

Access services for the Networking BroadBand Services architecture

by C. P. Immanuel
G. M. Kump
H. J. Sandick
D. A. Sinicrope
K. V. Vu

This paper describes how Networking BroadBand Services (NBBS) provides multiple protocol support over emerging high-bandwidth networks. It describes the problems associated with handling these multiple protocols that are seen now more than ever in the emerging asynchronous transfer mode (ATM) infrastructure. The paper addresses the access services of NBBS and how the NBBS transport functions are used to provide these access services. A detailed example of one access service component is given to provide the reader with a feeling for typical access services issues. Brief examples of access services are given to show how the concepts described are being used.

The phenomenal growth of multivendor, multiprotocol enterprise networks has given rise to the difficult problem of finding an architecture to transport multiple protocols over wide area networks. This architecture must support voice, data, and multimedia applications (including those with real-time transmission requirements) in an efficient and integrated manner while taking full advantage of the characteristic properties of the wide area network.¹ The network must have flexible mechanisms that can handle the addressing, bandwidth, and quality-of-service characteristics demanded of each protocol it supports. For example, major problems such as broadcast storms, which occur when too many query messages are generated and sent to all other nodes in the network, must be eliminated or greatly reduced.²

The primary purpose of a transport system from the viewpoint of a network service provider is to

satisfy the communications requirements of users such that the required quality of service is optimally matched with the underlying transmission facilities. For example, some networks may support real-time applications that require tight bounds on packet delay and delay variation. Other networks support applications that require reliable delivery but are tolerant of delay, whereas most datagram networks support only "best-effort" delivery and can tolerate some level of packet loss, delay, duplication, or resequencing. A typical relationship between applications, required quality of service, and access services is summarized in Table 1.

Access services provide connectivity to multiple types of networking equipment using the protocols designed to meet the needs of that equipment. For example, a data terminal could use X.25, local area network (LAN), or frame-relay protocols. These protocols are designed to transfer data without much concern for delay but with high concern for data loss and errors. In contrast, a private branch exchange (PBX) or video camera would use a protocol that is delay-sensitive but does not consider minor bit losses or errors as a problem. These ap-

©Copyright 1995 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.

Table 1 Applications, required quality of service, and access services

Application	Required Quality of Service	Access Services
Voice (phone circuits)	Low delay, low delay variance, bandwidth reservation, low and medium error rates	Support of standard interfaces such as V.24, X.21, ISDN, ATM (voice)
File transfers	Medium-to-high delay, larger delay variance, reserved or nonreserved bandwidth, low error rates	ATM (data), frame relay, SNA, X.25, TCP/IP
Interactive data terminal	Medium-to-low delay, medium-to-high delay variance, reserved or nonreserved bandwidth, low error rates	ATM (data), frame relay, SNA, X.25, TCP/IP
Video	Low delay, low delay variance, bandwidth reservation, low-to-medium error rates	ISDN-PRI, SONET, DS-3, ATM
Imaging	Medium-to-high delay, larger delay variance, reserved or nonreserved bandwidth, low error rates	ATM (data), frame relay, SNA, TCP/IP

plications have conflicting requirements that a high-bandwidth, multiprotocol backbone must meet.

Over the last several years, a number of technologies have evolved to efficiently meet transport requirements for a variety of heterogeneous applications such as voice and low-speed data, video and high-speed data, and analog and digital sources. The advent of higher-speed digital processing technologies as well as the establishment of gigabit-per-second fiber optic communication systems have made it possible to assume that a single wide area network could meet all of the varying needs of an enterprise as illustrated in Figure 1.^{3,4} The requirement to support multiple protocols over a common transport network is thus much stronger today than it was ten years ago.


Historically, several mechanisms have been used with limited success to meet this requirement, including tunneling, routers, bridging, and protocol interworking:

- An example of the tunneling approach was providing X.25 service over IBM's Systems Network Architecture (SNA). Since SNA did not have any mechanisms for access services to support other protocols, it was necessary to tunnel X.25 packets inside the logical unit (LU) sessions of SNA. Although the LU sessions of SNA have proved to be excellent for SNA traffic, it was not a suitable choice for other protocol packets, where throughput and performance requirements were not well met with the SNA tunneling approach.⁵

- Routing and bridging approaches encapsulate data streams in packets of the backbone protocol. Although both approaches are elegant, neither offers a standardized way to guarantee fairness between encapsulated data streams in the backbone network. In addition, bridged networks use network-wide broadcasts to transmit network control messages and, therefore, cannot scale to a global level. Although bridging variants, such as smart bridging, can reduce the flow of network broadcasts, the scaling problem still exists.
- A third approach, which has been gaining in popularity, is to map various protocols to a single high-bandwidth transport, such as asynchronous transfer mode (ATM). This approach is known as protocol interworking. An example of this approach uses two frame-relay-ATM interworking units in an ATM network, one at the entry point of the ATM network and another at the exit point, to provide frame-relay bearer service to a user. In this approach, variable-length frame-relay packets are segmented into fixed-length ATM cells, transported across the ATM network as short fixed-length cells, and then reassembled back into variable-length packets on the other side of the ATM network.

The ATM technology has been developed to meet the quality-of-service requirements of multiprotocol networks. However, business constraints do not allow users to convert all their systems and applications to ATM overnight, and, in many cases, the required ATM applications are still being devel-

Figure 1 Multiprotocol switched wide area network



oped. Networking switch equipment that utilizes the access services framework provides the support necessary for users to maximize their current investments and yet takes advantage of the services of high-speed backbone networks such as ATM, with minimum investment.


This paper outlines the approach that has been taken in Networking BroadBand Services (NBBS), which combines the best aspects of all the previously discussed methods of service. Access services use the NBBS transport services to provide transport service for multiple protocols that already exist in customer premise networks across wide area networks such as ATM, without regression in terms of performance. This allows customers to maximize their existing investments and yet be able to take full advantage of the high-speed networking backbones for their emerging multimedia applications.

The remainder of this paper focuses on access services and how they meet the requirements of existing protocols. We first define the access services and put them into perspective within the broadband network structure. We then examine the access services components and explain the functions provided by the *protocol agent* component and how it utilizes the other components of the *access agent*. Next, we look in detail at the address resolution function as an example of a key function provided by the access services. The paper concludes with a short description of some protocols currently supported.

A multiprotocol network solution

Access services are used to provide entry into the NBBS network. They combine the best features of tunneling, routing, bridging, and protocol interworking transport mechanisms without degrading network performance.

Figure 2 Components of a high-bandwidth transport network



Access service placement and definition. In NBBS, components of a network may be classified into three categories. The first category is at the user level, where users have workstations and routers that communicate with others via a multiplicity of network interfaces and protocols such as frame relay, integrated services digital network (ISDN), Internet Protocol (IP), and ATM. The second category allows this set of interfaces to access a wide area high-speed network, such as those provided by a specialized common carrier or private network. The subject of this paper is the second category as shown in Figure 2. The third category is a set of services provided by the transport network, such as bandwidth management, multicast services, congestion control, and address resolution.

Access services are the collection of *access agents* that allow external protocols (e.g., frame relay, ATM, SNA) to obtain transport services from the high-speed NBBS network.


Access service components. To better understand access services it is useful to outline their components. As stated in the definition above, access services are provided by *access agents*. An access agent is a set of distributed support functions that

allow a user to interface to the NBBS high-speed network using a "native" protocol. The support functions are: a *protocol agent*, a *directory agent*, and a *connection agent*. Figure 3 illustrates the support functions within the access agent.

The protocol agent is the entity that is specific to the external protocol and matches the requirements of the external protocol to the NBBS backbone services. The connection agent establishes network connections across the NBBS network using the connection parameters provided by the protocol agent. The directory agent is the entity that locates target access agents using the destination user's address as provided by the protocol agent. To ensure that each access protocol is serviced according to its needs, NBBS combines the access agent structure with a multiprotocol backbone structure based on the ATM switch.

Functions performed by access agents. An access agent provides a variety of functions. It interprets external protocols, taking the best possible advantage of the NBBS features to establish, maintain, and clear connections requested by the user. Access agents also shield end users from the intricacies of backbone networks like NBBS. Entities con-

Figure 3 Access agent decomposition



nected via access agents do not typically require any special knowledge of functions available in the NBBS network but still can exploit them using their native protocol. In addition, interfaces to some NBBS value-added functions are offered through the access agent.

Connection establishment. An access agent may receive a connection request either from network management for a permanent virtual connection or via protocol signaling for a switched virtual connection. The access agent is responsible for interpreting the connection request and triggering the NBBS network to establish and maintain a connection across the NBBS network that meets the parameters of the connection request.

In general, there are four steps to service a connection request.

1. Find the access agent servicing the destination user—The protocol and directory agents work together to perform this function as explained in detail in the next subsection of this paper.
2. Calculate the parameters of the NBBS network connection—The protocol agent calculates the NBBS network connection parameters from those supplied by the external protocol in the connection request. The protocol agent ensures that the user will have a connection across the

NBBS backbone that meets the minimum bandwidth and quality-of-service (QOS) requirements specified in the connection request.

3. Determine the transfer mode of the network connection—The protocol agent must also select the best transfer mode to service the access protocol connection request. The NBBS architecture and implementations to date use several different transfer modes for data transport. In addition to the native ATM transfer mode, where the data are transported in 53-byte cells, NBBS offers a packet transfer mode and a source-routing automatic network routing (ANR) transfer mode. For example, for ATM traffic, the access agent would select the ATM cell transfer mode. For frame relay, a variable-length label swap transfer mode would be most efficient, as label swap switching is native to frame relay. For Advanced Peer-to-Peer Networking* (APPN*) or subarea SNA traffic, a variable-length source-routing mode such as the one used in IBM's high-performance routing (HPR) would be most desirable, since HPR natively utilizes a similar transfer mode.
4. Submit the connection request to the NBBS transport services—The protocol agent calls upon the NBBS network connection services to establish and maintain the network connection across the NBBS backbone.

One problem that users and network service providers face is ensuring fairness between multiple protocols and between different traffic classes. A large batch file sent from one user could delay or block interactive traffic sent from another user to the same destination. Protocol agents take advantage of the traffic management capabilities of NBBS, ensuring that each connection gets a fair share of network resources. Thus, users can request virtual connections to handle their different classes of traffic. For further information on NBBS network traffic management, see *Networking BroadBand Services-Architecture Tutorial*.⁶

For some protocols, it is necessary for the protocol agent to adjust the network connection to compensate for dynamic changes in connection requirements that must be made in real time. For example, a protocol agent that supports compressed voice, continuous bit stream, and data/fax capability, must be able to (1) reserve minimum bandwidth when compressed voice is in use, and (2) signal the backbone network to increase bandwidth dynamically whenever facsimile or modem traffic is detected.

In order to accomplish many of the tasks described above, it is necessary for the originating protocol agent to communicate with other protocol agents. For example, to ensure that a user can communicate with a destination user via a voice circuit, it is essential that matching services be invoked on egress access agents. Several mechanisms are available for communication between access agents for usage as required, such as the Inter-Access Agent Communication (IAAC) mechanism.⁶ These mechanisms use messages sent between access agents to signal changes in the connection characteristics or status. In addition, access agents utilize an interface with network management to provide customers with network statistics, vital accounting, and other network management information.

Access agents ensure that the services provided by the backbone network meet the contractual agreements required by the end terminal equipment, including quality-of-service parameters, delay requirements, and guarantee of bandwidth. In the case of some protocols such as frame relay and ATM, the protocol agent can provide value-added services, since it can take advantage of some functions provided by the NBBS network that are not available in other networks. A few examples of value-added services would include NBBS adaptive

bandwidth, equivalent capacity, and nondisruptive path switching.⁶ The value of NBBS adaptive bandwidth is exposed to frame-relay users in the following manner:

The frame-relay protocol agent allows the user to specify connection request bandwidth requirements in terms of minimum and maximum values. This feature gives the user the luxury of not having to determine busy-hour characteristics and yet being able to design a network that meets the budget constraints of his or her organization. The protocol agent can set up the connection to meet the user requirements and ensure that the NBBS network connection adapts to the rate of traffic received within defined limits.

We have already given one example of value added in the context of connection request parameters. However, value added is not limited to connection establishment. After a connection across the NBBS network is established, the protocol agent can allow users to increase or decrease the bandwidth reservation for the connection across the NBBS network. A request for such a change is triggered by the protocol agent requesting more bandwidth in response to a change in requirements, such as the user wanting to send a fax in the middle of a telephone conversation when voice compression is being utilized to minimize costs. A feature of the protocol agent services for voice is modem signal detection, which triggers bandwidth increase requests from NBBS to support facsimile transmission requirements.

Directory services and address resolution. Access services use external addresses (such as an E.164 address or an SNA name) to locate the called number or target node. This function is NBBS *directory services* (DS), a set of distributed services designed for locating destination users and for obtaining the characteristics associated with attaching devices or terminals.⁶ For example, DS may determine a user's ability to sustain reserved or nonreserved network connections. Locating a user means determining the network addresses of the access agent representing the user and through which the user can be reached. Directory services provide real-time address resolution service for high-speed, general-topology backbone networks, such as ATM networks.

A key directory service design objective is to minimize the overhead related to the transmission and

processing of network control messages for locating users. If the backbone network supports point-to-multipoint protocols, the directory service algorithms should take advantage of such services to "scope" search requests by broadcasting location requests to the smallest possible number of nodes. The scoping could be based on matching protocol type, address prefix matching, or a number of other schemes. With interworking services being defined in various forums such as the ATM Forum and the Frame Relay Forum, the scoping of searches based on protocol type may not always be possible. Scoping may need to be based on address prefix type, where the address prefix is part of the called address or telephone number. For example, in North America an address prefix could be 01-919, consisting of the North America code (01) concatenated with an area code (919). This scheme would allow an access agent to find a user normally associated with a different protocol. Therefore, one can achieve protocol independence with regard to directory searches. This means, for example, that frame-relay access agents could find partners that are frame relay or ATM attached.

The term address resolution means the ability to find the physical location or network address of a user or resource in the network based on the user's address (in this case the user's ATM address as defined by the ATM Forum UNI [user-to-network interface] 3.1⁷). NBBS support for the address resolution of ATM addresses satisfies three requirements. The first is support of user mobility and multihomed networks (an address connecting to two or more points in a network; ATM addresses cannot be tightly bound to the physical topology of the NBBS network). The second is address resolution efficiency in terms of time and network resources. The third is that the method must scale to large networks.

It is important to note that the address resolution method uses a combination of two NBBS base services. This ability to combine two or more base NBBS services to create a new service is an important feature of the NBBS architecture. It creates a powerful and flexible environment for designing the support of access services. To illustrate this point, the NBBS method for address resolution is described next in detail.

Address resolution overview. In order to understand the address resolution method selected, a review of three major ways for resolving addresses

follows (a more thorough analysis can be found in Reference 2). The broadcast query approach to perform address resolution is the first method. This method requires the source station to broadcast a

The address resolution method uses a combination of two NBBS base services.

query to every station in the network. The query contains the user's address of the target. If the target station receives the query, it sends a reply back to the source. This reply contains the information necessary to establish network connectivity between the source and target. A good example is the IP ARP (address resolution protocol⁸) on LANs.

The second approach requires the resource address to have two parts: (1) a prefix that contains, in Open Systems Interconnection (OSI) terms, a routing hierarchy, and (2) a host-specific portion. (The addresses adopted by the ATM Forum follow these rules.) The routing hierarchy in the address prefix is used to define the topographic location of the resource in the network.⁹ The E.164 international telephone numbering plan provides a good illustration of this approach. For example, based on the E.164 numbering plan in North America, the telephone number "1 919 543 5221" has the following topographic significance: "1" specifies North America (Canada and the United States), "919" specifies the area in North America (Eastern North Carolina), and "543" specifies an exchange within Eastern North Carolina (Research Triangle Park).

Because the address reflects the network topology, the location of resources in a large area of the network can be summarized by using the small address prefix that is common to all of the resources of the area. Based on the E.164 example above, calls between the United States and France can be routed using only the country code. (Of course, this case is the simplest. Usually other policy concerns are involved in call routing.) This ability to

summarize location information is a powerful way to control network-control overhead, i.e., routing table size and distribution of routing update information.

The third method for resolving an address is the creation of a hierarchy of distributed directories, for example, the IP Domain Name System or the International Organization for Standardization (ISO) X.500 directory.¹⁰ In a general environment, these distributed directories resolve source station address queries by using hop-by-hop hierarchical search protocols. Since the external network addresses may incorporate only very few hierarchical layers (typically only two: network and station identifier), address resolution procedures would in very many cases have to be navigated through the first, second, or third level of the gigantic worldwide directory tree.

Scoped query rationale and structure. Although all of the address resolution approaches described above are extremely useful, none was able to meet all of our initial objectives for fast address resolution, user mobility, and ability to scale to large networks. The broadcast query approach is fast but does not scale.¹¹ The routing hierarchy is fast and scales, but in order to control network control overhead it restricts user mobility, e.g., when a person moves to another location his or her telephone number usually must change. The distributed directory technology as it exists today scales and allows user mobility, but based on the hop-by-hop search protocols discussed earlier, it does not offer the subsecond response times required for high-speed interconnection (for example, see Reference 12).

To solve our problem we decided not to view the three approaches described above as discrete functions but as three points on a continuum of address resolution methodology. From this perspective, other options or approaches are possible. Using this line of reasoning, we were able to meet our initial requirements by creating a hybrid address resolution method, the scoped query, which incorporates positive elements from the three approaches. A brief description of the important features of the scoped query follows.

To meet our objectives for user mobility we decided to use the query mechanism. The question was how to keep the search fast without overwhelming the network. To control the search, the


scoped query capitalizes on the routing hierarchy contained in the ATM address prefixes. These address prefixes are used to form groups of AAs supporting resources with a common address prefix. These groups of AAs, which are not required to have topological proximity, are called directory groups and have hardware multicast connectivity or "trees" between the group members. A network function exists so that given a target resource, the multicast tree associated with the appropriate directory group can be obtained. Once obtained, the multicast tree is used to send queries to the members of the directory group supporting the address prefix of the target resource. This ability to control the distribution of query is the origin of the name "scoped query." In addition, the level of scoping is defined on a network basis and is flexible to the degree that the address allows it, i.e., E.164 addresses are aligned on nibble boundaries, whereas the ATM private address (NSAP, or network service access point, form) can be aligned on bit boundaries.

The scoped query as outlined above meets our three initial requirements. First, the use of hardware-assisted multicast allows fast distribution of queries, avoiding the hop-by-hop traversal of a name server hierarchy. Second, the query is not broadcast to every node in the network but is multicast only to the directory group supporting the relevant address prefix. This method minimizes network control overhead and therefore allows the approach to scale to large networks. Third, the directory function is separated from the network topology. This separation permits user mobility without invoking substantial network control overhead.

In terms of implementation, the scoped query is built on top of NBBS services and functions already available. Set manager (SM) is used to control the membership of a directory group and to form and maintain multicast connectivity between group members; directory services (through the directory agent) are used to send and receive the query; and the remote-access-to-tree-transfer mode, an NBBS function implemented in hardware, is used to transmit the query to the directory group.⁶ This last function allows all members of the group to receive the query at approximately the same time.

Scoped query example. The scoped query function has two basic parts. First, each access agent is system-defined with the resource address prefixes it supports (this must be done as part of the ATM Fo-

Figure 4 Multicast groups formed of directory agents serving the same address prefixes



rum UNI 3.1⁷ address registration procedure). The directory agent of the access agent then uses NBBS set management to join the group(s) of directory agents associated with each address prefix it supports. The set manager uses a network-defined function to be sure that all directory agents (DA) are placed in the appropriate directory group, thus ensuring that all resource addresses are reachable. The SM makes sure that the directory group is registered with the registrar of the network, the centralized set database function of set management. Finally, as part of the set manager function process, each of the directory agents in the directory group for a given prefix is placed on a multicast tree. As mentioned earlier, this tree will be used later for the scoped queries.


In Figure 4, directory groups for the DAs supporting the prefixes 726 and 725 have been formed. It is important to note that in this figure the groups have been topologically segregated for ease of graphic depiction. One of the access agents (AA) supporting 725 could have been located between the two AAs supporting the 726 area.

The second part of the scoped query involves the actual address resolution process. As an example,

how the scoped query works in support of an ATM SVC (switched virtual circuit) setup procedure is presented. First, a user initiates the setup of a virtual channel connection (VCC) to another user by sending a Q.2931 SETUP message to its network interface, an NBBS ATM AA. The SETUP message contains the called party (or target) address, which in this case is 725-4564. The DA checks its local directory and is unable to find the target address. The DA goes to its local SM in order to obtain the multicast tree of the directory group supporting the target address. The set manager has either cached the tree (keyed off the address of the target) or it must consult the registrar to obtain the tree associated with the directory group serving the target address.

The tree is returned to the DA, which in turn sends an undirected query that travels point to point through the network until it reaches the multicast tree of the directory group. At this point, the query becomes a multicast query sent to each member of the directory group. Again, it is important to note that only the DAs that have registered for this group receive the search. As part of the query procedure, the path back to the originating DA (DA serving the initiating user) source is collected in the header of the query. This path allows the destination DA (DA

Figure 5 A query is sent to the directory group serving the 725 address prefix



serving the target resource) to reply to the query request using a point-to-point flow. The reply terminates the query operation, returning the information necessary to set up the VCC across the NBBS network. See Figure 5.

Examples of access agents

The relevance of an architecture is best demonstrated by showing its applicability to a representative sample of real-world scenarios. Examples of possible access agents for various protocols include real-time voice, ISDN, ATM (permanent virtual circuit, or PVC, switched virtual circuit, or SVC, and multicast), frame relay (PVC, SVC, multicast), high-level data link control (HDLC), X.25, LAN Bridging, Fibre-Channel Standard, switched multimegabit data service (SMDS), IP, and Resource Reservation Protocol (RSVP). In this paper, we limit our illustrations to ATM, frame relay (with the unique application of frame relay-ATM interworking), and voice-PBX. It should be noted that these examples are for illustrative purposes only, and although the principles outlined in this paper can be extended to any layer-2 or layer-3 protocol, not all may be actually implemented in any given product.

Access services essentially terminate the signaling protocol with subsequent data transport being performed at the physical layer.

ATM access agent. ATM has been defined by many as the base technology for the next generation of global communications. It is well suited for high-speed networking and the transport of a wide range of different traffic types, including voice, data, and video.

The NBBS ATM access agent provides a fully standards-compliant (i.e., conforming to International Telecommunication Union-Telecommunications, or ITU-T Q.2931-1994¹³ and ATM Forum UNI 3.1⁷) ATM user-to-network interface (UNI), providing the network side of the interface for user access to the NBBS network, as well as the user side when the NBBS network is attached to other ATM bearer services. Both permanent and switched virtual circuits (PVC and SVC) are supported.

Although the ATM standards concentrate on the physical, ATM, and ATM adaptation layer interfaces between the user and network, the standards do not address many difficult networking issues. It is here where NBBS excels, offering a total network-

ing solution for ATM, with a suite of required networking services, including efficient bandwidth management, congestion control, multicast services, nondisruptive fault recovery, and network management. The ATM access agent takes full advantage of these services as it provides seamless access between the user and network, guaranteeing quality of service and efficient management of multiple traffic classes.

Though fully standards-compliant, the ATM access agent allows access to value-added services, which enables users to take advantage of additional NBBS capabilities, such as adaptive bandwidth (network adjusts bandwidth reservation on an as-needed basis) and nonsymmetrical connections (ATM bidirectional connections may be routed by different paths through the network).

Frame-relay access agent. Essentially, frame relay is statistical multiplexing over a shared network. It is intended for use between intelligent endpoints and is implemented over high-quality transmission facilities that connect programmable switches or nodes. Its principal usage is LAN interconnection over wide areas. However, usage is expanding to include interconnection of routers and network subareas. Frame relay is also seen by industry experts as the leading "feed" technology to ATM broadband networks.

Conceptually, a frame-relay virtual connection is comprised of bandwidth allocated on a path through the frame-relay network from one user to another. Each node in the path has a table entry that maps the address of the frame coming into the node to the frame address going out of the node toward the next node in the path. As a frame makes its way through the path, its address is replaced, or "swapped," at each node until it reaches the destination user. Frames traveling in the opposite direction do the address swapping in reverse.

The frame-relay specifications define the formats and procedures at the interface between the user and the network.^{14,15} How the frame-relay service is supported across the network backbone is up to each individual network service provider. In NBBS, the frame-relay access agents utilize NBBS services to support standards-compliant, switched, and permanent virtual connections at both the user and the network sides of the interface.^{16,17}

When a frame-relay access agent establishes a frame-relay virtual connection across an NBBS network, many NBBS services are utilized to create and support the connection. Directory services functions are used to locate the destination user of the virtual connection. Once the destination user is found, the frame-relay access agent uses the network connection and congestion control services of NBBS to reserve and regulate the virtual connection. The originating and destination frame-relay access agents may also signal between themselves to negotiate various connection options and to clear the connection across the NBBS network when the users have completed their data exchange.

The NBBS frame-relay access agent provides several enhancements over and above standard frame-relay service, including bandwidth adaptation, nondisruptive path switch, rate-based congestion control, and integrated frame-relay-to-ATM interworking:

When requested, the frame-relay access agent establishes a virtual connection that takes advantage of the NBBS bandwidth adaptation feature. The bandwidth of the virtual connection increases and decreases according to the offered traffic load within limits provided by the user.

The frame-relay access agent enables the NBBS nondisruptive path switch feature to ensure virtual connection reliability.

The NBBS congestion control services used by a frame-relay access agent avoid network congestion by triggering frame-relay congestion control procedures at the periphery of the NBBS network based on the rate of the user's offered traffic.

To support the move to ATM technology, the NBBS frame-relay access agent provides an integrated frame-relay-to-ATM interworking function that allows the interconnection of existing frame-relay equipment to NBBS ATM products. The frame-relay-to-ATM interworking function converts the frame-relay packets to ATM cells that can then be sent as though they were generated by ATM access agents for communication to other ATM users.

Voice and PBX access agent. The voice and PBX access agent has the special features that were mentioned in the other two examples, as well as some others. For example, as mentioned earlier, the ac-

access agent exploits the fact that NBBS has the capability of dynamically increasing or decreasing bandwidth reservations. Since voice support in NBBS-supported products includes the use of "voice servers" (which provide voice-related functions such as compression, silence removal, and echo cancellation), the access agent allows close coupling and integration of these voice server functions with those of NBBS. For example, the in-band signaling capability of the access agents allows the connection to be set up across the backbone with the assurance that the matching features are properly invoked at both ends. If a certain type of voice compression is used at one end, the corresponding piece is properly invoked at the far end. The access agent also has the ability to reduce the bandwidth reservation whenever a PBX customer hangs up the telephone or when any type of voice compression is invoked by the voice server.

For isochronous interfaces (i.e., where the source and destination equipment require clock synchronization), packet-switching broadband networks must compensate for the delay jitter of packet transmission through them, and they must limit the packetization delay for voice and video transmission. Access services ensure that end-to-end delay requirements are met and communicate enough information to the destination access agent to size playout buffers. They may also provide enough information to the destination to ensure invocation of various compatible support modules such as echo cancellation, compression, and A-law/ μ -law conversion for voice transmission. NBBS architecture provides an in-band mechanism for access agents to communicate with one another for this purpose.

Summary

In this paper we discuss how an access agent model provides a robust structure for enabling the backbone NBBS network to deal effectively with a large number of diverse access protocols, to allow high throughput across the wide area backbone and to have unique protocol processing at the access agents. We have also shown how this approach provides additional value to the customer by exploiting the features of NBBS. Finally, through some examples of access agents, we have shown how the access agents allow customers to use NBBS services without absorbing the cost of understanding the internal details of NBBS.

Acknowledgments

The authors would like to thank the other members of the NBBS team for their comments and suggestions on this paper. In particular, appreciation is expressed to Nick Huslak, Didier Giroir, and Raif Onvural for their thorough reviews. We would also like to recognize the contribution that Willi Doeringer, Jeff Derby, and Doug Dykeman made to the early development of the access services.

*Trademark or registered trademark of International Business Machines Corporation.

Cited references

1. P. F. Chimento, J. E. Drake, L. Gün, E. A. Hervatic, C. P. Immanuel, G. A. Marin, R. O. Onvural, S. A. Owen, and T. E. Tedijanto, *Broadband Network Services for High Speed Multimedia Networks*, IBM Corporation, Dept. BUAA/B664, P.O. Box 12195, Research Triangle Park, NC 27709 (1993).
2. W. A. Doeringer, H. D. Dykeman, M. L. Peters, H. J. Sandick, and K. V. Vu, *Efficient, Real-Time Address Resolution in Backbone Networks, Local Area Network Interconnection*, R. O. Onvural and A. A. Nilsson, Editors, Plenum Press, New York (1993).
3. M. de Prycker, *Asynchronous Transfer Mode—Solution for Broadband ISDN*, Ellis Horwood, New York, London, Toronto, Sydney, Tokyo, Singapore (1991).
4. *High Speed Networking Technology: An Introductory Survey*, GC24-3816-00, International Support Center, IBM Corporation, Research Triangle Park, NC (1992); available through IBM branch offices.
5. *X.25 Guide*, GC24-3458-01, International Technical Support Center, IBM Corporation, Research Triangle Park, NC (1991); available through IBM branch offices.
6. *Networking BroadBand Services—Architecture Tutorial*, GG24-4486-00, International Technical Support Center, IBM Corporation, Research Triangle Park, NC (1995); available through IBM branch offices.
7. ATM Forum, *ATM User-Network Interface Specification Version 3.1*, ATM Forum Inc., Mountain View, CA (1994).
8. D. Comer, *Internetworking with TCP/IP: Principles, Protocols and Architecture*, Prentice-Hall, Inc., Englewood Cliffs, NJ (1989).
9. *Network Service Definition*, ISO/8348, International Organization for Standardization, Geneva (1988).
10. J. Postel and J. Reynolds, *A Standard for the Transmission of IP Datagrams over IEEE 802 Networks*, RFC 1042, Internet Engineering Task Force (IETF), Geneva (1988).
11. S. Zatti, "Naming in OSI: Distinguished Names or Object Identifiers?" *Proceedings of the CompEuro Conference*, V. A. Monaco and R. Negrini, Editors, Bologna, IEEE Computer Society Press, Los Alamitos, CA (May 13–16, 1991).
12. R. P. Davidson and N. J. Muller, *LANs to WANs: Network Management in the 1990s*, Artech House, Boston, London (1990).
13. ITU-T Study Group XI, *Draft Recommendation Q.2931 (DSS2 Basic Call)*, International Telecommunication Union, Geneva (1994).
14. Frame Relay Forum, *User-Network Implementation*

- Agreement (FRF1)*, Frame Relay Forum, Foster City, CA (1992).
15. *Frame Relay Forum, Network-to-Network Implementation Agreement (FRF2)*, Frame Relay Forum, Foster City, CA (1992).
 16. *Digital Subscriber Signalling System #1 (DSS1)*, *Signalling for Frame Mode Basic Call Control*, Recommendation Q.933, ITU-T, Geneva.
 17. *Frame Mode Bearer Services*, Recommendation I.233, ITU-T, Geneva.

Accepted for publication June 19, 1995.

C. Paul Immanuel *IBM Networking Hardware Division, 800 Park Offices Drive, Research Triangle Park, North Carolina 27709 (electronic mail: immanuel@ralvm6.vnet.ibm.com).* Mr. Immanuel is a senior engineer at IBM in the Networking Architecture development organization. He is a team leader for access agent development for Networking BroadBand Services (NBBS) architecture and has designed access agent architectures for frame relay and voice. He also has interests and assignments in the areas of optical networking and intelligent networks. Prior to his present assignment, he was a lead architect for link services architecture and developed a variety of link-level enhancements for both Systems Network Architecture and Advanced Peer-to-Peer Networking Architecture, including development of link-level security algorithms for connections through switched networks. Mr. Immanuel joined IBM in 1984 after 18 years of design and development experience in analog microwave, digital, and optical transmission systems. His experience includes hardware design as well as system design and planning of these systems for a number of companies, including the ITT Corporation. Since joining IBM, he has received a number of awards, including invention awards and a corporate Outstanding Technical Achievement Award for his work on NBBS architecture. He has a B.S. degree in electrical engineering from North Carolina State University and is a senior member of the Institute of Electrical and Electronics Engineers (IEEE). He has also taken several graduate-level courses in digital signal processing, communications, and computer science.

Garry M. Kump *IBM Networking Hardware Division, 800 Park Offices Drive, Research Triangle Park, North Carolina 27709.* Mr. Kump is a staff engineer in the Networking Architecture organization, where he works with the Networking BroadBand Services group. After serving in the United States Air Force from 1976 to 1984 as an electronic communications technician, he received his B.S.E. degree in electrical engineering from the University of Nebraska at Lincoln, where he became a member of IEEE and was accepted into the Eta Kappa Nu Electrical Engineering Honor Society. Coming to IBM in 1989, he first worked in unit testing of Interconnect Communication products, moving shortly thereafter into product development and then to communication architecture. His main interests lie in software development, especially in the area of high-speed communications.

Haldon J. Sandick *IBM Networking Hardware Division, P.O. Box 12195, Research Triangle Park, North Carolina 27709 (electronic mail: sandick@vnet.ibm.com).* Mr. Sandick is an advisory programmer in the Networking Hardware Division's Networking Architecture (NA) group. He joined IBM in 1987 to work on the VTAM[®] design of APPN Directory Services. In

1990, he moved to the Networking BroadBand Services architecture team in NA where he has worked on the architecture required to support various access technologies. During the last few years he has focused on supporting ATM interfaces. He received an M.S. in education from Fordham University (1981) and an M.S. in computer science from New York University (1987). In 1994 Mr. Sandick received an Outstanding Technical Achievement Award for his work on the NBBS architecture. His current interests lie in the design of networking software for high-speed networks.

David A. Sinicrope *IBM Networking Hardware Division, 800 Park Offices Drive, Research Triangle Park, North Carolina 27709 (electronic mail: david@vnet.ibm.com).* Mr. Sinicrope is a staff programmer in IBM's Networking Architecture group currently working on the Networking BroadBand Services (NBBS) architecture and representing IBM in the Frame Relay Forum. After receiving a B.S. in computer science at Northeastern University, he joined IBM in 1988 as an associate programmer designing and developing protocol test systems, later leading the frame-relay activities of the IBM Conformance Test Center. His current interests lie in networking software development, especially in the area of high-speed communications.

Ken V. Vu *IBM Networking Hardware Division, P.O. Box 12195, Research Triangle Park, North Carolina 27709 (electronic mail: kenvu@vnet.ibm.com).* Mr. Vu is an advisory engineer in the Emerging Technology Department of the Networking Architecture group. He is currently responsible for traffic management in NBBS architecture. He joined the former IBM System Products Division at IBM Rochester, Minnesota in 1980, where he designed and developed the Computer Control Reactive Iron Etching System and Digital Image Inspection System. In 1985 he transferred to Research Triangle Park and worked on the Workstation Controller for the AS/400[®] and the IBM 3174 Establishment Controller. He joined the Networking Architecture group in 1989 and was the team leader for NBBS Access Architecture. His research interests include performance analysis of ATM traffic management, congestion control, and bandwidth management. He received a B.S. degree in electrical engineering from the University of Missouri at Rolla in 1979. In 1984, he received an M.S. degree from the University of Wisconsin at Madison. He is currently working on his Ph.D. degree in computer engineering from North Carolina State University.

Reprint Order No. G321-5587.