SNA Perspective On Internetworking

Volume 14, Number 9
September, 1993
ISSN 0270-7284

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

Broadband Network Services Fills in The ATM "Cloud"

In July, IBM announced its plans for supporting broadband networking. The announcements included a major commitment to support current and future Asynchronous Transfer Mode (ATM) standards. IBM also announced its Broadband Network Services (BNS) architecture which includes support for ATM standards while adding a range of value-added features.

What is the relationship between BNS and ATM? Almost everyone has seen ATM "cloud diagrams" with networked systems connected to an amorphous blob. Essentially, BNS defines the innerworkings of the "cloud". While standards have been defined for interfaces between user systems and ATM networks, standards which defined the innerworkings of ATM networks have not yet been completed. BNS is designed to fill this gap.

(continued on page 2)

Broadband Networking With Asynchronous Transfer Mode

Asynchronous Transfer Mode (ATM) is undoubtedly the most talked about technology in the networking business today. This attention is justified because ATM is not just another incremental improvement in networking technology, it has the potential to revolutionize enterprise networking. ATM provides a very high speed cell-based switching infrastructure which is capable of carrying a wide variety of traffic including data, voice, and video.

While ATM is a very new technology, we believe that it will have a substantial impact on the networking marketplace over the next year or two. Some vendors are already beginning to ship ATM products, while many others, including IBM, will begin deliveries in 1994. In this article we will provide an overview of the current state of ATM standards and analyze the capabilities of ATM.

Why ATM?

Today's networks are made up of combinations of LANs, such as Token-Ring and Ethernet, and various wide area networking technologies including dedicated circuits and packet switching networks based on X.25 and frame relay standards. All of these

(continued on page 10)

In This Issue:

(continued from page 1)

Role of BNS

BNS provides the potential for both technical and economic benefits. From a technical perspective, BNS will be used by IBM to bridge the "standards gap" which is currently faced by ATM product implementors. As we have described in the accompanying background article on ATM, there are many ATM networking standards that have not yet been defined. This means that IBM or any other vendor building ATM networking products must, at least in the short run, employ proprietary solutions to fill in the gap.

Over the longer term IBM wants to make BNS an open industry standard. IBM is working with the ATM Forum to set additional industry standards. As a part of this standards process IBM is contributing its technologies in the hope that some of the now proprietary elements of BNS will eventually become industry standards. The BNS architecture is almost certain to evolve over time because some of the standards that emerge may differ from the technologies now included within BNS. In some cases, these emerging industry standards might replace elements of BNS while in other situations BNS might support the standard along with alternative BNS technologies.

It is also important to keep in mind that BNS addresses a wide range of broadband networking issues including technologies other than ATM. For example, BNS networks are capable of switching variable length information packets in addition to the fixed-length cells defined by ATM.

From an economic point of view, the key objective of BNS is to reduce customers' recurring bandwidth costs. Even though the cost of a given amount of bandwidth will decline in the future, new applications supporting video, voice, and image data will require huge increases in the total amount of bandwidth in the enterprise network. The result is that the total recurring costs of wide area bandwidth will soar. IBM estimates that the bandwidth management capabilities of BNS can improve efficiency by 20% to 50% depending on the characteristics of the traffic.

Overview of BNS

The three major elements of BNS are Control Point (CP) Services, Access Services, and Transport Services.

CP Services

CP Services provides the network management functions for a BNS network. These management functions include bandwidth management, congestion control, computation of optimum routes across the network, and non-disruptive rerouting of connections across the network. A network directory service is provided to maintain information about network resources and users.

CP Services is also responsible for creating and maintaining a spanning tree which connects all of the network management points in the network. This spanning tree is used to distribute control information such as topology database updates.

Access Services

Access Services provides a wide variety of interfaces to the BNS network. Entities called Access Agents provide these interfaces to devices which attach to the BNS network. Interfaces are provided for data networking applications as well as for other types of information such as voice and video.

Note that the data networking interfaces are not limited to any particular protocol layer. The BNS network could be used to provide a low level (layer 1) circuit emulation service, a frame relay service (layer 2), or it could provide a full APPN Network Node emulation. We feel that this philosophy of the network adapting to the user rather than the other way around is one of the key value-added features of BNS.

Transport Services

The third major element of BNS is Transport Services. Transport Services is responsible for providing the end-to-end connections used to transport information across BNS networks. Transport Services employs a variety of routing/switching schemes and packet formats. The Transport Services protocols are selected to match the requirements of the communicating users.

Transport Services supports ATM's fixed-length cell format as well as a variable-length packet service which is called Packet Transport Mode (PTM). Support for native ATM switching protocols within BNS networks means that traffic between ATM users will require no conversion as it traverses the BNS network. The use of PTM and its variable length packets within the BNS network will allow packet-oriented data communications devices to communicate across a BNS network without the overhead of converting variable-length packets into fixed-length cells.

BNS supports standard ATM cell switching based on Virtual Path Identifiers (VPIs) and Virtual Channel Identifiers (VCIs). Other switching techniques are also natively supported by BNS switching systems. One of these switching techniques, Automatic Network Routing (ANR), should be familiar to SNA Perspective readers. ANR is the routing protocol used by APPN's new High Performance Routing (HPR) protocol (see SNA Perspective March 1993 for a review of HPR).

Just as BNS optimizes communications between ATM users by natively supporting ATM protocols, it provides the same type of native switching/routing support for APPN HPR users. This will provide APPN users with very high performance, low overhead connections across BNS networks.

The Structure of BNS Networks

Like any network architecture BNS defines networks which are made up of nodes which are interconnected by communications links. The communications links within BNS networks are unidirectional connections between adjacent pairs of nodes. Note that this is a logical view of the communications links because, of course, the actual communications links may be full-duplex facilities. BNS treats such full-duplex facilities as pairs of BNS logical links.

Nodes within BNS networks are categorized as either nodes or subnodes. A node is a group of one or more subnodes which share some or all of their network control functions. For example, the subnodes within a node might share a single network topology database. This node/subnode concept permits the design of networking products which contain multiple, distributed switching elements without requiring the implementation of the full network control function on each switching system. For simplicity in this article we will generally refer to BNS switching systems as nodes.

BNS Node Structure

The BNS functions which are implemented in nodes and subnodes can be described by the functional layering model shown in Figure 1.

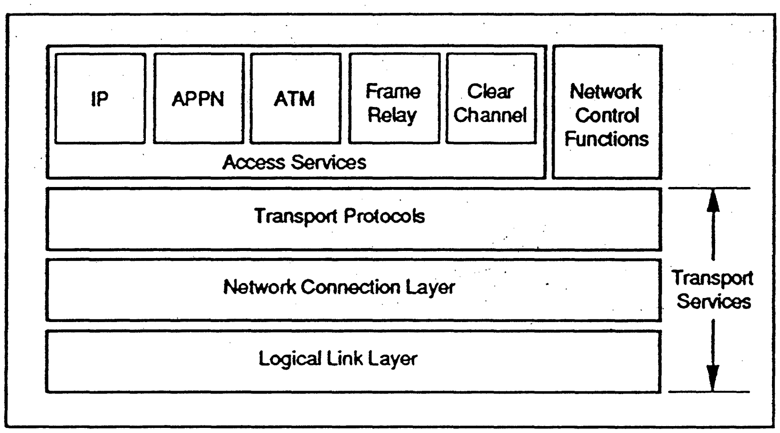

Figure 1

The Logical Link Layer (LLL) deals with communications between adjacent nodes. The LLL can support a wide range of communications facilities including T1, T3, SONET, and LANs such as Token-Ring and Ethernet.

The Network Connection Layer (NCL) uses the services of the LLL to provide end-to-end connectivity across BNS networks. NCL is responsible for packet/cell switching. NCL handles transmission scheduling and buffering based on quality of service requirements, as well as the implementation of congestion control procedures. NCL supports the three types of end-to-end network connections (NCs) shown in Figure 2.

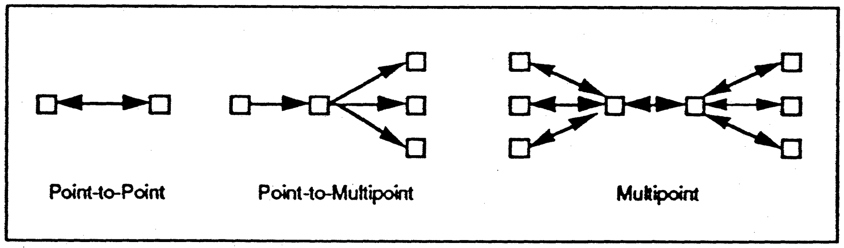

Figure 2

Point-to-point NCs provide a pair of unidirectional end-to-end paths, one in each direction, while point-to-multipoint NCs are unidirectional. Multipoint NCs provide bidirectional any-to-any communications among a group of users by employing either a single bidirectional tree connection or multiple unidirectional trees.

Note that the NCs provided by NCL are unreliable connections. If reliable message delivery is required, it must be provided by a transport protocol running on top of NCL.

The transport protocols provide end-to-end logical connections called transport connections. The transport protocol layer handles segmentation and reassembly of the information stream as well as providing such optional functions as reliable delivery of information and end-to-end flow control.

BNS Packet Formats

BNS defines a standard format for packets of data flowing through the network. Actually, BNS networks are capable of handling packets which conform to the BNS format as well as native ATM packets (cells). As we mentioned earlier in this article, BNS provides native ATM switching support. ATM performance is maximized by eliminating the need to perform any cell/packet conversions as ATM information passes through a BNS network.

The format of BNS packets is shown in Figure 3. A Network Connection header and a Transport Protocol header are added to the front of a variable length information payload. The length and format of the TP header is determined by the transport protocol being used.

The NCL header contains a routing mode identifier which is used by the switching nodes to determine the type of routing/switching to be used for each arriving packet. Note that when ATM packets enter a BNS network the first four bits of the ATM cell header, the Generic Control Field, are replaced by the BNS routing mode identifier. This is possible due to the fact that the Generic Control Field only has significance on the User-to-Network interface and not for cells flowing between switches within the network.

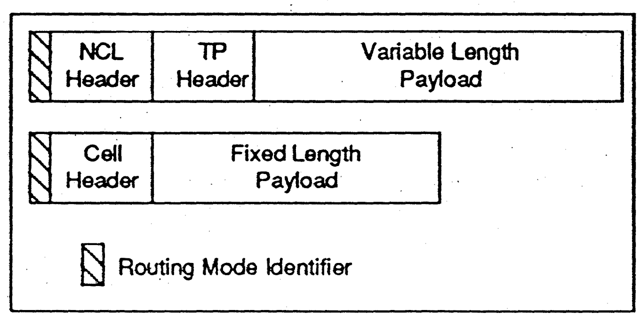

Figure 3

The NCL header also includes a Routing Information Field (RIF) which contains the identifier/address which the switching nodes use to determine where each packet should be forwarded. The format of the RIF is different for each of the routing modes supported by BNS.

BNS Routing Modes

BNS routing/switching is handled within the Network Connection layer. Each of the switching nodes in a BNS network is capable of supporting several different packet routing modes. Each packet of data flowing through a BNS network carries the routing mode identifier which we have already discussed. When a BNS switching node receives a packet of information it looks at the routing mode identifier to determine which type of routing is to be performed. The following is a list of the types of routing supported by BNS switches.

- ATM Routing
- Automatic Network Routing (ANR)
- Tree Routing
- Remote-Access-to-a-Tree Routing
- Label Swap Routing

BNS switches support the industry standard ATM routing scheme using the Virtual Path Identifiers (VPIs) and Virtual Channel Identifiers (VCIs) which are carried in the ATM cell headers. All ATM traffic entering the BNS network is routed natively. The ATM Overview article in this issue of SNA Perspective describes the methods used to route cells in ATM networks.

Automatic Network Routing

Automatic Network Routing (ANR) is the same routing technique used by APPN's High Performance Routing (HPR). ANR is a source routing scheme in which each packet carries a complete description of the packet's path through the network. The routing information field carries a list of ANR labels. As the packet moves through the network each switch in succession makes its routing decision based on the first label in the list and then it removes that label so that the next switch can again use the first label in the list to make its routing decision. At each switching node the first ANR label in the routing information field identifies the outbound link for that switch.

The ANR routing mode is used to set up connections in BNS networks and for other network control functions. ANR provides a datagram service within the connection-oriented BNS environment.

Tree Routing

The tree routing mode implements a multicast capability among the BNS switches. The routing information field carries a unique identifier which identifies a tree. The switches then forward the packet on all outbound links associated with that tree.

The tree routing mode is used to implement the BNS Control Point spanning tree as well as various types of multicast functions. Tree routing could be used, for example, to transmit a video signal to a group of viewers simultaneously.

Remote-Access-to-a-Tree Routing

Remote-access-to-a-tree routing combines ANR and tree routing to allow nodes which are not a member of a tree to send information to that tree's members. ANR routing is used to send a packet from the non-member node to a member node and tree routing can then be used to send the packet to all members of the tree. The routing information field in this case would contain the list of ANR labels needed to reach a tree member node and the last ANR label would be followed by the tree identifier.

Label Swap Routing

Label swap routing uses a technique which is similar to APPN's intermediate system routing (ISR). Each packet carries a label which is looked up in each intermediate switching node and then that label is swapped for the label which will identify the outbound link on the next switching node. The labels are set up in the intermediate switching nodes at call set-up time. Unlike APPN's ISR, BNS label-swap routing can support multipoint routing by associating one or more of the labels with multiple outbound links within the switching nodes.

Each of the routing modes that we have discussed, except for ATM routing, can implement a copy option which is used to support network management functions. When the copy option is in use, each intermediate switching node will receive a copy of the packet as it is sent through the network. This linear multicast function is used to inform intermediate switching nodes when a connection is being set up.

Managing BNS Network Performance

Bandwidth management in broadband networks is particularly challenging because of the wide range of traffic types. End user bandwidth requirements range from a few thousand bits per second for some types of data traffic to millions of bits per second for image data transfers and multimedia applications.

Bandwidth management complexity is increased by the fact that some connections operate at constant bit rates while other types of traffic may be very bursty. Some types of information such as voice and video also require connections which introduce very little transmission delay. Some applications, such as video, are also sensitive to packet loss. Effective network management must take all of these user requirements into consideration.

Congestion Control Techniques

Some connections across BNS network require specific levels of performance which is referred to as quality of service (QOS). When these types of connections are established, the network reserves the appropriate amount of bandwidth to meet the QOS requirements.

Bandwidth reservation is simple for connections operating at a constant bit rate, but it's quite a bit more difficult for traffic whose bandwidth requirements vary over time. The simplest solution for handling variable rate traffic is to reserve the maximum (peak) bandwidth that will be used. This approach is simple but very wasteful because the reserved bandwidth will go unused during times when traffic is not running at the peak level.

The BNS approach to bandwidth allocation is to compute the equivalent capacity for a connection. Equivalent capacity is the minimum amount of bandwidth required to ensure that the number of packet losses due to congestion do not exceed a maximum guaranteed value for the connection.

In order to calculate the equivalent capacity for a connection, the following parameters must be provided:

- Mean rate
- Peak rate
- Mean duration of bursts
- Maximum acceptable packet loss ratio

When the connection is initiated, the equivalent capacity bandwidth will be reserved on each link along the connection. If the required bandwidth is not available, the connection is not set up.

Note that when reserved bandwidth is not being used it is made available to nonbandwidth-reserved connections. These connections make no quality of service guarantees, they provide best effort delivery service.

Setting Up a Network Connection

Figure 4 shows a simple example of some of the steps involved in setting up a network connection. The connection request is initiated by a network user. The Connection Agent within the user's Access Agent receives the request which includes a

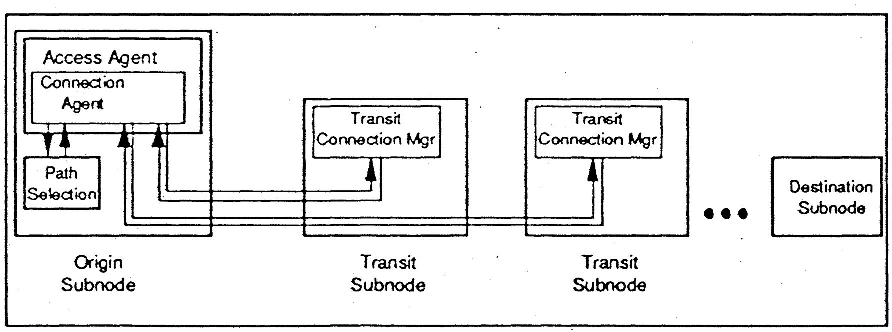

Figure 4

description of that user's bandwidth requirements. The Connection Agent uses these bandwidth requirements to compute the equivalent capacity required for the connection based on the quality of service required.

The Path Selection service uses information stored in a topology database to compute the best path to the destination. The topology database contains information about the bandwidth which is currently available in the network.

The Connection Agent then sends connection requests to each of the Transit Connection Managers (TCMs) in the intermediate Transit Subnodes. Each of these TCMs determines whether the requested bandwidth is available and sends a response to the Connection Agent. If the TCMs are able to allocate the requested bandwidth they broadcast a Topology Database Update (TDU) on the CP spanning tree. This notifies all of the CPs in the network of the change in bandwidth availability due to this connection request.

Enforcing Bandwidth Allocations

We have just seen how the requested equivalent capacity is initially assigned to a network connection. In order to ensure that quality of service commitments are met for all users, the network must enforce the bandwidth allocations that are made at the time that connections are started. These enforcement procedures are initiated when a connection is established and they constantly monitor the traffic entering the network. Whenever the agreed-upon traffic characteristics (peak rate, mean rate, and mean burst duration) are exceeded, the traffic policing mechanism acts to control the incoming traffic.

The controls that can be employed include delaying the traffic, marking packets as discardable in the event of congestion, and discarding packets before they enter the network. The peak rate is enforced by a packet spacing mechanism which discards packets which exceed the agree-upon peak rate. A leaky bucket is used to control the mean rate and the burstyness of the traffic. When these parameters are exceeded, packets can be marked as being "excess" packets. When congestion occurs in the network, those packets which are marked excess will be discarded before other packets.

Managing Transmission Delay

The techniques that we have just described are designed to manage the loss of packets in a BNS network. Another key performance characteristic of BNS networks is transmission delay. In order to manage delays, BNS defines three general categories of network connections based on their transmission delay requirements. These three categories of connections are:

- Real-time bandwidth-reserved
- Nonreal-time bandwidth-reserved
- Nonreserved

These three classes of traffic each have a logically separate queue within the switching nodes. Transmission priority is given first to real-time traffic and then to non-realtime bandwidth-reserved traffic.

The fact that BNS networks can handle variable length packets of information makes support for real-time traffic more challenging. Consider the situation where a packet of real-time information arrives in the output queue of a switching node, but a long packet is already in transit over the outbound link. For most data applications this situation would have little impact, but for real-time applications such as video this type of delay is intolerable.

ATM handles this situation by breaking up all information into fixed-length, 53 byte cells. This effectively time-slices the output links so that the maximum delay in transmitting a cell will be the length of time required to send 53 bytes of information.

How does BNS, with its variable length packets, handle this problem? BNS switches implement a technique which is called preemptive-resume. This mechanism allows real-time packets to interrupt the transmission of lower priority packets. This preemptive resume mechanism is the key to supporting both variable-length packets and real-time information transmission in BNS networks.

BNS Access Agents

Access Agents are one of the key elements of the BNS architecture. Users interface to BNS networks through the Access Agents. One of the strengths of BNS is that it accommodates a wide range of existing networking interfaces. In most cases, user systems will require no upgrades to interface to BNS networks.

Some examples of interfaces that will be provided by BNS Access Agents include:

- Clear channel services with industry standard interfaces such as V.35
- An ATM UNI interface
- A frame relay interface
- An access service which functions as an IP router
- An access service which functions as an APPN Network Node

The Access Agents are made up of several major elements. The Protocol Agent provides the interface to the end user system and initiates requests for services to other elements of the Access Agent. The Directory Agent interfaces with the BNS directory system to locate users and resources in the network. The Connection Agent establishes network connections with remote users via their local Access Agents.

Note that Access Agents provide interfaces between like systems as shown in Figure 5. In this example, a native ATM user is connected to a remote ATM user while three APPN users are connected through APPN Access Agents. Note that the APPN Access Agents would appear as standard APPN Network Nodes to their users.

IBM's BNS/ATM Products

Along with the announcement of the BNS architecture, IBM discussed future plans for BNS/ATM products. These products include VLSI ATM chip sets, and ATM products for both the LAN and WAN environments.

The basic building block of IBM's ATM products is a VLSI ATM switching element which is capable of switching at speeds in excess of 8 gigabits/second. That equates to the ability to switch 32 million packets/second.

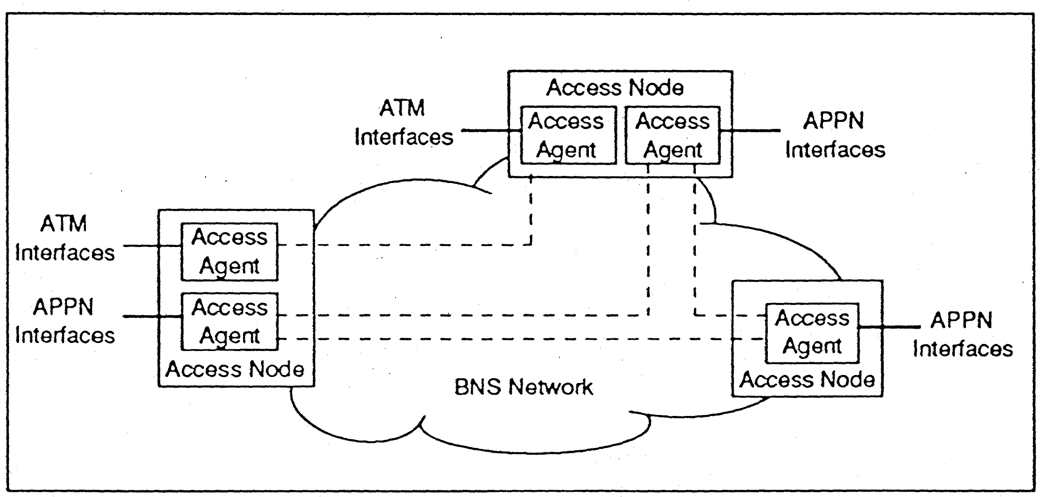

Figure 5

IBM has stated that ATM upgrades for the IBM/Chipcom 8250 hub will be delivered during 1994. This ATM support will include the LAN emulation services of BNS to provide interoperability between ATM-attached systems and existing LAN-based servers and workstations.

ATM hub products will provide the type of local area connectivity shown in Figure 6.

IBM will also deliver an ATM wide area switching product during 1994. The WAN switch is called the Transport Network Node (TNN). The TNN switch will include BNS bandwidth management support which is particularly important in the wide area environment where bandwidth is a recurring cost. TNN wide area networks will be capable of carrying both fixed-length ATM cells and variable-length Packet Transfer Mode (PTM) packets simultaneously. Figure 7 shows a network which is configured using both LAN and WAN BNS/ATM products.

How will existing IBM networking products interface with these new ATM LAN and WAN products? There are several possibilities. First, any systems which are currently attached to Token-Ring, Ethernet, or FDDI LANs could interface via the 8250 ATM switching support which will provide transparent interoperability between LAN-attached and ATM systems. Systems which require a WAN interface to ATM networks can use the frame relay access agent support which will be included in the TNN product. The following are some of IBM's products which support the frame relay interface:

- 3745
- 3174
- 6611
- AS/400
- RouteXpander/2

An ATM interface will be added to the 3172 which will provide the mainframe connection to ATM networks. Some of the other interfaces that will be provided by the TNN product include:

- HDLC
- ISDN
- Fiber Channel
- Constant bit rate voice interface

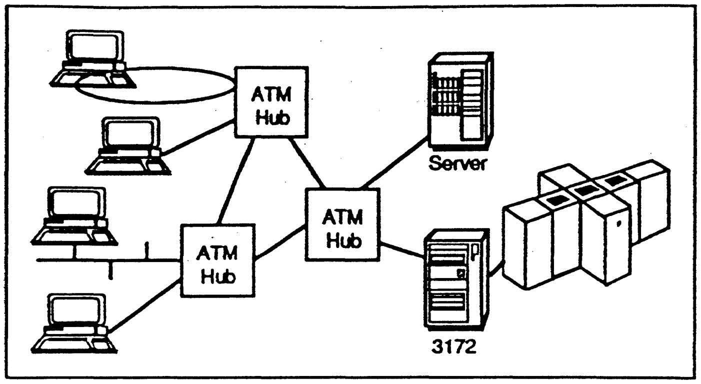

Figure 6

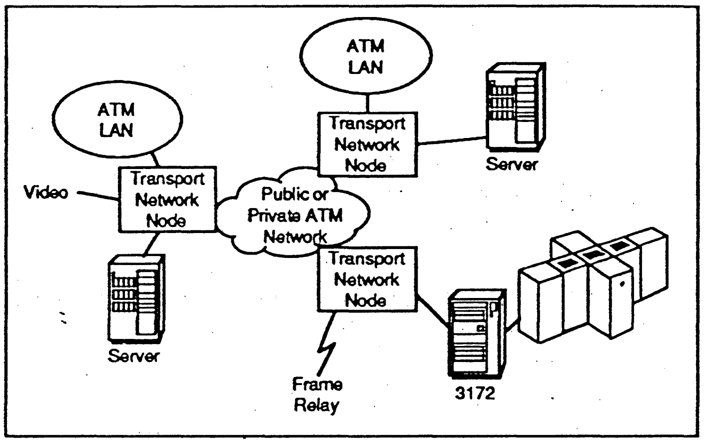

Figure 7

The bottom line is that most existing networking products, whether they are from IBM or other vendors, will be able to interface with the new BNS/ATM networks even if they don't have native ATM interfaces.

Summary

BNS appears to be a very comprehensive architecture for building broadband networks. Other ATM vendors, including Adaptive and others, have developed their own techniques for dealing with at least some of the issues that are addressed by BNS. As all of these vendors move forward with their own designs, the ATM Forum continues work on various aspects of ATM networking. Most of the ATM vendors, including IBM, are very active in the ATM Forum and their current proprietary solutions are developed with an eye toward compatibility with the industry standards as they emerge.

It remains to be seen whether IBM is successful in having elements of BNS incorporated into the industry standards, but in any case BNS could, at the very least, coexist with almost any standards that ultimately emerge. The bottom line is that IBM's ATM product plans reflect the "new IBM's" very positive position on support for industry standards while keeping the door open to add substantial value to broadband networking products. ■

(continued from page 1)

LANs and WANs are then tied together with internetworking devices such as bridges and routers, or by dedicated, single protocol devices such as the 3745/NCP.

The result is a crazy quilt of technologies and products that is very difficult to configure and manage. As networks scale up, the problems grow even more complex. One of the benefits of ATM is that a single technology can be used in both LAN and WAN environments, and ATM networks do not rely on devices such as bridges and routers to tie the various segments of the network together. Internetworking devices such as routers will still have a role in ATM networks, acting as firewalls and access devices, but they are less likely to be the "glue" that holds the network together.

ATM networks are also capable of providing the high bandwidth connections required to support image-based applications and video. ATM links and switches can operate at gigabit speeds and beyond. This bandwidth is also scalable which means that each user can access the network at a speed that is appropriate for its applications. By contrast, today's shared media LANs require that all users operate at the same speed. LANs also share a fixed amount of bandwidth among all users. As the number of users on a LAN segment increases, the average bandwidth available to each user decreases.

Realtime (isochronous) traffic such as voice or video require very low and very consistent transport delays across networks. Providing such low latency connections is also a strength of ATM. ATM's short, fixed-length cells and high speed switches can ensure that connections with very short transit delays can be established. Packet switching technologies such as X.25 or frame relay cannot effectively support isochronous traffic. Shared media LANs are also not well suited to supporting realtime traffic because of contention among users for the fixed amount of bandwidth that is available.

We will now look at how ATM networks deliver these networking services.

Structure of ATM Networks

High speed cell switching systems form the backbone of ATM networks. ATM networks are made up of switching systems which are interconnected by physical transmission links. As information enters the network it is segmented into fixed length cells. All information including data, voice, and video is carried in identically formatted cells. This greatly simplifies the design of the switching nodes. The sole task of the switches, except for management functions, is to switch cells from inbound links to the appropriate outbound link. ATM is designed to allow these switching functions to be implemented in hardware which results in high performance.

Virtual Connections

ATM is a connection-oriented communications service which means that a logical connection must be established between pairs of users before they can begin to communicate. In ATM networks, a logical connection between a pair of users is called a Virtual Connection (VC) . Figure 8 (right) shows how Virtual Connections are used to provide a logical communications path between users.

In many respects Virtual Connections are similar to the virtual circuits used within X.25 and frame relay networks. One fundamental difference, though, is that VCs are unidirectional pipelines between users. If bidirectional communications is required, two VCs must be allocated. Because separate Virtual Connections are used to support communications in each direction, the amount of bandwidth allocated

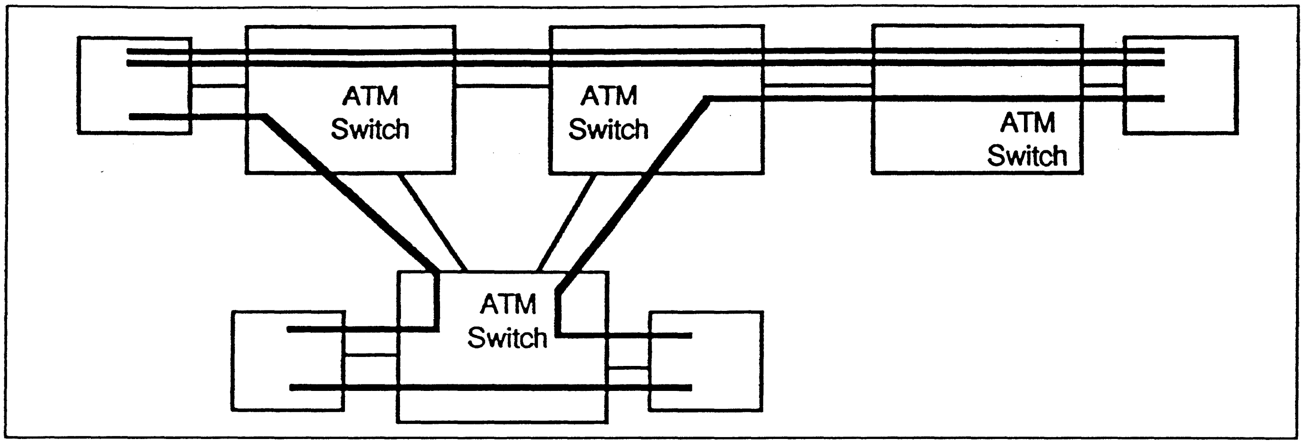

Figure 8

can be different in each direction. This is known as asymmetrical bandwidth allocation.

In data networking environments it usually makes sense to allocate bandwidth symmetrically when a connection is established between a pair of users, but the multimedia applications which ATM is designed to support can have very different requirements. An extreme example might be a video application which would require large amounts of bandwidth in one direction, but little or none in the opposite direction. Another benefit of the unidirectional allocation of bandwidth is that Virtual Connections in each direction might be routed over different physical paths across the network. Each path could have different cost or quality of service characteristics in addition to different bandwidth characteristics.

Routing Traffic Within ATM Networks

The end-to-end Virtual Connections are routed across the network through a series of concatenated logical connections known as Virtual Paths and Virtual Channels. Figure 9 shows the relationship between Virtual Channels and Virtual Paths. Each Virtual Path represents a collection of Virtual Channels which are switched as a unit. Multiple Virtual Channels can be multiplexed through a single Virtual Path.

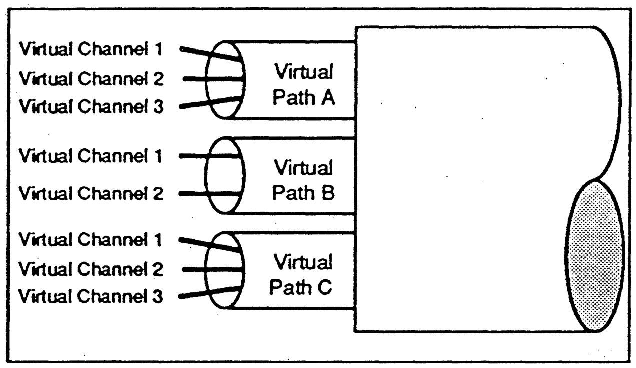

Figure 9

The header of each cell contains both a Virtual Path Identifier (VPI) and a Virtual Channel Identifier (VCI). These identifiers are assigned locally at each hop across the network and they only have local significance. Label swap routing is performed on each hop, similar to the technique used in APPN routing.

Figure 10 (page 12) shows how Virtual Connections between network users are supported by a concatenated series of Virtual Paths and Virtual Channels. Note that Virtual Paths can exist between adjacent nodes or they can traverse one or more intermediate switches. When a Virtual Path passes through an intermediate switching node, its Virtual Channels are not demultiplexed; they are all switched based on the VPI. When a Virtual Path does terminate in a switching node, the Virtual Channels are demultiplexed and then each of them is switched independently based on the VCI.

Virtual Paths can simplify network administration and management by allowing managers to deal with

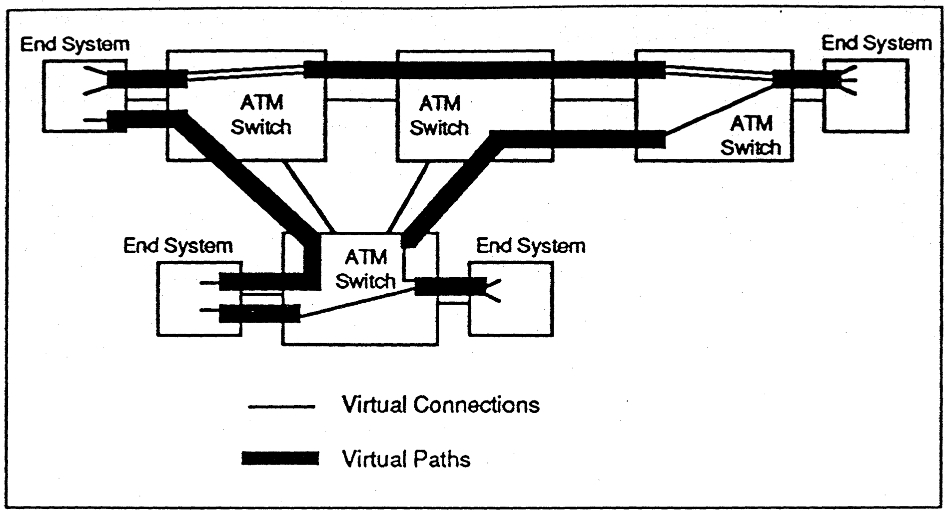

Figure 10

Virtual Paths rather than each of the individual Virtual Channels which might exist in a network. This is particularly convenient when configuring the network and describing route characteristics. The use of Virtual Paths can also improve the performance of the network by simplifying switching decisions. The set-up times for Virtual Connections can also be improved by activating Virtual Paths on a semi-permanent basis. When the required Virtual Paths are already active, the Virtual Channels needed to support a Virtual Connection can simply be mapped to the available Virtual Paths. In this type of environment the Virtual Paths would be relatively static while Virtual Channels are created and destroyed on demand.

Format of ATM Cells

All information flowing through ATM networks is segment into fixed-length packets which are called cells. ATM cells are 53 bytes (octets, technically) in length. Short, fixed-length cells are used in order to ensure very short transmission delays across the network and to simplify the design of switching equipment. Short and relatively consistent latency is particularly important in supporting real-time information such as voice and video. Figure 11 shows the format of an ATM cell.

Each cell is made up of a 5 byte header and a 48 byte payload.The header contains the Virtual Path Identifiers (VPIs) and Virtual Channel Identifiers (VCIs) that we have already discussed. The VCI is 16 bits in length. The VPI is 8 bits in length for cells which are flowing between user systems and the switch through which they access the backbone of the ATM network. The interface between the user system and the network is called the User-Network Interface (UNI).

Cells flowing between switches within the network have a slightly different VPI format. The interface between switches is called the Network-Network Interface (NNI) and cells flowing across this interface user a 12 bit VPI field. The 4-bit Generic Flow Control (GFC) field which is carried in the UNI cell format is replaced by an extended VPI field for cells flowing across the NNI interface.

All cell headers contain a payload type field which is used to identify the type of information being carried. There is also a Cell Loss Priority indicator which is used to identify cells which are eligible to

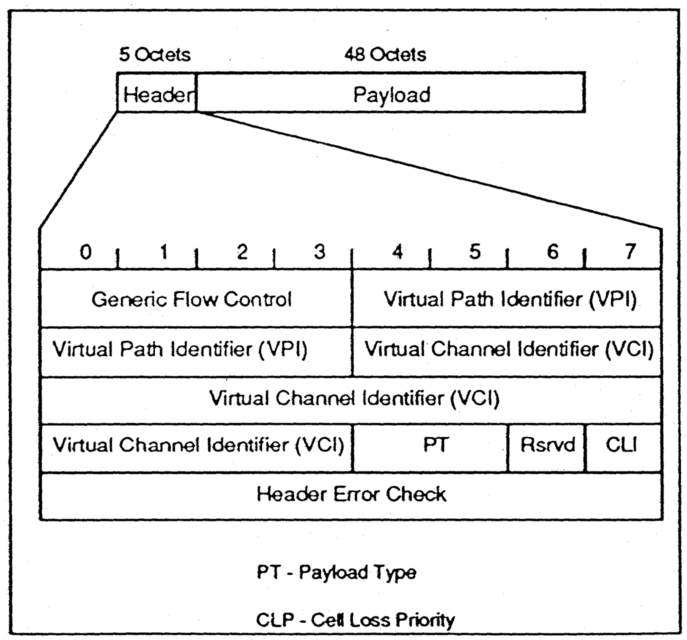

Figure 11

be discarded during times of network congestion. The entire cell header is error checked by the Header Error Control Field. The header is error checked to ensure that cells are not misdirected, but the payload is not error checked. The integrity of the information in the payload is the responsibility of the network users.

Interfacing to ATM Networks

Since ATM networks are capable of carrying a wide variety of information, including data, voice, and video, a standard structure must be provided for interfacing the users of these very different types of information to ATM networks. The overall structure is defined by the Broadband ISDN (B-ISDN) standard as shown in Figure 12.

When people talk about ATM they are usually refering to this entire structure of which ATM is a major component. The B-ISDN standards define a layer, called the ATM Adaptation Layer (AAL), which is responsible for mapping the various types of information to and from the ATM-defined cell formats. AAL-to-AAL protocols are also added to the information stream to support the transmission of various types of information across the ATM network. For example, information may be added to allow the multiplexing of multiple end user data streams over a single ATM Virtual Connection.

Note that when calculating the overhead involved in transmitting information over an ATM network, the overhead that may be added by the AAL is frequently overlooked. It is common for network designers to add only the overhead of the 5-byte ATM header for each 48-byte payload. It is important to remember that the payload may also contain overhead information which is added by the AAL.

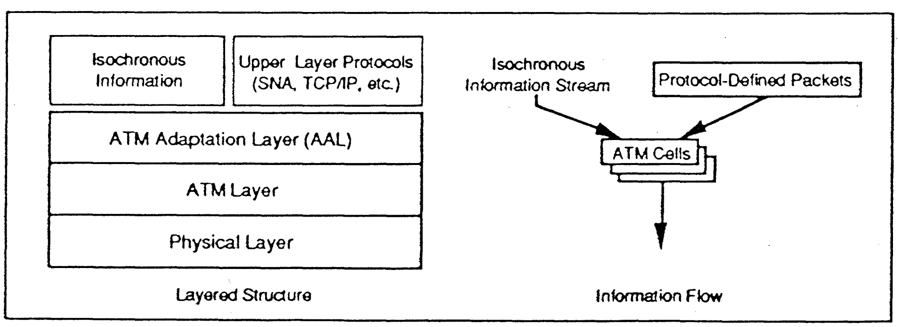

Figure 12

The B-ISDN standard defines four general classes of information that can be sent across an ATM cell-based network. Each of these service classes is supported by one or more types of AAL as shown in Figure 13.

	Class A Circuit Emulation	Class B Variable Bit Rate Service	Class C Connection-Oriented Data	Class D Connectionless Data
AAL Type	AAL 1	AAL 2	AAL 3/4 or AAL 5	
Source/ Destination Timing	Related		Unrelated	
Bit Rate	Constant	Variable		
Connection Mode	Connection Oriented			Connectionless

Figure 13

The service classes differ in their support for three key characteristics of information transfer across networks.

- Source/destination timing relationship - some types of information transfers rely on a synchronized timing relationship between the sender and the receiver.
- Bit rate - some types of information have a constant bit rate while others do not
- Connection mode - some types of information transfer are connection-oriented while others are connectionless

Note that the service classes A through D are also sometimes designated as classes 1 through 4. Some documents also omit the AAL Type 5. This is due to the fact that AAL5 was not part of the original CCITT standards for Broadband ISDN. AAL5 was added by the ATM Forum to provide more efficient support for high performance data communications.

Service class A is designed to support traffic which requires a close timing relationship between the sender and receiver, and which operates at a constant bit rate. A good example of this type of traffic would be a standard video signal. Class B is similar to A except that the bit rate of the traffic varies over time. Voice connections would fit into this category.

Classes C and D are designed to support data transmission. This traffic can tolerate considerable delays in network transit time because the senders and receivers are not synchronized with one another.

It is important to keep in mind that the ATM Adaptation Layer defines only general levels of service. There are no standard programming interfaces defined by the AAL.

Interfacing to ATM Networks

The currently-defined interfaces between network users and the ATM network are cell-level interfaces which operate over a variety of media and speeds. The ATM Forum has defined the following user-network interfaces (UNIs):

- DS3 - 45Mbps
- Multimode fiber - 100 Mbps
- Multimode fiber - 155 Mbps
- Sonet STS3c - 155 Mbps

Each of these interfaces uses the ATM cell format which was previously show in Figure 11.

The ATM Forum has recently finalized work on a different type of interface which is called the Data Exchange Interface (DXI). DXI allows the attachment of non-ATM systems to an ATM network. Figure 14 shows how the DXI interface can be implemented in a device such as a DSU. In this case the DSU implements the AAL functions including the conversion of user data into ATM cells and vice versa.

Summary

To be sure, ATM is not the only potential solution for applications which require high bandwidth and multimedia support. Emerging technologies and products such as switching hubs address many of the same issues. Some these solutions will also enjoy a cost advantage over ATM, at least in the short run. Despite all of the choices available to network designers we believe that ATM will be the clear standout as a strategic enterprise networking technology because of its scalability and manageability. ■

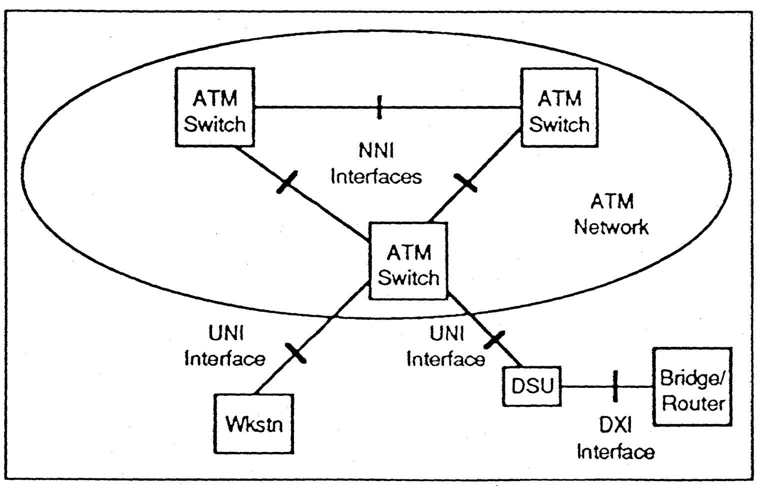

Figure 14

Consultant's Corner

Thumbs up Cisco!

This month I break with my long standing policy of focusing on technology and avoiding mention of specific companies (other than IBM of course).

This month's focus is Cisco.

In the history of technology evolution (and revolution), there are often notable turning points - "red letter" days. One of those days occurred in August. On that day, Cisco, the market leader in IP routers (and, I would add, also the market leader in tunneling SNA over IP backbones) joined the APPN camp and abandoned APPI!

What does this mean? To cisco? To IBM? To the market?First, confusion is lessened. The market leader in SNA (IBM) and the market leader in routers (Cisco) have converged their SNA routing strategies. Both vendors and users alike hear a clear message. The basis of SNA routing is APPN.

Second, Cisco catapults itself into the hearts of SNA users. SNA users want SNA, period. And, SNA means APPN. Cisco no longer finds itself in the position of missionary selling (which selling costs marketing $$$). More than anyone, cisco is the direct beneficiary of its own decision.

Third, Cisco (and other vendors) free up research and development budgets (and technical training and mind share). Rather than develop three technologies (including APPI), two can be developed (DLS and APPN). The most aggressive vendors canreduce the development budget to one - APPN. (After all, APPN with DLUR/DLUS enhancements does everything that DLS does plus more, yes? And DLS for NetBIOS? MPTN does NetBIOS.) Quicker time to market. Reduced support burden. Cash reserve for development of value added enhancements.

Fourth, IBM's credibility as an open company is enhanced. (Note also that IBM is the licensor of APPN technology to cisco.) Cisco (and APPI partners) agreesto work within the APPN Implementors Workshop (AIW) toward multivendor SNA consensus. IBM, learning from the APPN experience, is well positioned to broaden its open posture for APPN enhancements - HPR - , and then onward beyond SNA - Broadband Network Services (BNS).

Fifth, the APPN marketplace benefits big time. Vendors can develop products without wondering whether an alternative technology is preferred. Users can plan with a single solution in mind. All can train and develop expertise in a single technology - APPN. The whole is indeed greater than the sum of the parts. As evidence of this effect it should be noted that within a week of Cisco's APPN announcement, another major vendor, Well fleet, also announced an APPN product plan based upon licensed technology from Data Connection Limited (DCL). This brings to five (in order of announcement) the router vendors with announced APPN routing product plans - IBM, 3Com, Crosscom, Cisco, and Well fleet.

A year ago it would have been difficult to predict this outcome.

Looking back, when Cisco defined its first PU4-based SNA strategy, it was an understandable (though not excessively visionary) attempt by Cisco to give users more of what they already had (1974-1982 vintage subarea SNA backbones). When Cisco inaugurated its second SNA strategy with the APPI battlefield campaign, the goal was apparently to seize (some would say "muscle") the initiative for SNA routing marketplace dominance (along with cisco's APPI allies). Certainly Cisco's APPI attack loosened up IBM, though some would say (such as myself) that the opening up was coming anyway. APPI accelerated a natural evolving process. (Hindsight is 20-20?) Though APPI had been wavering of late, it certainly was not dead. It took courage on Cisco's part to reverse course, kill APPI, and confederate SNA strategies with IBM.

Was the net-out of APPI positive or negative? Will IBM hold steady and play it open from here out? Will Cisco resist the temptation to maneuver the

AIW in APPI'ish directions? Perhaps in the spirit of fostering ongoing peace it is best not to speculate.

Thumbs up Cisco.

And, congratulations to the APPN marketplace!.■

Given the acceptance of APPN by cisco and the APPI community - here are proposed new challenges for SNA development in the next two years.

- Finish DLS - keep it pragmatic and simple, don't get dragged into the not-a-bridge, not-a-router complexity pit. Look toward 802.1G and APPN/DLUR for mainstream technical solutions.
- Deploy APPN in routers, with statistical buffer reservation for APPN/ISR, then evolve to APPN/HPR.
- Upgrade end systems stacks to APPN-ENp
- Finish and deploy DLUR - DLUR completes APPN.
- Fix holes in APPN - standardize the APPN/IP encapsulation method. Define and standardize RTP data link layer - both PPP over RTP and ANR over PPP.
- Define and deploy MPTN, especially NetBIOS-MPTN (as an alternative to NetBIOS DLS caching/tunneling) and MPTN-3270 (an alterative to tn3270). ■