple layer-3 protocols over the same physical media, but where each layer-3 logical network is completely separate from the others. For example, an end system designed for attachment to a TCP/IP network can communicate with a like end system, but not with an end system designed to work with SNA or OSI transport protocols.

The IBM 6611 Network Processor, for example, is a multiprotocol router that allows the equal coexistence of SNA, TCP/IP, IPX/SPX, XNS, DECnet**, and AppleTalk layer-3 protocols. Figure 4 illustrates some of the many configurations of multiprotocol networks that are now possible. The media between network processors may be point-to-point T1 (1.544 Mbs transmission rate) links, X.25 networks, or frame-relay networks. The common approach in such network processors is to have independent, parallel routers, one for each protocol. Bridging and other data link connections are, however, also included. In the early releases of the IBM 6611, SNA and NetBIOS are encapsulated in TCP/IP. Reportedly, in a later phase, full SNA routing will be added using SNA/APPN. In

any case, such "ships in the night" approaches share transmission media among multiple protocols, but preserve separate logical networks and do not provide for communication among them.

Any application over any transport

Often the network user would like to write an application once and have it run on different networks, and to create a logical connection across heterogeneous networks that matching applications can use. We need to improve the return on investments in applications by simply using them

A critical ingredient of the multiprotocol network is a comprehensive and efficient system management process.

on more networks. Similarly, we need to be able to use existing workstations with more applications on more networks.

Existing wide area networks, using TCP/IP, SNA, or OSI protocols, are called single-protocol networks because they do not carry other layer-3 protocols without some form of layer-3 encapsulation or transport service conversion. Applications that run on one such network usually cannot run on other networks. However, transport service switching offers considerable promise for such network interoperability in key situations.

Transport service switching adds the ability for an application designed to work with one transport service (layers 1-4) to communicate with a like application using a different transport service. Layers 5-7 must still match.

This interoperability concept is illustrated in Figure 5. There, a pair of workstations, on any one of three different networks, uses different types of applications via a single transport protocol. For example, with such workstations on each network, one could have interoperation among

TCP/IP applications on SNA/APPN networks, SNA applications on TCP/IP networks, NetBIOS applications on TCP/IP networks, NetBIOS applications on SNA/APPN networks, and OSI applications on TCP/IP networks. Still other combinations can be developed for various applications over OSI, SNA, and TCP/IP transports.

By the addition of an internetwork transport service switch (or gateway), shown in Figure 6, these same applications could all interact with like applications across any pair of these networks. Workstations could then be used for more applications on more networks. In this scenario, existing applications would still be usable, as consolidation to fewer transport protocols proceeds, because of the use of transport service switching facilities in end systems.


Architectural studies have indicated the feasibility of such a high degree of interoperability without requiring uniform protocols and without requiring replacement of installed networks. IBM has already indicated that applications written to the UNIX** sockets specification will run on SNA subarea and APPN networks. Market demand, of course, will determine which of these other many combinations will also be provided.

In the following paragraphs, we examine the facilities in each of the layers of Figure 1, starting from the top and working our way down, and provide more specific illustrations of network evolution. This will provide the basis for finally discussing the evolution to the fast-packet multiprotocol network.

System management

An obviously critical ingredient of the multiprotocol network is a comprehensive and efficient system management process. The management process must apply to systems using a variety of communication protocols, including OSI, SNA, TCP/IP, and LAN-oriented protocols. Each of these protocols provides a somewhat different set of management information. For example, SNA alerts, TCP/IP traps, and OSI events must all be "caught" and processed. Even with the transition to the OSI Common Management Information Protocols (CMIPs), an interim accommodation of existing protocols is needed. A common management database drawn from multiple sources is the goal.

Figure 5 Any application on any transport



The IBM system management model allows hierarchical and distributed system managers, which may be determined by customer organization. For example, one could have management entities, in four tiers, as illustrated in Figure 7. These are the node agents, regional system managers (such as a manager of a LAN), domain system managers that have responsibility for a portion of the network (such as LANs and WANs), and the

enterprise system manager where multidomain decisions are made.

The system design must provide the appropriate portion of management applications at each management level in a distributed system. For example, fault management may be primarily at a lower level, with exceptional situations referred to a higher level. Also, a single logon may result in

Figure 6 Interoperation with single-protocol transport networks




Figure 7 Possible management structure, showing four tiers of system management




Figure published with permission of Addison-Wesley Publishing Co. (see Reference 2).

authentication assists over an entire domain of the network.

The expectation is that the OSI efforts in network management will ultimately provide the common denominator blueprint for the next stage. While interim management systems are necessary, the OSI object definitions and CMIP protocols are expected to be more effective and less expensive in the long run. This is partially because the generality of the OSI architecture avoids the necessity of creating and adding new modules for each new device. Other CMIP assets include emphasis on event-driven objects, rather than requiring a manager to poll an agent, rich object control func-

tions, the ability to apply a command to a large number of objects, and good automation potential.¹⁰

The systems management process is thus the integration of a wide variety of management applications that are part of network management facilities from multiple vendors and use multiple communication protocols.

Common programming interfaces

The SAA Common Programming Interface is a key aspect of SAA application services. There are distinct elements in the CPI, including the elements

for transaction communications, file transfer, distributed database, resource recovery, and information presentation. Moreover, some application services can both provide a component of the CPI and in turn use the CPI offered by some other application service. One can foresee the Common Programming Interface for Communications (CPI-C) to:

- Be consistent across all SAA and AIX/OSF computing environments and others that will implement the CPI
- Be accessed at least via all SAA high-level languages
- Enhance the portability of applications across all SAA and AIX/OSF operating environments, and others that support the CPI¹¹
- Offer network independence, as seen by the applications, for TCP/IP, OSI, and SNA communications subsystems (and any other protocols that might similarly be supported under the CPI with appropriate mappings);¹² this simplifies and preserves the usability of a wide range of programs across multiple communications service providers

Currently, CPI-C provides only conversational support. The known requirement is to support conversational, remote procedure call (RPC), and messaging APIs over a range of transports, with common directory, recovery, and security facilities.

At the transport services (TS) boundary is the X/Open Transport Interface (XTI). This provides usability of programs that are written to this lower transport API, across both OSI and TCP/IP transport services. As mentioned above, this usability could be extended to other transports as well.

Distributed data

Distributed information is becoming a practical reality, with the ability to transparently use both Structured Query Language (SQL) data and file data from local and remote sources, on LANs and on WANs. Today, as part of the IBM Information Warehouse* concept, the Enterprise Data-Access/SQL Server enables users to transparently and interactively access multivendor relational and nonrelational data from workstations and terminals, using TCP/IP or SNA.¹³

The sophistication of data service is illustrated by the IBM plan for the progression of styles from one

SQL statement on one remote database, to multiple SQL statements per unit of work with multiple databases, and still later adding multiple databases per SQL statement.¹⁴

Data integrity is, of course, essential, as data are distributed across the network. The two-phase commit procedure enables all parties to a distributed operation to gracefully return to their original state if any participant runs into difficulty. This procedure, incorporated into the Common Programming Interface (CPI) and supported in SNA's logical unit (LU) 6.2, proves to be a relatively simple way of insuring data integrity across the network.

The continuing direction is to allow the use of multiple communication protocols to permit transparent access to heterogeneous relational and nonrelational data on multivendor platforms, application access to any file data anywhere, and the ability to print and view any document anywhere.

Gateways at the application layer

When common multiprotocol application services are not yet available, gateways at the application level can be used. These are exemplified by the X.400/PROFS*/DISOSS (Professional Office System/Distributed Office Support System) interchange between OSI and IBM messaging systems. The IBM Open Network Distribution Services (ONDS) product is illustrated in Figure 8. Acting as a special Message Transfer Agent (MTA) within the X.400 Message Transfer System, the ONDS provides connection services to proprietary electronic mail systems.

ONDS uses the OSI/Communications Subsystem for the communication via X.25 or LANs to other OSI-based nodes. Users can exchange messages, through the ONDS gateway, with users attached to other X.400 systems on the network and users attached to PROFS or DISOSS systems, via PROFS or DISOSS connections.

In addition, the multiprotocol, application level mail switch made by Soft Switch is modular so as to accommodate the electronic mail systems of different subnetworks. Soft Switch's Central supports access to and from systems that include: DEC ALL-IN-1, Novell MHS, Dialcom, HP Desk-Manager, MCI Mail, Banyan Mail, Wang Mail-

Figure 8 Using a special MTA and user agent to connect DISOSS or PROFS to an X.400 message handling system




Figure published with permission of Addison-Wesley Publishing Co. (see Reference 2).

way**, Data General CEO, and 3COM 3+Mail, in addition to the IBM mail systems.

Universal directory

A universal directory service can also be foreseen. The X.500 standard's universal naming is a hierarchical structure, with naming authorities at each level and the ability to associate parameters with each name. In this directory, all items of interest can be given a hierarchical, globally unique distinguished name. This system can incorporate SNA, OSI, and other protocol naming methods (e.g., SNA LU names and transaction program names), so all applications can easily reference other applications on different platforms. A worldwide, multiprotocol naming and directory service is inevitable.

Then, the X.500 directory services at the application layer, in combination with dynamic network-layer directory services (such as those of APPN) and LAN discovery protocols, have a clear potential to serve the future mechanisms for destination location and route determination across multiple LAN and WAN subnetworks.

Data exchange facilities

At the next level, we have the end-to-end data-exchange architectures. Three mainline services are projected: conversational; remote procedure call (RPC), which extends the programmer's familiar call procedures to remote locations; and Queued Transaction Messaging (IMS-like), which uses the management of message queues as the primary control mechanism.

LU 6.2 is an implementation of the conversational mode. Advanced Program-to-Program Communications (APPC) is a common term for LU 6.2. LU 6.2 verbs are a superset of the CPI-C calls. LU 6.2 also supports the CPI for Resource Recovery (CPI-RR).

LU 6.2s have become an industry workhorse for conversational, interactive communications. The stable, well-defined verbs of LU 6.2 permit effective and robust interactions among cooperating application processes. These can be distributed throughout the enterprise on a broad variety of different systems, from different vendors, from workstations to mainframes.¹⁵

An example of APPC used in a client/server mode is the APPC/MVS server. In concert with the SAA Common Programming Interface CPI-C, the application developer uses program name references for transparent application-to-application connectivity. There are two distinct styles: (1) multiple instances of the server (one for each requesting application), and (2) single server instance that handles multiple concurrent requesting applications, where the operating system can handle multiple concurrent tasks.

APPC/MVS schedules conversations to multiple server applications in response to client's connection requests. Each server has access to MVS data and services. Each client may use multiple concurrent servers that are all different, or multiple instances of the same application, or a server within which there can be multiple transaction programs. Each server can also initiate conversations to other applications on the same or different systems. APPC/MVS is thus an example of the way in which client/server concepts can help to integrate the facilities and advantages of workstations and mainframes, on LANs and WANs.

APPC is also widely used on the Application System/400* (AS/400*), OS/2, and DOS.¹⁶ On the AS/400, client/server functions running over LU 6.2 include: bytestream file serving, SQL serving, print serving, a query server for flat files, data queues (a read/write interface for interprocess communications), a remote command facility, IBM 5250 terminal emulation, and messaging. The AS/400 also has Distributed Data Management (DDM) record-level clients and servers, as well as Distributed Relational Database Architecture (DRDA) clients and servers.

OSI/Communications subsystems

The IBM OSI/Communications Subsystem includes not only certain application services and the data exchange facility (layers 5-7) but also the transport and network layers below that. This subsystem can be inserted into many different system environments.¹⁷

Not every node needs to have such an OSI subsystem for communicating with other OSI systems. As shown in Figure 9, a remote API is provided whereby, through the services of an

LU 6.2s have become an industry workhorse for conversational, interactive communications.

intervening SNA network, using LU 6.2 sessions, an OSI application in one location, A, can use the OSI/Communications Subsystem at a different location, B, to talk to locations C and D. The SNA network thus carries both the OSI and SNA traffic on shared lines. This is another special case where one API (the OSI API) is made to be usable for more than one communication service provider (both OSI and SNA).

Transport service providers

Proceeding down to layers 4 and below, the primary transport service providers are in four major protocol "islands" at layers 3 and 4, with OSI, TCP/IP, SNA, and NetBIOS (or other LAN-oriented) protocols. Many other proprietary protocols, such as DECnet, Novell's IPX, Xerox's XNS, and Apple's AppleTalk may also need to be integrated in an enterprise information system.

TCP/IP continues to be upgraded. Recent advances include the Open Shortest Path First (OSPF) routing protocol, which will more intelligently select an optimum route, based on more complete information on alternate paths.¹⁸

Figure 9 Local and remote OSI applications using the OSI/Communications Subsystem




Figure published with permission of Addison-Wesley Publishing Co. (see Reference 2).

Each of the above transport service providers may involve subnetworks within it, operating at layers 1-3a. Figure 10 shows five sets of layer 1-3a subnetworks: SDLC WAN, X.25 PSDN, LANs and MANs, ISDN and frame relay, and lastly a subnetwork with bandwidth management between

voice and data. Switched Multimegabit Data Service (SMDS) and Asynchronous Transfer Mode (ATM) in Broadband ISDN (Integrated Services Digital Network) are also included. Still another subnetwork is that of mobile communications.¹⁹ Each of these can deliver information from one

Figure 10 Target capability of general interoperability




Figure published with permission of Addison-Wesley Publishing Co. (see Reference 2).

end system to another end system. Moreover, each may be used as a subnetwork within another network. All of these subnetworks are largely independent of the higher layer protocols and standards (layers 5-7) used for end-to-end communi-

cation. All of these subnetworks are currently or potentially part of SNA, OSI, and TCP/IP networks. It is feasible, therefore, for all of them to be further integrated into an overall, multiprotocol, OSI/SNA/TCP/IP network architecture.

APPN

At layer 3c, Advanced Peer-to-Peer Networking (APPN) is an architecture extension to SNA²⁰ that provides the base for easy-to-use and high-performance wide area networks of the future, which in turn can incorporate any of the above subnetworks. Currently, APPN services are used for two related purposes: (1) to reduce the need for coordinated network definitions, thereby simplifying network installation, additions or deletions of nodes, and reconfiguration; and (2) so that end systems can enter a network, become dynamically registered, and communicate across the network without the aid of any remote system services control point (SSCP).

Key attributes of APPN networks include:

- Dynamically updated network topology database services for automatically recording network topology changes and making these changes immediately known throughout the network
- Distributed, network-layer directory services for automatically registering and locating logical units; for example, the current location of the serving network node can be determined for any remote LU that is identified only by name
- Dynamic route selection services to select the best available route to the remote LU; this selection is based on the continuously up-to-date topology database and a class of service specified by the LU initiating the session
- Adaptive pacing and transmission priority to control the flow of traffic for each session and to jump ahead of lower priority traffic at queuing points in the network

The APPN systems also now offer improved performance in areas like session establishment, directory searches, and taking advantage of parallelism in multiprocessing systems. In addition, APPN significantly provides the architectural direction for later network evolution, such as the incorporation of fast-packet switching and multi-protocol transport, with bandwidth management and congestion control.

APPN will be available on all SAA platforms. Advanced Interactive Executive (AIX) implementations are also expected. The long-desired goal of any-to-any capability, across LANs and WANs, is brought closer by the inclusion of APPN in SAA

Networking Services/2, operating in the Personal System/2* (PS/2*), as well as having APPN in mainframes. This integration of the LAN and WAN layer-3c facilities thus facilitates a broad spectrum of interacting workstations, mainframe hosts, and LAN servers.

APPN and subarea networks

Most of the seven-layer architectural structure is common to APPN and subarea networks. This commonality includes all the upper layers, 5-7, and all of the links and subnetworks (in layers 1, 2, and 3a, including X.25, LANs, Fiber Distributed Data Interface [FDDI], BISDN, frame relay, SDLC, and channels).


Both subarea and APPN networks also stress the ability to have high line-utilization with mixed data types, without impacting interactive traffic. Both include advanced priority and congestion controls.


Now, with the support of APPN in Virtual Telecommunications Access Method (VTAM*), IBM System/390*s, AS/400s, and PS/2s can participate fully in Advanced Peer-to-Peer Networking. A combined subarea and APPN network, with IBM System/390s performing as APPN end nodes and network nodes, is illustrated in Figure 11. For the gradual extension of portions of subarea networks to APPN, with smooth, stepwise transitions, the following options are available:

- Mainframes can have both subarea capability and APPN capability, and are then called interchange nodes. These freely communicate with pure subarea nodes and APPN nodes.
- A mainframe and one or more communication controllers that it manages can together be designated as a composite node. That composite node functions as one APPN network node. Multiple composite nodes then appear to the APPN network simply as multiple APPN network nodes.
- Mainframes functioning as data hosts, without responsibility for communication controllers, can function as APPN end nodes.

Capabilities of the VTAM SNA/APPN thus include: completely meshed subarea/APPN directories; VTAM and the network control programs (NCPs) in communication controllers (like the IBM 3745) together functioning as a composite network node;

Figure 11 SNA/APPN





no impact to subsystems or applications—all APIs fully supported; total support for dependent LUs (e.g., LU 0, LU 2, and LU 6.2); and large network support of directory servers.

The directory server accumulates destination locations for a number of network nodes. This reduces broadcast traffic in the larger networks. If one directory server does not have the desired location, it first queries other directory servers. If that also fails, then a broadcast to all network nodes is initiated.

VTAM permits the addition to the network of intelligent workstations running LU 6.2 applications without requiring definitions to be added to VTAM or the network control program (NCP). Sessions may connect mixed pairs of LUs (one independent of a system service control point [SSCP] and one dependent on an SSCP). The independent LU can also initiate the session. This “self-definition” of devices will also be extended to include IBM 3270 displays attached to an IBM 3174 cluster controller on APPN as well as subarea networks.

The direction for dependent LUs, such as LU 2, is to allow them to connect to mainframes via any path in the APPN network. To achieve this, it is a requirement that serving network nodes will possess the subarea type of boundary function.

The SNA subarea networks are thus also participating on five evolutionary tracks: expanded application services, end-to-end system management, ease-of-use, incorporation of multiple communication protocols, and use of high-speed facilities. Some current examples are:

- Ease of use: full APPN network node (NN) and end node (EN) capability within VTAM; dynamic LU definition in the subarea network for independent LUs that are in connected APPN networks or attached low-entry networking (LEN) end nodes; and dynamic attachment of channel-connected control units
- Multiple subnetwork capabilities within subarea networks: X.25 PSDN; LANs 802.3, 4, 5; FDDI; and serial optical channel (ESCON*, the Enterprise Systems Connection*)
- Use of the subarea network as a backbone with multiple-protocol passthroughs for: X.25 end

systems and PSDNs, TCP/IP workstations, APPN networks, and NetBIOS on LANs

The more than 40 000 SNA subarea networks around the world are a highly evolved, dependable set of networks that carry a tremendous workload. They can be expected to continue to evolve and to provide dependable service for the foreseeable future. It is expected that the full ad-

**The addition of fast packets
promises a comprehensive
WAN/LAN/MAN network that is
entirely protocol-independent.**

vantages of further evolutions of APPN will be incorporated in SNA subarea networks in ways that will combine the advantages of each, further promote commonality and interoperability with other transport protocols, and at the same time preserve customers' earlier investments.

Common addressing

The future interoperation of multiple subnetworks that use different layer-3c protocols creates the problem of common addressing among these protocols. Various address-mapping schemes can be devised. Fortunately, however, the OSI addressing approach provides the basis for such commonality, and the SNA addresses can be incorporated in the OSI address scheme.

The approach for network addresses, taken by the International Organization for Standardization (ISO) and the Consultative Committee on International Telegraph and Telephone (CCITT), adds an extra level of hierarchy, so as to encompass the existing numbering, organizational, and geographically-based plans under a single umbrella.

A resulting picture is shown in Figure 12, where the headers are grouped to emphasize the layer-3c intersubnetwork header and the layer 1-3a link/subnetwork-access headers. In the OSI layer-3c header, there is the layer-3c Network Service Access Point (NSAP) address, on which we now focus.

SNA registered resources can be included in the OSI address space as part of an NSAP address. A registered AFI (Authority and Format Indicator) code is used, signifying that the next field, the IDI (Initial Domain Identifier), uses an ICD (International Code Designator) format. The AFI used by IBM is 47 and the ICD designator for IBM has been established as 0018. The network-qualified SNA name (Net_id.SNA_name) is then placed in the DSP (Domain Specific Part) of the address. TCP/IP addresses can similarly be included as a Domain Specific Part of the OSI NSAP.

LANs

LANs are the prime subnetworks now in many enterprises. Remarkably complex combinations of LANs can be interconnected with high performance. The design for the local area network²¹ at the IBM facility at Raleigh is a good illustration. Each floor of each building may be a distinct LAN segment. Dual mini-campus backbone segments (for high availability) then are bridged to each LAN segment. The multiple mini-campus backbones, in turn, are bridged to dual higher backbones. The main reason for the two levels of backbones is for problem isolation and backbone traffic reduction.

That IBM campus has grown to use a token ring bridged network consisting of 57 interconnected rings, supporting 11 000 terminals and 48 host computers. About half of the terminals are IBM 3270 displays and half are programmable workstations. The network appears as a single logical LAN. Users get essentially local response times and access to any host.

Multiple types of LANs now become the distributed access links to mainframes. Figure 13 illustrates a multiprotocol, multioperating system, multivendor configuration, with four types of LANs feeding multiple hosts via the IBM 3172. These LANs may be 4 and 16 Mbps token ring, Ethernet,²² PC network, FDDI, or Manufacturer's

Figure 12 The intersubnetwork and subnetwork addressing and headers




Figure published with permission of Addison-Wesley Publishing Co. (see Reference 2).

Automation Protocol (MAP) 3.0 token bus LANs. These LANs can all be channel-attached to hosts (e.g., AIX/370, MVS, or VM) using a single, shared, interconnect controller. The design of this 3172 bridge is such that multiple higher layer protocols, including TCP/IP, SNA, and OSI can be passed through from LANs to hosts.

Other protocols are also accommodated. With the upgraded IBM 8209 to also handle the IPX protocol, one can have a mixture of Novell and IBM servers, on IEEE 802.3 Ethernets and IEEE 802.5 token rings, which are connected by IBM 8209s. A client workstation then has access to servers from both vendors on token ring networks and Ethernets.

Figure 13 Multiprotocol, multioperating system, multivendor IBM 3172 configuration




Figure published with permission of Addison-Wesley Publishing Co. (see Reference 2).

High-speed LAN-to-LAN

Along with multiprotocol bridges and/or routers (described previously), a growing variety of options will be used at lower layers for LAN-to-LAN backbones. Besides the usual point-to-point dedicated lines between routers, the following high-speed options emerge:

- ISDN at 64 kbps and 1.544 Mbps in the United States, and 2.048 Mbps in Europe
- Frame-relay data interface and associated fast-packet networks at 1.544 Mbps and higher²³
- FDDI—high-speed LANs at 100 Mbps
- Private bandwidth management (time division multiplex and fast packet) at T3 (45 Mbps) and higher
- Metropolitan area networks (MANs) at 45 Mbps and higher
- SMDS (Switched Multimegabit Data Service), providing a backbone service for data-oriented MANs
- Broadband ISDN using Asynchronous Transfer Mode (ATM) at 150 Mbps and higher

All of these high-speed link/subnetwork-access facilities operate at layers 1–3a. All are commonly available to any of the layer-3c and higher protocol stacks, as illustrated in Figure 1. All, therefore, are able to serve as protocol-independent subnetworks. All but frame relay and SMDS look to the inclusion of isochronous traffic, for voice, image, and perhaps video, as they evolve.

The Aurora Gigabit Testbed

A key testbed for fast-packet and broadband service to multimedia applications is the Aurora Gigabit Testbed.²⁴ Its objective is the exploration of technologies that may be appropriate for use in a National Research and Education Network (NREN) operating near gigabit-per-second bandwidths.²⁵

One of the backbone technologies employed in Aurora is that of plaNET (a later version of PARIS²⁶), developed at IBM. This has been targeted to support heterogeneous traffic types, using packets of different sizes, priorities, and routing methods. These include variable-size PARIS packets using source routing, small fixed-size ATM cells using label-swapping routing, and packets being routed in multicast. The latter employs

a spanning tree algorithm, and topology information is broadcast through hardware over the spanning tree. The source-routing packets use Automatic Network Routing (ANR) with hardware assist to achieve very rapid connection setup. The design thus can support both frame relay and ATM standard interfaces.

Bandwidth reservations are guaranteed. After a user agreement on bandwidth is reached, the user is monitored to insure that the agreement is not exceeded. If that occurs, the excess is accepted only at a lower discard priority, and only if extra bandwidth is available.

The multimode plaNET design will also be evaluated in reference to a pure ATM, fixed-cell design conceived at Bellcore.

Aurora includes experiments with actual applications. Configurations will be tested with loads that simulate expected traffic classes such as video, voice and bulk, and interactive data transfer. Performance will be judged, not merely of the network, but of the entire distributed application process, including the network.


Such testbeds will determine the technologies used in tomorrow's high-speed, high-performance, multiprotocol, multimedia transport systems, to which the SNA/APPN architecture is evolving.

The fast-packet multiprotocol network

The combination of the fast-packet addition to APPN, the various LANs (802.3, 802.4, 802.5, and FDDI), and metropolitan area networks (802.6) promises a comprehensive WAN/LAN/MAN network that is entirely protocol-independent. By this, one usually means that all routing within the network is done at layer 2, and any layer-3 protocol and its data can be carried simply as information, without any special layer-3 type of encapsulation.

In the later APPN, illustrated in Figure 2, fast-packet, multimedia transport facilities will perform hardware-assisted routing at layer 2, and will efficiently carry any layer-3 protocol (e.g., TCP/IP, OSI, SNA, IPX, AppleTalk) as data. Customers will be able to consolidate to use fewer protocols (or one protocol), at the rate they desire. The transport's internal routing and topol-

Figure 14 Multiprotocol network and workstation



ogy features will be independent of the end-system interfaces it must provide. It therefore will be adaptable to change in the appearances it must provide. A single routing and topology update system within the network will reduce installation, reconfiguration, and recovery workloads, and reduce routing, topology, and recovery traffic in the network.

Such a protocol-independent network is illustrated in Figure 14. There, the combination of many LANs, all interconnected by protocol-independent WANs, is called the Equal-Coexistence Multiprotocol (fast-packet and LAN) Network.²⁷

A multiprotocol workstation, shown in Figure 14, could then send data using more than one layer-3

protocol. In the example, it could be OSI, SNA, TCP/IP, NetBIOS, IPX, or AppleTalk. The recipients of that data could be workstations on that same multiprotocol network or on three attached networks: SNA workstations on an SNA/APPN single-protocol network, socket workstations on a TCP/IP single-protocol network, or a LAN-oriented workstation on a LAN.

The fast-packet feature might be a frame-relay, ATM, or IBM plaNET variety. In particular, it can be the coming extensions to APPN, which are targeted to be able to use or provide standard frame-relay or ATM (BISDN) services. The number of protocols actually handled by this multiprotocol network would be a customer option.

Full interoperation between the fast-packet APPN version and earlier APPN versions is required. Accordingly, it is expected that it will be possible for the SNA/APPN single protocol transport network (SPTN) to be gradually absorbed within the more advanced version, as indicated by the arrow in Figure 14.

Migration to a target transport protocol. In theory, all or part of the projected IBM fast-packet network could also be consolidated to handle only one target protocol (this, e.g., could be OSI, TCP/IP, or SNA). This is illustrated in Figure 15. All the workstations then would need to support only that one target transport protocol. Preservation of the application investment could then be achieved by the same techniques cited earlier using transport service switching in the end systems. By that means, the end system could serve more than one type of upper layer. For example, a workstation's upper layers could support OSI, SNA, TCP and/or LAN-oriented applications.

Only one common protocol in an entire enterprise is, however, an unlikely or distant prospect for most large corporations, because of the heavy investment in existing networks. Also, the continued use of LAN-oriented protocols will persist in LANs (as well as the more general TCP/IP, OSI, and SNA protocols used in both WANs and LANs). Some pairs of transport protocols (layers 1-4) are thus very likely. The target may very well, then, be coexistence protocols of Figure 14, plus only a partial movement toward that target protocol of Figure 15. For many others, the foreseeable future will also include the stages shown in Figures 4, 5, and 6, as well.

In this environment, the projected evolutionary approach aims to allow each customer to first consolidate only portions of the network, and to consolidate to the small set of protocols that is most appropriate for each portion. Different cus-

**The architecture facilitates
a stepwise movement
toward common services
at all levels.**


tomers will consolidate to different sets of protocols. The approach is thus designed for network stability and durability through a basic flexibility that permits orderly change.

Summary of unification tools

Various approaches are outlined above for the orderly integration of OSI, SNA, TCP/IP, and LAN-oriented protocols like NetBIOS, along with growing subnetworks like frame relay, BISDN, etc. These approaches can fit together in various option sets, depending on the customer network history and desired evolutionary path. A composite of these tools follows.

- Application layer tools
 - Key applications and application services (subsystems) can be common to underlying communication protocols.
 - Common programming interfaces (CPI-C and XTI) can support the above by helping to avoid the necessity of writing programs to multiple communication APIs. They can also help to reduce the number of different communication protocols that must be maintained in the network.
 - Application layer gateways can sometimes provide interoperation between different application services, like mail systems designed for different operating systems.

Figure 15 Migration to a target transport protocol



- Automated, distributed, dynamic network services
- Ease-of-use services, like distributed directories, dynamic topology update, and dynamic route selection, make complex network integration feasible.
- Multiprotocol coexistence
 - Layer-2 bridges can often provide simple, highly efficient, and multiprotocol connections of very large LAN clusters.

- Layer-3c routers can meet the need for a degree of isolation and control among very large LAN clusters, and among LANs whose protocols use extensive broadcast techniques.
- Multiprotocol routers in which multiple layer-3c routers are used in parallel currently add an important capability.²⁸ These routers are particularly advantageous when used in conjunction with subsystems that are independent of network layer protocols, such as LANs,

frame relay, and others cited above. The combination of multiprotocol routers and protocol-independent subnetworks provides an attractive vendor-independent transport.

- Any application over any transport
 - Switching near the transport service level among SNA, OSI, TCP/IP, and NetBIOS is also an attractive option. Then, an upper-layer stack can be switched to one of several lower-layer stacks. For example, by utilizing a transport level switching boundary, TCP/IP-based applications using TCP/IP upper stacks can be connected via an SNA lower-stack network, and vice versa.
 - Transport-switch gateways, then, are an option for interoperation among SNA, OSI, TCP/IP, and NetBIOS transport service providers. For example, it has been shown that gateways (using transport service switching) could interconnect NetBIOS LANs with OSI, SNA, or TCP/IP networks.^{9,29}
- Any subnetwork under any network protocol
 - Link/subnetwork-access facilities, including all the LANs, X.25, BISDN, SDLC, MANs, and frame relay, can be commonly available to all applications and their upper communications protocol stacks.
- Multimedia integration
 - Fast-packet WANs (as indicated in Figure 2) will provide both a protocol-independent transport and the integration of multimedia (voice, text, image, video, and data).

Thus, a very comprehensive family of complementary integration, coexistence, interoperation, and consolidation techniques is available. Combinations of these techniques, under appropriate system management, congestion control, and worldwide naming and addressing conventions, are essential parts of the strategy to address the myriad of circumstances found in today's multi-protocol world.

The above is part of a careful, continuing, evolutionary process, as advancing technology steadily adds new options. The approach involves the acceptance of heterogeneity as a normal and often necessary occurrence. It promotes commonality of language and function as seen by the user. It fosters common application services and cooperation among remote systems of different sizes and architectures. It shields the user from much of the underlying complexity, and provides a consistent set of user operations. The process

finds commonality of boundaries and interfaces at multiple levels. All this has progressed rapidly, so that one can appropriately refer to the development process for an open systems network architecture that has been the subject of this paper.

Acknowledgments

While the work from which parts of this paper were drawn (see Reference 2) acknowledges the help of 64 individuals from many IBM locations, the author would like to specially acknowledge the insights provided by James P. Gray, Phillippe Janson, and Ellis Miller, and the support of John Hunter, director of architecture and technology, and Rick McGee, manager of communications architecture, at IBM. Any errors in the paper, however, are solely the responsibility of the author.

*Trademark or registered trademark of International Business Machines Corporation.

**Trademark or registered trademark of X/Open Co., Ltd., Apple Computer, Inc., Digital Equipment Corp., UNIX Systems Laboratories, Inc., Wang Laboratories, Inc., Microsoft Corp., Information Builders Inc., or Sun Microsystems, Inc.

Cited references and notes

1. Presented as IBM's telecommunications vision by Ellen Hancock, Vice President and General Manager of IBM Networking Systems, at COMNET 89, in Washington, D.C., February 8, 1989.
2. This paper includes material, with permission, from the author's recent book, *Communications for Cooperating Systems*, published and copyrighted 1991 by Addison-Wesley Publishing Co., Reading, MA. Figures 1, 3, 7-10, 12, and 13 published with permission.
3. The term *fast packet* in this paper refers to (1) data-only frame relay, (2) multimedia Broadband ISDN (BISDN), using Asynchronous Transmission Mode (ATM), or (3) a combination of these two.
4. The term *end-to-end data-exchange facilities* refers to the upper-layer programming support at each end system that helps to manage the conversation or dialog between two end users.
5. The XTI is an interface that is close to the older sockets and Transport Layer Interface (TLI) interfaces that have been common in the TCP/IP community. It is also close to the NetBIOS end-user interface (NetBEUI).
6. The OSI network layer is subdivided into three parts: the Subnetwork Independent Convergence layer (3c), the Subnetwork Dependent Convergence layer (3b), and the Subnetwork Access layer (3a).
7. In a broad sense, openness pertains to the ability of a customer to make choices and to achieve interoperability. Openness usually involves adherence to international standards. However, other aspects include adherence to

- industry standards and those of multivendor consortia, publicly available specifications, and code licensed at industry practice.
8. AIX and OSF are distinct. The term AIX/OSF is used here to project the incorporation of OSF/1 into AIX as indicated by IBM. The AIX/OSF term is not, however, an official IBM term.
 9. K. Britton, "Introduction to Multi-Protocol Transport Networking," paper submitted by IBM to X/Open, dated December 2, 1991.
 10. R. C. Williams, "Enterprise Management Directions," document number RCW0005, internal IBM presentation dated December 18, 1991.
 11. IBM statements indicate the intent to provide CPI-C on DOS and Windows** as well.
 12. IBM has reportedly indicated that programs written to CPI-C will run over TCP/IP networks. IBM statements also indicate that when the OSI-TP functions are adopted in OSI, IBM will implement OSI-TP and CPI-C over OSI-TP. This will enable customers to write single applications that operate over SNA and OSI networks. In addition, the CPI-C has been licensed to X/Open to be incorporated as part of their industry standards.
 13. EDA/SQL** is developed by Information Builders Inc., a member of the IBM International Alliance for Information Warehouse. Databases and files are accessible on IBM MVS, VM, and OS/2, DEC VAX** and MicroVAX, UNIX, and Hewlett-Packard systems.
 14. The IBM Distributed Relational Database Architecture (DRDA) has been adopted by a growing number of database vendors, including Borland International Inc., Computer Associates, Gupta Technologies Inc., Informix, Lotus Development Corp., Micro Decisionware, Novell Inc., Oracle Corp., and Sybase Inc.
 15. APPC is now supported in a wide variety of environments, including DOS, Windows, DEC VAX, AIX, Mac/OS, UNIX, Wang VS, OS/2 SE & EE, AT&T 382, UNISYS-A Series, and Sun** workstations. Reportedly, more than 88 vendors and 178 products already use LU 6.2, facilitating communications among IBM and OEM facilities.
 16. There are reportedly over one million personal computers using the DOS APPC facility.
 17. OSI capabilities are now on all SAA systems and the AIX RISC System/6000*.
 18. TCP/IP mail, file transfer, and remote-logon facilities are now on all SAA systems, DOS, and AIX.
 19. The IBM 9075 PCradio*, for example, is a battery-powered computer terminal that can transmit, via cellular radio, the public switched telephone network (using an integrated modem), or an external radio transceiver system.
 20. IBM has stated its intent to make available APPN end-node and network-node specifications.
 21. R. G. Campbell, "A Communications Strategy for Growth," *IBM Personal Systems Technical Solutions*, G325-5005, IBM Corporation; available through IBM branch offices.
 22. Ethernet connected to an IBM 3172 may be version 2 standard and thin wire, 802.3 10base5 wire and 802.3 10base2 wire.
 23. IBM reportedly intends to support frame relay as both a DTE (end system) and a DCE (services provider). The IBM 3745, for example, will be a frame-relay DCE and will natively multiplex both SNA traffic and frame-relay traffic.
 24. D. D. Clark, B. S. Davis, D. J. Farber, I. S. Gopal, B. K. Kadaba, W. D. Sincoskie, J. M. Smith, and D. J. Tenenhouse, "The Aurora Gigabit Testbed," submitted for publication in *Computer Networks and ISDN Systems* (1992).
 25. The principal participants in Aurora are Bellcore, IBM, MIT, and the University of Pennsylvania. Collaborating carriers are Bell Atlantic, MCI, and NYNEX Corp.
 26. I. Cidon and I. S. Gopal, "PARIS: An Approach to Integrated High-Speed Private Networks," *International Journal of Digital and Analog Cabled Systems* 1, No. 2, 77-85 (April/June, 1988).
 27. For simplification, all network-access facilities, including network processors, routers, and bridges, are assumed (in Figures 14 and 15) to be part of the Equal-Coexistence Multiprotocol Network.
 28. The IBM 6611 router, for example, supports TCP/IP, XNS (Xerox), IPX (Novell), DECnet, AppleTalk, SNA, and NetBIOS protocols, and also provides source route bridging. It connects Ethernet, token ring, SDLC, serial (up to T1), frame relay, and X.25.
 29. An ingredient of such gateways is the ISO standard Interdomain Routing Protocol (IDRP). Although this standard currently refers to forwarding connectionless protocol data units, it can readily be used by gateways between networks that use other routing protocols, both connection-oriented and connectionless.

General references

- P. E. Green, R. J. Chappuis, J. D. Fisher, P. S. Frosch, and C. E. Wood, "A Perspective on Advanced Peer-to-Peer Networking," *IBM Systems Journal* 26, No. 4, 414-428 (1987).
- M. T. Rose, "Transition and Coexistence Strategies for TCP/IP and OSI," *IEEE Journal on Selected Areas in Communications* 8, No. 1 (January, 1990).

Accepted for publication February 12, 1992.

Rudolph J. Cypser 101 Young Road, Katonah, New York 10536. Following several years at the Ames Aeronautical Research Laboratory, Dr. Cypser was an instructor and then assistant professor at MIT. At IBM he was responsible for input/output equipment on the SAGE air defense project, a mobile computer prototype for the U.S. Army, and the technology for a large computer system for the U.S. Air Force. After an assignment on the IBM corporate staff, he spent several years in Europe at the headquarters for the IBM laboratories there. Subsequent activities included product planning for communications-oriented hardware and software, and program director of IBM's Systems Research Institute. Before retiring from IBM in 1989, he was director of technical communications for IBM. He is the author of an earlier book, *Communications Architecture for Distributed Systems*, and a recent book, *Communications for Cooperating Systems*. He has taught this subject at both the IBM Systems Research Institute and Syracuse University. Dr. Cypser is currently with Kim Pathways, a consulting and publishing firm.

Reprint Order No. G321-5468.