SNA over Frame Relay
- The Frame Relay Forum -
- The APPN Implementers Workshop -

September 30, 1997

David Sinicrope
IBM Corporation
Networking Architecture and System Design
P.O. Box 12195
Research Triangle Park, NC 27709 USA
david@raleigh.ibm.com

Ralph Case
IBM Corporation
eNetwork Architecture
P.O. Box 12195
Research Triangle Park, NC 27709 USA
rcase@vnet.ibm.com

Marcia Peters
IBM Corporation
eNetwork Strategy
P.O. Box 12195
Research Triangle Park, NC 27709 USA
mlp@vnet.ibm.com

Contents

Acknowledgements

The authors would like to thank all members of the APPN Implementer's Workshop and the Frame Relay Forum who contributed to the review and publication of this paper. In particular the following people dedicated an exceptional amount of time and effort:

- Marc Bernstein, Bay Networks
- Jim Cobban, Nortel
- Lori Dreher, Lucent Technologies
- Gary Dudley, IBM
- Larry Greenstein, Nuera Communications
- Mark Kaplan, Newbridge Networks
- Michelle Olesiejuk, Frame Relay Forum
- Harry Silverstone, Bay Networks
- Peter Tam, Nortel
- Ed Tremblay, IBM

Terms and Definitions

Terms and Abbreviations	*Definitions*
ANR	
	Automatic Network Routing - In High-Performance Routing (HPR), a highly efficient routing protocol that minimizes cycles and storage requirements for routing network layer packets through intermediate nodes on the route.
APPC	Advanced Program-Program Communications
APPN	
	Advanced Peer-Peer Networking - An extension to SNA featuring (a) greater distributed network control that avoids critical hierarchical dependencies, thereby isolating the effects of single points of failure; (b) dynamic exchange of network topology information to foster ease of connection, reconfiguration, and adaptive route selection; (c) dynamic definition of network resources; and (d) automated resource registration and directory lookup.
APPN End Node	Node that hosts applications at the edge of an APPN network
APPN Network Node	Node that provides network control and packet forwarding in an APPN network
CP	Control Point
DLCI	Data Link Connection Identifier
FRAD	Frame Relay Access Device
HDLC	High-level Data Link Control
HPR	High Performance Routing
	An addition to APPN that enhances data-routing performance and session reliability.
Logical Unit	The entity in a node that enables users to gain access to network resources and communicate with each other.
LEN node	Low-Entry Networking node
NCP	Network Control Program
PU	Physical Unit
PVC	Permanent Virtual Circuit
RTP	Rapid-Transport Protocol
	A connection-oriented, full-duplex transport protocol for carrying session traffic over High-Performance Routing (HPR) routes.
Session	A logical connection between two network accessible units that can be activated, tailored to provide various protocols, and deactivated, as requested.
SNA	Systems Network Architecture
SSCP	System Services Control Point

SDLC	Synchronous Data Link Control
SVC	Switched Virtual Circuit
Transmission Group	A connection between adjacent nodes (a logical link)
Type 2.0 Node	A node that attaches to a subarea network as a peripheral node and provides a range of end-user services but no intermediate routing services. (e.g. cluster controller (3174), workstation/PC, router or FRAD)
Type 2.1 node	A general category of node that includes APPN network node, an APPN end node, or a LEN end node It can also attach as a peripheral node to a subarea boundary node in the same way as a type 2.0 node.
Type 4 Node	A node that is controlled by one or more type 5 nodes. It can be a subarea node, or, together with other type 4 nodes and their owning type 5 node, it can be included in a group of nodes forming a composite LEN node or a composite network node. (subarea SNA Communications Controller, e.g. 3745 w/ Network Control Program (NCP))
Type 5 Node	A node that can be configured to be any one of the following: • APPN end node • APPN network node • LEN node • Interchange node • Migration data host (a node that acts as both an APPN end node and a subarea node) • Subarea node (with an SSCP) Together with its subordinate type 4 nodes, it can also form a composite LEN node or a composite network node. (subarea Host Node, e.g. System 390/Virtual Terminal Access Method (VTAM))
VC	Virtual Circuit
VTAM	Virtual Telecommunications Access Method

Introduction

With the recent growth of frame relay as the data networking industry's leading wide area technology and the strength of System Network Architecture (SNA) as the industry's leading protocol suite for mission critical applications, it seems a natural match to use these two technologies together. To date many SNA networks have been modified to use frame relay as a substitute for analog and digital leased lines. The primary motivation has typically been network hardware and transmission line cost reduction. However, as frame relay continues to develop, with many implementations or services supporting features such as quality of service and switched virtual circuits, it is important to note the specific requirements and features of the frame relay payload protocols and how frame relay can best meet the requirements and enhance the features.

This white paper is intended to assist those who wish to consider Frame Relay as a transmission protocol for their SNA data traffic. The document gives an overview of the issues to be considered and decisions to be made. The intent of the document is to provide a high level overview that can be used as a starting point to focus on issues of particular concern.

Systems Network Architecture (SNA) and its successors Advanced Peer-Peer Networking (APPN) and High Performance Routing (HPR), are extremely broad topics which stretch well beyond the scope of this white paper. This paper focuses on aspects of SNA related to using frame relay as a transmission medium to carry SNA traffic. It attempts to explain just enough about SNA control, routing, error recovery, congestion management and class of service to understand the associated frame relay issues. For more details on SNA components, internal workings, or SNA in general, please refer to Systems Network Architecture Technical Overview [2] or the APPN Overview [1]. These documents provide a high level starting point for those interested in learning more about SNA and contains references to the related architecture references that are the doorway to the world of SNA algorithms, GDS variables and control vectors.

This document assumes the reader already has a reasonable understanding of frame relay and data networking. See the other Frame Relay Forum white papers and education modules for information on frame relay.[1]

The remainder of this paper gives an overview of the SNA market and a brief history of the SNA technologies. The paper briefly describes each of the SNA technologies highlighting the distinguishing features. The document then goes through the parts of SNA that interact with frame relay and the standards formats used to transport SNA over frame relay. Finally, the document discusses the issues significant to SNA traffic and how these can be addressed when using frame relay.

[1] See "References" on page 22 and ***http://www.frforum.com/*** for details on the frame relay technologies.

The SNA Market

While the growth and glamor of the Internet and its associated TCP/IP protocol suite eclipses SNA in the trade press daily, SNA unobtrusively continues its vital role as the workhorse of enterprise networking. SNA applications and networks exist in enormous numbers today, running critical business and governmental applications worldwide. New SNA applications and networks are being deployed, and SNA will continue to grow for a long time.

It is estimated that over twenty trillion dollars have been invested in SNA applications in over 40,000 enterprises worldwide. According to current surveys, SNA accounts for 61% of wide area network enterprise traffic and 68% of enterprise WAN budgets. Contrary to the image portrayed by some of the trade press, SNA is alive and well. Fifteen years of annual surveys find no decrease in SNA penetration or any significant plans to convert SNA applications. SNA remains a vital solution for customers in their mission-critical applications. In fact, it continues to grow, with a reported 4.7 million units of SNA client software shipped in 1995 and an estimated 5.38 million in 1996. Existing single-enterprise SNA networks may have as many as one million terminals and logical units and an average of 435,000 active sessions. [4]

Customers have come to depend on the stability, predictability, security, reliability, dependability, interoperability, and high resource utilization that SNA networks provide, and they increasingly want the high availability and performance provided by APPN/HPR.

SNA provides a base that promotes reliability, efficiency, ease of use, and low cost of ownership; enhances network dependability; improves end-user productivity; allows for resource sharing; provides for network security and resource management; protects network investments; simplifies problem determination; accommodates new facilities and technologies; and lets independent networks communicate. SNA can be very frugal with expensive networking resources such as links. With careful tuning, link utilizations as high as 98% have been reported. SNA also allows for extremely large networks: enterprises with tens to hundreds of thousands of attached terminals and applications are not uncommon. All these features make it a favorite for mission-critical corporate and governmental applications.

Traditionally, SNA networks have been connected across the wide area using leased lines or switched facilities running SDLC or X.25 protocols. As networks grow, the cost of link facilities and hardware grows tremendously. Also the expense to manage multitudes of links is quite large.

One of the key factors driving the installation of frame relay connections is the conversion of SDLC leased lines to frame relay PVCs. One frame relay access link can multiplex several connections that required many adapters. The savings on access hardware and leased line charges can be dramatic. Frame relay is usually added as a software-only upgrade to products that already support SDLC leased lines. Due to the efficiencies of frame relay, it is usually the preferred data link technology.

Another factor is that it is more feasible, technically as well as economically, to fully mesh an SNA network with frame relay than with leased lines. This gives more paths to the network and consequently better reliability.

With the advent of multiprotocol networks, frame relay became essential, allowing the multiplexing of SNA and other protocol traffic over the same virtual circuit, allowing savings in hardware and leased line cost.

History and Background

IBM's first release of SNA in 1974 did for networking what System/360 had done for IBM computing a decade earlier. It brought order by providing commonality and structure through a single architecture for data communications, and ended the anarchy of the multitude of disparate methods and link protocols then in use for connecting devices to host systems. Originally designed for the "glass house," subarea SNA's hierarchical structure connected many simple devices to one powerful mainframe. IBM added multiple-host networking in 1977 and transmission priority in 1980. Priority allowed more important (e.g., interactive) traffic to proceed before less time-critical (e.g., batch) traffic, improving link utilization. In 1982 IBM introduced Advanced Program-to-Program Communication (APPC) so applications could embrace the new distributed transaction programming paradigm.

While APPC let programmers write distributed programs, the original hierarchical SNA network structure inhibited any-to-any connectivity, since all data had to flow through one or more host-controlled subareas. To address this, IBM introduced SNA's second generation, Advanced Peer-to-Peer Networking (APPN) in 1986.

APPN is an open data networking architecture that has decentralized control with centralized network management, allows arbitrary topologies, has connection flexibility and continuous operation, and requires no specialized communications hardware. It replaced the coordinated system-definition required in subarea SNA with automatic configuration definition, and fully embraces the peer-to-peer and client-server paradigms. It provides sophisticated route selection, dynamic topology updates, and accommodates existing subarea networks. In 1994 IBM added the Dependent Logical Unit Requester (DLUR), allowing APPN networks to carry subarea SNA traffic. Recognizing that customers were best served by an open architecture, in 1993 IBM sponsored the first APPN Implementers' Workshop (AIW), a consortium of networking vendors sharing an interest in APPN. As the standards body for SNA technologies, the AIW continues to meet three times a year. The latest APPN standards can be found on the World Wide Web at ***http://www.networking.ibm.com/app/aiwhome.htm***.

To improve APPN availability and performance, IBM developed High-Performance Routing (HPR). This third-generation SNA is a fully-compatible upgrade to APPN. Building upon APPN's topology and directory services, HPR adds nondisruptive rerouting, improved routing performance, and rate based congestion control, while reducing memory and processor use in intermediate nodes.

In 1992 and 1994, IBM developed Peripheral and Extended Border Nodes for partitioning very large networks into smaller subnets.

In 1996, the AIW approved "HPR Extensions for ATM Networks." This standard lets users exploit Asynchronous Transfer Mode Quality of Service from existing SNA applications.

In 1997 IBM added native multi-link transmission groups to HPR products. This popular feature from subarea SNA tunes network capacity by aggregating low-speed links, dials extra bandwidth on demand, and maintains the integrity of a transmission group despite individual link failures.

We are starting to witness universal access to the corporate network from any client or browser. New linkages to the corporation's most valuable information resources, the corporate MIS databases, are enabling electronic commerce to thrive. Even as companies take advantage of the explosion of internet-based services, SNA preserves the continuing immense value of their mission-critical applications.

Different Types of SNA

As mentioned above, SNA has progressed into three distinct models. The three models are subarea SNA, APPN, and HPR. Each is described briefly below. In all models, all communications are connection-oriented. A session is established between partner applications for data to flow through the network.

Subarea SNA

Subarea networks are the most classic form of an SNA network. They are characterized by hierarchical network roles. They typically involve a host node (Type 5 - e.g. System 390 mainframe with VTAM), communications controllers (Type 4 - e.g., 3745/NCP) and several peripheral nodes (Type 2 or 2.1 - e.g., cluster controllers or workstations/PCs, etc.). All communication is mediated by the Type 5 nodes which contain a System services Control Point (SSCP).

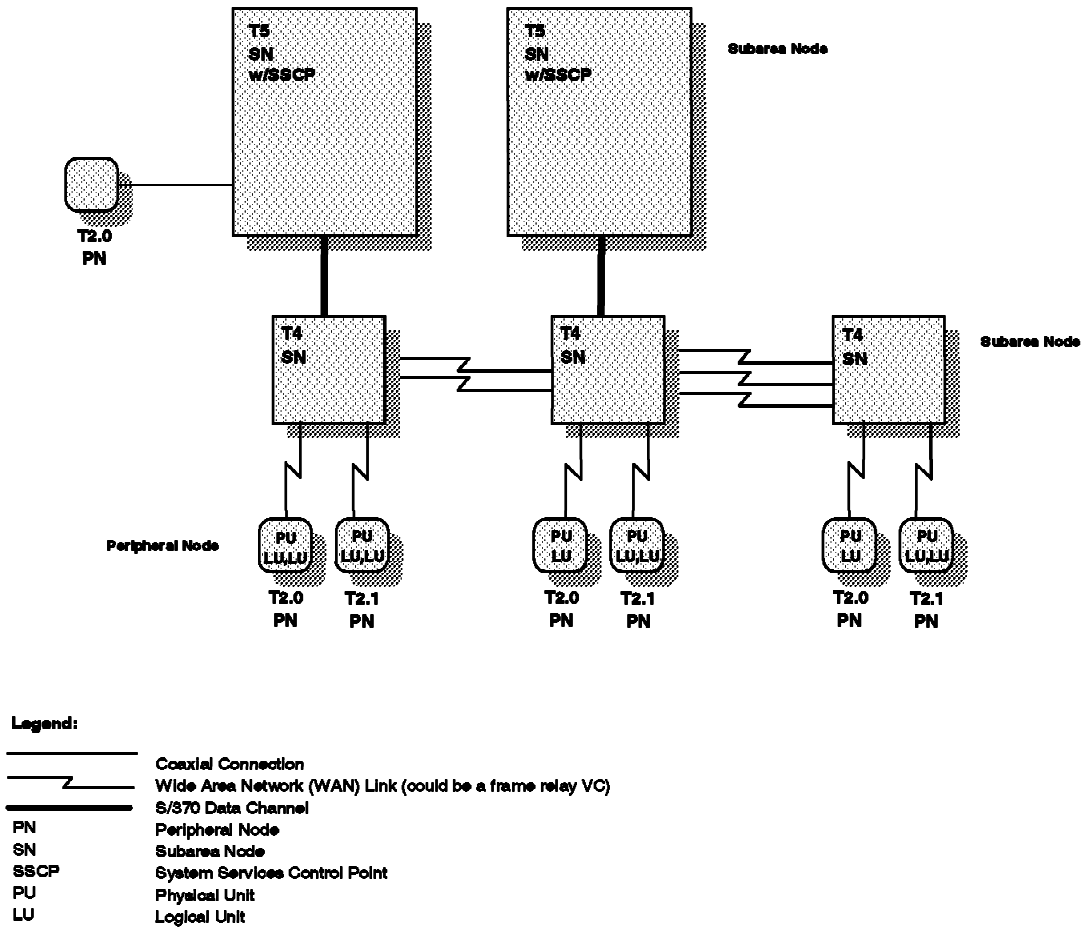

Figure 1. Subarea Hierarchical Network

Each node in a subarea network contains a Physical Unit (PU). The Physical Unit is responsible for performing local node functions such as activating and deactivating links (e.g., frame relay connections, or leased lines) to adjacent nodes. To do this, the PU must exchange control information with the controlling

SSCP. Once the necessary links are activated, the programs or terminals can exchange data using sessions between Logical Units (LUs). These sessions are also controlled by the SSCP in the Type 5 nodes.

The hierarchical nature of subarea networks has some disadvantages in that the centralized control of communications results in static routes and a great deal of configuration.

Subarea SNA depends on the data link control layer (typically SDLC or LLC2) for reliable delivery of packets from one node to the next. The DLC provides window flow control on an individual linkbetween nodes. SNA also uses flow control end to end to ensure that intermediate nodes do not get congested and applications are not overrun.

Routes are predefined as a series of "hops" between any two nodes in a subarea SNA network. Route definitions are static, and ordered according to desirability. When a session is initiated, the first available route from the list of predefined routes is selected for the session traffic. The list of available routes, and transmission priorities for those routes, depends on the Class Of Service (COS) for the new session, so sessions of different COS may be assigned to different routes. If any link or node along the route fails, the session is terminated and the user or application can start a new session, if desired.

Advanced Peer-Peer Networking (APPN)

As the name suggests, APPN enables nodes to communicate without requiring mediation by a Type 5 node. This gives the network better connection flexibility, scalability, and reliability.

The APPN extensions to SNA distribute the network control into many Control Points (CPs). Each Control Point has a partial responsibility for many of the same functions of the SSCP, and others as well. The functions include locating and discovering routes to partner nodes and selecting which particular routes to use. This relieves network personnel from having to configure locations and routes. Control Points exchange topology and directory information among themselves using Control Point-to-Control Point (CP-CP) sessions.

Unlike subarea SNA, there are only two types of APPN nodes: end nodes and network nodes. End nodes support the APPN protocols through connection to a network node and are located at the periphery of the APPN network. Network nodes and their interconnecting links form the intermediate routing network. They are responsible for interconnecting the end nodes. For example, they perform route selection based on:

1. the information given to them by the end nodes,
2. directory information accumulated through searches and
3. topology information exchanged among themselves and other network nodes.

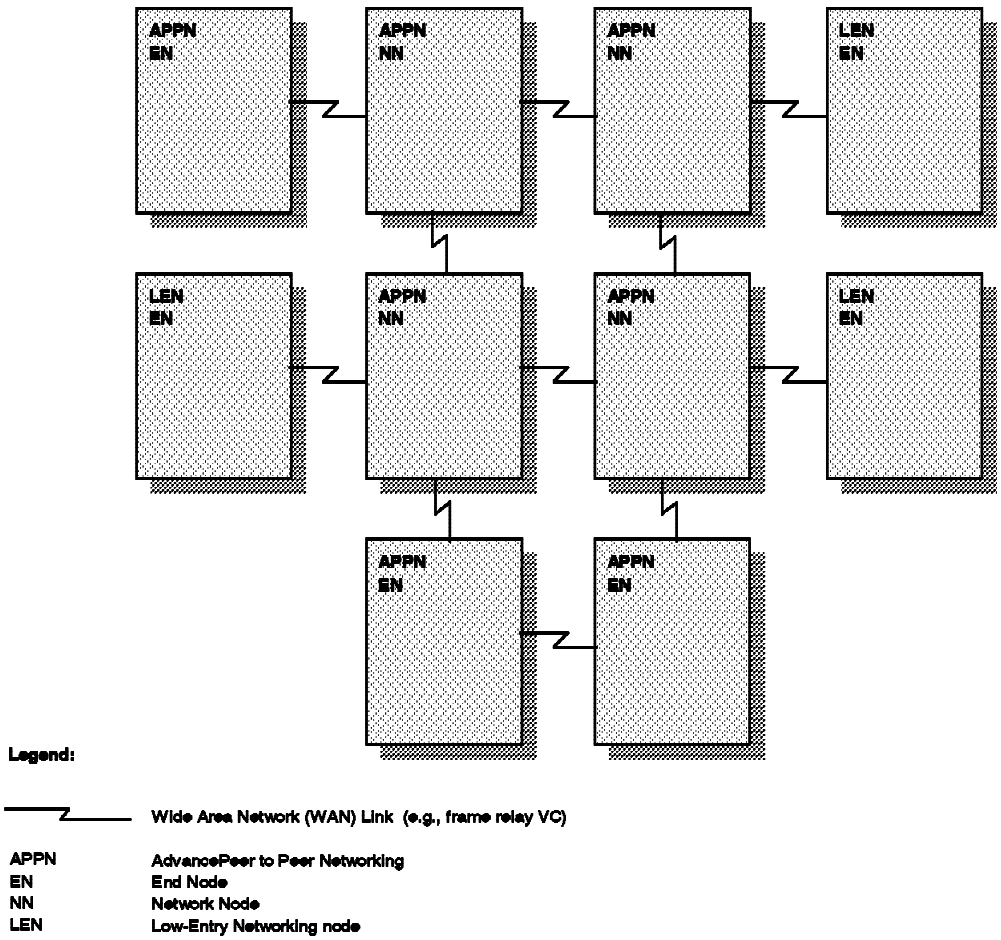

Figure 2. APPN Network

APPN/ISR network nodes use Intermediate Session Routing (ISR) protocols to forward packets along a pre-determined path through the network. This path, a series of “hops” from node to node, is determined dynamically when the session is set up but does not change for the duration of the session. If any link or node along the route fails, all sessions using that resource are terminated and must be restarted by the end user or application. Like subarea SNA, APPN with ISR depends on the data link control layer (DLC) between adjacent nodes over each individual “hop” to ensure reliable delivery of packets. The APPN/ISR frame formats are compatible with subarea SNA frame formats used to attach peripheral nodes. Addresses on each link are only locally significant, and APPN/ISR performs label swapping at the intermediate nodes.

APPN/ISR uses IEEE 802.2 Logical Link Control type 2 (LLC2) to ensure error recovery above frame relay. This protocol is similar to other HDLC based protocols (e.g., SDLC) in that it sends and receives acknowledgments for some number of frames before continuing to send more. It also supports retransmission of lost or corrupted frames. If a frame is lost or corrupted, LLC2 asks that the sender retransmit the lost frame and all the frames sent after it. Although very effective for older, less reliable transmission lines, this means of retransmission is inefficient for today's high speed, high quality lines.

APPN/ISR uses the same high level flow control used for pacing sessions in subarea SNA, but it is used on each hop to obtain maximum utilization of each link while avoiding congestion at any node. This requires that buffers be allocated at each intermediate node, and the number of required buffers increases relative to the number of sessions, the bandwidth, and the propagation delay of the links.

APPN selects session routes based on Class Of Service (COS) requirements. Data with different attributes can be sent over different paths with different attributes. For example batch traffic may be sent over a satellite link with low cost, high bandwidth and long propagation delay, while credit card transactions are sent over a secure path with low propagation delay.

High Performance Routing (HPR)

The High Performance Routing (HPR) extensions to APPN use the existing APPN control algorithms for locating resources and selecting routes. It also adds additional features for transporting the data. These features exploit today's high-speed, high-quality links, more powerful end systems and standard protocol switched backbones.

HPR minimizes the processing required in intermediate nodes using Automatic Network Routing (ANR) and Rapid-Transport Protocol (RTP). RTP provides end-to-end transport between any two end points in the network. Intermediate nodes are not aware of SNA sessions or RTP connections. ANR is a source-routing protocol. Each network-layer packet is routed based on the information in the packet. See Figure 3.

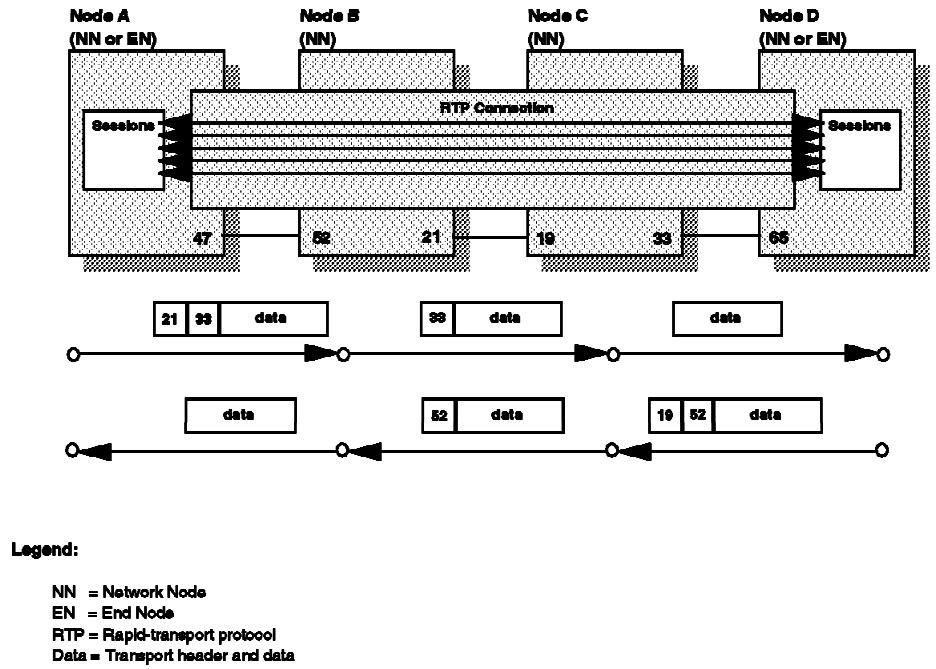

Figure 3. Automatic Network Routing, with Rapid Transport Protocol for SNA sessions

Rapid-Transport protocol is used above ANR to ensure error recovery and retransmission end to end. This eliminates the intermediate nodes from having to do route look-ups, participate in hop-by-hop error

recovery, and reduces the processing and buffers needed in the intermediate nodes. RTP uses selective retransmission, meaning that only lost frames are retransmitted, making more efficient use of network bandwidth.

Another feature of HPR is Adaptive Rate-Based (ARB) flow and congestion control. To ensure that the sender is not sending data to a congested network or receiver, the Adaptive Rate-Based algorithm exchanges information between the two end points of an RTP connection. This information tells about the state of the network and other end system so the sender can regulate the amount of data it sends accordingly and avoid congestion. Since ARB handles flow control end-to-end, there is no need for LLC windows; data packets are sent using connectionless services (e.g., LLC Type 1).

APPN with HPR uses the same COS-based route selection algorithm as APPN/ISR. The route is determined when the session is set up, but if there is a failure along the path, e.g. PVC status indicates inactive, the RTP end points dynamically find and switch to a new path without disrupting the sessions, as opposed to dropping the sessions as in subarea SNA and APPN.

SNA and Frame Relay protocols

Relative to the SNA protocols, Frame Relay is used as a transmission medium, (see Figure 4). It is used to transport the SNA traffic of one or more *transmission groups or TGs* between SNA nodes. A TG is a connection between two adjacent SNA nodes that is identified by a transmission group number.[2]

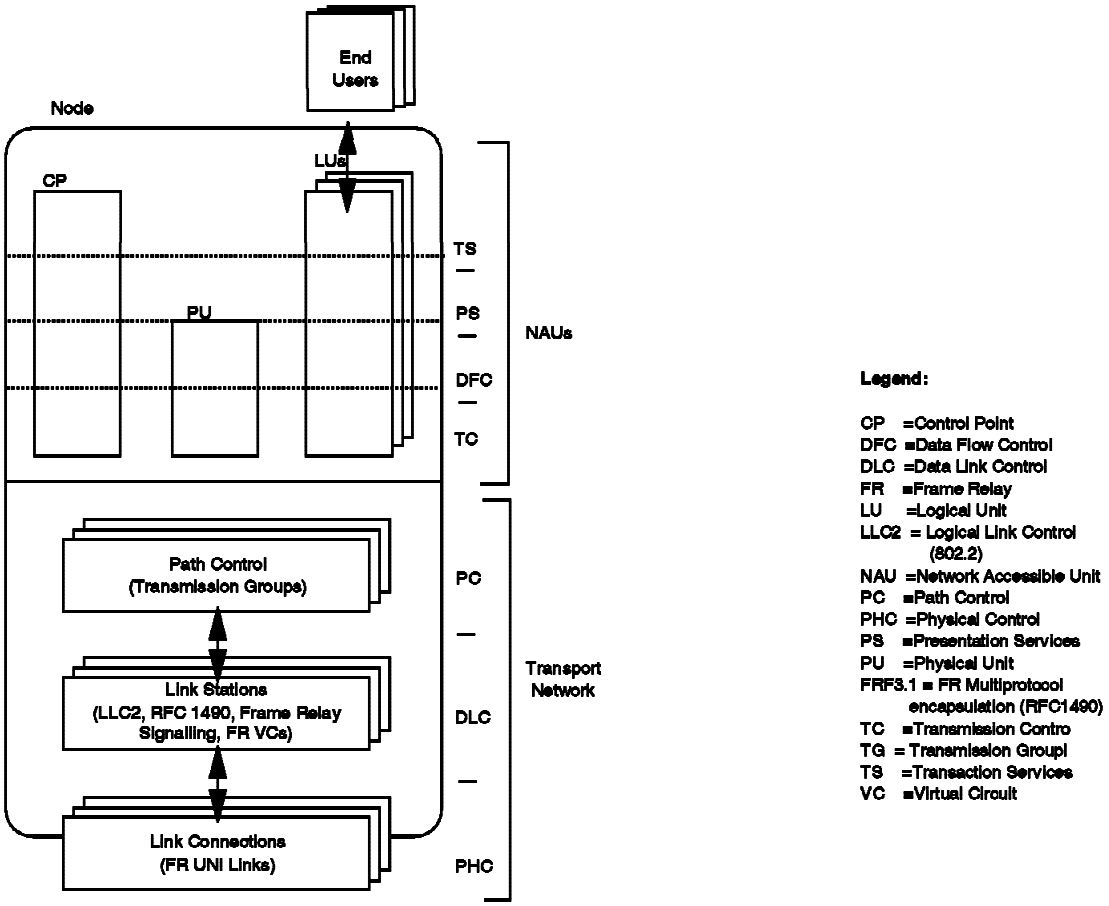

Figure 4. SNA and Frame Relay Protocols

One of the features of frame relay is that it exploits today's high-speed, high quality transmission lines by not performing hop-by-hop error recovery as in X.25. The frame relay network puts the responsibility for error recovery on the end equipment if it needs it. Because SNA typically carries mission critical data, error recovery is needed over frame relay. The layer directly above frame relay is typically Link Layer Control

---

[2] A TG also collectively refers to the hardware and software in each of two connected SNA nodes that control the particular connection(s), along with the transmission media between the nodes. More than one connection may connect the same two SNA nodes. In this case, each of the connections may be used for a distinct TG, or the connections may be combined in parallel to form a single TG. This is referred to as a *multi-link TG (MLTG)*.

For example, when two SNA nodes are connected using frame relay, the nodes refer to the physical adapters and the logical connection between them as a TG. Each of the physical components in a TG may also be shared with other TGs.

(IEEE 802.2) type 2. This layer provides reliable retransmission of any lost or corrupted frames. For APPN/HPR this function is performed by the Rapid Transport Protocol (RTP).

Above the layer providing reliable transmission are the the SNA layers responsible for control (CP-CP and SSCP-PU) sessions. These layers control link status, routing and topology, and data sessions. Above these layers are the data sessions (LU-LU) that are used to transfer data from user to user or program to program. For more information on these layers please see [1] and [2].

Multiplexing SNA Connections over Frame Relay

Normally networks are configured such that a router or Frame Relay Access Device (FRAD) is attached between a frame relay network and a local area network. TGs from SNA nodes on the LAN to partners across the frame relay network are concentrated over frame relay virtual circuits (VCs). There are several choices for how the TGs are actually related to frame relay links, and the choices have implications on the function and cost of the network design. The main choices are described in this section.

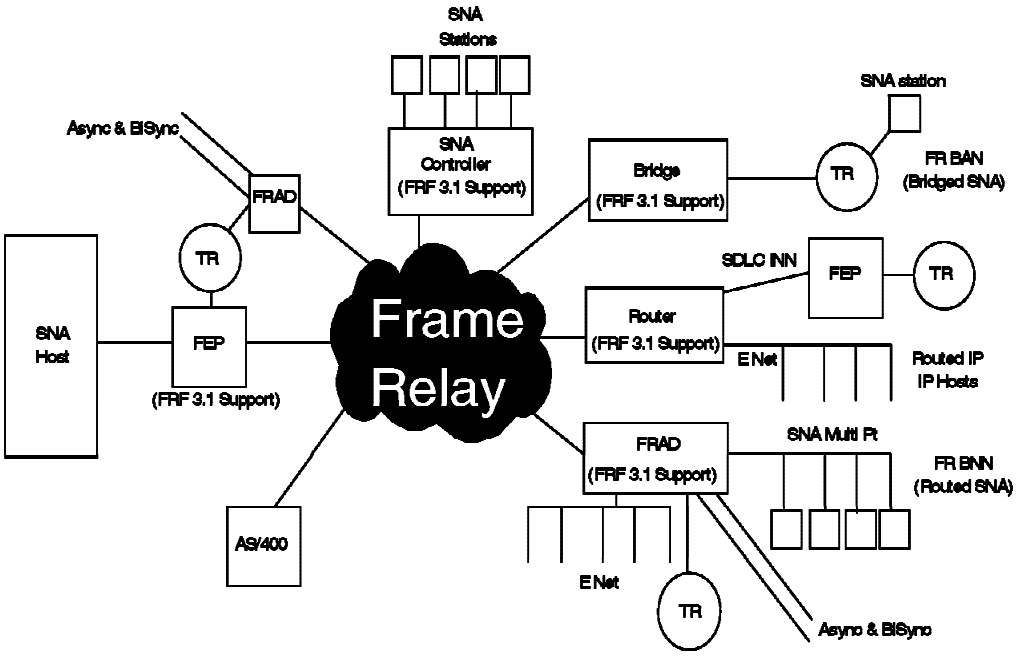

Figure 5. Frame Relay Network Transporting SNA

One VC per Transmission Group

It is possible to configure one frame relay Permanent Virtual Circuit (PVC) per TG. However, this is not practical except in small, private network scenarios, as the cost of one PVC per TG in a public network would be prohibitive. In addition, the amount of system definition required for a larger or growing network may quickly become overwhelming.

SAP Multiplexing

Service Access Point (SAP) Multiplexing is typically used with the Boundary Network Node (BNN) format described below. Using the frame relay multiprotocol routed frame format, there are 127 valid 802.2 SAPs that can be used for multiplexing on any frame relay virtual circuit. Each TG is identified by a unique pair of SAPs. The advantage of this technique over one VC per TG is that larger capacity frame relay VCs may be better utilized; versus multiple smaller capacity frame relay VCs that may be under utilized. However, for LAN campuses that have a large number of workstations, there may not be enough SAPs available to represent all of the SNA TGs. In addition, the SAP-TG associations are typically made through system definition, which becomes complex in a large campus frame relay device and at a host or central site. Since the Media Access Control(MAC) addresses of the workstations are not used on the frame relay VC, network management has no way to know the MAC addresses of the various workstations. They are hidden from the remote network management.

MAC Multiplexing

MAC Multiplexing uses the Boundary Access Node (BAN) format described below. The LAN MAC address as well as the SAPs can be used to differentiate between different TGs. The advantage to this approach is that the MAC address used to identify the TG on the LAN is the same MAC address used to identify the TG on the frame relay connection. Because all MAC addresses are unique and are carried over the frame relay connection, little to no mapping definition is needed. This means that BAN scales well to large campus configurations. In addition, because the MAC address is known to the communications controller or host-end router, the network management also has knowledge of this device and can individually address the device. However, the MAC addresses within the BAN format also add overhead to each frame.

DSPU Support

DownStream Physical Unit (DSPU) support is another technique used by some vendors to multiplex traffic from multiple LAN TGs over a frame relay connection. In this technique the device with access to the frame relay network contains a single SNA device. All traffic from the LAN devices is made to look like Logical Units (LUs), i.e., program and terminal traffic, within the access device. The benefit to this approach is that it overcomes the SAP limitation, while still using the low overhead BNN format. However, the relationships between LAN addresses and access device LUs must be defined. Also, because the access device is mapping LAN addresses to its own LUs, the network topology behind the access device is hidden from network management.

Formats and Functions for SNA over Frame Relay

When frame relay is used as a TG or link, the SNA packets are typically encapsulated in one of the formats defined in the Frame Relay Forum Multiprotocol Encapsulation document (FRF.3.1). These formats are commonly known in SNA terms as *Boundary Access Node (BAN)* and *Boundary Network Node (BNN)*. As an alternative, SNA packets may be encapsulated in other network layer protocols such as TCP/IP to be transmitted across the frame relay network as in Data Link Switching (DLSw). Each of these methods is described briefly below. More detail on the SNA formats and functions can be seen in the IBM Frame Relay Guide. [7]

Aside from the standard formats and functions below, several vendors have developed proprietary methods for transporting SNA over frame relay. The proprietary methods typically have some enhanced features over and above the methods listed below. Details on these methods can be obtained from the individual equipment vendors listed on the Frame Relay Forum World Wide Web site.

Boundary Network Node (BNN)

Frame relay Boundary Network Node (BNN) uses the routed frame formats given in Frame Relay Forum Multiprotocol Encapsulation (FRF 3.1). See Figure 6 on page 14. The codepoints in these formats identify the layer 2 protocol (e.g. LLC2) as well as the SNA frame format (e.g. subarea SNA, APPN or HPR) encapsulated in the frame.

For subarea SNA and typically for APPN, LLC2 is used for error recovery across a frame relay connection. When using SAP Multiplexing, one LLC2 connection typically spans the LAN and the frame relay link to ensure error recovery. For APPN/HPR, Rapid Transport Protocol (RTP) is used for end-to-end error recovery.

The strength of BNN encapsulation is its low frame overhead. It uses the fewest bytes per frame of any of the other formats discussed in this paper. The drawback of BNN is its limited TG multiplexing capability.

This encapsulation is typically used for SAP multiplexing, DSPU and APPN devices because of its low frame overhead. The SAP multiplexing and DSPU techniques above address the limited TG multiplexing capabilities of this format. APPN devices route traffic to individual SNA devices and do not typically need TG multiplexing support.

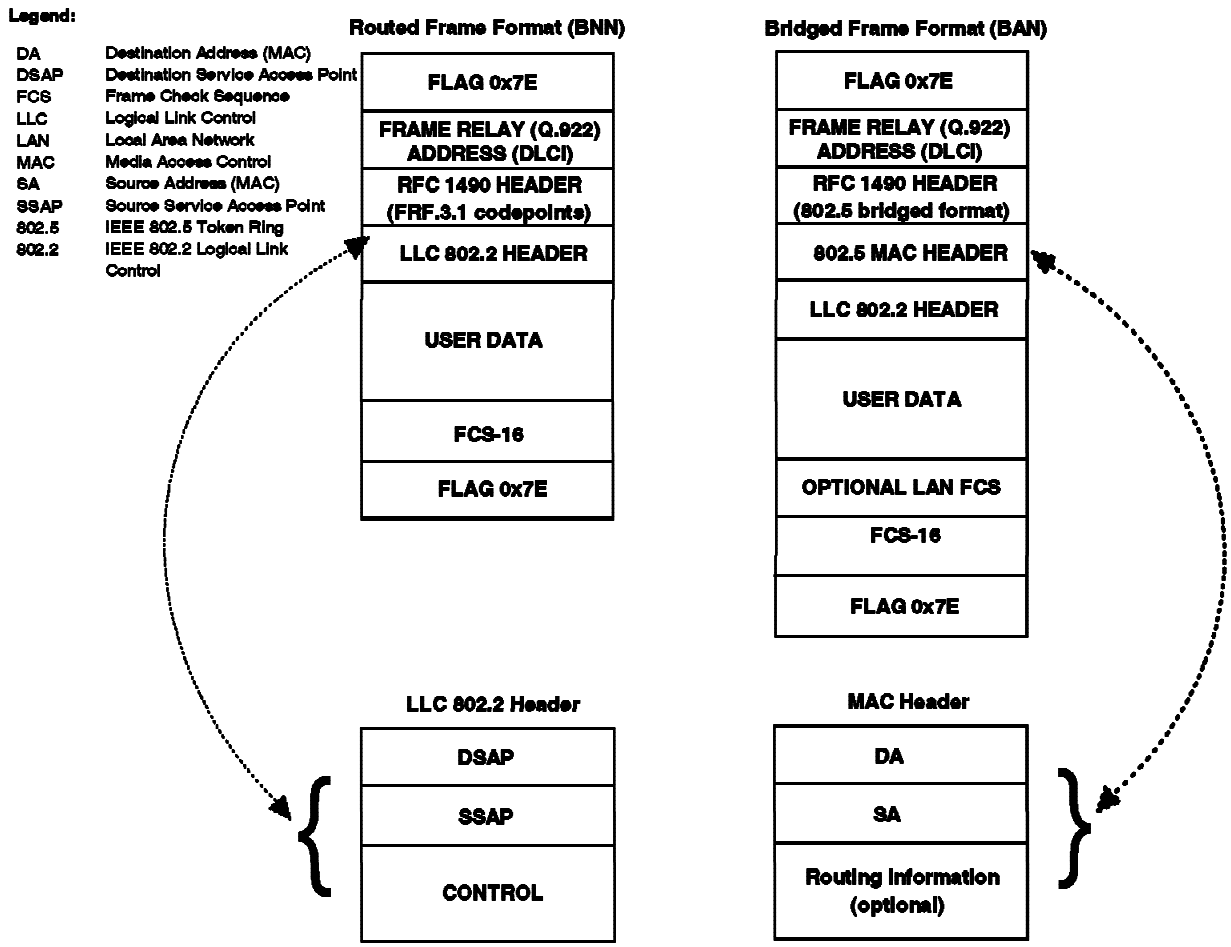

Figure 6. BAN and BNN Frame Formats

Boundary Access Node (BAN)

BAN (Boundary Access Node) is essentially LAN bridging and uses the bridged frame formats in the Multiprotocol Interconnect over Frame Relay specification (RFC 1490). See Figure 6. The codepoints within the RFC 1490 header identify 802.5 bridged frames.

One LLC2 connection may used across the LANs and frame relay to ensure error recovery. In this case the LLC2 connection might time out if the delay across the LANs and frame relay networks is too long or the LLC2 timers are set incorrectly.

Alternatively, a BAN device may terminate LAN LLC2 connections and multiplex the SNA traffic across frame relay networks using different LLC2 connections. This function alleviates the timing constraints of the LLC2 connections. This is similar to DLSw, but does not require the overhead of TCP/IP encapsulation.

BAN is used to support MAC multiplexing described in “MAC Multiplexing” on page 12. BAN was designed to address the limitations of BNN when supporting large numbers of LAN devices. There is no limit on the number of LAN stations supported for a VC. No configuration is needed when adding a LAN station. Network management alerts from LAN stations are fully identified using the MAC address and frame relay DLCI.

Because BAN can be done with LAN bridging, it is typically used to connect a LAN to frame relay when the function of a full SNA protocol stack is not needed.

The strengths of BAN encapsulation are its flexibility and scalability, typically making it the preferred choice.

The disadvantages of BAN are a slightly larger frame size and the lack of support in older versions of the Network Control Program (NCP).[3]

Data Link Switching (DLSw)

Data Link Switching (DLSw) is a technique for carrying SNA over a TCP/IP network. Frame relay comes into play when the TCP/IP network uses frame relay to interconnect its routers. In this way the SNA traffic is transported over TCP/IP over frame relay. This technique uses the IP encapsulation format described in RFC 1490 [14], as opposed to those defined for SNA in FRF3.1.

For error recovery, DLSw breaks the path of an SNA packet into segments. For example, the local LAN, the TCP/IP/frame relay network used to span the wide area, and the remote LAN. In the local LAN an LLC2 connection is used for error recovery and is terminated in the router connected to the frame relay network. The local router uses a TCP/IP connection to transmit the frame and ensure error recovery across the frame relay circuits and the intermediate TCP/IP routers. The remote router terminates the TCP/IP connection from the frame relay circuit and uses a another LLC2 connection to transfer the data across the remote LAN to the destination.

[3] Versions prior to NCP V7.3.

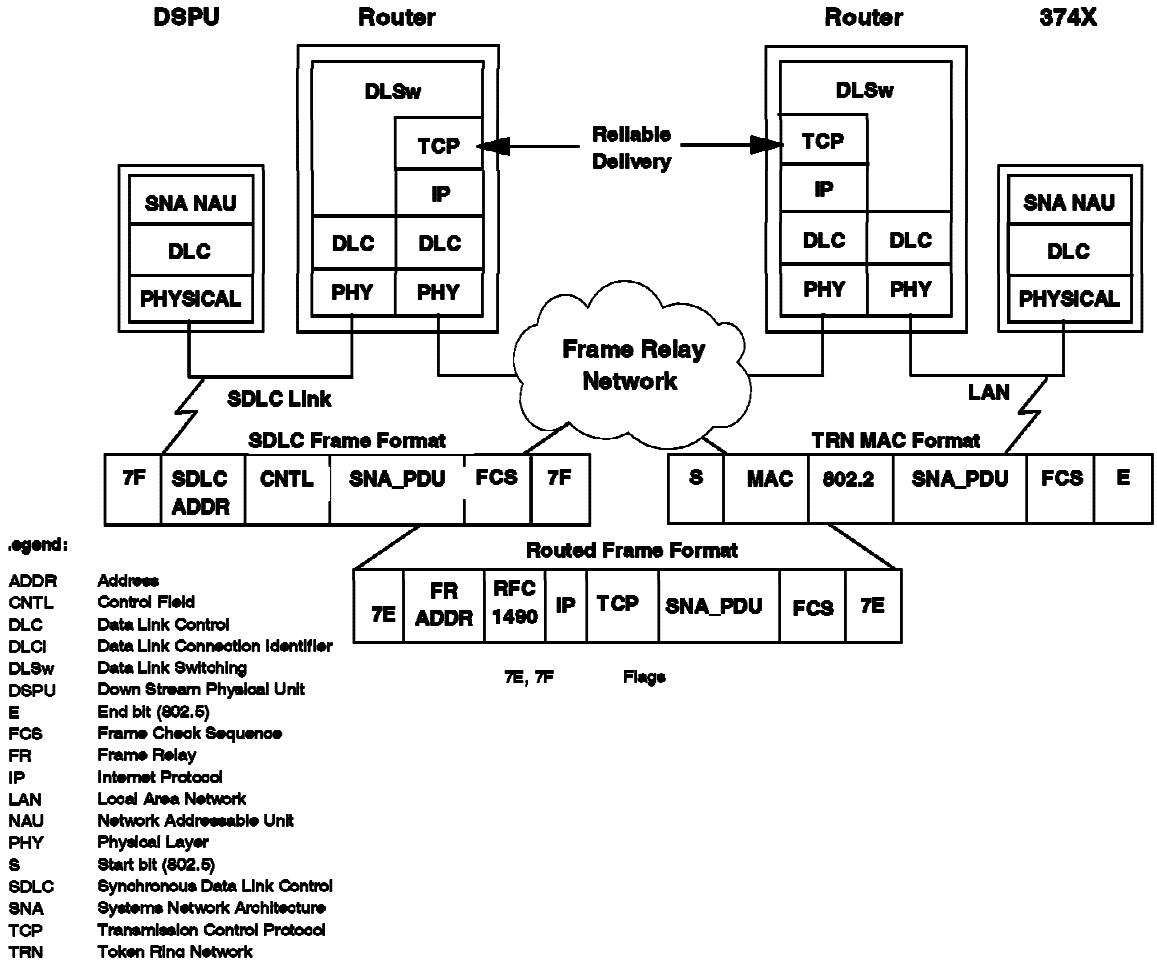

Figure 7. Data Link Switching (DLSw)

DLSw is beneficial in a predominately IP routed network where the percentage of SNA traffic is small. Its main strength is that it terminates the LLC2 connections at the edge of each LAN, reducing the amount of LLC2 acknowledgements across the frame relay network. Although not described in the DLSw specification, it is possible to encapsulate multiple SNA packets into one TCP/IP frame.

The major drawback of DLSw is the amount of per frame overhead, the highest of the formats presented here. Also, both SNA and IP traffic are tagged with the same multiprotocol encapsulation header. To the frame relay layer, all the traffic appears as IP traffic and differentiating between the IP encapsulated SNA traffic and the true IP traffic becomes difficult. This is important if the SNA traffic carries mission-critical data and is to be given priority over other IP traffic. See “Fairness with SNA and other Network Layer Protocols” on page 18.

In addition, the individual SNA priorities cannot be respected by intermediate routers. All SNA traffic to a destination is put onto the same TCP/IP connection regardless of whether it is an interactive terminal session or batch file transfer traffic. Although this causes no problems with the SNA protocols, the end-user response time may suffer.

SNA over FR Issues

Quality of Service

Delay

Subarea SNA and APPN/ISR are delay sensitive because they use LLC2. LLC2 was designed for point-to-point reliability and its timer defaults are set accordingly. LLC2 typically demands an information frame acknowledgement within 1-2 seconds. However, this timer is generally configurable (timer ACK_TIMER) and may be made longer. Increasing the acknowledgement timer is typically recommended when using one LLC2 connection across both the local and wide area media.

In general the LLC2 protocol timers should be configured based on average time for a frame to be sent to the remote site and to get an acknowledgment back (the round trip delay). The LLC2 protocol will recover in the abnormal cases where the actual delay exceeds the average delay.

Momentary degradation in delay will not generally be enough to drop an SNA session, although it may be noticeable to the end user. The LLC2 protocol will recover any frames that are not acknowledged in the allotted time period, usually transparent to the end user.

Since HPR uses RTP as a reliable transport instead of LLC2, its protocol timing constraints are considerably more flexible than subarea SNA and APPN. The amount of time it takes before RTP declares a frame lost is based on the round trip time it measures and is not configured or fixed as with LLC2.

Using priority frame relay VCs should be based on reducing the user response time and not satisfying SNA protocol constraints. For example, a high priority VC is desirable when it will carry predominately interactive (e.g., 3270) traffic. A high priority VC would not necessarily be used for non-interactive SNA traffic just to satisfy the LLC2 timers. By increasing the timer values a normal VC will work well. It is also important to note that putting all SNA traffic on a high priority VC does not guarantee low response times. See “SNA priorities” on page 18.

Loss

SNA is very well suited to frame relay with respect to loss. Subarea SNA, APPN and HPR all have robust error recovery procedures and will recover in nearly all cases of lost frames without any noticeable difference to the user in terms of response time or session loss. Subarea SNA and APPN use LLC2 which relies on go-back-N recovery. This means that if a receiver detects a lost frame (missing sequence number) it tells the sender to retransmit all frames sent starting with the one that was lost or corrupted. This is done transparently to the application protocol.

HPR/RTP uses a more sophisticated form of recovery called selective retransmission. In this method the receiver tells the sender to retransmit just the frame that is missing or corrupted. This method is more suited to frame relay in that higher speed, higher quality connections, and lower buffer memory cost[4] make recovery without unnecessary retransmissions a more desirable design point. DLSw also uses this form of retransmission.

[4] Selective retransmission requires a device to store all frames that arrive after a rejected frame until the rejected frame is retransmitted. This requires more memory for receive buffers than go-back-N which discards all frames arriving after a rejected frame and asks that they all be retransmitted. Selective retransmission was cost prohibitive until memory prices dropped enough to make it affordable.

These error recovery mechanisms allow the applications to continue functioning in cases where the network is heavily congested and is discarding frames.

SNA priorities

SNA has 4 traffic classes or priorities, network control, high, medium and low. These are set by the application based on its requirements. Network control is used for transport of network control traffic. High priority is used for interactive traffic like 3270 terminal traffic, and low for batch traffic like file transfers. The priorities must be respected when running multiple classes over the same frame relay VC and when running multiple protocols over the same frame relay VC. See “Fairness with SNA and other Network Layer Protocols.” For example, obtaining a priority PVC from a network provider may not improve response time if batch traffic is sent at the same priority on the PVC as interactive traffic. Even when giving SNA priority on a multiprotocol PVC, the internal SNA priorities must be respected or response time sensitive traffic may find itself queued behind a batch transfer. Only subarea SNA, APPN, or HPR nodes fully respect SNA priorities. Other devices (e.g., DLSw or FRAD bridges) that do not examine the SNA priority information embedded in the frames cannot respect the SNA priorities.

Fairness with SNA and other Network Layer Protocols

It is important to note that frames will be delivered in the order in which they are transmitted on the frame relay VC. They will not be reprioritized on the VC within the frame relay network. For frames to have special treatment within the frame relay network they must be sent on separate VCs which have distinguished service characteristics. This gives the distinguished VCs precedence for network resources above those that are not distinguished.

There are many implementation or service specific schemes for giving SNA traffic priority on the same VCs as other multiprotocol traffic. Most products use the FRF 3.1/RFC1490 header as a way to discern protocols and prioritize one over the other. This generally gives SNA traffic, which is often interactive, priority over traffic such as IP or IPX which is typically not as mission-critical. This also generally satisfies the LLC2 timer requirements of subarea SNA and APPN.

Ideally prioritization considerations should be based on the application as well as the protocol. For example, problems may arise by queuing TCP/IP Telnet traffic behind SNA batch traffic. However, even in this case, it is still important to provide SNA traffic with some minimum guaranteed bandwidth because of the LLC2 timer sensitivity. Also, in the event of network congestion, SNA will typically react sooner than TCP/IP. See “Congestion Control.” If the SNA traffic is not given some minimum bandwidth it could be "starved" by the TCP traffic which typically does not react until a packet is dropped.

Congestion Control

Congestion control is important to keep throughput up. If the frame relay network is congested, it makes more sense to reduce the amount of traffic offered to the network than to offer traffic that the network will discard and have to be retransmitted. There are two ways for SNA devices to control congestion.

Explicit Controls - FECN/BECN

When the network sets the FECN and BECN bits in the frame header, most SNA devices will reduce the LLC2 window size to control the amount of traffic offered to the network. In HPR, which does not use

LLC2, a FECN causes the Adaptive Rate Based (ARB) flow control to slow down.[5] Both of these mechanisms allow the network to recover and then traffic ramps back up to the contracted CIR and beyond, if allowed.

Implicit Controls - Adaptive Rate Based Flow Control (ARB)

HPR is uniquely suited to frame relay in that it periodically measures the round trip delay across the network and adjusts the offered traffic to the network accordingly. This avoids congestion rather than waiting for a FECN or a lost packet to react. It also maximizes throughput in the event that the frame relay network is lightly loaded.

Selective Retransmission vs. Go Back N

As mentioned above, HPR uses a selective retransmission method of recovery versus a go-back N. In a congested network this helps to further reduce congestion by not retransmitting frames that were not lost. Only those frames that are lost or corrupted are retransmitted.

Security

Security is frequently an issue when choosing frame relay equipment. Unlike leased line networks the media of a frame relay network is typically shared between many users subscribing to a frame relay service. Meeting security requirements can be done in several ways depending on the level of protection desired. For example, end devices may encrypt the data on a virtual circuit or the path to the destination may be chosen unpredictably to avoid security attacks. The level of security is dependent on individual network requirements.

HPR Extensions for ATM and FR/ATM Interworking

HPR extensions for ATM were developed to let SNA applications run directly over ATM networks, without an intermediate layer such as LAN emulation. APPN/HPR is given direct access to ATM signaling. ATM addresses are stored and distributed by the normal APPN topology algorithms and APPN link definitions can specify ATM throughput and QoS. “Smart” applications may be able to specify needed throughput and QoS for APPN to request from the underlying ATM subnetwork. In addition, APPN route selection understands ATM characteristics so that optimal routes can be chosen.

SNA over ATM can be interworked with frame relay PVCs using either of the two types of frame relay/ATM interworking defined: Network Interworking or Service Interworking.

Network Interworking

Network Interworking with ATM (FRF.5) [11] is transparent to the SNA protocols and applications. To SNA, all traffic looks like it is running over Frame Relay. The same congestion bits (FECN, BECN, DE), DLCI and multiprotocol encapsulation headers are maintained in both the frame relay and ATM segments of the connection.

[5] BECN is not used in ARB.

Service Interworking

For Service Interworking, the SNA multiprotocol encapsulation headers in FRF3.1 are translated by the ATM Service Interworking Function (FRF.8) [12] to the ATM equivalent (RFC 1483 [13]). The ATM congestion information is also translated to and from frame relay. In addition, the *HPR Extensions for ATM Networks* specification was also designed to handle service interworking with frame relay.

Connection Network Support and Switched Virtual Circuits (SVCs)

APPN/HPR is designed to take advantage of ATM's Switched Virtual Connection (SVC) capability. APPN's "connection network" function allows nodes that are connected to a common shared (e.g., LAN) or switched (e.g., ATM) facility to bring up direct connections with each other. These connections need not be defined in advance; they are dynamically created when needed for a new session.

In the same way that APPN and smart applications use ATM SVCs, they can also be adapted to use frame relay SVCs. This allows them to request wide area services tailored to the traffic and needs of the application. The use of frame relay SVCs is an important work item to be considered by both the AIW and the Frame Relay Forum.

Similarly, the SVC support can be extended to include FR/ATM SVC Interworking. The FR/ATM SVC Interworking specification is currently under development within the Frame Relay Forum. This feature is also an important work item under consideration by both the AIW and the Frame Relay Forum.

Summary

SNA is the predominate protocol of choice for mission critical applications and will continue to be in the foreseeable future because of the large investments in capital equipment, application development, training, and operational procedures. Frame Relay has established itself as the WAN protocol of choice that has the facilities to support SNA and its manageability, reliability and efficiency features most effectively and economically. The advent of QoS and SVCs in Frame Relay strengthens this bond by adding further support to the SNA features and flexibility in network configuration that SNA is uniquely qualified to advantage.

This paper has only given a brief overview on the topic of SNA and providing SNA services over frame relay. There is a wealth of information contained in the references listed below and is also available from those vendors who support and implement these protocols.

References

1. *Networking With APPN: An Overview* (IBM G325-0204-00).
2. *SNA Technical Overview* (IBM GC30-3073-04).
3. *White Paper - SNA over Frame Relay*, M. Bernstein, Bay Networks June 1996.
4. *Inside APPN: The Essential Guide to the Next-Generation SNA* (IBM SG24-3669-03).
5. *APPN Architecture Reference* (IBM SC30-3422-04).
6. *APPN High Performance Routing Architecture Reference Version 3.0* (AIW-HPR8).
7. *IBM Frame Relay Guide* (IBM SG24-4463-01).
8. *Frame Relay Forum User to Network Implementation Agreement*, FRF.1.1, Frame Relay Forum, January 1996.
9. *Frame Relay Forum Multiprotocol Encapsulation*, FRF.3.1, Frame Relay Forum, June 1995
10. *Frame Relay Forum User to Network SVC Implementation Agreement*, FRF.4, Frame Relay Forum, 1994.
11. *Frame Relay Forum Frame Relay/ATM PVC Network Interworking Implementation Agreement*, FRF.5, Frame Relay Forum, December 1994
12. *Frame Relay Forum Frame Relay/ATM PVC Service Interworking Implementation Agreement*, FRF.8, Frame Relay Forum, April 1995.
13. RFC 1483 *Multiprotocol Encapsulation over ATM Adaptation Layer 5.* , Juha Heinanen, July 1993.
14. RFC 1490 *Multiprotocol Interconnect over Frame Relay*, T. Bradley, C. Brown, & A. Malis, July 1993.
15. RFC 1795 *Data Link Switching: Switch-to-Switch Protocol AIW DLSw RIG; DLSw Standard Version 1.0*