SG24-4777-00

IBM Nways 2219, 2225, 2230 WAN Switches: Services and Technology

October 1996

International Technical Support Organization
Raleigh Center

International Technical Support Organization

SG24-4777-00

IBM Nways 2219, 2225, 2230 WAN Switches: Services and Technology

October 1996

Take Note!

Before using this information and the product it supports, be sure to read the general information in Appendix C, "Special Notices" on page 141.

First Edition (October 1996)

This edition applies to the IBM Nways WAN switches (2219, 2225, 2230) running Release 4.1 software or below.

Comments may be addressed to:
IBM Corporation, International Technical Support Organization
Dept. HZ8 Building 678
P.O. Box 12195
Research Triangle Park, NC 27709-2195

When you send information to IBM, you grant IBM a non-exclusive right to use or distribute the information in any way it believes appropriate without incurring any obligation to you.

Contents

Preface

This redbook covers in detail the newest members of the IBM Nways WAN broadband switch family. It provides information about the IBM 2219 frame relay switch, IBM 2225 multiservices switch, and the IBM 2230 ATM switch.

This redbook was written for networking specialists, customers, and business partners. Some knowledge of broadband WAN networking is assumed.

How This Redbook is Organized

This redbook is organized as follows:

- Chapter 1, "Wide Area Networking Requirements"

  This chapter starts by outlining the objectives of this publication. This is followed by a brief discussion of the key wide area networking issues facing network managers, administrators and operators. Finally, the IBM Nways 2219, 2225 and 2230 family of products is positioned in relation to the services and technologies that are nowadays driving wide area networking.

- Chapter 2, "Multiservices Concepts and Implementation"

  This chapter focuses on the frame relay, SMDS and ATM capabilities of each platform. Each type of service offered is briefly outlined. The prime objective is to explain the implementation of each service and, in some cases, a brief technical description of the service is given.

- Chapter 3, "Hardware Overview"

  This chapter contains a fairly detailed description, switch by switch, of the hardware characteristics of each switch and, if applicable, related software.

- Chapter 4, "Managing Networks Using Nways Wide Area Element Manager"

  This chapter is dedicated to network management. First, the general management framework is described; then the role of the Nways Wide Area Element Manager is detailed. Finally, there is a configuration section, where the main parameters for each network object are described, to illustrate the services discussed in Chapter 2, "Multiservices Concepts and Implementation."

- Appendix A, "Release 4.2 Preview for the Nways 2225"

  This appendix describes the enhancements planned for Release 4.2 of the IBM Nways 2225.

- Appendix B, "Frame Relay Supported Standards"

  This appendix describes the frame relay standards that are supported.

The Team That Produced This Redbook

This redbook was produced by a team of specialists from around the world working at the Systems Management and Networking ITSO Center in Raleigh, NC.

Donna Fox is an Advisory ITSO Specialist at the Systems Management and Networking ITSO Center.

Petr Havlik is a Networking Marketing Consultant for IBM Czech Republic.

Vasco Almeida is a Networking Solutions Specialist for IBM Portugal.

Thanks to the following people for their invaluable contributions to this project:

Thom Langlois and Ron Keohane
Cascade Communications Corporation
Westford, Massachusetts

Randy Coffelt, Greg Valicka, Pat Quinn, Robert Burton
Advantis
Schaumburg, Illinois

Christopher Metz
IBM New York

Tony McNeill, Bill Caldwell, Matt Trzyna
Networking Hardware Division Marketing, RTP

Tim Kearby
International Technical Support Organization, Raleigh Center

Philippe Blaevoet
Networking Hardware Division
IBM La Gaude

Patti Bullen
Networking Marketing Specialist
IBM Canada

Comments Welcome

We want our redbooks to be as helpful as possible. Should you have any comments about this or other redbooks, please send us a note at the following address:

redbook@vnet.ibm.com

Your comments are important to us!

Chapter 1. Wide Area Networking Requirements

This chapter starts by outlining the objectives of this redbook. This is followed by a brief discussion of the key wide area networking issues facing network managers, administrators and operators. Finally, the IBM Nways 2219, 2225 and 2230 family of products is positioned in relation to the evolving services and technologies of wide area networking.

1.1 Introduction

In this redbook, we describe the IBM Nways 2219, 2225 and 2230 Wide Area Networking Switches family of products. The objective was not so much to describe the products, but to describe them from a practical viewpoint. Thus the emphasis was put on the services provided by these products, which took precedence over architectural or implementation considerations; these considerations, when described, are intended to give practical information and illustrate how they contribute to the implementation of the required services.

The pace of change in networking hardware and software is such that one of the primary concerns of network managers and administrators has become how to build, manage and expand a network in a scalable, yet nondisruptive, fashion. In the WAN environment this need is even more strongly felt, since bandwidth does not come for free; indeed it can be very expensive and therefore has to be efficiently utilized.

The following sections outline some of the key networking issues and related consequences as briefly as possible, in order to not lose sight of the practical inclination of this redbook.

1.2 The Traditional Data Network

The overriding concern in this environment was to save cost on what was then expensive, low-speed communication lines; the immediate by-product of this concern was that communication lines had to be as fully utilized as possible.

This, in turn, drove the assumptions on which network exploitation had to be performed. To minimize retransmissions that would impact line utilization, hop-by-hop recovery was the rule; routing in the network could be handled by software, since the bottleneck resided in the link and not in the processing nodes; congestion and flow control mechanisms were key to keeping the network stable, while providing high link utilization; applications could not expect throughput rates higher than a few thousands of packets per second.

Although the technological assumptions that legitimized traditional data networks are irrevocably outdated, the fact remains that many low-speed lines are still in operation today, supporting mission critical applications that have to be run and networked for years to come.

1.3 The Evolution of Transmission Technology

The spectacular evolution in this area means that applications are no longer confined to transmission rates of 9.6 Kbps, 64 Kbps, or even 2.048 Mbps. Two almost unrelated advances have profoundly influenced available bandwidth expectations.

First, coding and compression techniques have rendered economic feasibility to digital processing and transmission, thereby enlarging the types of traffic and applications that users could realistically require.

The second was the emergence of fiber as an economically viable infrastructure on which networks could be built. So the speeds that we could expect rose by orders of magnitude, and at the same time, the error rate decreased by orders of magnitude. Concomitantly, technological improvements in the area of transmission over copper wires cannot be overlooked either, since they enabled higher bandwidth availabilities without incurring in the cost of infrastructure replacement.

The cost, but not necessarily the price, of transmitting a given piece of information has decreased dramatically and will probably keep decreasing in the years to come.

1.4 The Rising Demands of the User Community

The new technological developments folded nicely with long-standing user expectations. These expectations spread in four distinct but complementary directions.

The first was that new types of applications became possible (for example, voice, video, imaging, etc.).

The second was that it became technically possible and sometimes economically advantageous to integrate the different traffic types generated by these new applications seamlessly onto a common network infrastructure.

The third was that applications of the traditional type, using different protocols, could be run concurrently, since bandwidth availability could be bought to compensate for protocol inefficiencies or conflicts.

The fourth was perhaps the most important one, which was the transformation of the characteristics of the user workstation, opening up a world of possibilities in applications and related design issues.

In the old paradigm, the network constrained the user; in the new one, the user is driving the network in the high-speed direction.

1.5 The Characteristics of High-Speed Networks

If user requirements of today and tomorrow are to be satisfied, and efficient exploitation of technological advances is to be achieved, the characteristics of networks must be different from those in the past.

First, the dramatic rise in the range of available transmission speeds meant that the bottleneck moved from the links to the network processors. As a result, in a high-speed network, moving information from node to node has to rely on hardware-driven switching instead of software-driven routing.

Second, a network node must be very fast to cope with the flood of information that any number of medium and high-speed links can input into or extract out of it.

Third, this network node′s operation has to be based on an architecture designed for the support of traffic types of conflicting characteristics and diverse demands as far as quality of service is concerned.

1.6 How IBM Nways 2219, 2225 and 2230 WAN Switches Fulfill Requirements

The changes outlined in the previous sections had a profound influence on wide area networking technologies, techniques and architectures. Additionally, the increasing exchange of information and information types between the wide area and the local area environments spells even more change in the years to come.

Nevertheless, the networking community has, over the past few years, managed to come to an agreement on a number of standards, applicable to wide area network environments, that give network managers and administrators the ability to evolve their networks in a stable and cost-effective manner over an acceptable period of time. Those standards apply mainly to frame relay, switched multi-megabit data service (SMDS), and asynchronous transfer mode (ATM). In this redbook, it is clear that the IBM Nways 2219, 2225 and 2230 switches support these network disciplines and are fully compliant with the related standards.

Frame relay, as a standardized technique for interfacing to a packet network, has gained practically universal acceptance. Its popularity derives mainly from the fact that a network node can use a single physical port to maintain many virtual connections, thereby facilitating implementation of meshed small to medium sized networks in an economical way. Both the IBM Nways 2219 and 2225 switches support frame relay switching in full compliance with the relevant standards.

Switched multi-megabit data service (SMDS) comprises the definition of a network interface and related functions, enabling the implementation of metropolitan area networks (MANs). It offers connectionless transfer of messages, closed user groups with services such as address screening and validation, and broadcast services across low or high-speed lines. It is gaining increasing acceptance with network service providers. Again, both IBM Nways 2219 and 2225 support the applicable standards.

Asynchronous transfer mode (ATM) is the technology adopted to deliver Broadband ISDN services. Its acceptance among operators, equipment vendors and users stem from the fact that ATM supports all types of traffic over medium or high-speed lines, can be used both in wide and local area environments, and is designed to operate in the technological environments of today. IBM Nways 2225 supports ATM access services, as well as frame relay to ATM interworking functions (IWF), while IBM 2230 is a high-capacity ATM switch.

Chapter 2. Multiservices Concepts and Implementation

This chapter gives detailed information on the services that are supported by the current release of software for the IBM 2219, 2225, and 2230 WAN Switches. The purpose of this chapter is to provide an overview of features implemented in the switches.

The IBM 2219, 2225 and 2230 WAN switch platform is designed to to meet all requirements for high-speed, high-capacity, multiservice WAN data networks. The hardware platform allows the building of scalable private and public WAN multiservice networks in a very efficient way. IBM has implemented the following types of access and interworking services:

- Frame relay services
- ATM services
- SMDS services
- ISDN services

2.1 Frame Relay Services in IBM Nways 2219 and 2225

Frame relay is a fast packet switching technology well suited for bursty LAN traffic. More and more customers are moving from X.25 and leased-line networks to frame relay, which is most efficient for transmission of data traffic in WANs today. You can find a detailed description of frame relay in *Frame Relay Guide*, GG24-4463.

Frame relay switching is implemented in the IBM Nways 2219 and 2225 Switches. IBM Nways 2219 supports a total of 350 PVCs per unit independently of the number of ports. IBM Nways 2225 supports up to 975 PVCs originating per module. IBM frame relay implementation strictly adheres to industry standards. The following features are supported:

- Frame relay UNI-DCE interface
- Frame relay UNI-DTE interface
- Frame relay NNI interface
- Direct FRAD
- Translated FRAD (PPP to RFC 1490)
- Frame relay trunk
- Direct line trunk

In addition to these, IBM Nways 2225 implements the following interworking functions:

- Frame relay/ATM network interworking
- Frame relay/ATM service interworking

2.1.1 Logical Port Configuration

For each physical port or n x DS0 channels, you can define a logical port. Because IBM Nways switches use a symmetrical architecture, you can define by software any logical port as an interswitch-trunk, network-to-network link or user-to-network link. Symmetrical architecture lets you mix any type of connections on any single I/O module or on any contiguous or noncontiguous block of n x DS0 channels, when using 24/30 bundle T1/E1 I/O module. For non-frame relay devices you can define a direct FRAD or translated FRAD logical ports. The following sections describe each logical port configuration.

2.1.1.1 Frame Relay UNI-DCE

Frame relay User-Network Interface (UNI-DCE) performs frame relay DCE link management functions, and the frame-relay DTE device can be attached to it as illustrated in the following figure. For example, you can connect IBM 2210 or IBM 6611 routers to the IBM Nways 2219 or IBM Nways 2225 UNI-DCE ports. In such a case, IBM Nways switches will provide the frame relay backbone for the IBM routers.

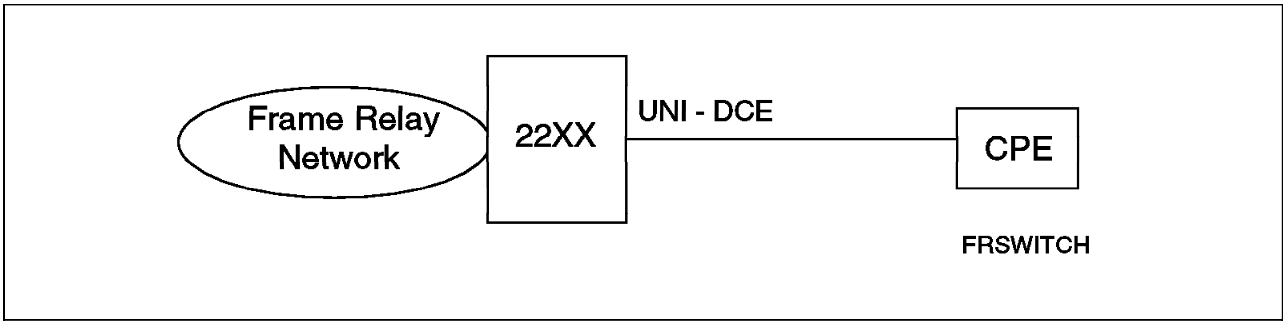

Figure 1. Frame Relay UNI-DCE Logical Port

2.1.1.2 Frame Relay UNI-DTE

Frame relay UNI-DTE (feeder function) performs frame relay DTE link management functions. The IBM Nways Switch acts as a DTE device connected to the frame relay DCE switch. This function is useful when interconnecting IBM Nways switches over public frame relay networks using the OPTimum trunking. OPTimum trunking is discussed later.

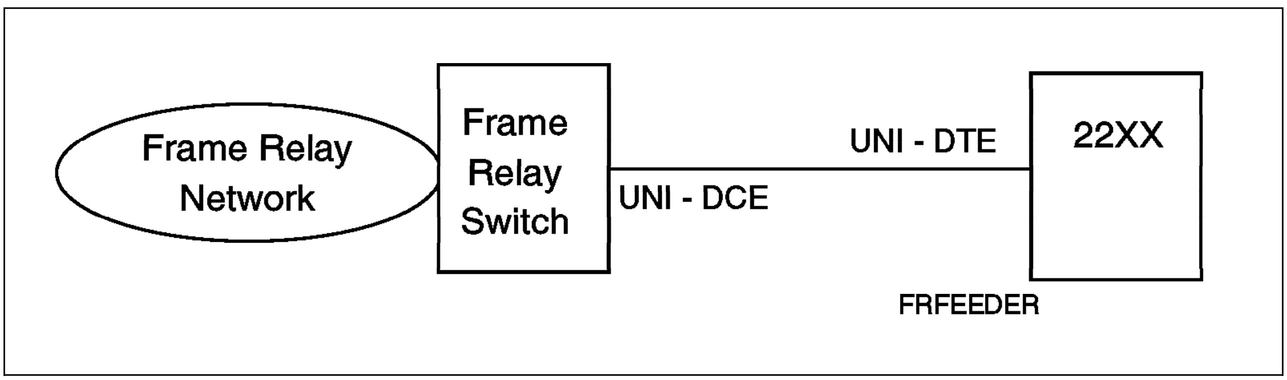

Figure 2. Frame Relay UNI-DTE Logical Port

2.1.1.3 Frame Relay NNI

The frame relay network-to-network (NNI) interface performs both frame relay DTE and DCE link management functions. An NNI logical port is useful for connecting two different types of switches or two frame relay networks. It performs bidirectional messaging, which enables both networks to share the status of PVCs in each network.

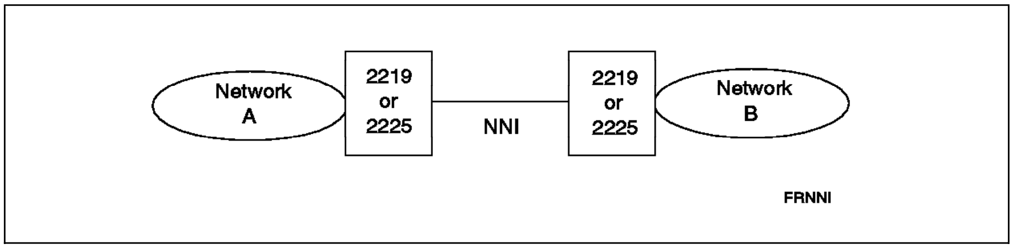

Figure 3. Frame Relay NNI Logical Port

2.1.1.4 Direct FRAD

A logical port defined as frame relay direct FRAD performs frame relay encapsulation/de-encapsulation for HDLC/SDLC-based protocols. The FRAD function encapsulates incoming HDLC/SDLC traffic to the IBM Nways Switch and de-encapsulates it upon exiting the network. As the IBM Nways Switch does not examine the HDLC/SDLC header, this configuration is limited to one point-to-point PVC only. This function is useful when you are interconnecting devices that support only the HDLC/SDLC-based protocol over the IBM Nways network. For example, many old X.25 switches use the HDLC protocol on their trunk side.

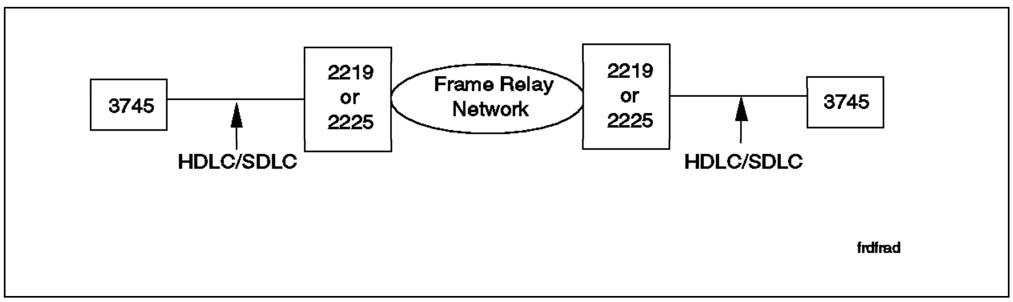

Figure 4. Frame Relay Direct FRAD Logical Port

2.1.1.5 Translated FRAD

The frame relay translated FRAD logical configuration enables communication between a DTE device using PPP (point-to-point protocol) and a DTE device using frame relay to access the IBM Nways network. The frame relay DTE device has to be configured for frame relay RFC 1490 Multiprotocol encapsulation. This function is very useful when you are connecting branch office routers that support PPP protocols to a central frame relay router using the frame relay backbone. Using this technique, you can interconnect more PPP routers to the frame relay router. IBM Nways switches also enable the interconnection of a PPP router with an ATM-connected device using the multiprotocol translated from RFC 1490 to ATM RFC 1483. For more information, refer to 2.2.5.1, “Frame Relay/ATM Service Interworking” on page 31.

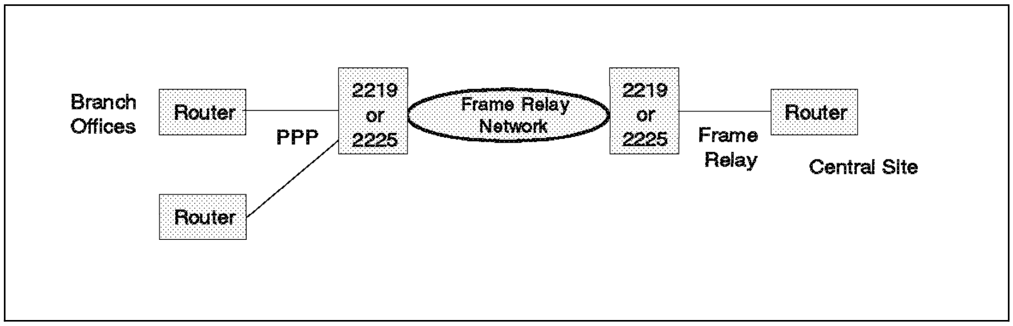

Figure 5. Frame Relay Translated FRAD Logical Port

2.1.1.6 Direct Line Trunk

The direct line trunk allows for a direct trunk connection to another IBM 2219 or 2225 Switch. The trunk carries traffic destined for other switches in the network using a special type of protocol called IBM Nways trunk protocol. This protocol provides users with IBM Nways added value features for rate enforcement (such as red, amber and green frame descriptors). For more information on rate enforcement see 2.1.3.2, "Rate Enforcement" on page 10. Direct trunk can be used only on the leased line or on any other type of transparent data channel providing the same service as leased line. IBM Nways switches reserve 5% of trunk bandwidth for control traffic. For more details about bandwidth reservation, refer to 2.1.3.6, "Virtual Trunk Bandwidth" on page 15. IBM Nways 2225 or 2219 allows you to specify which type of traffic can traverse a particular trunk; you can use one of the following options:

- Management data. This trunk will carry only management data.
- Frame relay, SMDS and management data.

2.1.1.7 Frame Relay OPTimum Trunk

Frame relay trunk provides a switch-to-switch IBM Nways trunk through a public frame relay network. You can interconnect two separate networks based on IBM Nways switches over a public frame relay network. The switch uses tunneling to maintain the IBM Nways trunk protocol headers. The IBM Nways trunk protocol is encapsulated into the standard frame relay packets and sent over the public frame relay network. This function is also referred to as OPTimum trunking . To use this function you have to define your logical port either as a UNI-DTE feeder or a frame relay NNI to allow link management exchange between IBM Nways and the public frame relay network. You can specify if the trunk will be used for management data only, frame relay and management traffic, or frame relay, SMDS and management traffic.

Note: You cannot use the OPTimum trunk feature on the T1/E1 channelized interfaces.

Figure 6 on page 9 illustrates direct and OPTimum trunks inside the IBM Nways network.

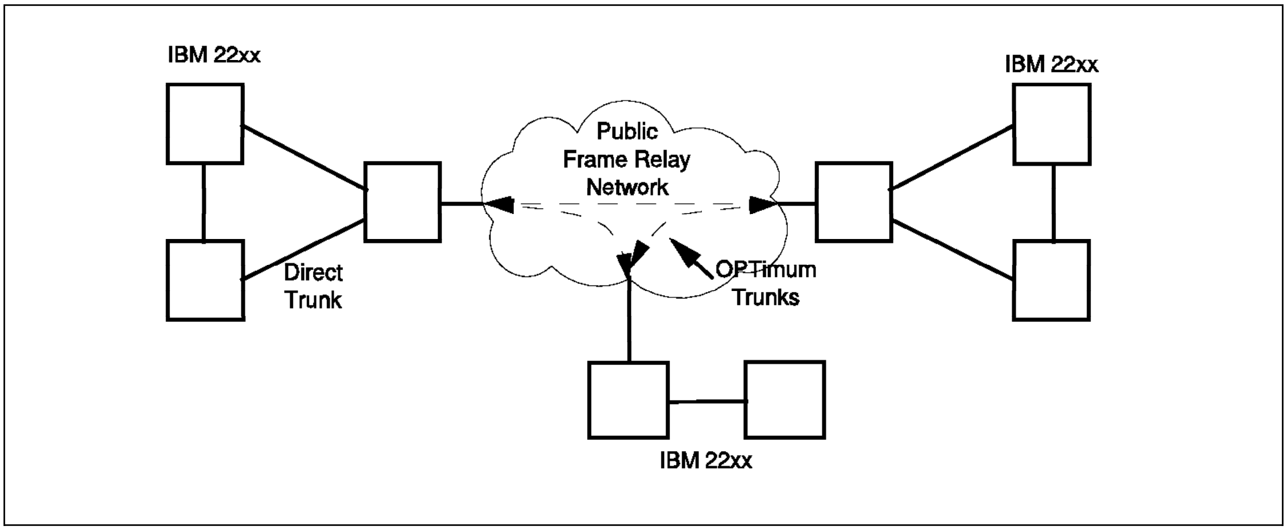

Figure 6. IBM Nways Direct and Frame Relay OPTimum Trunks

2.1.2 Frame Relay Implementation Overview

IBM Nways 2219 I/O modules support different types of interfaces at transmission rates from 2.4 Kbps to 4 Mbps. IBM Nways 2225 ranges from 9.6 Kbps to 45 Mbps. The maximum supported frame size is 8192 bytes. IBM Nways 2219 and 2225 Switches strictly adhere to all frame relay standards. For a list of the frame relay standards, see Appendix B, "Frame Relay Supported Standards" on page 139. However, IBM also implements some extensions that enable you to build more reliable and efficient frame relay networks.

2.1.2.1 Supported Frame Format

IBM Nways switches use standard frame relay format. Extended addressing is not currently supported. This limits the maximum number of user VCs to 975 per physical port. The maximum frame length is 8192. However, the maximum length of 4096 octets is recommended.

2.1.2.2 Link Management Interface

Link management interface provides dynamic notification to the user of addition and deletion of PVCs. It also monitors each connection to the network through the periodic heartbeat keep-alive polling process.

There are three types of link management protocols that can be used in the frame relay network. Their functions are very similar, but they do not interoperate; both ends of line (DTE-DCE or NNI-NNI) have to use the same type of LMI. IBM Nways Switch supports all three types of link management interfaces:

- Local Management Interface (LMI) developed by StrataCom, Digital Equipment Corp., Northern Telecom and Cisco Systems. LMI was created in the early days of frame relay.
- ANSI T1.617 Annex D.
- CCITT Q.933 Annex A.

Additionally, the Nways 2219, 2225, 2230 WAN Switches could be configured for autodetection or have Link Management Interface disabled. When you configure a port for autodetection, the IBM Nways Switch automatically senses the type of the protocol employed by the user devices and adjusts accordingly.

2.1.3 Frame Relay Value-Added Features

The IBM Nways 2219 and 2225 Switches comply with frame relay standards; for a list of supported standards refer to Appendix B, “Frame Relay Supported Standards” on page 139. Frame relay added-value features extend the standard frame relay implementation. These features provide for:

- Improvement of network throughput
- Support delay sensitive traffic
- Traffic management
- Fault-tolerant design

2.1.3.1 IBM Nways 2225 High-Speed Frame Relay

The frame relay standard defines 2 Mbps access speed to a frame relay network. IBM Nways 2225 Switch high-speed frame relay enables you to connect a frame relay device using higher-speed access lines. High-speed I/O modules, such as HSSI or E3 cards, provide for connection to IBM Nways frame relay network with access speeds of up to 44 Mbps. The IBM Nways 2225 Switches can be connected using high-speed trunks.

2.1.3.2 Rate Enforcement

Rate enforcement is implemented on a per DLCI basis on user links at the ingress of the switch. As data is received over time interval T, a determination is made as to whether the frame is under the committed burst size (Bc), over the committed burst size but under the excess burst size (Be), or over the excess burst size.

Any PVC in the frame relay network is subject to a contract between user and network in terms of traffic characteristics. The network knows what traffic to expect from the user, and the user should behave according to the contract. The standard traffic descriptors for each frame relay PVC are:

Committed Burst Size (Bc)

The maximum number of bits accepted by the network under normal conditions during time interval T.

Committed Information Rate (CIR)

The rate at which the network agrees to transfer data end-to-end under normal conditions.

Excess Burst Size (Be)

The maximum number of uncommitted bits, during time interval T, the network accepts above the committed burst size (Bc).

Committed Rate Measurement Interval (T)

The length of time over which data is received before user adherence to the contract is confirmed. The input rate is averaged over an increment of time T. T is calculated according to the following formula: T = Bc/CIR.

Contract enforcement in the IBM Nways Switch implementation conforms to ANSI T1.606 Addendum 1 and is based on the leaky bucket algorithm. IBM Nways switches keep track of every DLCI on user links in ingress ports. Green, amber and red are descriptors that are used to identify packets as they travel through the network. As data is received over time interval T, determination of the type of frame is done according to the following:

Green frames are used to identify packets where the number of bits received during the current time, including the current frame, is less than or equal to Bc. Green frames are never discarded by the network, except under extreme circumstances such as a node or a link failure.

Amber frames are used to identify packets where the number of bits received during the interval T, including the current frame, is greater than or equal to Bc but less than Bc + Be. These frames are forwarded with the Discard Eligible bit (DE) set and are eligible for discarding when passing a congested node.

Red frames are used to identify packets where the number of bits received during the current time, including the current frame, is greater than Bc + Be. Red frames are forwarded with the DE bit set only when the graceful discard feature is enabled, otherwise they are discarded.

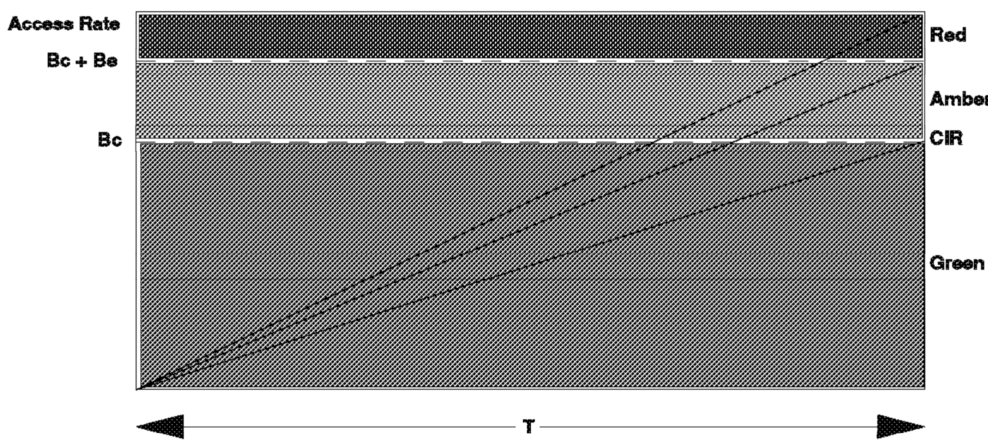

Figure 7. Green, Amber and Red Frame Tagging

Note: Release 4.1 only allows for a fixed interval of T = 1 second.

2.1.3.3 Congestion Avoidance and Management

Congestion avoidance and management are very important functions in any type of high-speed packet-switched network. They influence data throughput and network stability. Nways 2219, 2225, 2230 WAN Switches have implemented robust algorithms for congestion avoidance and recovery. Congestion control is designed to implement ITU-T and ANSI frame relay standards.

The Nways 2219, and 2225 WAN switches use standard congestion notification bits, FECN and BECN, to inform users about network congestion. The user device should slow down transmission, when notified.

Data traversing the IBM Nways Switch network is queued for transmission in the egress I/O adapters. The state of each transmit queue is checked for pending congestion. A time-averaged algorithm, Time Averaged Queue Length (TAQL) is executed each time a new frame is queued for transmission. TAQL is calculated and compared with the predetermined threshold. The default thresholds are not changeable or configurable.

If TAQL exceeds the first threshold, the link is considered as *mildly congested* and the following steps are taken to avoid congestion on the link:

1. All red frames are discarded.
2. All frames transmitted on the mildly congested link are marked with the FECN bit.
3. All frames received on the mildly congested link are marked with the BECN bit before being forwarded out on another link.

To avoid possible oscillations, the TAQL algorithm uses time averaging and hysteresis calculation. In addition to monitoring the average queue length, Nways 2219, and 2225 WAN Switches monitor the actual depth of each transmit queue each time a new frame is queued. If the actual size of the queue exceeds a second predetermined threshold, the link is considered *severely congested* and the following steps are taken:

1. FECN and BECN bits are set the same way as for the mildly congested links.
2. All incoming red packets are discarded unless graceful discard is enabled.
3. All incoming amber packets are discarded.

If the previous procedure does not improve the state of the transmit queue, and queue length reaches the third predetermined threshold, there are no more buffers available for additional packets. In this *absolutely congested* state, the Nways 2219 and 2225 WAN Switches do the following:

1. BECN bits are set.
2. All incoming packets are discarded.

If the network is designed with no oversubscription of trunks (for example, there is enough bandwidth to accommodate all CIRs), this state should be reached only in the event of link or node failures.

The size of the queues are 56 byte buffers, and this is the smallest block of memory that can be assigned to a frame. Thus, a 57 byte frame will end up using two buffers or 112 bytes of memory in the queue. Each logical port has its own queue and the number of buffers in that queue varies, depending on the card type and speed of the logical port. A short summary is below:

- Channelized E1 - one DS0/lport 174 buffers = 9744 bytes
- Channelized E1 - two DS0/lport 340 buffers = 25,200 bytes
- Channelized E1 - three DS0/lport 520 buffers = 29,120 bytes
- Channelized E1 - >3 DS0/lport 588 buffers = 32928 bytes
- UIO port - 588 buffers = 32,928 bytes
- HSSI/ATM ports - 17,070 buffers = 955,920 bytes

There is only one physical queue per logical port. However, inside the physical queue, there are four logical queues, one per priority. Thus, congestion is based on the total amount of data for all priorities that is currently in the queue. When discards are required due to congestion, frames of all three priorities are equally likely to be discarded.

Inside the physical queue, there are three congestion thresholds: mild, severe and absolute. When the time-averaged queue length hits the mild threshold, all red frames arriving at the queue will be discarded and will not be queued (regardless of priority). However, all red frames already in the queue will be left in the queue and will eventually be transmitted. When the time-averaged queue length hits the severe threshold, all arriving amber frames will be discarded instead of queued. Again, all amber frames already in the queue are safe and are not discarded. The absolute threshold means the entire queue is full and no more frames will fit. When this threshold is reached, all arriving green frames will be discarded.

Prior to Release 4.1, the queue sizes and the congestion thresholds are fixed. (Mild congestion = 25% of queue length and severe congestion = 50% of queue length.) In Release 4.2, the user will be able to configure the queue sizes and the thresholds. Since the current congestion scheme is not using the full amount of memory on the IOPs, these queue sizes can be increased. Queueing studies have been conducted, and we feel that the values we have already set will optimize the data throughput. Nevertheless, customers require flexibility so we are providing it. Whether they increase or decrease their throughput by playing with these queuing values remains to be seen.

A single 8 Kbps frame will cause the queue to be approximately 25% full (that is, the mild congestion threshold) for ports up to E1 speed. Just hitting the mild threshold will not invoke discards, however, since the congestion algorithm works according to a time-averaged queue length. Also, the simultaneous or near-simultaneous arrival of four such 8 Kbps frames will push the queue to absolute congestion, potentially causing green frames to be discarded. In order to determine the expected value of the queue size and the probability that four frames will be queued at once at the egress port will require more detailed knowledge of your traffic distribution (as well as more queuing knowledge than we currently have). It is a possible to consider that if the maximum frame size were decreased to 4096 bytes, eight simultaneous or near-simultaneous maximum size frames would be required to hit the absolute threshold.

2.1.3.4 Frame Relay Circuit Priority

A priority queue mechanism has been implemented, enabling the selection of different priorities for frame relay virtual circuits. This added-value feature allows you to differentiate between time-sensitive traffic, such as voice over frame relay and the data traffic. IBM Nways 2225 and 2219 give you the capability to either prioritize mission critical data or to preserve service levels for delay sensitive traffic.

Hardware implementation uses two types of queues. The hardware queue is used by the physical link drivers, and the software queue is used for the buffering of user frames. The software queue is divided into four priority queues PQ0, PQ1, PQ2 and PQ3. PQ0 has the highest priority level and is used only for IBM Nways Switch management traffic. When you configure a frame relay PVC, you have to choose one of the three priorities, which define the type of queue used for that PVC.

If the hardware queue has enough space to accept frames, IBM Nways Switch does not use the prioritization and immediately submits frames for transmission. As soon as the hardware queue is full, the IBM Nways Switch starts to buffer the frames in the software queues. The software queues are emptied in strict order, starting with the highest priority (PQ0) and completing with the lowest priority (PQ3). The software will begin to dequeue data packets from a particular queue only if all higher queues are completely empty. Note that this algorithm does not differentiate between frame types, so amber frames with priority 1 will be transmitted before green frames with priority 2.

All four priority queues form a queue called Instantaneous Queue (I_QUEUE), which is used by the congestion avoidance and recovery mechanisms. The I_QUEUE length is calculated using the following formula:

```
len(I_QUEUE)=len(PQ0)+len(PQ1)+len(PQ2)+len(PQ3)
```

Nways 2219 and 2225 WAN Switches congestion avoidance and recovery algorithm works with the length of the I_QUEUE. For more information, refer to 2.1.3.3, "Congestion Avoidance and Management" on page 11. The following figure illustrates the relationship between the PQs and the I_QUEUE.

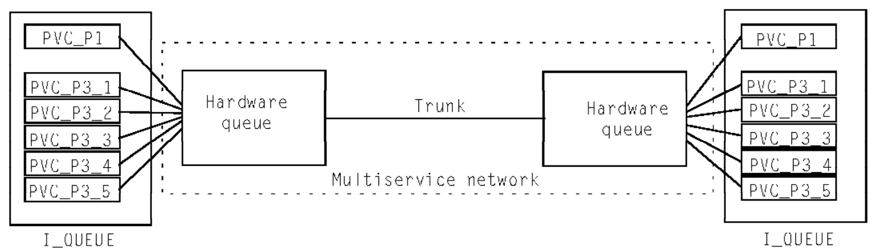

Figure 8. The I_QUEUE and Hardware Queue

Priority Queues, together with the graceful discard option and support for CIR=0, enables you to configure different frame relay classes of service.

2.1.3.5 Zero CIR PVCs

Zero CIR PVCs are defined such that CIR=Bc=0. Using Zero CIR PVCs, it is possible to divide the PVCs into three groups, according to the importance of transported data:

Class 1 - most important data - CIR>0, Bc>0 and Be>0. The user data will be transported using green, amber and red frames.

Class 2 - less important data - CIR=0, Bc=0 and Be>0. The user data will be transported using amber and red frames.

Class 3 - expandable - CIR=0, Bc=0 and Be=0. The user data will be transported only via red frames.

Since the network discards frames based on the level of congestion starting with red frames (see 2.1.3.3, "Congestion Avoidance and Management" on page 11) the Class 3 PVC will be transported without discarding any frame only when there is no congestion on its path. The amber frames of Class 2 PVC will traverse the mildly congested nodes, but will be discarded in the severely congested node.

Class 1 to 3 PVCs can use different frame relay circuit priority. Using different frame relay circuit priorities and classes enables you to create up to nine different classes of service.

2.1.3.6 Virtual Trunk Bandwidth

You are able to oversubscribe (use total of all CIRs greater than the real bandwidth of the trunk) to a frame relay network to take advantage of the statistical nature of packet switching. Virtual bandwidth is calculated in order to allow for oversubscription, using the trunk utilization factor K. To determine the virtual bandwidth, use the following formula:

Virtual Bandwidth=0.95*(trunk configured bandwidth)/K

Coefficient 0.95 represents 5% of the trunk bandwidth reserved for IBM Nways Switch management traffic. The IBM OSPF algorithm then assigns metrics to these trunks according to the following formula:

Available Virtual Bandwidth=Virtual Bandwidth-(sum of PVC CIR)

The trunk utilization factor K (sometimes referred to as the oversubscription factor) is set during the configuration of the trunk using management applications. Using the trunk utilization factor, together with the trunk administrative cost, allows you to adjust the OSPF algorithm.

2.1.3.7 Routing in the Nways Switch Network

To be able to determine the path for a particular PVC in the network, some type of routing algorithm has to be used. Nways 2219, 2225, 2230 WAN Switches use the OSPF, open shortest path first, routing protocol to determine the best path for the PVC through the network. The OSPF link state routing protocol is the Internet standard for the Interior Gateway Protocol (IGP). For details, refer to Internet RFC 1131.

Each switch in the network maintains a link state database based on OSPF link state updates. The Dijkstra algorithm runs on top of the link state database, which results in a tree of the best paths through the network. Nways 2219, 2225, 2230 WAN Switches use bandwidth availability and user-defined administrative costs as metrics for each trunk. As each PVC is mapped to a trunk, bandwidth availability decreases, which causes OSPF link state metrics for that route to increase.

Note: The best path is not necessarily the most direct path (minimum hops) through the network.

When all switches in the network are running software Release 4.1 or higher, you are able to assign for each trunk an administrative cost. When you first define a circuit, IBM Nways Switch looks to its link state database for a path that has enough virtual bandwidth available to handle requested CIR for the PVC. If more than one path exists, the switch chooses the path with the lowest administrative cost. If there is more than one path with the same administrative cost, the one with the largest available bandwidth is chosen. If there is not enough bandwidth in the network, new PVCs cannot be built. However, this does not mean that existing PVCs will not be rerouted in case of trunk failure. The following list summarizes the key OSPF routing functions:

> **CIR Guaranteed Bandwidth TOS (Type of Service) Routing**: A metric associated with the available, uncommitted bandwidth is kept for each interface. The cost for routing through the trunk is inversely proportional to the uncommitted bandwidth. The best path for a new PVC is determined using the Dijsktra algorithm.

Route Recovery: When a tandem node or trunk fails, new paths for affected PVCs are re-calculated immediately at the ingress (origin) nodes. The PVCs are rerouted to the new route with recovery time typically under 4 seconds.

Note: In this case, the available bandwidth can become negative (if actual bandwidth is oversubscribed).

Load Balancing: New PVCs are distributed equally over all routes with the same costs. The cost of each route is re-calculated periodically to reflect used CIR bandwidth for the existing PVCs. A PVC is rerouted, when a less costly path is discovered. For more information, see 2.1.3.8, "Reroute Tuning."

The Nways 2219, 2225, 2230 WAN Switches allow you to manually define PVCs. If you do so, you are bypassing the OSPF algorithm to make a routing decision for the defined PVC. Manually defined paths give you additional control over routing inside the network. You are still able to use the *Alternate Path Option* to reroute this PVC when a predefined path fails. For future enhancements, refer to Appendix A, "Release 4.2 Preview for the Nways 2225" on page 127.

2.1.3.8 Reroute Tuning

IBM has implemented an algorithm for the rerouting of existing PVCs over a less costly discovered path. The reroute tuning feature enables you to customize the rate of reroute requests per switch. You can define how many PVCs can issue reroute requests during the single reroute batch request as well as the delay between each rerouting batch.

For example, if the switch has four cards, each with 50 PVCs, and you set the reroute count for 5 PVCs and the reroute delay for 10 seconds, the switch will perform a batch reroute, consisting of the first five circuits on each card (for a total of 20 circuits). The switch then waits for 10 seconds before it begins to reroute the next batch of 20 circuits.

It is recommended that you set the reroute count and reroute delay no lower than one PVC every ten seconds. A lower setting may cause network instability. However, to quickly balance a set of circuits after the trunk failure, you can use the previous example.

You can adjust the load balancing algorithm not only via reroute count and reroute delay, but you can also specify the type of circuit. The following options are available:

Negative-to-Positive - Only circuits traversing trunks with negative virtual bandwidth are considered. Switch reroute these only when a path with sufficient bandwidth exists. This option provides the least aggressive search for an alternate path.

Unrestricted - All circuits are considered. If a path with a greater bandwidth exists, the circuit is rerouted. This is the most aggressive rerouting option.

Negative-to-Anything - If the current circuit path involves one or more trunks with negative bandwidth and a path with less negative or sufficient bandwidth exists, the switch reroutes the circuit.

Disable - Disable the load balancing algorithm.

2.1.3.9 Multicast Frame Relay Service

Multicast frame relay is useful in applications where the stations are routers or bridges or for video broadcast applications. The Frame Relay Forum defines three types of multicast services over the frame relay network:

One-Way Multicast Service - All frames are sent from the root station to the members (leafs) of the multicast tree.

Two-Way Multicast Service - Packets originating in the root station are sent to all members of the multicast group. Frames sent by leaf stations are transmitted to the root over the multicast DLCI.

N-Way Multicast Service - All transmissions are duplex and are multicast; members of the multicast group are transmission peers.

The Nways 2219, 2225, 2230 WAN Switches currently support only One-Way Multicast Service.

To define a multicast group, you must first create individual PVCs from the logical port, which performs broadcast (root) to all other members of the multicast group (leafs). Then you select a number for broadcast DLCI. (You can select any valid DLCI number.) As a last step, you define which DLCI numbers are participants of the multicast group.

Frames addressed to the multicast DLCI will be replicated by the multicast server to all leafs of the multicast group. Frames are delivered using the DLCI numbers defined in the first step.

The multicast server has been implemented in the IBM 2210 and 2225; there are no special hardware or software requirements to support this function. Today implementation of a multicast server is centralized. The root switch performs all multicasting functions.

Figure 9 on page 18 explains the multicast function. The left side of the diagram shows the main site as the root of the multicast group and has PVCs to leafs B, C and D, which are leafs. The DLCIs associated with these PVCs are 100, 101 and 102, respectively.

Note: You do not have to use the same DLCI numbers on both ends, because DLCIs are only locally significant. The root has configured all three DLCIs as part of the multicast group with DLCI number 200. When the root transmits a single frame on DLCI 200, that frame is delivered to all leafs on DLCI 100 (site B), 101 (site C) and 102 (site D). The switch at site A performs the multicast server function.

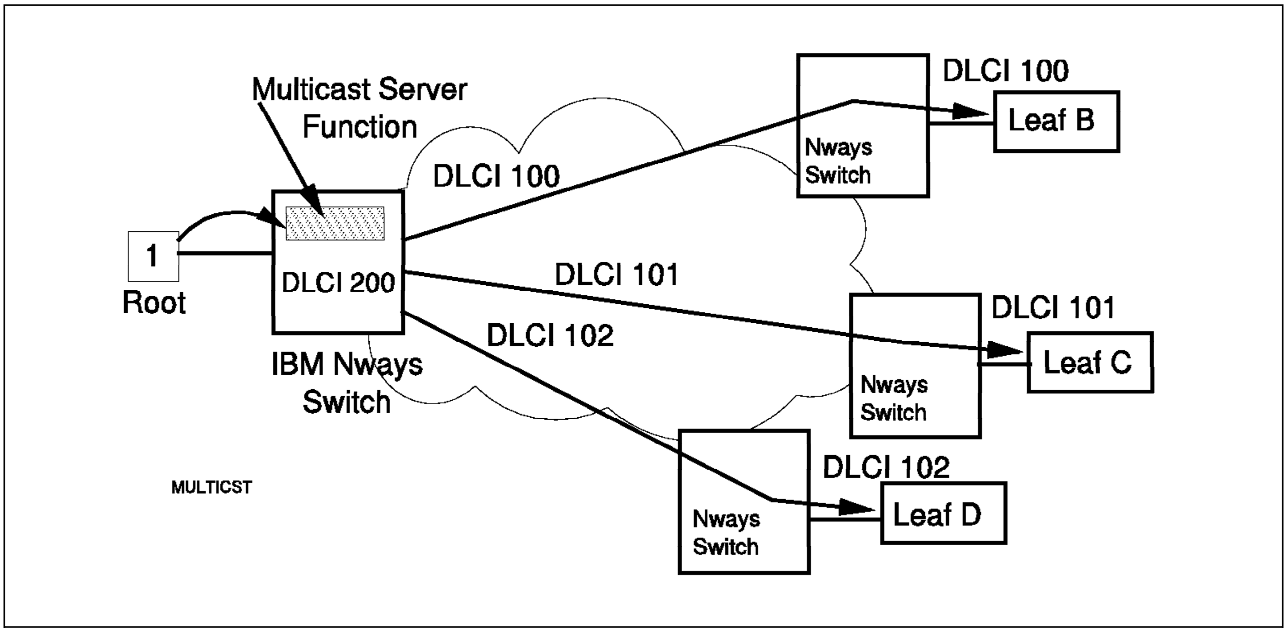

Figure 9. Multicast DLCI Example

You can distribute the multicast functions using the NNI interface inside the network as shown in Figure 10.

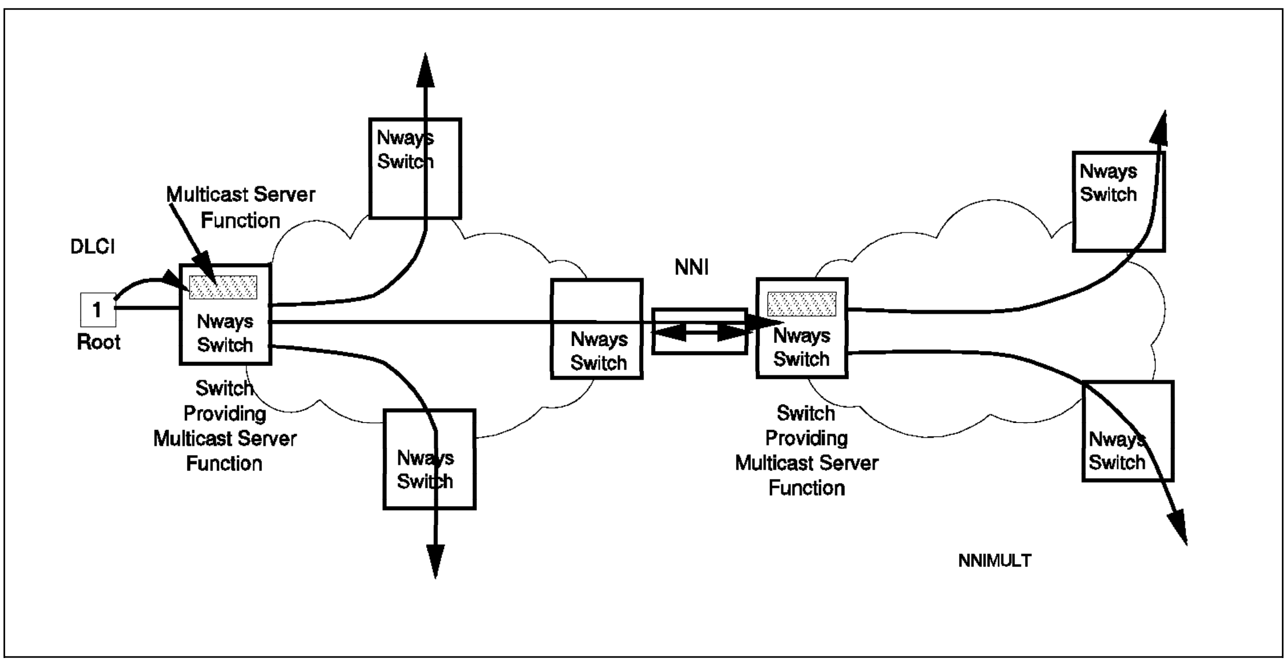

Figure 10. Distributed Multicast Server Function

The NNI interface divides the multicast tree into two separate parts. The left multicast server sends the data to the NNI interface as it is a leaf of the multicast tree. The egress leaf NNI DLCI is configured as a root of the right multicast tree. The major disadvantage of this configuration is that you have to divide your network into two different routing domains.

You can implement the distributed multicast server in the same OPSF domain using two frame relay ports on the same switch. In this case, one port will be configured as a leaf of one part of the mutlicast tree, the other one will be configured as a root for the other part of the mutlicast tree. The following figure illustrates this type of configuration:

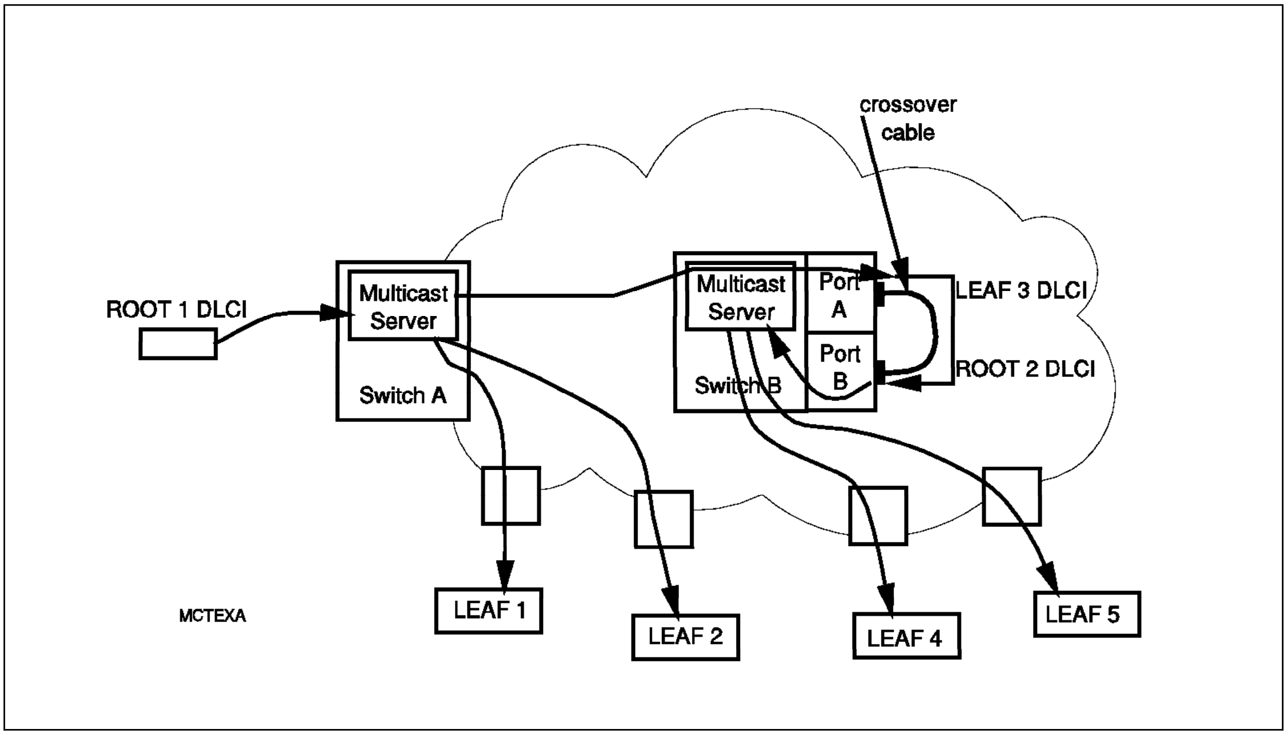

Figure 11. Distributed Multicast Server Example

Root 1 sends the multicast frame to the multicast server in switch A, which sends a copy of the frame to leafs 1, 2 and 3. Switch B has configured A and B ports, which are connected to each other using a simple crossover cable. Leaf 3 DLCI on port A is used as a root 2 DLCI on port B. The multicast server in switch B receives the leaf 3 DLCI frame as a multicast frame and sends the copy to leafs 4 and 5.

One IBM Nways 2225 Switch can provide multicast server functions for 32 multicast groups. The CIR of the ingress port, the size of the incoming frame, the number of leafs in a group and the location of ingress and egress ports will affect performance of the multicast functions. Currently, there is no performance data available.

2.1.3.10 Fault-Tolerant PVC

Nways 2219, 2225, 2230 WAN Switches have a built-in industry-unique feature, which allows for backup of logical frame relay ports. Using this option, you enable easy recovery from a catastrophic disaster in the network.

Frame relay UNI-DCE, UNI-DTE and NNI logical ports can be configured to provide backup service for any logical port in the network. When the backup port is not used, then its status is idle with no DLCIs on it, and it does not use the IBM Nways Switch resources. In case the primary port fails or you need to take the primary port offline, activate the backup port, and the OSPF and PVC manager will reroute all frame relay PVCs from the primary port to the designated backup port. This is currently done by issuing a command. Once the primary logical port is restored, you are able to manually disable the backup ports, and all of the PVCs will be rerouted to the primary port. The key features of fault-tolerant PVCs are:

- Fault-Tolerant PVCs can be implemented with no impact on customer CPE devices. Only the Nways 2219, 2225, 2230 WAN Switches are aware of PVC rerouting to the disaster recovery center.
- IBM OPTimum trunk can be used together with the Fault-Tolerant PVC feature in the public data networks, maximizing the reliability of the network while keeping cost low.
- When using the Fault-Tolerant PVC feature, you do not have to reconnect the PVCs from the failed physical port one by one: a manually entered switch command will reroute them all at once.
- All actions of Fault-Tolerant PVCs are initiated by the administrator of the network.

The following figure illustrates the backup port functionality.

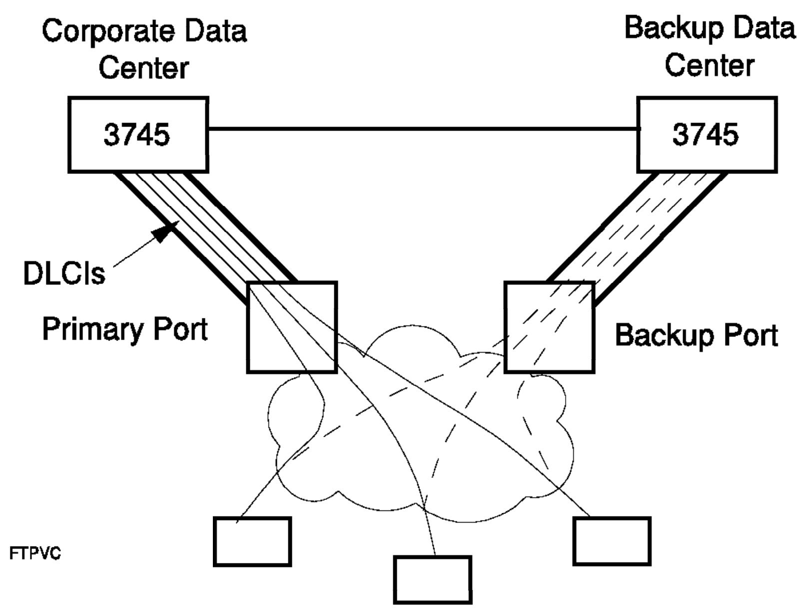

Figure 12. Fault-Tolerant PVC Example

2.1.4 Frame Relay Network Troubleshooting

Nways 2219, 2225, 2230 WAN Switches allow you to troubleshoot network frame relay traffic in a very efficient way. You can use several powerful diagnostic functions for layer 2 frame relay diagnostics:

- PVC loopback functions
- Last invalid DLCI

2.1.4.1 PVC Loopback

The PVC loopback function is supported on a per-PVC basis; you can specify either endpoint of the PVC to be looped back. This function does not affect traffic on other PVCs on the same logical port and lets you diagnose the frame relay circuits.

You can use three types of diagnostic tests on any PVC in the network:

PVC Local Loop - All traffic is looped against locally connected equipment. Statistics show frames received and sent between the switch and the locally attached device.

PVC Remote loop - All traffic is looped back to the far end endpoint of the PVC.

Both Loops - Both local and remote data will be looped back.

The following figure shows the local and remote PVC loops.

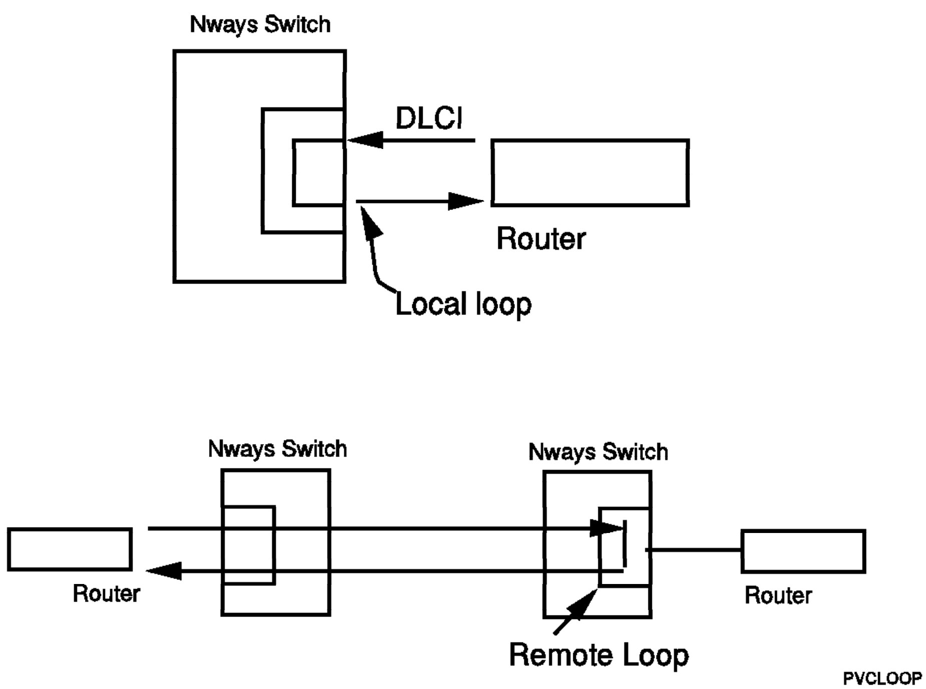

Figure 13. Local and Remote PVC Loopback

2.1.4.2 Last Invalid DLCI

This function displays the last invalid DLCI. It provides you with information on the last invalid DLCI received on the current port. When a user sends frames with a DLCI that is not defined for the port, the frame is said to be an invalid DLCI. Nways 2219, 2225, 2230 WAN Switches have a MIB object (lportLastInvalidDLCI) for each logical port; this value reflects the last invalid DLCI received on a particular logical port.

2.2 IBM 2225 ATM Services

Asynchronous transfer mode is the latest networking technology designed for transport of voice, video and data traffic. ATM technology can be deployed in the local and wide area networks, targeted to provide seamless interworking between the WAN and LAN environments. ATM is not only providing high-speed connections, it also supports quality of services (QOS), because it is connection-oriented. ATM allows implementation of many value-added features, such as VLAN support, efficient bandwidth management, flow control, etc. IBM has implemented ATM technology in a wide range of products including the campus ATM machines, the 8260 and 8285 as well as the wide area network switches the 2220, 2225 and 2230.

The IBM Nways 2230 is a high-end backbone ATM switch, and its features are discussed in 2.8, "IBM Nways 2230 ATM Overview" on page 52. The IBM Nways 2225 is not designed as a high-speed ATM backbone switch, but IBM has implemented a rich set of features including frame relay, SMDS and ATM, focusing on providing ATM DXI and ATM UNI interworking services with frame relay. There are not many users who can afford high-speed ATM services at each location using a WAN public network. With the 2219, 2225 and 2230 switches you can build scalable data networks using frame relay technology to connect small branches and ATM technology for high-speed backbones and big data center connections.

This section does not describe ATM technology, it only describes the specific ATM features. For more details about ATM technology, please refer *Asynchronous Transfer Mode Overview: Architectural Tutorial*, SG24-4525.

2.2.1 IBM Nways 2225 ATM Access and Interworking Functions

IBM Nways 2225 set of ATM functions enables you to take advantage of an ATM backbone wide area networks. The IBM Nways 2225 Switch was designed to support both high-speed cell-based and low-speed frame-based ATM access. ATM/frame relay interworking allows you to integrate current frame relay services and convert these services to use the advantages of high-speed ATM backbones. The following access and interworking functions are available:

- ATM frame-based access services
- ATM cell-based access services
- ATM cell and frame-based interworking functions

2.2.2 ATM Frame-Based Access

ATM frame-based access was developed to provide ATM services over low-speed access lines. The frame-based access does not send user data packets using small ATM cells (53 bytes); it uses HDLC frames (up to 9 KB). This approach provides improved efficiency of low-speed access line bandwidth when compared to the cell-based ATM access, because when the frame size increases and overhead stays constant, efficiency increases. The frame-based access does not support delay sensitive traffic, due to the long frame sizes.

IBM 2225 supports the following frame-based access services:

- ATM DXI
- ATM FUNI

2.2.2.1 ATM DXI Overview

The ATM Data eXchange Interface (DXI) was standardized by the ATM Forum to provide ATM support for users over low-speed lines. The ATM DXI is supported on all frame-based modules such as UIO, HSSI and DSX-1 I/O modules at speeds from 56 Kbps to 45 Mbps. ATM DXI standard specifies the interface between the user DTE device, such as a router, and the DCE equipment called DSU. The DSU provides conversion between the ATM DXI frames and ATM cells and could be either external to the ATM network or integrated inside the switch.

The ATM Forum defines three major types of DXI interfaces, whose characteristics are summarized in the following table.

Table 1. ATM DXI Interfaces

Characteristic	**Mode 1a**	**Mode 1b**	**Mode 2**
DFA Header Octets	2	2	4
Maximum number of circuits	1023	1023	16777215
AAL 5 Support	Yes	Yes	Yes
AAL 3/4 Support	No	Yes	Yes
DTE SDU Length (AAL 5 assumed)	9232	9232	65535
Bits in the Frame Check Sequence (FCS)	16	16	32

Figure 14 shows the format of the ATM DXI frame for Modes 1a and 1b that is used between the DTE and DCE. The DSU units use either the AAL 5 or AAL 3/4 adaptation layer to fragment/reassemble the DTE SDU (Service Data Unit) into ATM cells. The DTE has to know if AAL 5 or AAL 3/4 is to be used, as these two adaptation layers use different Common Part Convergence Sublayer (CPSC) PDUs.

Note: The convergence sublayer is the part of AAL which sends and receives data to and from the higher layers of the device.

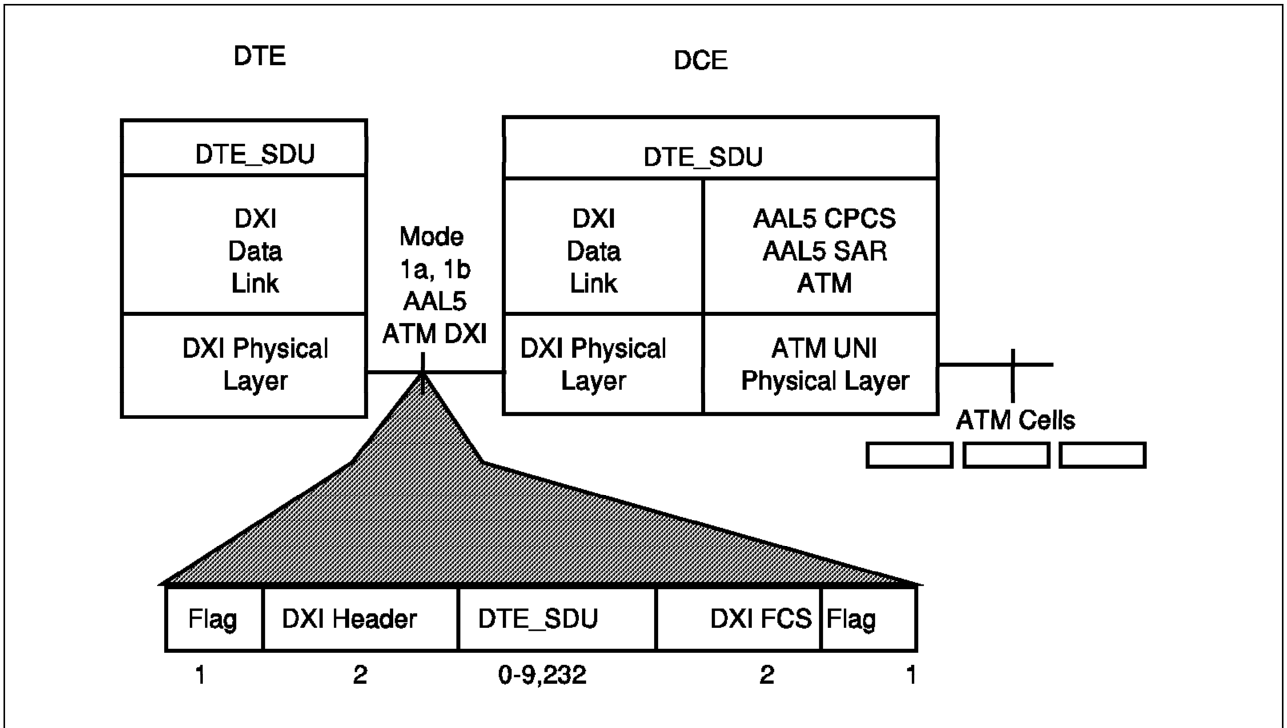

Figure 14. ATM DXI Frame

The DCE maps the DXI frame address (DFA) to the low-order bits of the VPI/VCI fields in the ATM cell header. The two-octet DXI header shown in the following figure is used by the DXI interface in Modes 1A and 1B; Mode 2 uses a 4-octet header.

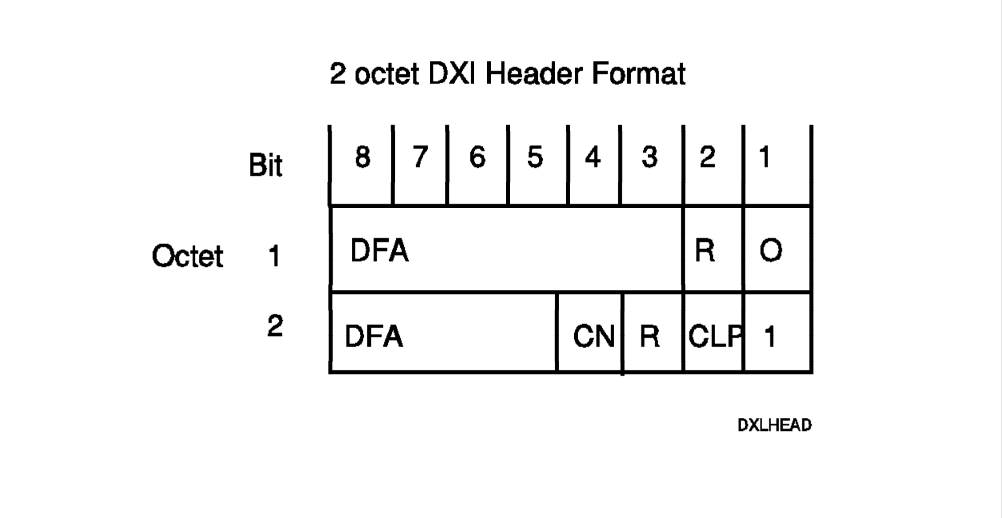

Figure 15. Two-Octet DXI Header Format

IBM Nways maps the flow control fields to and from ATM cells according to the following DXI specifications:

- Congestion Notification (CN) is set when the ATM cell′s Payload type is set to 01X, which means congestion has been experienced in the ATM network; otherwise, CN is cleared. In the DTE to DCE direction, the CN field is always cleared.
- Cell loss priority (CLP) bit is set based on the DXI header In the DTE to DCE direction, the CLP field is always set to zero.
- Bits ″R″ are reserved for future use.

The ATM DXI interface between DTE-DCE is managed by the DTE device using the Local Management Interface (LMI), which allows the DTE to obtain the type of DXI interface, as well as to set and query the AAL type assigned per VCC. The DSU will pass the Interim Local Management protocol to the DTE device.

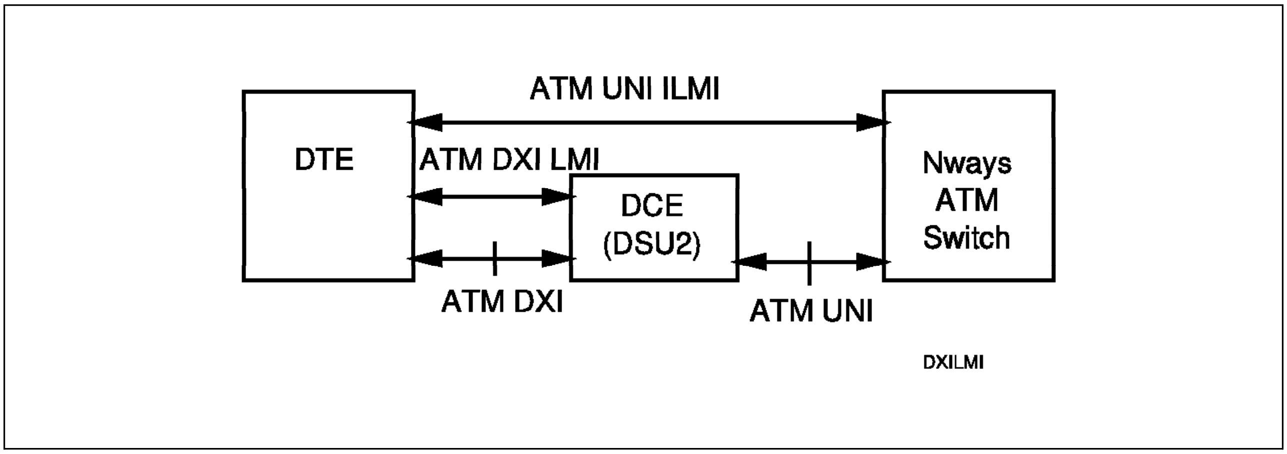

Figure 16. ATM Data eXchange LMI Implementation

2.2.2.2 FUNI Overview

A frame-based user-to-network interface (FUNI) was standardized by the ATM Forum to provide frame-based access to an ATM network. FUNI is based on the ATM DXI interface and has very similar characteristics to ATM DXI. One of the differences between ATM DXI and FUNI is that the ATM DXI interface requires an external unit to make the conversion into cells which are sent over an access line. The following figure illustrates the different accesses via the FUNI and ATM DXI standards:

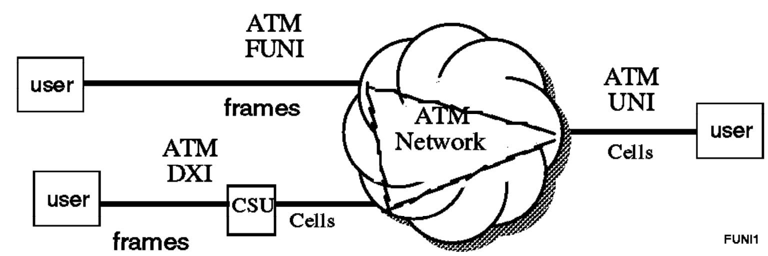

Figure 17. FUNI and ATM DXI Difference

Note: The IBM Nways 2225 Switch integrates DSU functions in the frame-based adapters; therefore, both ATM DXI and FUNI use the frame-based access to the switch. The IBM implementation of the DXI and FUNI interfaces use the same type of algorithm, since the major difference between these two interfaces (external ATM DXI DSU) does not apply in the case of IBM Nways 2225 Switch.

The FUNI interface was designed to support frame-based access to ATM cell networks using low-speed lines. The ATM forum standard defines access line speeds of up to 2048 Kbps; E1, T1, V.35 and fractional n x DS0 channels are supported.

As the FUNI interface uses frames for data transmission, it is not suitable for time-sensitive traffic, such as CBR. The interface was designed to support only the VBR and UBR classes of service, using either AAL 5 or AAL 3/4 adaptation layers.

The IBM FUNI implementation supports up to 975 virtual connections; direct mapping of the Frame Address (FA) to the VPI/VCI fields is provided. FUNI interface supports up to 128 virtual circuits and 16 virtual paths as defined by the ATM Forum standard. The following figure illustrates mapping between the FA and VPI/VCI fields. Note that the same type of mapping is used by the ATM DXI interface.

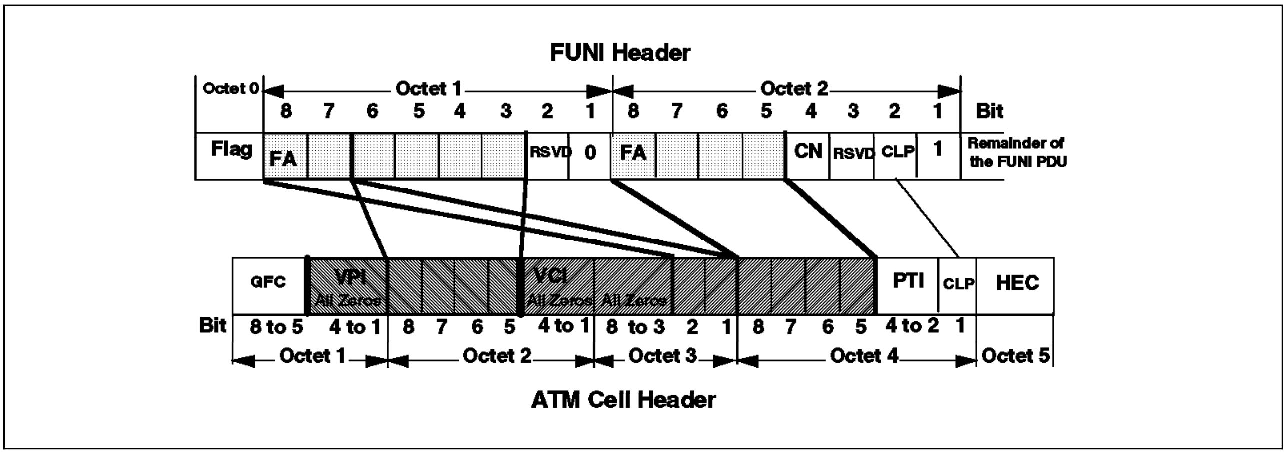

Figure 18. FUNI FA to ATM VPI/VCI Mapping

The FUNI interface uses modified HDLC frames. IBM Nways 2225 provides necessary mapping and frame segmentation. Note that segmentation to the ATM cells is done only when you have installed the ATM UNI adapter, otherwise the FUNI frames are transmitted using the IBM Nways trunk protocol. The following figure illustrates the FUNI reference model.

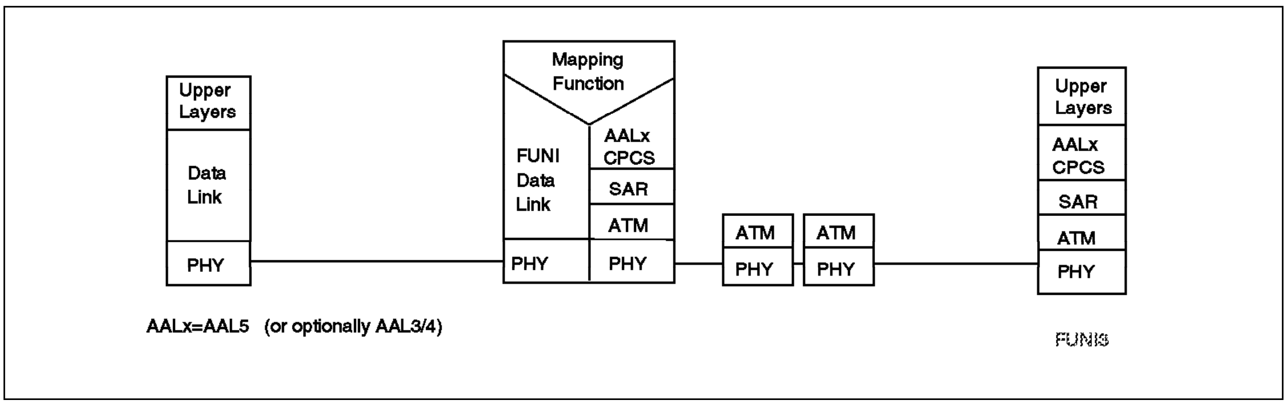

Figure 19. FUNI Reference Model

ATM DXI, FUNI and frame relay frames have a very similar structure. As shown in the following figure, DXI and FUNI headers are identical; however, they are different from the frame relay header. The FUNI is divided into two different types of interfaces, identical to modes 1a and 1b of the ATM DXI standard. For more information, refer to 3.2.5, "IBM Nways 2225 Components" on page 82.

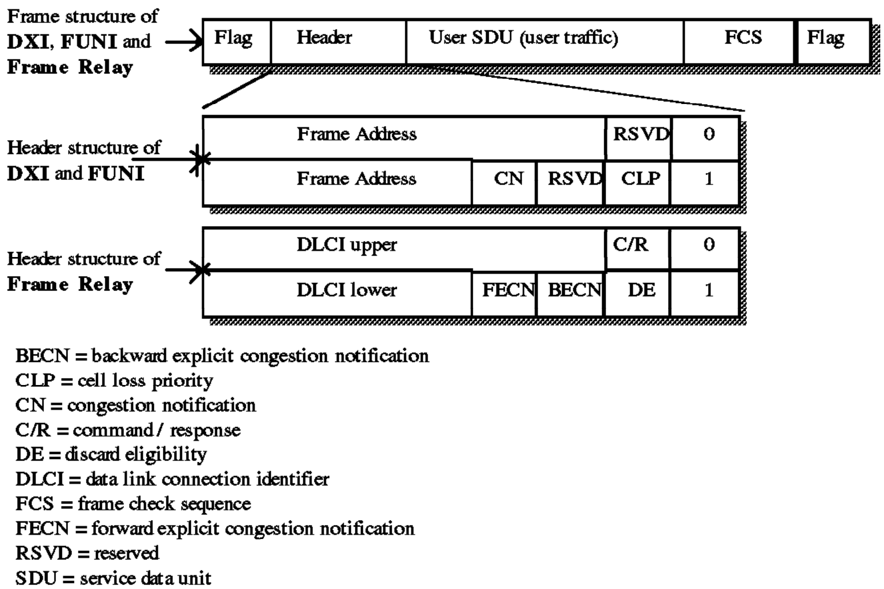

Figure 20. FUNI, DXI and Frame Relay Frame Headers

2.2.3 IBM Nways 2225 Frame-Based ATM Interfaces

The IBM Nways 2225 Switch supports the following types of frame-based interfaces:

ATM DXI Interface - IBM Nways Switch currently supports ATM DXI interface Mode 1a as specified in the ATM Forum document (ATM_Forum/93-590R1). The following features are supported:

- Up to 975 virtual connections
- DTE SDU frame size up to 9232 octets
- 16-bit FCS between the DTE and DCE
- DXI LMI support

IBM 2225 FUNI Interface - The IBM implementation of the FUNI interface is fully compliant with the ATM Forum standard. Both Modes 1a and 1b are supported with the following characteristics:

- Up to 975 virtual circuits
- AAL 5 for mode 1a, both AAL 5 and AAL 3/4 for mode 1b
- DTE SDU size up to 9232 octets, AAL 5 is used
- 16 bit FCS

As the IBM Nways 2225 integrates the DXI CSU in the switch, implementation of both FUNI and ATM DXI interfaces is identical. These two types of logical ports are selectable under the same configuration menu item. The following sections briefly describe the possible configurations of the ATM DXI/FUNI interfaces. The IBM Nways Switch supports these types of configurations:

- FUNI-to-FUNI switching
- FUNI-to-frame relay interworking
- FUNI-to-Cell ATM; this configuration is supported only when you use the ATM UNI I/O modules inside your network.

The IBM Nways 2225 Switch supports the following frame-based ATM logical ports:

- ATM DXI/FUNI DCE interface
- ATM DXI/FUNI DTE interface

2.2.3.1 ATM DXI/FUNI DCE Interface

ATM DXI DCE interface can be defined on any frame-based I/O module. Using this type of interface, you can connect to the IBM Nways Switch network any type of ATM DXI/FUNI DTE device, such as a router, configured as ATM DXI/FUNI DTE. The IBM Nways 2225 DXI/FUNI DCE implementation allows you to configure PVCs between different DXI/FUNI ports or between the DXI/FUNI and UNI ports; in this case, the ATM UNI I/O module performs the fragmentation and reassembly functions. Figure 21 shows possible configurations using the ATM DXI DCE port on the IBM Nways 2225.

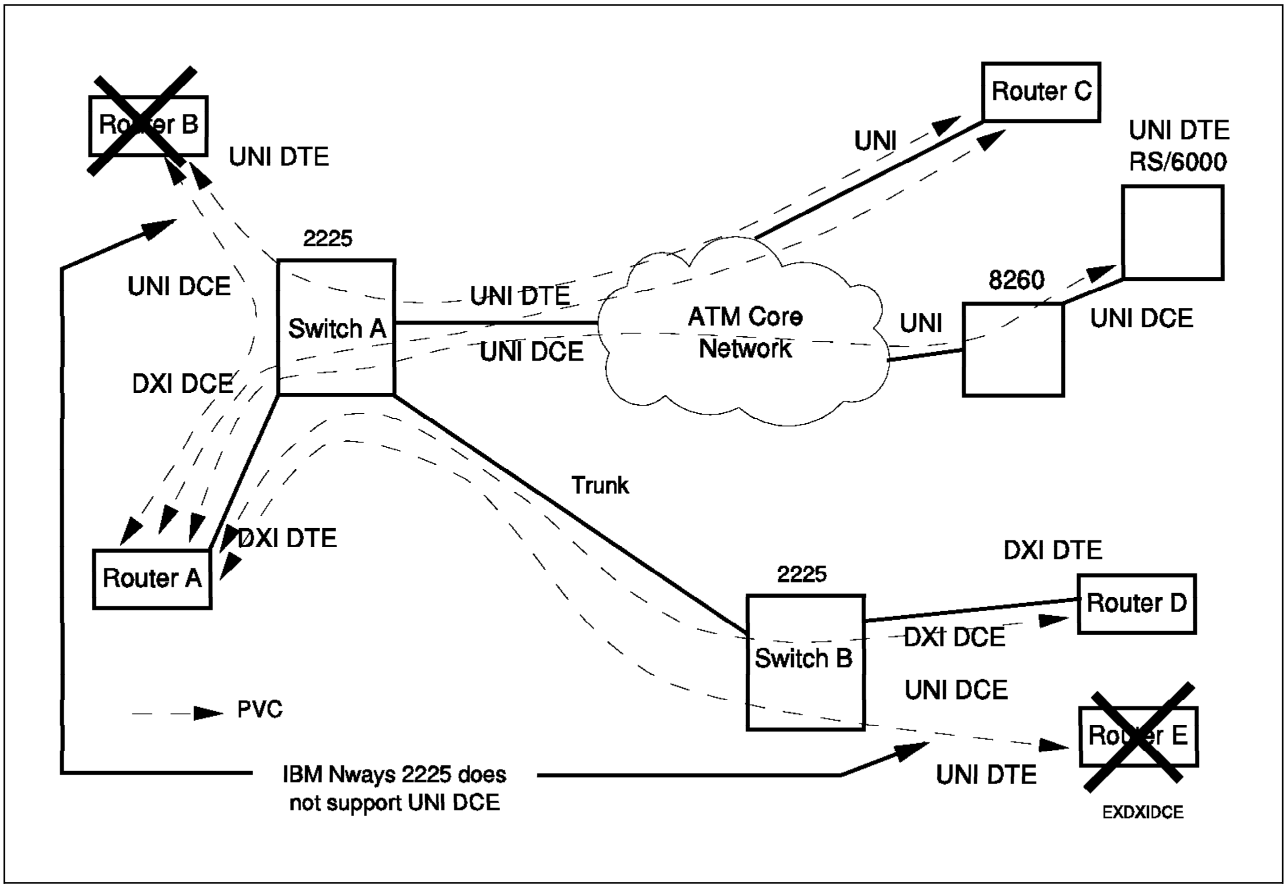

Figure 21. IBM Nways 2225 DXI-DCE Interface Usage Example

Router A is connected to switch A using the ATM DXI DTE interface. Switch A is connected to the ATM core network, which could be based on the IBM Nways 2230 Switches, and to switch B using one of the trunks (direct, frame relay, SMDS or ATM OPTimum trunk). Router C and the 8260 are directly connected to the ATM core network. This type of configuration allows you to define PVCs between all of the ATM UNI and ATM DXI end devices: Routers A, B, D, E and RS/6000. (The same type of higher-level protocol encapsulation is assumed.) In the current release of software, routers B and D are not supported using the UNI DTE interface because the 2225 does not support UNI DCE. This will change with the next release, see below.

The main objective in this example is to show the services provided by the ATM DXI DCE interface in switch A, enabling router A to communicate with:

- Router B, which is attached to the same switch over an ATM UNI interface (available with Release 4.2 on cell-based modules)
- Router C, which is reached over an ATM network
- RS/6000, 8260-attached via an ATM UNI interface
- Router D, which is attached to switch B in a manner equivalent to that of router A
- Router E, which is attached to switch B over an ATM UNI interface (available with Release 4.2 on cell-based modules)

2.2.3.2 ATM DXI/FUNI DTE Interface

This type of configuration is useful when interconnecting IBM Nways 2225 Switches using a low-speed line over the ATM public network or the customer′s high-speed ATM network. An ATM DXI/FUNI DTE interface could be used for ATM DXI/FUNI OPTimum trunking. Figure 22 shows example configurations of an ATM DXI/FUNI DTE interface.

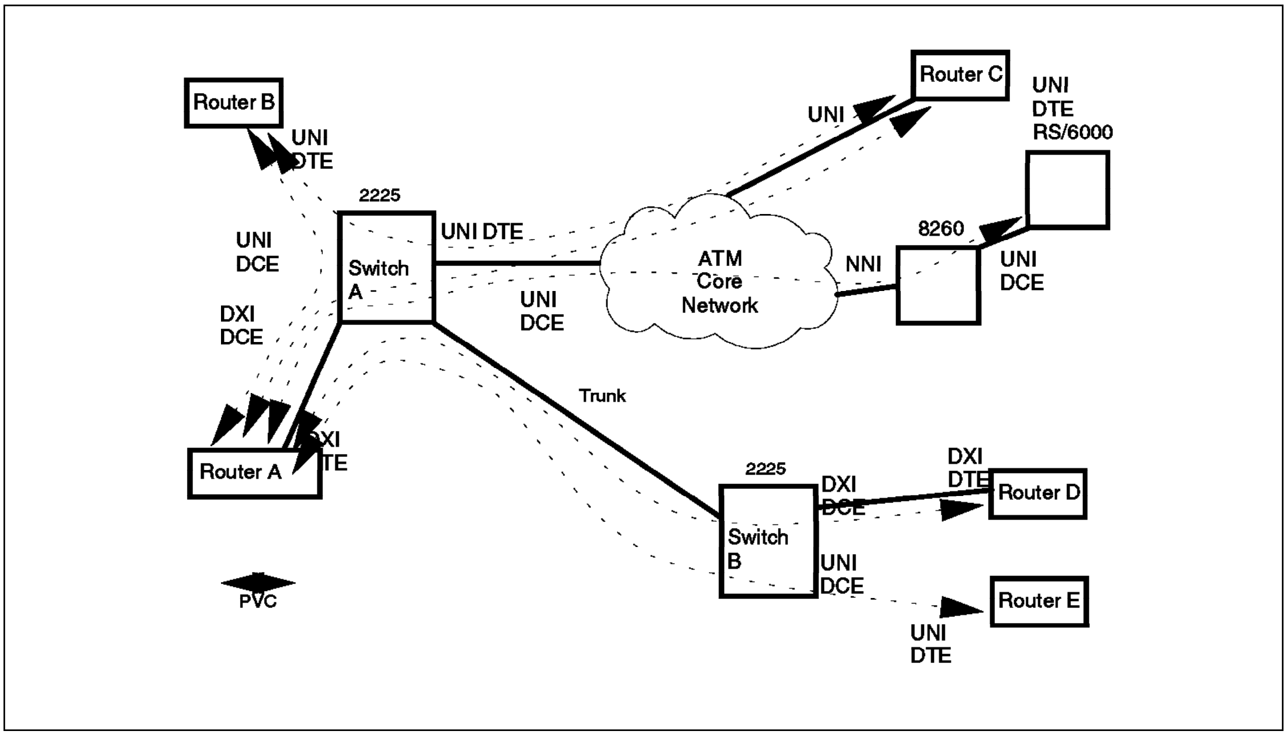

Figure 22. IBM Nways 2225 DXI-DCE Interface Usage Example

Switch A is connected to the ATM core network using the DXI DTE interface. There are two PVCs defined for OPTimum trunk connection to switch B and switch C and one PVC carrying the data traffic between routers A and B.

2.2.4 ATM Cell-Based Interfaces

The IBM Nways 2225 ATM cell-based interface was designed to support IBM Nways 2225 Switch connection to an ATM network. The current version of the ATM cell-based interface supports an ATM feeder, and interworking functions. For future enhancements, refer to Appendix A, “Release 4.2 Preview for the Nways 2225” on page 127.

The following ATM cell-based logical ports are supported:

- ATM UNI DTE interface
- ATM OPTimum trunk interface

2.2.4.1 ATM UNI DTE Interface

IBM has currently implemented the ATM UNI interface on the E3 and DS3 I/O module. You can have multiple modules in the switch to support redundancy or to provide multiple connections to the ATM core networks. However, multiple ATM I/O modules do not implement a switching fabric inside the IBM Nways 2225 Switch, since switching is provided on the level of network interworking and service interworking. This means that IBM Nways 2225 does not perform switching at the ATM cell level. The ATM UNI I/O module supports up to 500 circuits for the frame relay/ATM interworking functions. When configured for OPTimum trunking, up to 123 OPTimum trunks are supported per module; these trunks can support up to 4000 circuits. This interface was developed mainly for interconnection of the IBM Nways 2225 Switches with ATM networks. Current implementation of the ATM UNI interface supports only ATM DTE functions.

Note: We can connect to any device that is an ATM DCE which would include routers or LAN switches.

2.2.4.2 ATM OPTimum Trunk Interface

ATM OPTimum trunk allows you to interconnect IBM switches using the ATM core network (core network could be based on the IBM Nways 2230 ATM backbone switches or you can use public ATM network). This feature allows you to optimize performance and throughput of the network and lower the cost using a single access line to support multiple trunks. Prior to software Release 4.1, only frame relay and SMDS OPTimum trunks were supported.

IBM Nways trunk protocol is fragmented and reassembled using AAL 5 adaptation layer and sent over the predefined PVC to the remote IBM Nways 2225. This configuration multiplexes multiple frame relay PVCs and SMDS SVCs over the one ATM virtual circuit. An ATM UNI I/O module is required for support of the ATM OPTimum trunk feature.

Figure 23 on page 31 shows a simple configuration for this type of trunk over an ATM network.

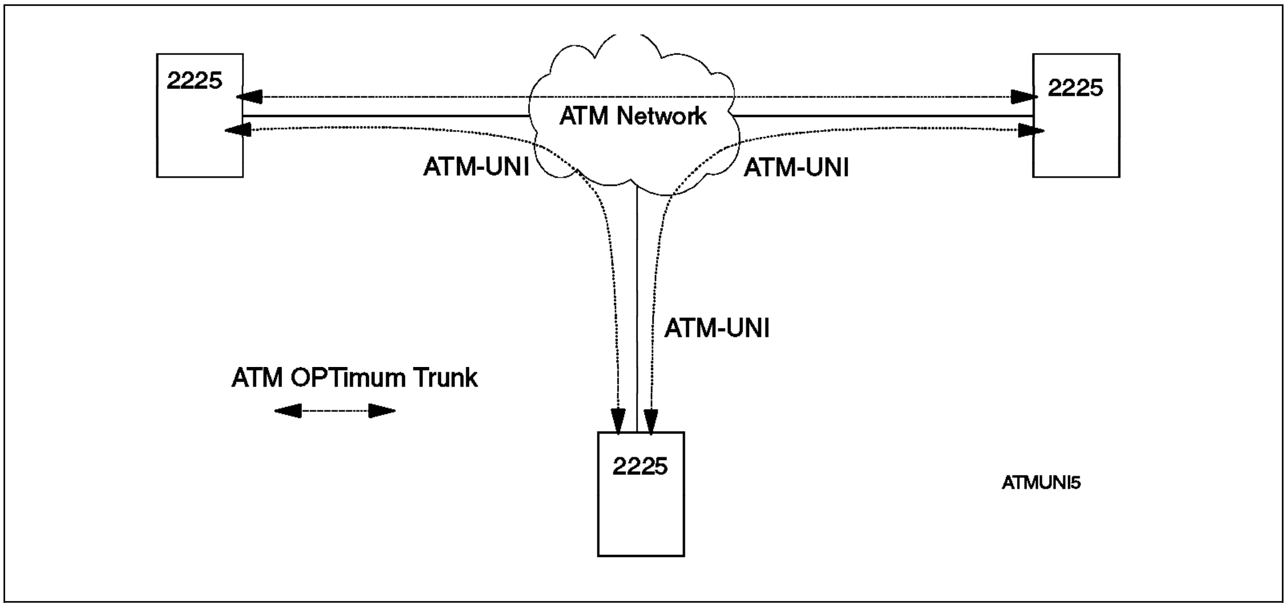

Figure 23. OPTimum Trunk over ATM UNI Interface

2.2.5 Frame Relay and ATM Cell-Based Interworking Functions

Interworking functions enable integration of frame relay and PPP devices to high-speed ATM networks providing a customer with smooth migration to ATM high-speed services. ATM interworking can be provisioned in a number of ways, which can fall into two types of service:

1. Network interworking, enabling two frame relay-based devices to communicate via an ATM service.
2. Service interworking, enabling a frame relay or PPP device to communicate with an ATM device.

IBM Nways 2225 Switch implements the following ATM interworking functions:

- Frame relay/ATM service interworking
- Frame relay/ATM DXI/FUNI service interworking
- PPP/ATM and ATM DXI/FUNI service interworking
- Frame relay/ATM network interworking
- Frame relay/ATM DXI/FUNI network interworking

2.2.5.1 Frame Relay/ATM Service Interworking

Frame relay/ATM service interworking was designed to enable communication between frame relay devices and ATM users; the ATM device performs no frame relay functions and the frame relay device performs no ATM functions. All interworking is performed by the Interworking Function (IWF), which is implemented in the IBM Nways 2225 switch. IBM has implemented frame relay/ATM PVC service interworking according to the Frame Relay Forum Agreement FRF.8. This function enables you to interconnect any ATM device using RFC 1483 encapsulation with a frame relay device using RFC 1490 encapsulation. Figure 24 on page 32 shows the service interworking reference model.

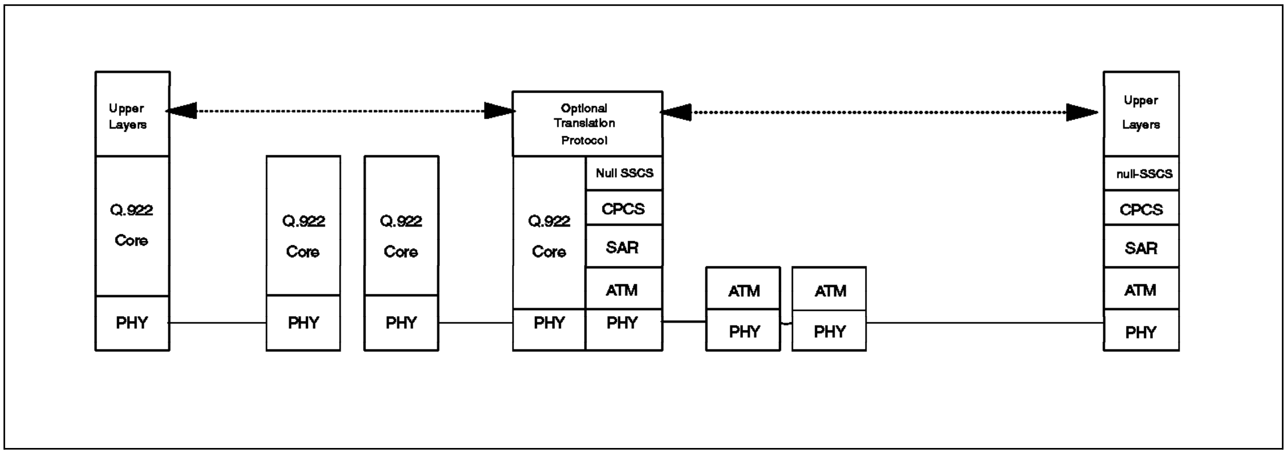

Figure 24. Frame Relay/ATM Service Interworking Protocol Stacks

In the frame relay to ATM direction, the frame relay packets are mapped to the AAL 5 PDUs. IBM Nways 2225 IWF strips the frame relay Q.222 header, maps some of the control fields to the ATM header, and segments the data into ATM cells using the AAL 5 adaptation layer. At the same time, RFC 1490 (Multiprotocol over Frame Relay) translated to RFC 1483 (Multiprotocol over ATM) takes place (in the optional translation protocol part).

In the ATM-to-frame-relay direction, IBM Nways 2225 uses the delineation function of the AAL 5 adaptation layer, optional RFC 1483 to RFC 1490 translated and frame relay encapsulation. IBM Nways 2225 Switch supports translation of IP, IPX, SNA, DECNET AppleTalk and NetBIOS protocol.

The following standard frame relay/ATM interworking flow control mapping is performed by IWF and implemented in the IBM Nways 2225:

- **DE to CLP** - Frame Relay Forum defines two modes of operation for mapping discard eligible (DE) to cell loss priority (CLP):

  Mode 1: DE bit is copied to the ATM CLP field of every generated cell.

  Mode 2: CLP field contains preconfigured value (either 0 or 1). This value is configured by the network administrator.

- **CLP to DE** - In the ATM-to-frame-relay direction, the following modes of operation are defined:

  Mode 1: If one or more cells of AAL 5 PDU has the CLP bit set, then IWF sets the DE bit.

  Mode 2: The DE bit is set to one preconfigured value (either 0 or 1).

IBM Nways 2225 Switch with Release 4.1 is supporting only mode 2 of operation. Mode 1 will be added in a later release which which will include both DE to CLP and CLP to DE mapping.

When the frame relay/ATM Service or Network Interworking PVC is configured, the ATM traffic contract parameters (PCR, SCR and MBS) are stored internally as frame relay PVC descriptors using the (CIR, Bc and Be) to minimize the changes to the PVC manager of IBM Nways 2225 Switch. If you enter the frame relay parameters, software calculates the corresponding ATM values for you using following formulas:

SCR=CIR/(Octet Per Cell * Bits Per Octets)

MBS=(Bc+Be)/(Octet Per Cell * Bits Per Octets)

PCR=(CIR*(1+Be/Bc))/(Octet Per Cell * Bits Per Octets)

CIR=SCR*(Octet Per Cell * Bits Per Octets)

Bc =SCR*(Octet Per Cell * Bits Per Octets)*(MBS/PCR)

Be =(PCR-SCR)*(Octet Per Cell * Bits Per Octets)*(MBS/PCR)

When using adaptation layer AAL 5, the Octet Per Cell is equal to 48, and the Bits Per Octet is equal to 8.

The following figure shows an example of using the frame relay/ATM Service interworking.

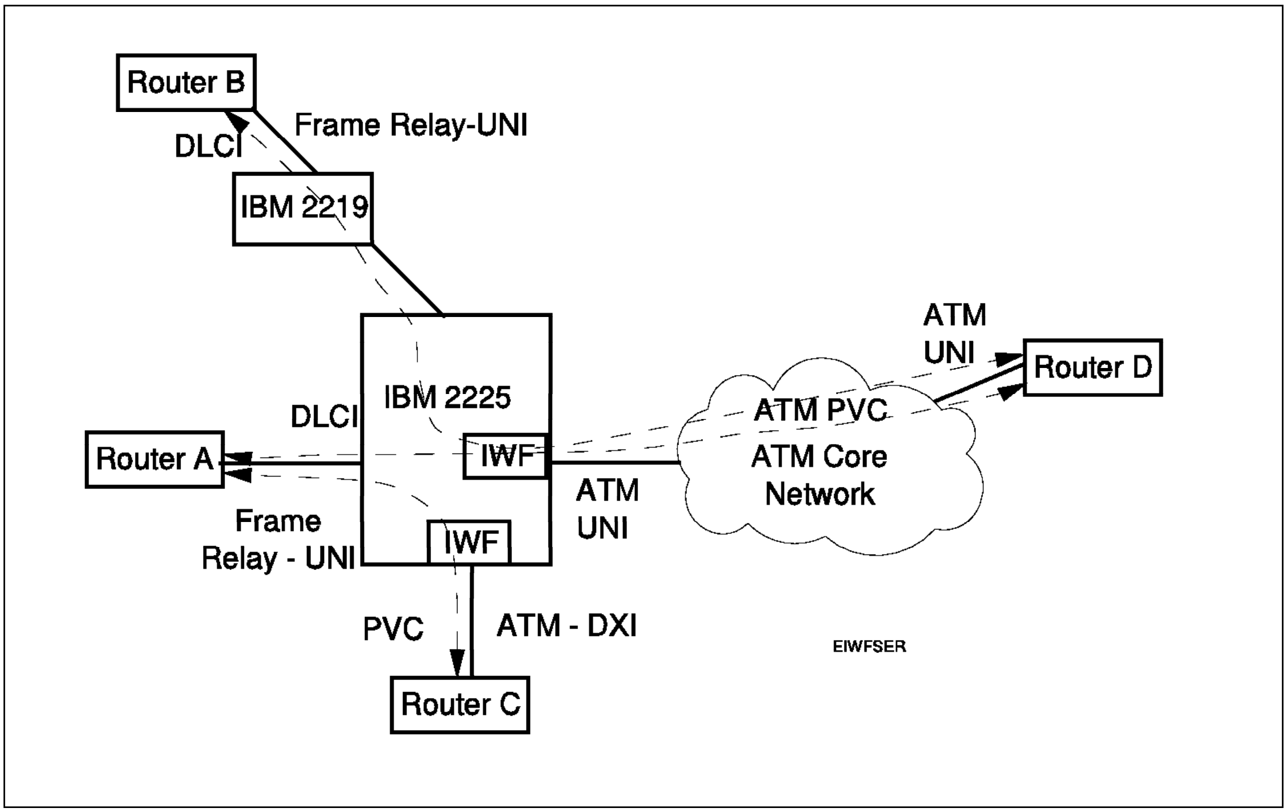

Figure 25. Example of FR/ATM Service Interworking Function

Figure 25 shows the IBM Nways 2225 IWF functionality. Router A and router B are connected to the frame relay part of the multiservice network. Using the IWF function they can communicate with the ATM DXI-connected router C and with the ATM-connected router D. More specifically, routers B and D are communicating according to the protocol diagram depicted in Figure 24 on page 32. This feature allows you to build very scalable multiservice high-speed data networks.

2.2.5.2 Frame Relay/ATM DXI/FUNI Service Interworking

IBM Nways 2225 Switch enables you to configure frame relay/ATM DXI/FUNI service interworking PVC. The functions performed by the switch are very similar to the previous frame relay/ATM service interworking.

2.2.5.3 PPP/ATM and ATM DXI/FUNI Service Interworking

This scenario uses both the translated FRAD function and the service interworking function. The ingress port configured as a translated FRAD translates the PPP frames into frame relay packets. The service interworking function provides the translation of the frame relay either to the ATM cells or to the ATM DXI packets. For more information about the translation FRAD, refer to 2.1.1.5, “Translated FRAD” on page 7.

2.2.5.4 Frame Relay/ATM Network Interworking (Scenarios 1 and 2)

Frame relay/ATM network interworking was put in place to enable communication between two frame relay devices using the ATM network. In contrast with the frame relay/ATM service interworking, both ATM devices have to perform frame relay specific functions running the FR-SSCS (Frame Relay - Service Specific Convergence Sublayer) within AAL 5. This means that both devices (interconnected over the ATM network) only fragment/reassemble frame relay packets using FR-SSCS within AAL 5. Both end devices have to be configured to interoperate at the frame relay protocol layer; no protocol translation exists. Figure 26 shows the frame relay/ATM network interworking Scenario 1, where two switches are connected over the ATM network.

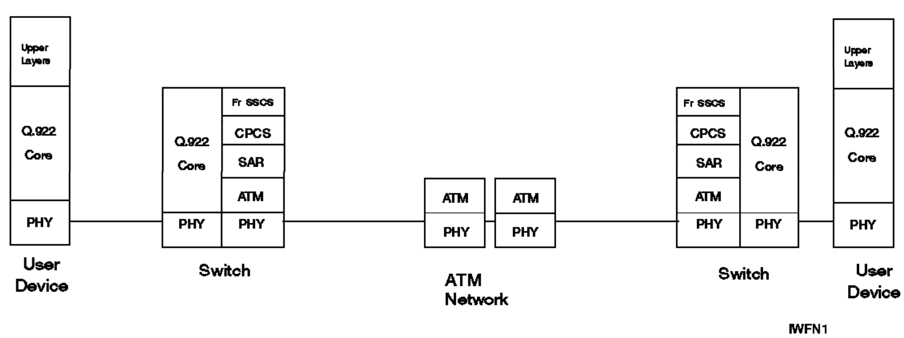

Figure 26. ATM/Frame Relay Interworking over ATM Network

Scenario 2 is shown in Figure 27 on page 35. In this case, the user device performs FR-SSCS/AAL 5 functions and is connected to the frame network over the ATM network using the IWF.

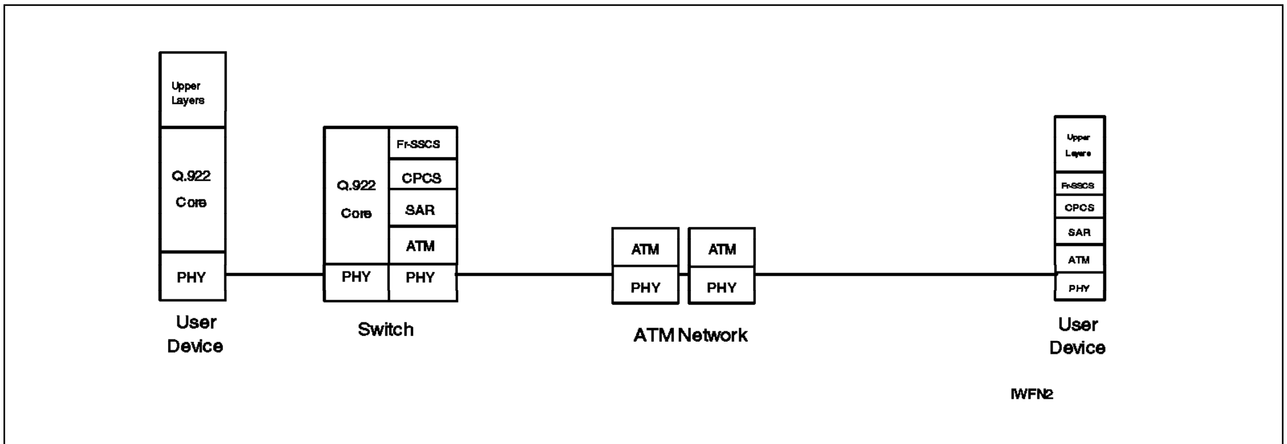

Figure 27. ATM/Frame Relay Interworking with an ATM End System

IBM Nways 2225 implementation of frame relay/ATM network interworking is based on the Frame Relay Forum standard FRF.5. Other functions, such as CLP/DE mapping, are implemented in the same way as described in 2.2.5.1, "Frame Relay/ATM Service Interworking" on page 31. Figure 28 shows a configuration using the frame relay/ATM network interworking.

Note: When you interconnect two IBM Nways 2225 networks over ATM, you should use the ATM OPTimum trunk instead of the frame relay NNI over ATM to preserve the added value of IBM Nways trunk protocol.

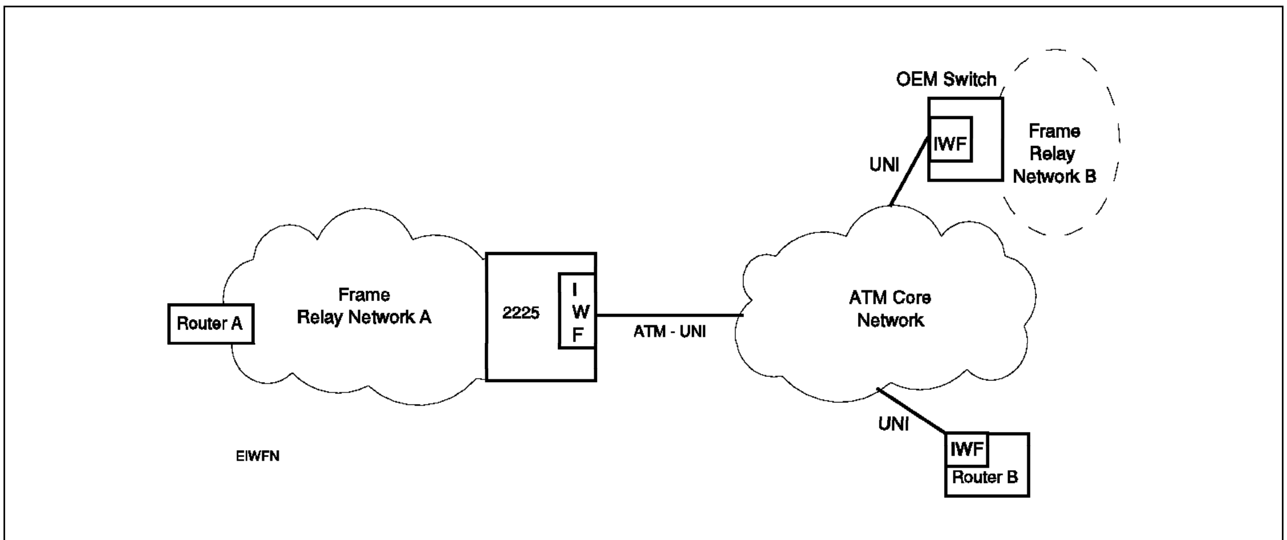

Figure 28. Frame Relay/ATM Network Interworking Example

The above example shows two frame relay networks (network A based on the IBM Nways 2219 and IBM Nways 2225 switches, and network B based on OEM switches) interconnected over an ATM network using the UNI interface. Any frame relay device connected to frame relay network A can communicate with any device on frame relay Network B (scenario 1). Routers A and B can exchange frame relay frames (scenario 2). Note that router B has to support FR-SSCS (that is, a Service Specific Convergence Sublayer function for frame relay); the reader can contrast this with the frame relay/ATM service interworking section above.

2.2.5.5 Frame Relay/ATM DXI/FUNI Network Interworking

IBM Nways 2225 Switch enables you to configure frame relay/ATM DXI/FUNI network interworking PVCs. The functions performed by the switch are very similar to the previous frame relay/ATM network interworking; however, the switch performs the frame fragmentation/reassembly and traffic shaping functions.

2.3 IBM 2225 Traffic Shaping

Traffic shaping is typically implemented at the port level when interconnecting two different networks. Bursty traffic can be handled internally inside the IBM Nways Switch network. However, other connected networks might not be able to accept these bursts of data. IBM ATM UNI modules provide you with strong traffic shaping features. The traffic shaping control has no influence on the source of the bursty traffic, and it does not modify its behavior. The ATM UNI I/O module has significant buffers for burst smoothing purposes. In extreme cases, cells may be lost due to buffer overflow.

ATM UNI I/O modules support up to eight separate traffic shaping queues, each queue being configurable for the following ATM traffic parameters:

Peak Cell Rate (PCR) - Peak cell rate is the maximum cell transmission rate. It defines the shortest time between two cells. However, cells can be sent at the peak cell rate only if there are available credits. One credit is consumed for each transmitted cell. The maximum number of credits is defined by the maximum burst size. When all credits are used, the user end device can transmit at the sustainable cell rate. New credits are accumulated when no traffic is transmitted.

Sustainable Cell Rate (SCR) - Sustainable cell rate is the average cell rate, which is guaranteed by the network under normal conditions.

Maximum Burst Size (MBS) - Maximum burst size determines the maximum number of cells that can be transmitted at the peak cell rate. In other words, the MBS specifies the maximum amount of credits that can be accumulated.

The ATM UNI I/O module supports eight rate queues, which are used for frame segmentation and traffic shaping. The rate queues are organized in two banks of four queues. Bank 1 has higher priority than bank 2; traffic assigned to a higher priority group can preempt the low priority group traffic when oversubscriptions occur. The PCR is user configurable for each of the eight queues.

Each rate queue maintains an internal timer that expires at the configured PCR rate (for example, for a PCR of 64 KB, timer is 1/64 seconds). When the timer expires, the hardware will service any frames linked to that rate queue. Rate queues in the bank (for example, Q1, Q2, Q3, etc.) are serviced sequentially. The rate queue data structure is two-dimensional. For the same rate queue, all frames with the same VC are linked horizontally, while all frames from different VCs are linked vertically. The following figure illustrates the implementation of the rate queues.

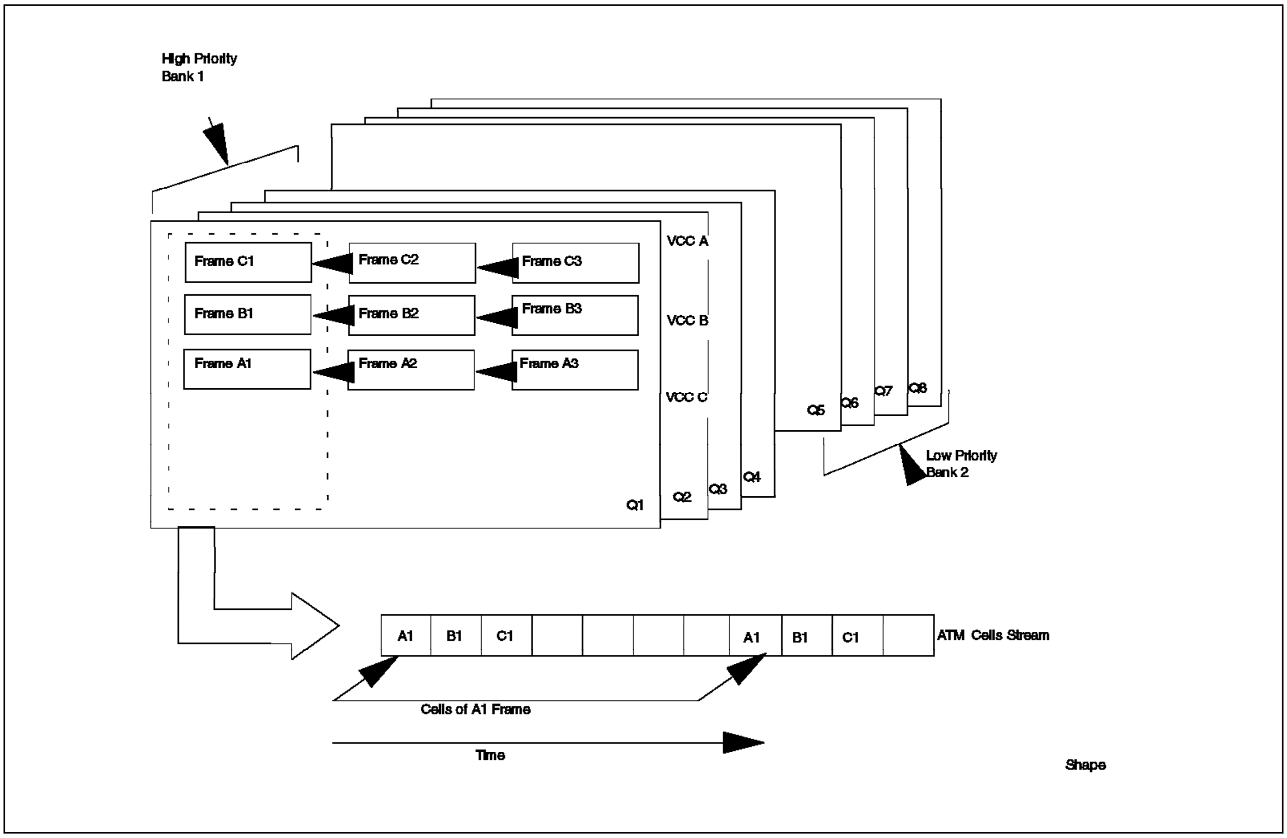

Figure 29. ATM UNI I/O Module Traffic Shaping Queues

Frames related to interworking circuits are placed in the rate queues according to their priority. Setting circuit priority to 1 or 2 selects PCR queue bank 1; priority 3 circuits use a queue in bank 2. To select the rate queue for a particular PVC, the switch chooses the queue with a configured PCR greater than or equal to that of the circuit. If you are configuring an ATM OPTimum trunk, you can choose one of the eight rate queues at configuration time.

If you define a frame relay PVC, as Bc equal to CIR and Be equal to zero, the ATM traffic shaping for that connection (assuming traversing ATM Core Network) is turned off, and MBS is forced to be equal 32. In reality, MBS does not affect traffic shaping, as the PCR is equal to the SCR. Note that configuration with Bc=CIR and Be=0 turns off the SCR/MBS traffic shaping, but the PCR traffic shaping is still in use. Frames sent over this connection will be segmented to cells using the PCR rate defined for this queue.

2.4 IBM Nways 2225 PVC Manager

IBM does not implemented CAC functions in the IBM Nways 2225 switch. However, a very similar approach was taken to determine if the requested connection should be accepted by the network and which route in the network should be used. These tasks are performed by OSPF and the PVC manager, which takes care of any type of connection inside the network. OSPF runs in each IBM Nways 2219 and IBM Nways 2225 Switch and is responsible for selecting the optimum path for circuits.

The implementation of the PVC Manager is based on the virtual bandwidth (see 2.1.3.6, “Virtual Trunk Bandwidth” on page 15). For ATM type connections, conversion of the ATM connection descriptor (PCR, SCR and MBS) to corresponding frame relay parameters is used. This allows for reusage of the PVC manager for ATM type connections.

2.5 SMDS Overview

Switched Multimegabit Data Service (SMDS) is a connectionless, cell-switched data transport service developed by Bellcore mainly for metropolitan area networks (MANs). SMDS provides the flexibility needed for distributed computing and bandwidth hungry applications. Since SMDS is connectionless, it provides any-to-any services similar to those implemented in local area networks. Network operators do not have to predefine paths between each end-user device. Meshed networks, where each site requires connectivity with all other sites, can be built in a very efficient way. Significant savings can be achieved because less access lines are required, thereby minimizing distance charges associated with predefined circuits. SMDS began many trends, which were later adopted in ATM technology. IBM has implemented SMDS in IBM Nways 2225 Switches. Please note that you can use IBM Nways 2225 for SMDS switching, frame relay switching, The following types of services are available:

- SMDS DXI-SNI (subscriber-to-network interface)
- SSI (SMDS-to-access-server interface)
- SMDS OPTimum Trunk

2.5.1 SMDS Architecture

SMDS architecture uses the three-layer model based on the Distributed Queue Dual Bus (DBDQ) protocol defined in the IEEE 802.6 standard. This three-layer model includes functions for addressing, framing, error detection and physical transport:

- **Level 3** - Contains the user data, source and destination address and other header information.
- **Level 2** - A 53-byte long cell contains segmented layer 3 PDU including a level 3 header.
- **Level 1** - The physical layer with its physical interface specifications.

The following figure shows the layer 3 and 2 relationship.

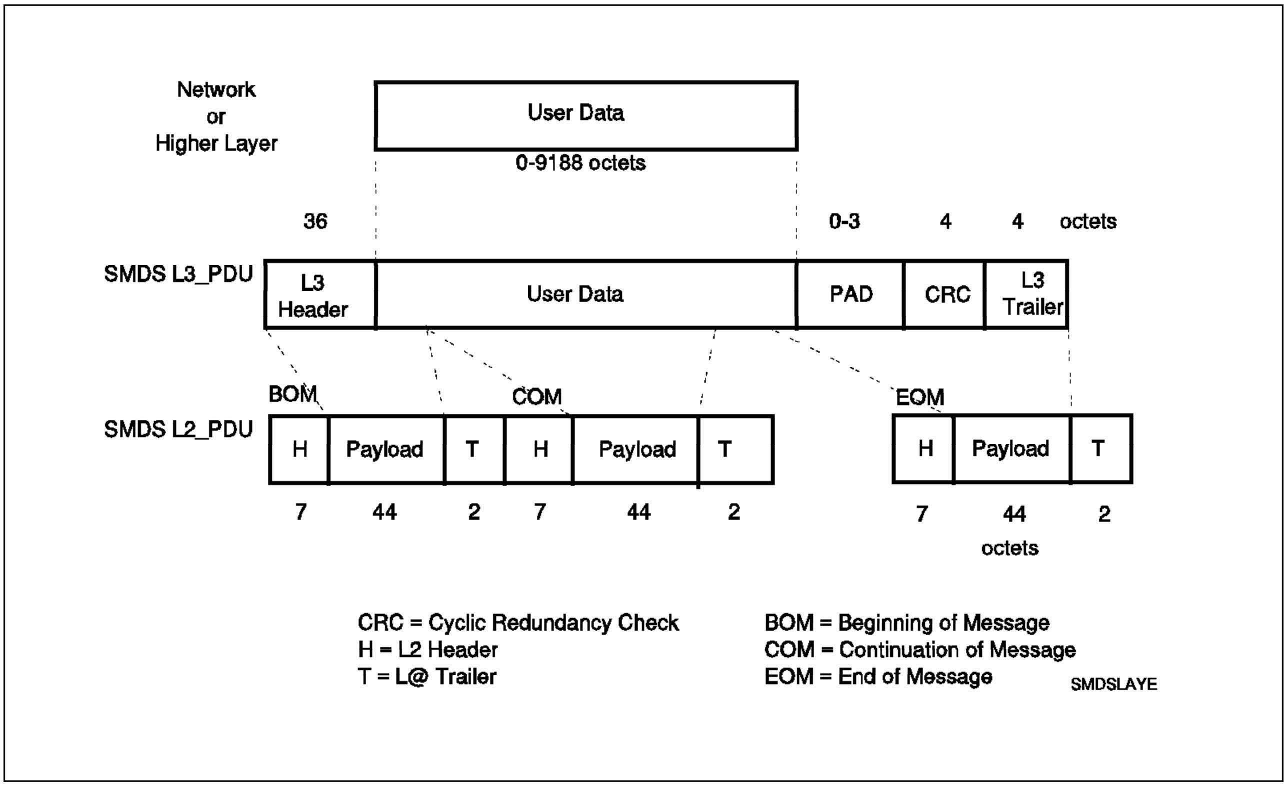

Figure 30. SMDS Layer Relationship

Layer 3 PDUs are segmented into 53-byte cells and sent to the SMDS Switching System (SMDS SS). The first cell (BOM - Begin of message) contains the full address of the destination, so the switch needs to examine only the first cell; other cells (COM - Continuation of Message and EOM - End of Message) are switched based on the information in the BOM cell. SMDS SNI is based on the DBDQ architecture, which is outside the scope of this redbook. If you need more information about DBDQ, refer to *High-Speed Networking Technologies*, GG24-3816. IBM Nways 2225 and IBM Nways 2219 do not implement the SNI interface.

2.5.1.1 SMDS Addressing

SMDS uses an E.164 address scheme (E.164 address is represented up to 15 binary coded decimal (BCD) digits). A user device (Customer Premises Equipment) can send a datagram to a fixed destination using point-to-point service or to multiple destinations using the point-to-multipoint service, defined as multicast address. The multicast address is defined for each closed user group by the network administrator. The SMDS architecture supports a screening function, which ensures that data will be delivered only to specific destinations. Screening can be used to create a closed user group; the following features are available:

- SMDS source address is validated for access lines to prevent other devices from connecting to the network.
- End device can block or accept specific blocks of data.
- User can limit destinations, which can be reached from a specific access line. IBM 2225 Switch allows you to define up to 260 E.164 destination addresses per I/O module.

2.5.1.2 SMDS Low-Speed DXI-SNI Interface

As in the case of ATM DXI, low-speed SMDS interfaces enable the CPE to access the SMDS services over the low-speed lines starting from 56 Kbps or 64 Kbps to T1/E1 speeds. Currently, there are three types of low-speed interfaces to access SMDS network:

- SMDS Interface Protocol over Frame Relay Interface (SRI), which allows the CPE to exchange SMDS layer 3 PDUs using the ANSI T1.618 frame relay frames at layer 2.
- Frame Relay-to-Access-Server Interface (FSI) enables the CPE to connect to SMDS over the frame relay network.
- Data Exchange protocol Interface/Subscriber-Network Interface DXI-SNI.

IBM Nways 2225 implements only the DXI-SNI interface. Native 802.6 cell format can be supported by various types of external OEM DSUs. Instead of using DBDQ level 2 and level 1 interfaces, DXI-SNI uses standard HDLC protocol over synchronous links. SMDS layer 3 PDUs are not segmented to the 53-byte cells using the DBDQ algorithm; instead, they are encapsulated into HDLC frames and sent to the IBM 2225 SMDS access server.

The SMDS access server enables you to combine the high-performance T3/E3 connections to the SMDS switching system with low-speed SMDS offering to traditional WAN users at speeds from 56 Kbps or 64 Kbps to T1/E1. The following figure illustrates SMDS access server functionality.

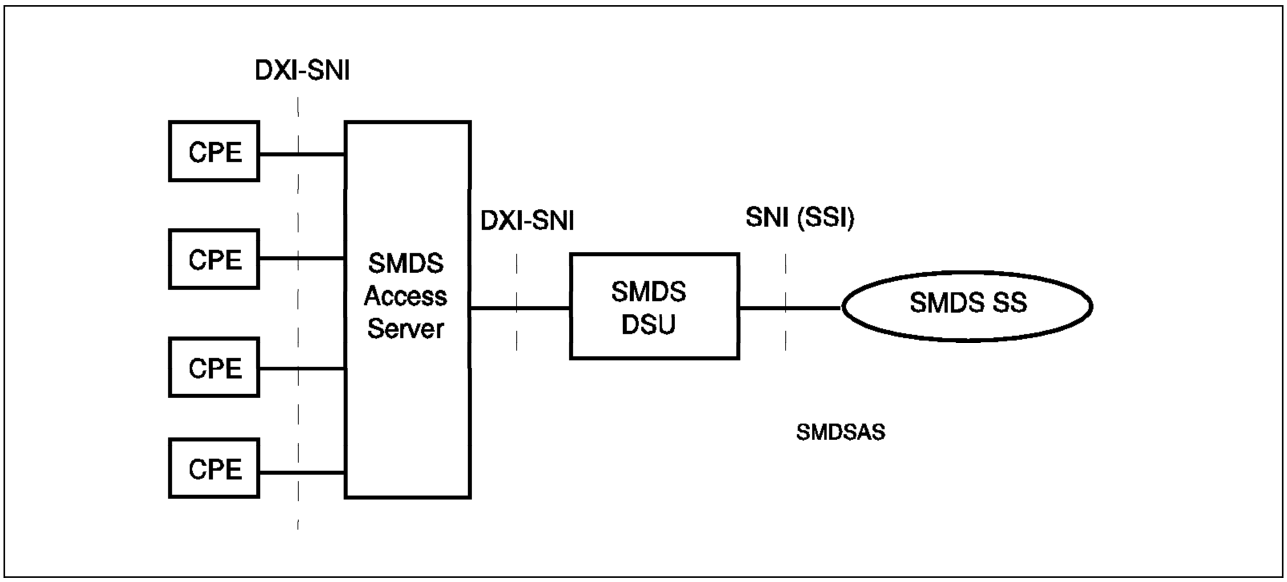

Figure 31. SMDS Access Server - Low-Speed Access to SMDS Network

IBM has implemented SMDS access server in the Nways 2225 Switch.

2.5.2 IBM SMDS Services Implementation

The rest of this chapter describes IBM implementation of SMDS services in the IBM Nways 2225 Switch. Currently, the IBM switching platform provides low-speed SMDS Access services and enables routing of SMDS datagrams between IBM switches implementing a functionality called SMDS access server/switching. The following key functions are implemented:

- **Switching of SMDS Datagrams** - IBM switches implement SMDS switching functionality between different types of DXI-SNI interfaces.

- **Network Management Functions** - You can manage remote IBM switches using the SMDS services.
- **System Functions** - The following system functions are provided:
  - Individual Address support
  - Group Address support
  - Source Address Validation
  - Address Screening
  - Heart Beat Poll
- **Operational Functions** - The support includes:
  - Memory Administration
  - SMDS Network Maintenance
  - Network Traffic Management
  - Network Data Collection
  - Customer Network Management
- **Performance and Quality of Service** - Provides support for SMDS OPTimum Trunk and IBM SMDS Management Information Base (MIB).

For more information about the management part of the SMDS network, refer to Chapter 4, "Managing Networks Using Nways Wide Area Element Manager" on page 111.

2.5.2.1 Routing in SMDS Network - Area IDs and Subscriber Numbers

IBM switches reuse the OSPF routing protocol implemented for the frame relay protocol. OSPF maintains and distributes the routing tables for adjacent switches in the SMDS network, treating the E.164 addresses as IP addresses; mapping of E.164 to IP addresses is performed (see Table 2 on page 42). Using the OSPF algorithm, IBM switches are able to calculate the path over which to send SMDS datagrams. The following types of trunks can be used to transport SMDS datagrams:

- SMDS, frame relay or ATM OPTimum trunks
- Direct trunks

When you define an SMDS network based on IBM Nways switches you have to divide the E.164 address into two parts: the *area number (area ID)* and the *subscriber number*. area ID is defined using the area mask; the area mask has to be the same for the whole SMDS network. The Area ID can start at any digit in the E.164 address, and the length can be up to 8 digits.

The rest of the address is used for identification of the end systems. Area ID and the subscriber number together use 15 BCD digits of the E.164 address. IBM Nways Switch defines area IDs automatically by applying the area mask to the individual addresses. As the area ID is used for the routing decision inside the network, defining the same address for two switches will lead to indeterminate routing inside the network. You can define more area IDs per switch but you can use only one switch per area.

The following example illustrates functionality of IBM Nways switches routing algorithm in an SMDS network. A 6-digit area mask was used.

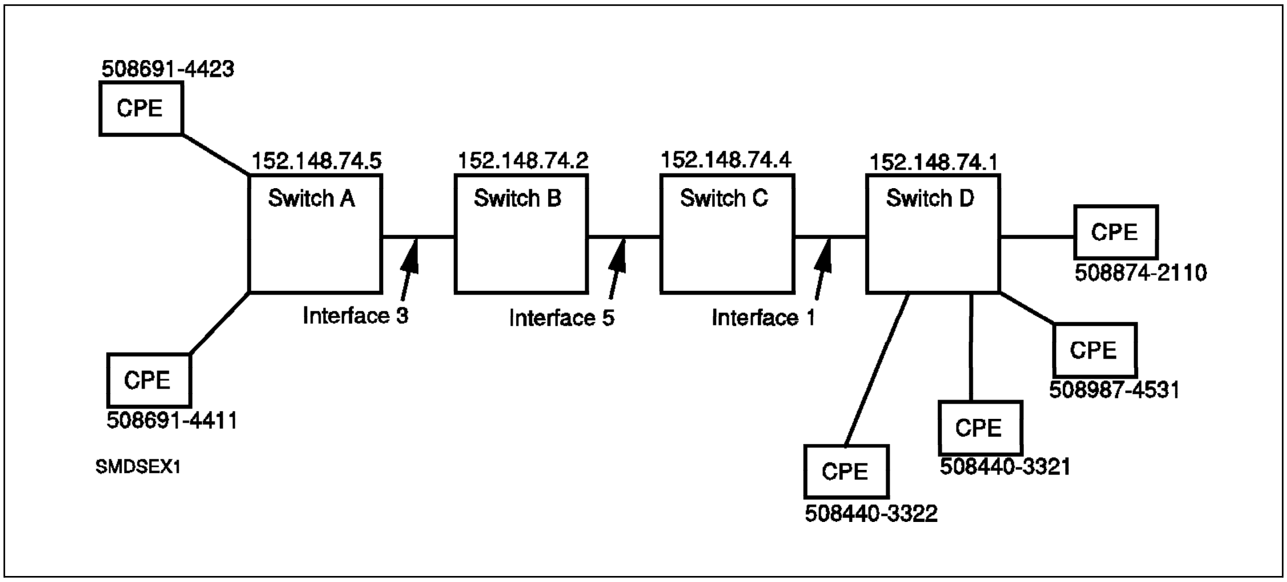

Figure 32. Example of SMDS Network Setup

The network setup shown in the figure results in the following routing table in switch A:

Table 2. Switch A Routing Table

Destination Area IDs	Next Hop Interface Number	Route Interface Number	Forwarding Node
508440	3	3,5,1	152.148.74.1
508987	3	3,5,1	158.148.74.1
508874	3	3,5,1	158.148.74.1
152.148.74.2	3	3	-
152.148.74.4	3	3,5	-
152.148.74.1	3	3,5,1	-

2.6 IBM Nways Switch SMDS Interfaces

The symmetrical architecture enables you to define on the same I/O modules different types of logical ports. One port of the I/O module can be used for frame relay service, while the other ports can be defined for SMDS or ATM DXI connections. Note that channelized E1 and T1 ports have the same capability per n x DS0 channel. IBM has implemented the following types of SMDS interfaces:

DXI-SNI Interface - Using this interface, you can connect user devices to the SMDS network implemented by the Nways switches.

SSI Interface - SSI (SMDS to access server interface) is used to connect the IBM switches to the SMDS cell-based switching system. IBM switches use the HDLC frames to transport the user datagrams to the DSU interface, DSU performs segmentation and sends the cells to SMDS SS using the SMDS cell-based interface. This interface enables IBM Nways Switch users to communicate with the devices connected directly to the SMDS SS cloud using the DBDQ protocol. You can of course interconnect Nways 2219, 2225, 2230 WAN Switches over the SMDS SS cloud as if they were SMDS devices, but using OPTimum trunking is a better way to do it. IBM OPTimum trunking provides value-added features and keeps the Nways 2219, 2225, 2230 WAN Switches in the same IBM OSPF domain. Figure 33 illustrates the use of the SSI interface.

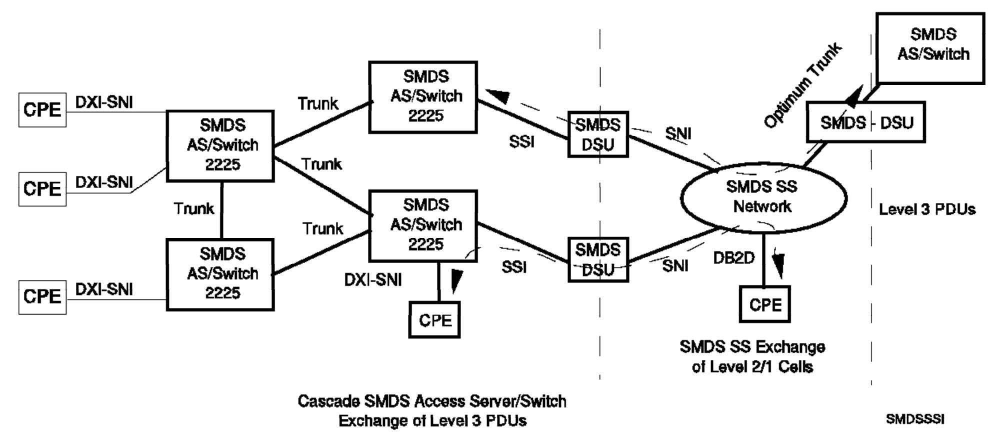

Figure 33. Usage of SSI Interface

SMDS OPTimum Trunk - When you connect an IBM Nways Switch to the SMDS SS cloud using the SSI interface, you are able to define SMDS OPTimum trunks over the SMDS SS network. SMDS OPTimum trunks use IBM Nways trunk protocol. Both sides of the trunk have to be Nways 2219, 2225, 2230 WAN Switches. The same physical port can be defined for one or more logical ports with the SMDS OPTimum trunk option. This capability allows you to use single physical ports for multiple destination trunks. Note that SMDS datagrams can be sent over any trunk supporting the IBM Nways trunk protocol. As an SMDS OPTimum trunk uses the IBM Nways trunk protocol, it can transport frame relay, SMDS, and managements. During trunk configuration, you are able to specify which type of protocols will be transmitted over this trunk. You have the following possibilities:

- Management traffic only
- SMDS datagrams and management traffic
- Multiservice traffic including the SMDS, management, and frame relay

2.6.1.1 SMDS Local Switching and Multiplexing

When you configure an SMDS DXI-SNI logical port, you can specify either a local switched logical port or a multiplexed logical port. When using the local switched option, all SMDS datagrams received on that particular port will be switched inside the IBM Nways OSPF domain; in other words, the SMDS datagrams will not be able to traverse an SSI interface.

The multiplex option allows you to associate a logical port with an SSI port. This association allows the switch to understand which SSI port should be used for the group addressed datagrams and for the datagrams with unknown area IDs. The multiplexed option does not disable local switching functions; the multiplexed logical port should, in fact, be considered a locally switched one with added SSI capability.

2.6.1.2 Delivering Individually Addressed Datagrams

Nways 2219, 2225, 2230 WAN Switches use virtual circuits to transport the SMDS datagrams through the network. These virtual circuits are established automatically between the nodes which provide SMDS services. The circuits are created only between switches at the time you configure the SMDS service on the switches. These virtual circuits are node-to-node connections, not port-to-port connections. Current implementation handles all SMDS datagrams as green packets with priority equal to 3. For future enhancements, refer to Appendix A, "Release 4.2 Preview for the Nways 2225" on page 127.

When the individual addressed packet is queued in the ingress SMDS access server, the following routing decisions are made:

If the destination area ID is in the local switch, then the datagram is delivered to the local DXI-SNI interface.

If the area ID is not in the local switch, but is in the IBM Nways OSPF switch domain, the datagram is delivered to the switch connecting the destination area ID. Local switching based on full SMDS address is then performed.

If the area ID lies outside the IBM Nways Switch OSPF domain, then the appropriate SSI interface is used to deliver the datagram through the public network.

2.6.1.3 Delivering Group Address Datagrams

One of the major advantages of the SMDS architecture is the ability to send a datagram to multiple destinations using the multicast function. You are able to define multiple group addresses per closed user group. When the CPE sends a datagram with the destination address field equal to one of its multicast group addresses, IBM switches deliver the datagram to all closed group destinations using the following process:

- Ingress switch delivers the datagram to all DXI-SNI interfaces which are in the same group. If the ingress DXI-SNI port is multiplexed to an SMDS SS cloud, then a copy is sent to the SMDS SS cloud using the SSI interface.
- Sends a copy of the packet to each switch in the same IBM OSPF domain.
- Other switches use step 1 to distribute group addressed datagrams.

2.7 Integrated Services Digital Network (ISDN)

ISDN technology was designed to improve the quality of voice transmissions and enable networking services at higher speeds. ISDN standard describes only the interfaces, it does not include the implementation inside the telecommunication switching system. There are two types of ISDN services:

- Narrowband ISDN
- Broadband ISDN

This section covers implementation and usage of the narrowband ISDN services inside the IBM Nways 2225 Switch. Narrow ISDN implements the TDM type of connections (64 Kbps and n x 64 Kbps) over the switched network.

ISDN provides a low-cost alternative to leased lines to access IBM Nways 2225 Switch. IBM has implemented the following features inside the IBM Nways 2225 Switch:

- ISDN Backup Trunk
- ISDN DTE inbound access to the IBM Nways 2225 Switch

Note: IBM Nways 2225 does not provide ISDN switching capability.

ISDN defines two basic types of access interfaces:

> **Primary Rate Interface (PRI)** - PRI is the only interface which is implemented on the IBM Nways 2225 Switch; it defines 23B+D channels over T1 for US standards; European standards use 30B+D channels over E1. You could use PRI for connecting the IBM Nways 2225 to the public ISDN network to enable customer′s access over the ISDN services.
>
> **Basic Rate Interface (BRI)** - BRI interface is defined as a 2B+D channel and will probably be used most by the user devices to access the IBM Nways 2225 multiservice network. Currently, IBM does not provide PRI interface on its switching platform.

For more details about the ISDN technology, please refer to *High-Speed Networking Technologies: An Introductory Survey*, GG24-3816.

2.7.1.1 IBM Nways 2225 ISDN Signalling

In the 2225, signalling messages are sent over the D-channel using the LAPD protocol. Packets encapsulated inside the LAPD protocol carry various types of information including called-up destination numbers. There are several types of signalling protocols describing the messages format, which are not compatible.

The ISDN switching system exchanges messages with the IBM Nways 2225 Switch over the D-channel. The PRI T1 I/O module allows you to define one D-channel per physical interface or define one super D-channel per whole I/O module. Standard D-channels always use the timeslot 24 (US version) or timeslot 16 (European E1 version).

When using the super D-channel (sometimes referred to as Non-Facility Associated Signalling), you define only one D-channel per whole I/O module. Super D-channel enables you to use more timeslots for user data. Please note that you can configure a I/O module to use either standard D-channel signalling or the super D-channel signalling. During configuration of the D-channel, you select the type of ISDN signalling used by the ISDN network provider.

IBM plans to extend the super D-channel capability to enable support of signalling associated with more than one I/O module.

IBM has implemented ISDN interface for the four-port T1 I/O module providing a 23B+D Primary Rate Interface. An E1 interface will be supported in the next release of IBM Nways 2225 software. IBM has implemented all necessary protocol layers including the data link layer Q.2931, network layer Q.2931/Protocol Control and application layer Q.2931/Call control. The following types of signalling are currently available for D-channel:

- AT&T 5ESS
- AT&T 4ESS
- Northern Telecom DMS 100

Other types of signalling will be implemented in the next releases of software.

2.7.1.2 ISDN Access Services

ISDN Access enables users to connect to the IBM Nways Switched network over the public ISDN network. The following types of services are supported via the ISDN interface on the IBM Nways 2225 Switch:

- **PPP over ISDN to frame relay** - ISDN user connects to the B-channel logical port which translates PPP protocol to the RFC 1490 and sends the data using the frame relay transport protocol to the remote destination, connected via frame relay UNI-DCE logical port. In such a case, B-channel logical port provides function of the Translation FRAD.
- **PPP over ISDN with ATM interworking** - Accessed B-channel logical port provides the PPP RFC 1490 translated. The far-end ATM device is connected over the PVCs configured for frame relay/ATM Service Interworking.
- **Frame relay over ISDN to frame relay** - Both logical ports (ISDN B-channel logical port and frame relay logical port) are defined for frame relay UNI-DTE services.
- **Frame relay over ISDN to ATM** - Accessed B-channel logical port is configured for the frame relay services. Communication with the ATM device is accomplished via PVC configured either for FR/ATM service or network interworking (this depends on the protocol stack executed in the ATM device).
- **SMDS over ISDN to SMDS** - B-channel is configured for the SMDS SNI-DXI DCE services to connect to the SMDS user closed group.

2.7.1.3 ISDN Trunk Backup

ISDN trunk backup provides a fault-tolerant networking service. Backup trunks are activated automatically when trunk failures occur or you can schedule activation of the trunk (for maintenance purposes). Immediate manual activation is also possible.

This function is supported over any type of I/O module with V.35 or X.21 interfaces, connected to an ISDN terminal adapter (ISDN-TA). IBM Nways 2225 control signals of the physical interface inform the ISDN-TA to initiate the ISDN call procedure. This setup is completely ISDN-transparent to Nways 2219, 2225, 2230 WAN Switches; all ISDN logic and signalling is performed in the ISDN-TA. IBM switches use the modem control signals to manage the ISDN-TA adapter. Both ends of the backup trunk have to be equipped with ISDN-TA as is illustrated in the following figure.

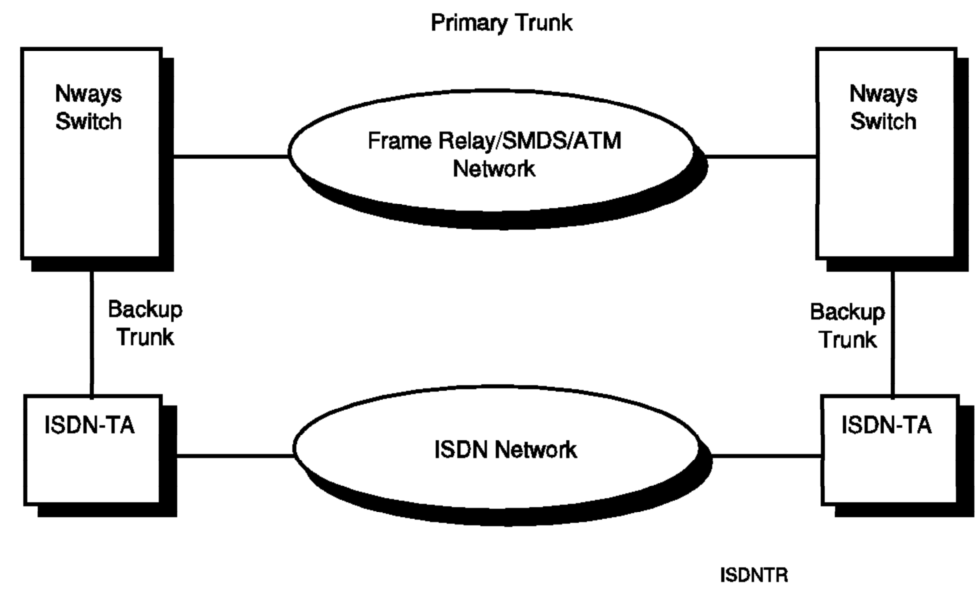

Figure 34. IBM Nways ISDN Backup Trunk

You can back up any type of IBM Nways Switch trunk (Direct Trunks, SMDS, ATM interworking, frame relay, OPIimum trunks). However, ATM virtual circuits are not supported due to the low speed of the backup trunk (typically less than E1).

When the backup trunk is activated, OSPF and the load balancing algorithm are used to reroute PVCs from their primary path. When more than one path exists, PVCs could be routed over a different path, not using the backup trunk services. In case of primary trunk restoration, backup trunk is automatically brought down. PVCs do not have to be rerouted over the original primary path if an alternate one exists.

Total bandwidth of the backup trunk could be less than the primary trunk; however, this does not have to be a problem, since you can provide users with additional bandwidth using multiple backup trunks for single primary trunk. You can define up to eight backup trunks per primary trunk; however, at least one dedicated backup trunk is required to provide fault-tolerant service for one primary trunk. The rerouting of the calls takes place even if the backup trunk does not provide enough bandwidth. In such a case, congestion may occur on the backup trunk, as the sum of the CIRs could be greater than the total available bandwidth. The following steps are taken when the primary trunk fails:

- OSPF and the load balancing algorithm redistribute affected PVCs using existing alternate paths, if they exist.
- Backup trunk is established and PVCs are restored, if there was no alternate path available.
- OSPF redistributes PVCs over the backup and alternate paths using the load balancing algorithm.

The trunk backup procedure is initialized upon receiving the `trunk down` message from the link level layer. To avoid rerouting oscillation, IBM Nways Switch monitors the trunk status for a while. If the link-up message is not received during this time, the backup procedure is initialized. The amount of time is configurable from the management station. A similar procedure is used when the primary trunk changes its status to link up.

The backup trunk is initialized using the physical interface signals between the IBM Nways Switch and the ISDN-TA. You have to configure one side to be originating, while the other one is answer only. The following signalling is used between the switch port and ISDN-TA:

- **Originate side** - IBM Nways Switch raises the DTR and RTS pins on its physical interface.
- **Answer side** - IBM Nways Switch monitors the ER and RS pins.
- **Call completion ISDN-TA side** - ISDN-TA is expected to raise the CS, DR and CD control signals.
- **Call completion switch side** - IBM switch monitors for activity of the CTS, DSR and DCD control signals.

If the ISDN call setup is not completed in five seconds, the IBM Nways switch enters the call setup retry cycle. After the time period called retry interval, IBM Nways Switch tries to establish the backup trunk again; when the maximum number of retries is reached, then IBM switch waits for a period of time called cycle interval. After cycle interval it enters the setup retry cycle again. All values (retry interval, retry number and retry cycle) can be defined from the management station.

2.7.2 ISDN Remote Access

IBM Nways 2225 is the first multiservice switch that provides remote access and switching functions in one chassis. Users are able to access the SMDS, frame relay and ATM services using the ISDN public network. Users can choose ISDN access either as a backup of the primary connection to the multiservice network, or as a low cost alternative to the leased line connection. The 2225 connects to the ISDN network via PRI interface.

Typical applications could be telecommuting, backup of primary connection or bandwidth on demand. Customer CPE is connected to the ISDN network using either ISDN-TA adapter or BRI interface. When needed, it requests the connection over the ISDN network to IBM Nways Switch, which provides the connectivity to the computer center using the frame relay, SMDS or ATM services. Figure 35 on page 49 illustrates usage of the ISDN access to the IBM Nways multiservice network.

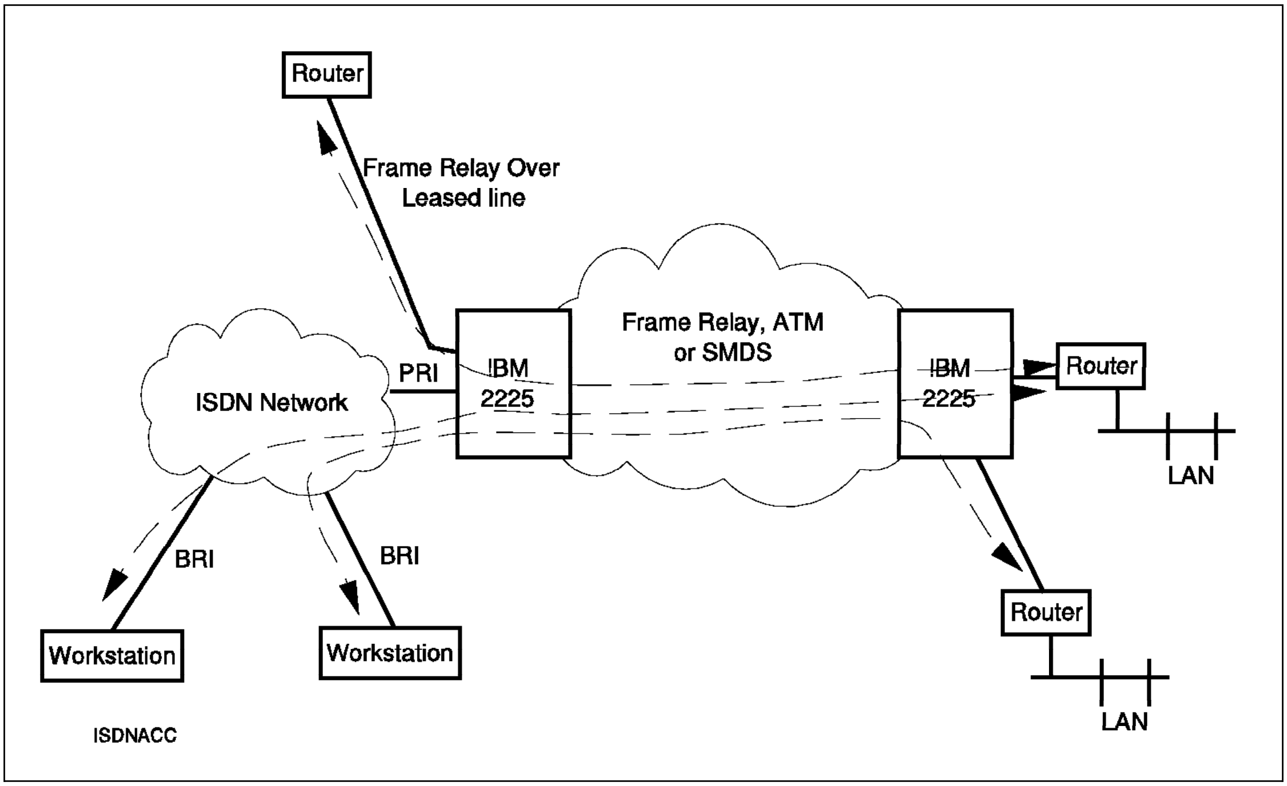

Figure 35. ISDN Access to the IBM Nways Switch Network

2.7.2.1 ISDN Logical Ports

The IBM Nways 2225 uses the logical port symmetrical architecture to associate the type of service with DS0 channels (timeslots). DS0 channels carrying the user data are referred as bearer channels (B-channels). You are allowed to configure for each B-channel one logical port with then associated type of service. Multiple types of service can be used per physical port; however, one B-channel can support only one protocol at a time. IBM Nways 2225 *does not* support different types of access service per B-channel via calling user identification. For available services, refer to 2.7.1.2, "ISDN Access Services" on page 46. When using standard D-channel signalling, you can define up to 23 logical ports (T1 version) or 30 logical ports (E1 version). Super D-channel allows you to define up to 95 B-channels per T1 I/O module or 119 logical ports per E1 module.

Each logical port has its associated *E.164 called address.* This number is used by the calling user, in order to access multiservice switch network over ISDN. You can assign a unique E.164 address to each B-channel or you can use one E.164 for more logical ports. In this case, you create a hunt group. For more details about hunt groups, please refer to the following section.

IBM Nways 2225 uses two methods for establishing calls between the switch and ISDN network:

- **Exclusive call establishment** - This mode of operation allows the ISDN public network to assign the B-channel when the request for call is received.
- **Preferred call establishment** - The IBM Nways 2225 Switch decides which B-channel will be used to establish the call.

2.7.2.2 B-Channel Hunt Groups

A hunt group consists of a group of B-channels, which can be accessed using the same telephone number (E.164 address). When you assign the same E.164 address to more than one B-channel, these are considered to be members of the same hunt group. Standard D-channel allows you to define 23 for T1 and 30 for E1 B-channels in one hunt group. When using the super D-channel, the maximum number of ports in a hunt group is equal to the number of time slots on the I/O module minus one.

When using the hunt group option, all of the B-channels have to be used for the same type of service (frame relay, SMDS, etc.), and when using a connection-oriented protocol, they have to provide connectivity to the same logical destination port. You define a pool of PVCs to the destination, when the device calls up the network; IBM Nways Switch automatically assigns and activates one PVC per destination. Figure 36 illustrates the hunt group port usage for the frame relay protocol:

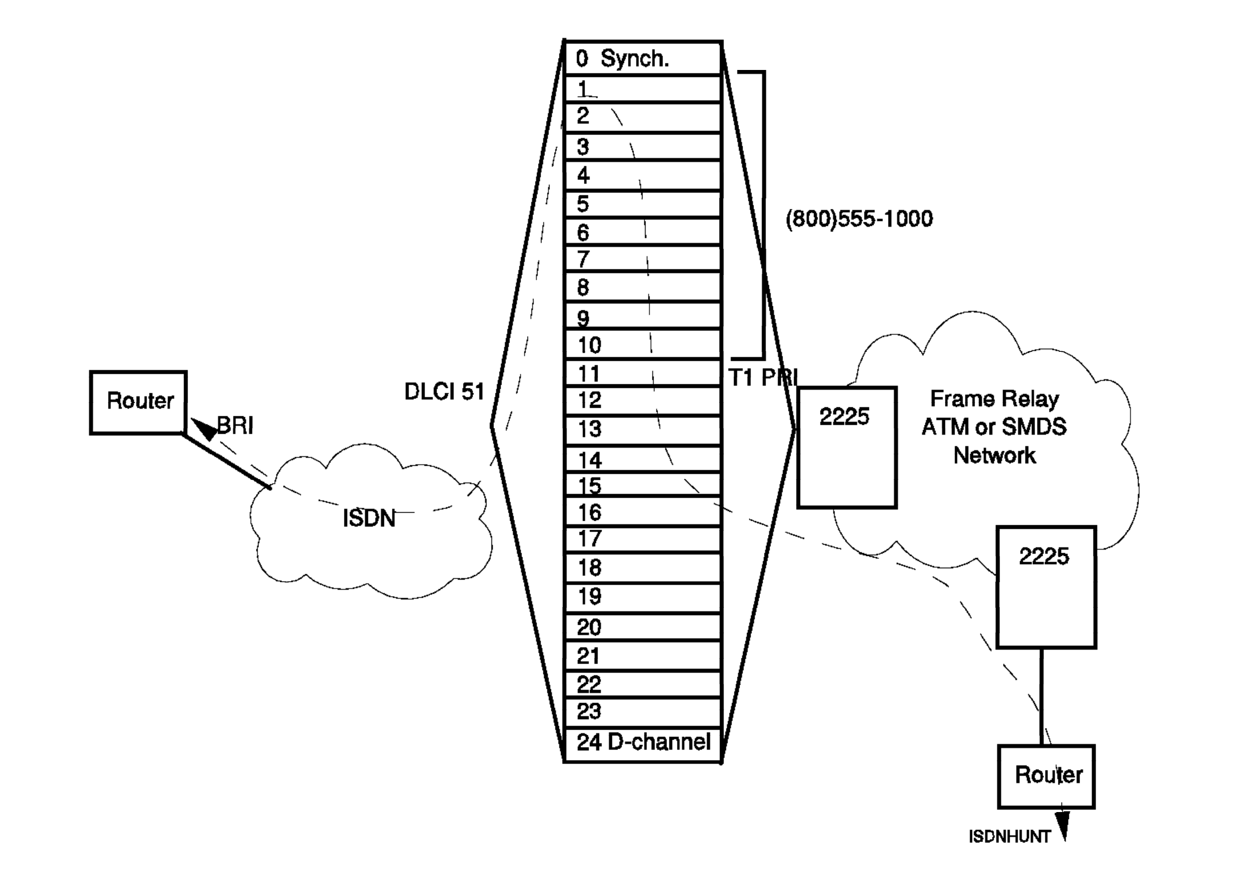

Figure 36. Hunt Group Access to Frame Relay Network

Table 3. Definition of the Hunt Groups Inside IBM Nways 2225

B-Channel	Physical port num.	Phone number	DLCI	Egress Point
1-10	1	(800) 555-1000	51-60	Port A
11-23	1	(800) 555-1001	101-113	Port B
1-23	2	(800) 555-1002	201-223	Port c

When the user router calls the telephone number (800) 555-1000, the ISDN public network will establish the connection to one of the free B-channels on the physical port 1. A free DLCI on the egress port A will be activated by IBM Nways 2225 Switch to provide the requested service to the calling user.

2.7.2.3 Caller ID Screening

This feature provides security by ensuring that only allowed users can access multiservice network over the ISDN public network. You can define up to 15 E.164 *calling addresses* per logical port. The calling address is the telephone number of the authorized calling user (user ISDN address). Therefore, end stations cannot access the multiservice network from unauthorized numbers; the call is denied by the IBM 2225 Switch. When you are using the hunt group option on the ISDN card, then you can define maximum (10*number of the B-channels in hunt group) calling addresses per hunt group.

2.7.2.4 Dynamic IP Address Assignment

This feature is very useful when the dial-in access device uses the PPP protocol and requires a dynamic IP address assignment. There are many cases where the PPP device does not have an assigned fixed IP address, and they request one during the establishment of connection. IBM Nways 2225 can provide these devices with the IP address during the call setup. You are able to define for each physical port a continuous block of pool addresses. When you are using the standard D-channel signalling, you define for each physical interface a block of pool addresses with the address range x.y.z.n to x.y.z.n+22; when using the super D-channel, the address range is for each physical interface (except the one with the signalling channel) x.y.z.n to x.y.z.n+23. IBM Nways 2225 does not provide the end station with a subnet mask; this has to be configured manually in each dial-in access workstation.

A router connected to the IBM multiservice network has to know which IP address was assigned to the dial-in user station. IBM Nways 2225 supports a function called Inverse ARP (I-ARP). When the router requests the IP address for the activated PVC using the I-ARP, IBM Nways 2225 replies with the IP address, which has been assigned to the dial-in user station.

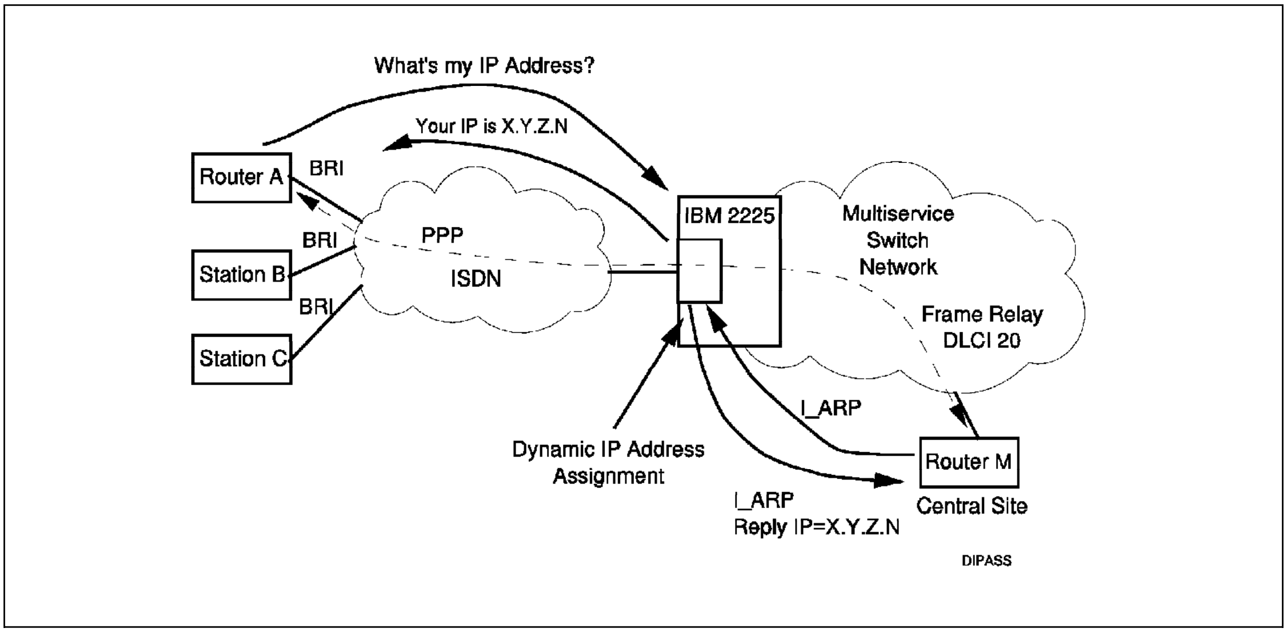

Figure 37. Device IP Address Assignment

Figure 37 shows the usual configuration of the dial-in access to the IBM multiservice switch network. Router A sets up the call using the ISDN network to access the multiservice network; it asks for the IP address assigned to it by the ingress 2225 switch. At the same time the DLCI to router M located in the central side is activated. Router M issues an I_ARP request for the remote IP address associated with the activated DLCI. IBM Nways 2225 replies with the same IP address as the one that was sent to router A.

2.8 IBM Nways 2230 ATM Overview

IBM Nways 2230 has been designed as a high-speed ATM backbone switch for public and private networks. Currently, IBM Nways 2230 provides you with pure ATM switching capabilities with high-density or high-speed ports. IBM Nways 2230 is a third-generation switch with distributed intelligence, high throughput and strong implementation of ATM QOS. IBM is implementing all of the latest ATM standards which, in combination with unique traffic management and distributed routing functions, provide next generation switching capabilities. For more details about the hardware architecture, please refer to 3.3, "IBM Nways 2230 ATM Switch Models 600 and 650" on page 96.

2.8.1.1 IBM Nways 2230 Logical Ports

IBM Nways 2230 implements logical ports architecture in a way similar to the IBM Nways 2225 and IBM Nways 2219 switches. Logical ports are associated with the physical ports. In this way, you can define the services that will be provided by the associated physical port. You can mix different types of services on one module. However, only one type of service is available per logical port.

2.8.1.2 IBM Nways 2230 Supported Class of Services

IBM Nways 2230 supports all classes of service defined by the ATM Forum standard. Implementation is done in the hardware as well as in the software architecture of the product. The following list briefly describes each of the supported classes of service:

Constant Bit Rate (CBR) - This class is designed to handle all traffic represented as a synchronous continuous bit stream such as uncompressed digitized voice, video stream or any other proprietary synchronous protocol. CBR traffic is delay-sensitive and requires guaranteed throughput rate.

Variable Bit Rate (VBR) Real Time - VBR is used for transmitting special types of delay-sensitive applications, which require low cell delay variation between the end points. Such an application could be a packetized video stream using a special type of video codec equipment. IBM Nways 2230 supports four different subclasses of the VBR-RT.

Variable Bit Rate (VBR) Non-Real Time - Handles packaging of long, bursty data streams. This type of service could be used for frame relay connections over the ATM network or similar bursty non-real-time traffic.

Unspecified Bit Rate (UBR) - This class was designed to handle primarily the LAN type traffic. User devices have to be able to accept any delay or loss of cells during the transmission over the network. IBM Nways 2230 handles the ABR traffic in the same way as UBR.

2.8.2 Supported ATM Interfaces

IBM Nways 2230 supports the following types of interfaces:

- ATM User Network Interface (UNI) 3.0 DTE and DCE logical ports for T3, E3, and OC3c physical interfaces.
- ATM User Network Interface (UNI) 3.1 DTE and DCE logical ports for T3, E3, and OC3c physical interfaces
- ATM Interim Inter-Switch Protocol (IISP)/Phase 0 DTE and DCE side for T3, E3, and OC3c interfaces
- ATM Direct Trunk
- ATM OPTimum Trunk

2.8.2.1 IBM Nways 2230 User Network Interface (UNI) 3.0 and UNI 3.1 DCE

This type of the service provides connectivity for ATM CPEs (user devices such as routers and workstations). IBM has implemented both interface types (UNI 3.0 and UNI 3.1) fully compliant with the ATM standard, both SVC and PVC connections are supported. You are able to mix the UNI 3.0 and UNI 3.1 interfaces on the switch. The IBM Nways 2230 Switch supports conversion between these two standards; the switch is able to set up the SVC connection between a CPE UNI 3.0 and UNI 3.1. The same function is implemented on the IBM 8260 and IBM 8285 switches. IBM Nways 2230 implements the following signalling and management protocols on the UNI interface:

- ILMI - Management information protocol based on SNMP provides status and configuration information. It is used for keep alive protocol to verify if the CPE device is active.
- Q.93B signalling - UNI 3.0 signalling protocol used for creation of the ATM SVCs over the ATM network.
- Q.2931 signalling - UNI 3.1 signalling protocol use for creation of of the ATM SVCs over the ATM network.

- OAM cells support - You can enable logical port to generate and process operation, administration and management cells.

A logical port configured as UNI DCE performs the following ILMI functions:

- Logical port status determination using polling. If the CPE does not respond within a given time period, the port is considered down.
- Send node and port prefixes to the end station, while accepting end station address ESI (End System Identifier).

Figure 38 illustrates usage of the ATM UNI DCE interface.

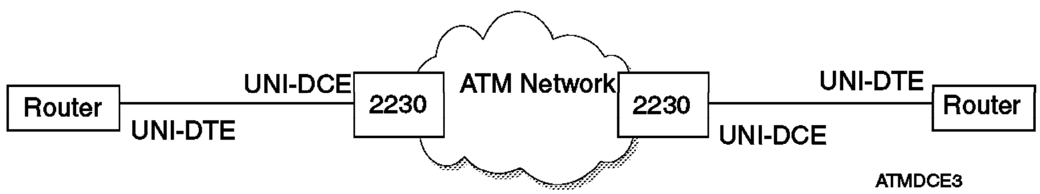

Figure 38. IBM 2230 ATM UNI DCE Interface

2.8.2.2 IBM Nways 2230 User Network Interface (UNI) 3.0 and UNI 3.1 DTE

A logical port configured as a UNI DTE service enables the IBM Nways 2230 to communicate with an ATM switch using ATM PVCs and SVCs. In this case, the IBM Nways 2230 works as access switch feeding multiple ATM PVCs and SVCs via the ATM UNI DTE logical port to the connected ATM network. UNI DTE logical port is mainly used for connection of the IBM Nways 2230 Switch to another ATM network when OPTimum trunks are used. Configuration allows you to choose whether to support standard signalling protocols (ILMI and Q.93B).

A logical port providing the UNI DTE service performs the following ILMI functions:

- Respond to the polling requests from the UNI DCE interface. If no polling request is received over a certain time period, the interface is considered to be down.
- During the address registration, UNI DTE accepts the network prefix and optionally sends its ESI identifier.

Figure 39 on page 55 illustrates usage of the ATM UNI DTE interface.

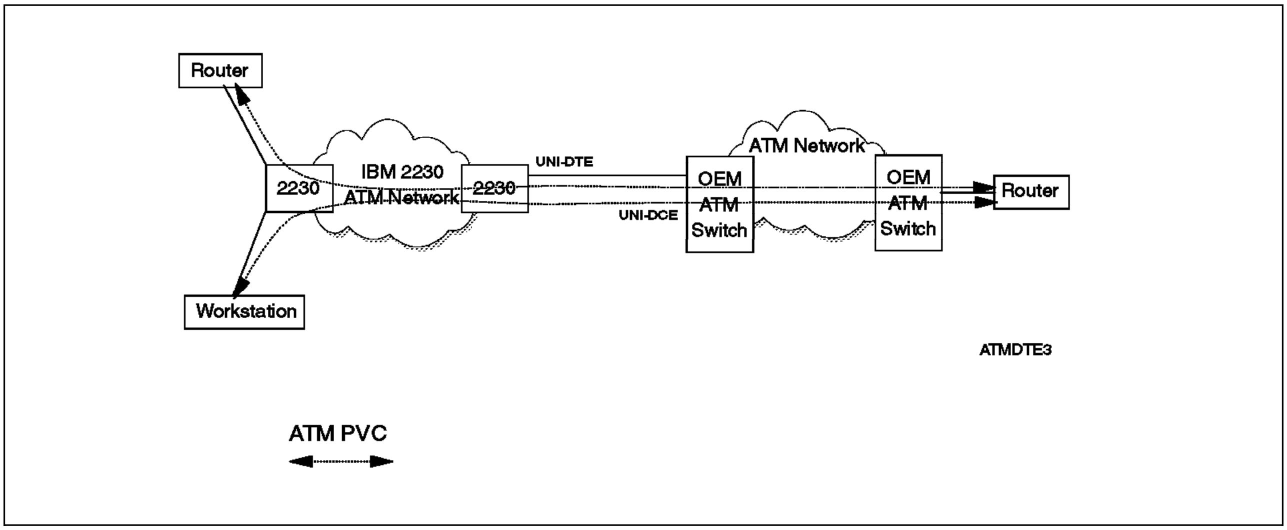

Figure 39. IBM 2230 ATM UNI DTE Interface

2.8.2.3 IBM Nways 2230 IISP Interface DCE and DTE

IBM Nways 2230 fully supports the ATM Forum Interim Interswitch Protocol (IISP). You can configure either the DCE or DTE side of the line. This logical port type enables interconnection with ATM networks based on other vendors′ switches. IBM Nways 2230 supports static routing for SVCs. The IISP logical port performs the ILMI functions; The DCE interface polls the peer node and waits for responses, while the DTE interface looks for poll requests and replies to them.

You are allowed to tune the parameters for the IISP logical port. For more details about the configuration of Maximum Valid VP/VC bits, UPC, CAC, ILMI, Q.93B and OAM support on the IISP interface, refer to 2.8.2.5, "IBM Nways 2230 UNI DTE and DCE Options" on page 56. IBM Nways 2230 implementation of administration costs for the IISP logical ports enables you to balance the load over the different IISP ports. Adjustment of the QOS parameters enables you to adjust the bandwidth and routing metrics for various classes of service. For more details, refer to 2.8.4, "Logical Ports Quality of Services Parameters" on page 60. An example of usage of the ATM IISP logical port is shown in Figure 40.

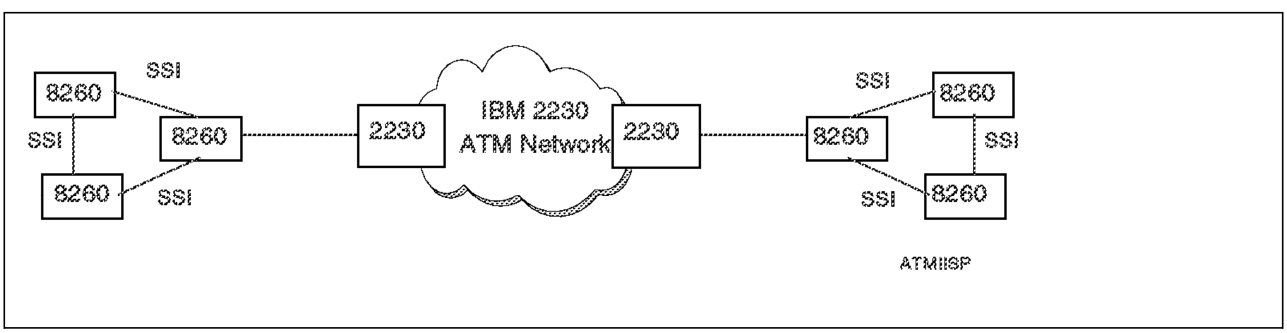

Figure 40. IBM 2230 ATM IISP Interface

2.8.2.4 IBM Nways 2230 OPTimum Trunk

This feature enables peer IBM Nways 2230 Switches to communicate over an ATM network, while remaining in the same OSPF routing domain. In this case, a permanent virtual path has to exist between the IBM Nways 2230 Switches.

Before you configure the logical port to support OPTimum trunking, you have to define an ATM UNI DTE logical port. You can configure more than one OPTimum trunk per logical port. The following figure illustrates the usage of the ATM OPTimum trunk.

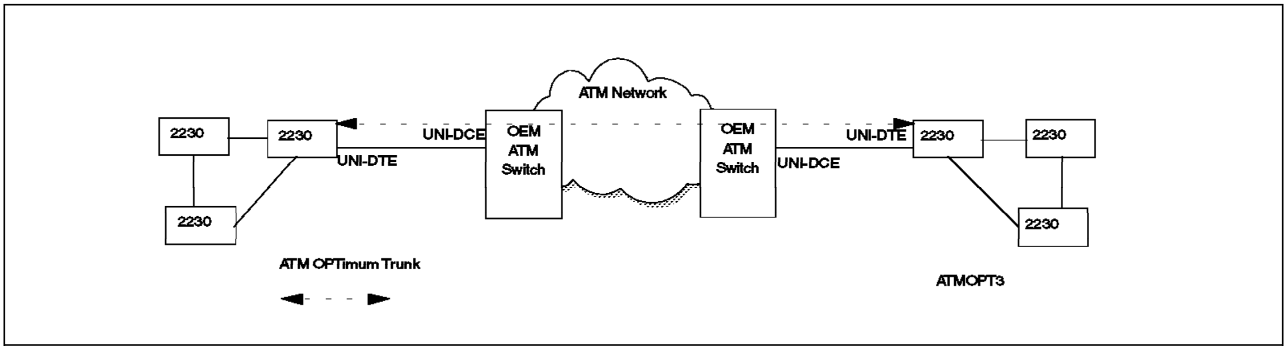

Figure 41. IBM Nways 2230 OPTimum Trunk Usage

2.8.2.5 IBM Nways 2230 UNI DTE and DCE Options

You can specify many options during the configuration of the UNI and IISP logical ports. These features increase flexibility when connecting to devices that do not fully support ATM Forum standards. The following parameters are configurable:

Number of Valid Bits in VPI/VCI - You can configure the number of VP and VC bits used in the ATM cell header. The interface can support either more virtual paths with less virtual circuits or vice versa. IBM Nways 2230 supports VPI fields from 0 to 8 bits and VCI fields from 5 to 15 bits; the sum of VCI and VPI cannot exceed 14 bits, as defined by the ATM Forum. This feature is useful in the following cases:

- To match the type of connected equipment settings.
- When you are configuring OPTimum trunk over the DTE logical port, you can either choose more OPTimum trunks with less VCIs supported, or less OPTimum trunks per logical port with more VCIs on it. To determine how many trunks and VCIs will be supported, use the following formula:

  Maximum OPTimum trunks per logical port = 2 power V - 1

  Maximum virtual channels per trunk = 2 power 32 - C

  where V is the number of VP bits and C is the number of VC bits used inside the ATM cell header.

ILMI Support - ILMI support could be disabled or assigned to specific virtual VP/VC channels (standard defines ILMI over the virtual channel with VP=0, VC=16). Polling interval and the loss threshold can be adjusted according to the requirements.

Q.93B/Q.2931 Support - Q.93B/Q.2931 protocol could be either disabled or enabled. When you disable the Q.93B/Q.2931 protocol, you can use that port only for the PVC support. You have the ability to set up the timers and threshold for the Q.93B/Q.2931 and Q.SAAL protocols.

UPC Function - IBM Nways 2230 has implemented the Usage Parameter Control (UPC) based on the leaky bucket algorithm. You are able to either enable or disable this function for the whole logical port. However, IBM implementation allows you to disable the UPC function per virtual circuit; in such a case, UPC has to be enabled for the logical port. In a disabled state, all traffic passes the port, including the nonconforming cells. As soon as you disable the UPC function on one of the ports, IBM Nways 2230 does not guarantee quality of service in the network.

DTE Prefix Screen Mode - This option is valid only for the DTE logical port and enables you to specify which networks can be connected to that interface by limiting the group of network prefixes accepted by DTE.

OAM Support - You can disable generation of the OAM alarms per physical port.

Call Admission Control (CAC) - IBM Nways 2230 performs the Call Admission Control function to determine whether the network has enough resources to fill the requirements for the requested QOS. If you disable that function, then the ingress port attempts to build the circuit, even if there are not enough resources to do so (applies only for VBR non-real-time and UBR traffic).

2.8.3 IBM 2230 Routing Implementation

An ATM system is a connection-oriented system. User data is sent over pre-established VCs. The VCs are established between two endpoints either on administrator request or on the end station request. The ATM network decides which path should be used and establishes a VC. All user data is transported over the same path unless VC is rerouted over a different path. Path selection in ATM networks uses a complex algorithm. The selected path has to meet all QOS requirements while the efficient usage of the network bandwidth has to be ensured. The IBM Nways Switch path selection algorithm consists of the following components:

- IBM Nways 2230 ATM routing - Virtual Network Navigator
- ATM link metrics
- Call Admission Control
- Reroute tuning

2.8.3.1 IBM Nways 2230 ATM Routing - Virtual Network Navigator

Determining the path inside the ATM network while providing the requested QOS is one of the most problematic issues of today's ATM network. The routing algorithm employed in the ATM network has to handle not only the physical topology problems, but also solve all issues related to providing the QOS for different traffic types (data, voice, video or circuit emulation) as well as effective load balancing of the traffic.

IBM has implemented a routing algorithm called Virtual Network Navigator which is based on the link state algorithm OSPF. IBM Nways 2230 re-uses the OSPF implementation for the IBM Nways 2219 and IBM Nways 2225. Some additional features were added to cover the specific issues of the ATM technology. Each switch builds its own topology database on the network based on the link state updates received from its neighbors. The powerful distributed routing intelligence allows you to build large, scalable networks. Each switch is able to determine the best path for the requested VC across the network while guaranteeing requested QOS.

IBM Nways 2230 Virtual Network Navigator algorithm uses a very similar approach to that which will be used in the ATM Forum PNNI standard. Due to the complexity of routing in the ATM network, the ATM Forum has not yet completed the PNNI standard, and it seems to be a long-term effort.

As described earlier, the ATM network has to provide different QOS types for different types of applications. IBM Nways 2230 has a strong implementation of QOS based on dedicated buffer switch resources for each QOS (see 2.8.4, “Logical Ports Quality of Services Parameters” on page 60). You are able to allocate resources in the network for each type of QOS; dynamic, static and static/dynamic resource sharing is available. To fully gain advantage from that approach, IBM implementation of the VNN algorithm builds separate topology maps for each class of service as shown in Figure 42.

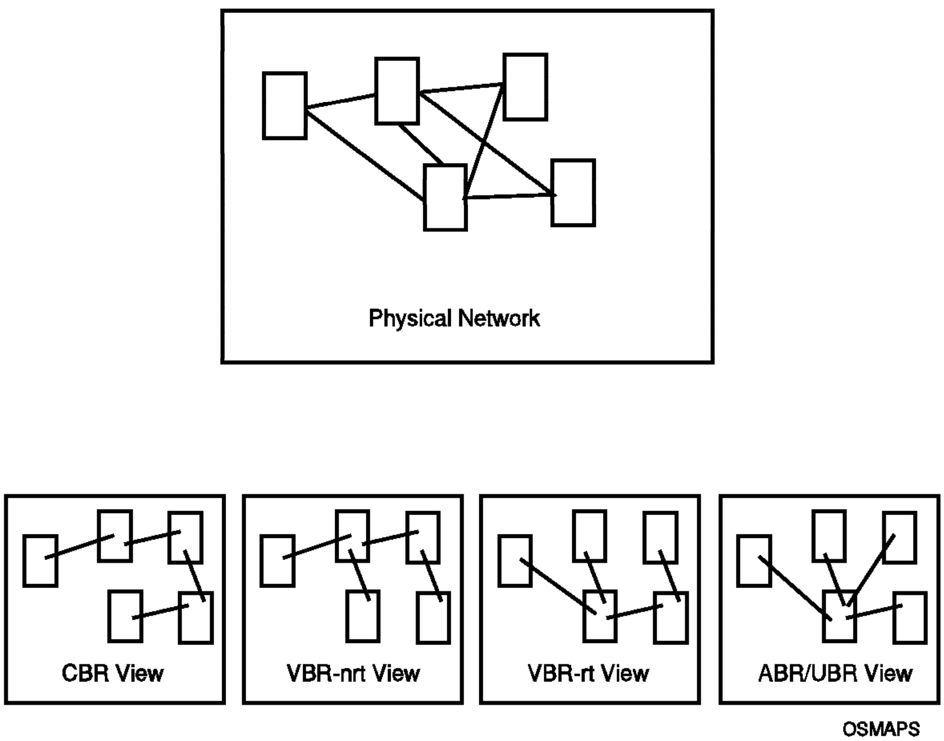

Figure 42. VNN Routing Topology for Different Classes of Service

The VNN algorithm is implemented in the IBM Nways 2225 and IBM Nways 2230 products, providing the best path determination using all packet-based, cell-based and OPTimum trunks. The algorithm interoperates with the IISP to enable interconnection of the IBM Nways 2230 with ATM networks based on OEM switches. VNN path determination ensures provision of the right QOS (bandwidth, end-to-end delay, cell delay variation and cell loss ratio) for each requested connection. Additionally, VNN includes policy enhancement, which allows users to select or avoid traffic over selected trunks, create virtual private networks and dedicate trunks only for traffic control. VNN consists of the following basic components:

- Topology database including the physical topology of the network and QOS capabilities for ATM trunks

- Topology database distribution algorithm based on OSPF
- Best-path determination for different QOS

VNN cooperates with the Virtual Circuit Manager (VCM), which runs in each node in the network. VCM is responsible for the setup of each circuit from the input to the output port in the local switch. VCM also performs the circuit acceptance decision based on the examination of the current resources of the local node. This means that each node in the network independently verifies the validity of QOS for the requested circuit, before accepting it. If the ingress node has made its decision based on the topology database, which was changed in the meantime, verification in each node ensures that the chosen path is valid. If the local node does not have enough resources to build the virtual circuit, the switch rejects the call and the original node recalculates a new path over the network.

2.8.3.2 ATM Link Metrics

IBM implemented enhancements to the OSPF routing algorithm to be able to cover the QOS specific details. New metrics are used to calculate the best path in the network. To improve performance of the whole algorithm, the best-route calculation is implemented by each I/O module. This distributed approach greatly improves performance during circuit initialization. The administrator of the network is able to predefine reserved resources for each QOS. This implementation together with the switching architecture guarantees end-to-end QOS for ATM connections over the IBM Nways 2230 network. For more details about the QOS reservation on the trunks, refer to 2.8.4, "Logical Ports Quality of Services Parameters" on page 60. The IBM Nways 2230 link metrics have the following characteristics:

- Available bandwidth per QOS
- Delay: IBM Nways 2230 measures propagation and transmission delays on the trunk. However, delay caused by congestion and queueing of cells is not included.
- Trunk capability indicates if the trunk is cell or frame-based.
- Version identifier provides the information about the version of the routing software currently executing in the switch. This is used for backward compatibility.

In addition to described metrics, IBM Nways 2230 Switch employs policy-based routing, which enables you to make additional tuning of the best-path decision algorithm. The following options are implemented:

- **Administrative path control** - You can assign the administrative cost to each trunk in the network. Trunks with a higher cost will be used less than trunks with a lower one. If the VNN finds two paths, which are providing sufficient resources to satisfy requested QOS, then the one with less administrative cost is chosen.
- **Management trunks** - You are able to define trunk only for management purposes. IBM Nways 2230 will not establish any data circuit over it.
- **Virtual Private Network trunk** - Some trunks can be dedicated only to particular users of the network, allowing you to build ATM virtual private networks.
- **Virtual Path Capability** - You can define whether the trunk can support VP switching.

2.8.3.3 IBM Nways 2230 Call Admission Control

Before the route decision can take place, the requested parameters for the virtual circuit have to be translated into metrics used by the VNN algorithm. This part is implemented in the CAC algorithm using the equivalent bandwidth. Requested QOS (for example, PCR, CDV, CLR, etc.) are converted to the equivalent bandwidth, which is compared with the unreserved bandwidth on available paths. IBM Nways 2230 implements both an enhanced equivalent bandwidth CAC algorithm and tunable oversubscription factor algorithm. CAC is a part of the VNN implementation.

IBM Nways 2230 tuning of the CAC algorithm enables you to achieve a desired cell loss ratio across all physical ports in the network. You can either tune the IBM Nways 2230 CAC algorithm, or you can configure your own CAC. Tuned or configured CAC is used only for the VBR-RT and VBR-NRT traffic types.

Note: If you tune or configure CAC, the IBM Nways 2230 network will not guarantee QOS in the network.

Tuning CAC

You are able to specify the allowed cell loss ratio objectives, which should be met across the network for all VBR-RT and VBR-NRT traffic. Cell loss ratio could be configured in the range from 10E-1 to 10E-15. IBM Nways 2230 uses these values to calculate the appropriate equivalent capacity for the requested connections.

Configuring CAC

IBM Nways 2230 enables you to configure the amount of bandwidth, which is allocated based on the requested QOS for each VBR-RT and VBR-NRT circuit. You can configure the CAC equivalent capacity calculation based on the physical port type used for a particular connection and requested SCR.

2.8.3.4 IBM Nways 2230 Reroute Tuning

IBM Nways 2230 implements load balancing for all selected PVCs in the network. Switches reroute the circuits over the paths, which provide more bandwidth than the current one. You are able to adjust the number of PVCs in the reroute batch request as well as the time between the batches to optimize load balancing in the network. The following parameters are available:

Reroute Count - Number of PVCs included in one reroute batch.

Reroute Delay - The amount of time between the reroute batches.

The load balancing algorithm per VC can be disabled.

2.8.4 Logical Ports Quality of Services Parameters

You can specify the QOS parameters for each port. These parameters represent metrics for the IBM VNN routing algorithm and are used by the Virtual Circuit Manager during call establishment.

IBM Nways 2230 allows you to adjust reserved bandwidth per QOS. This means that you can define bandwidth allocation for all CBR, VBR-RT, VBR-NRT and UBR circuits; in other words, you are able to divide bandwidth of the logical port between the four supported classes of services. IBM Nways 2230 enables you to use either dynamic or fixed bandwidth allocation.

Dynamic - In this case, IBM Nways 2230 assigns the bandwidth dynamically to a selected class. Dynamic allocation pools the bandwidth from the pool of the remaining bandwidths for that particular local port.

Fixed (Static) - You can specify the percentage of bandwidth that will be used only for the type of traffic. If you configure fixed bandwidth allocation for all classes, be sure that the total value is 100%, otherwise the rest of the unspecified bandwidth will not be accessible by user circuits.

2.8.5 IBM Nways 2230 Oversubscription

IBM Nways 2230 enables you to configure the oversubscription factor for VBR-RT, VBR-NRT and UBR classes of service. You can optimize the number of PVCs and SVCs which can pass a logical port by specifying the oversubscription factor. CAC algorithm calculates the equivalent bandwidth for each circuit. Virtual Network Navigator together with Virtual Circuit Manager place the call over the path with enough available bandwidth for the particular classes of service. (You can statically divide bandwidth between different class of service; for more information, see 2.8.4, "Logical Ports Quality of Services Parameters" on page 60.) The circuit is placed over the port only if the port has enough available virtual bandwidth for the particular class of service to do so.

```
Available Virtual Bandwidth = Configured Bandwidth * K - Total Used Bandwidth
```

where:

`Configured Bandwidth` is either fixed configured bandwidth for the particular class of service or available logical port bandwidth for all classes, which are using dynamic allocation.

`K` is the oversubscription factor used for the particular class of service.

`Total Used Bandwidth` is the sum of all equivalent bandwidths of the circuits in the same class of service.

When you configure the oversubscription factor higher than K=1, you can put more PVCs and/or SVCs over the logical port, but IBM Nways 2230 no longer guarantees the QOSs.

In extreme circumstances (node or link failure inside the network), the Available Virtual Bandwidth could become negative. Existing PVCs can be rerouted over negative bandwidth trunks. However, new PVCs cannot be built over such an overloaded trunk.

2.8.6 IBM Nways 2230 Point-to-Multipoint Support

VNN algorithm implements point-to-multipoint support for PVCs and SVCs in a very efficient way. VNN understands not only the network physical topology, but it knows the hardware architecture of the switch as well. Cell replication is postponed for execution at the I/O card level. When two multicast leafs are connected to different physical ports on the same I/O module, the switch sends only one cell to the I/O module, which replicates the cell to the leaf ports. This implementation maximizes bandwidth savings.

Point-to-multipoint connections are set up in stages; each leaf is added in a separate step, while the established part of the tree is not torn down. VNN multipoint routing capability together with high-performance IBM Nways 2230 hardware implementation make IBM Nways 2230 suitable for video distribution and other multicast-based applications.

2.8.7 IBM Nways 2230 Permanent Virtual Circuit Support

IBM Nways 2230 supports point-to-point and point-to-mutipoint permanent virtual circuits. Both Virtual Channel Connections (VCCs) and Virtual Path Connections (VPCs) are supported. You can configure only the ingress and egress physical port and associated traffic descriptors, and the VNN algorithm will automatically determine the best path for a configured PVC. IBM has implemented a rerouting and load balancing algorithm, which ensures guaranteed QOS and efficient usage of network bandwidth.

IBM Nways 2230 enables you to configure different classes of service and traffic descriptors for each direction of the circuit. For more details, refer to 2.8.8, "IBM Nways 2230 ATM Traffic Descriptors."

IBM Nways 2230 determines the optimum path for the PVC using the VNN routing algorithm, or you can predefine the path manually during the configuration. When the predefined path fails, the VNN can optionally reroute the circuit over the alternate path, if available.

2.8.8 IBM Nways 2230 ATM Traffic Descriptors

A traffic descriptor is a part of the contract between the user and the network. It defines a combination of the parameters which are used for traffic control. A traffic descriptor is used by the UPC to determine how and which cells can enter the network. The ingress port can either discard or tag the nonconforming cells. IBM Nways 2230 supports all ATM Forum standard combinations of traffic parameters:

PCR CLP=0, PCR CLP=0+1

IBM Nways 2230 checks the traffic conformance using the peak cell rate for both streams (CLP=0 and CLP=0+1) with no tagging option. If the stream CLP=0 exceeds the PCR, the nonconfirming cells are discarded.

If the PCR of the CLP=0+1 stream exceeds the contract, the nonconfirming cells are dropped according to a ratio of CLP=0 to CLP=1 bandwidth reservation. If you configure bandwidth 8000 Kbps for the CLP=0 stream, and 5000 Kbps for CLP=1 stream, then IBM Nways 2230 will drop 8 CLP=0 cells for each 5 CLP=1 dropped cells. This traffic descriptor can be used for CBR and UBR traffic.

PCR CLP=0, PCR CLP=0+1, Tagging

IBM Nways 2230 checks the traffic conformance using the peak cell rate for both streams (CLP=0 and CLP=0+1) with the tagging option. If the stream CLP=0 exceeds the PCR, the nonconfirming cells are tagged with the CLP field set to one. If the PCR of the CLP=0+1 stream exceeds the contract, the nonconfirming cells are dropped according to a ratio of CLP=0 to CLP=1. You can use this option with UBR and CBR traffic.

PCR CLP=0+1

Traffic contract is based on the peak cell rate of the CLP=0+1 aggregate stream. If the input exceeds the contract, all nonconforming cells are discarded without differentiation between CLP=0 and CLP=1 cells. This option can be used with CBR and UBR class of service.

PCR CLP=0+1, Best Effort

Traffic contract is based on the peak cell rate of the CLP=0+1 aggregate stream. When the input streams exceed the contract, nonconforming cells are sent over the network on the best-effort basis. You can use this option only for the UBR class of service.

PCR CLP=0+1, SCR CLP=0, MBS CLB=0

This option can be selected for the VBR class of service. Contract conformance is based on the peak cell rate, sustained cell rate and maximum burst size of the CLP=0 stream, no tagging is performed, and nonconforming cells are discarded.

PCR CLP=0+1, SCR CLP=0, MBS CLB=0, Tagging

You can use this option for VBR class of service. Traffic conformance is based on the peak cell rate, sustained cell rate and maximum burst size of the CLP=0 stream. Nonconforming cells are tagged with CLP=1.

PCR CLP=0+1, SCR CLP=0+1, MBS CLB=0+1

This option can be selected for the VBR class of service. Contract conformance is based on the peak cell rate, sustained cell rate and maximum burst size of the CLP=0+1 stream. No tagging is performed, and nonconforming cells are discarded.

2.8.9 IBM Nways 2230 ATM SVC Support

ATM switches should support switched virtual circuits (SVCs) to be able to support large ATM networks. When providing any-to-any connections, the number of VCs grow according to the following formula:

```
X=N*(N-1)/2
```

where X is the number of VCs and N is the number of ATM end stations.

For example, if the network is to provide any-to-any VC connection for 100 end stations, the number of SVCs which would be defined by the network administrator is 100*99/2=4450.

Some of the PVCs could transport data only for a short time period; however, they could allocate network resources even when they are not used. SVCs are built only for the time the connection is needed.

IBM Nways 2230 Switch SVCs implementation consists of the following components:

- Different type of ATM addressing support
- Finding a path between the source and destination ATM address using the IBM VNN algorithm
- Address screening
- Calling address insertion
- Address translated and tunneling

2.8.9.1 IBM Nways 2230 SVC Support

IBM Nways 2230 has implemented strong support for switched virtual circuits (SVC). SVC support is a very important feature as soon as the ATM network starts to grow in terms of the number of connected users. IBM Nways 2230 supports both User Network Interfaces 3.0 and 3.1 according the standards specified by the ATM Forum. IBM Nways 2230 SVC implementation supports the following classes:

- Bearer Class A
- Bearer Class C

- Bearer Class X

Quality of service class parameters and QOS profile selection are supported as well. Strong support of the addressing features includes calling number validation, per-port mixing of different AESA and E.164 addresses, address translated, tunneling and many other features described in previous sections. SVCs support the same QOS features as the PVC implementation.

The total number of PVCs and SVCs supported per node is currently up to 126000 VCCs, with each I/O module supporting up to 8000 VCCs. However, the IBM Nways 2230 design will allow up to 16000 VCCs per I/O module to provide a total capacity of 224000 VCCs per switch.

An addressing scheme is used for the SVCs setup decision. Users communicate over connections, which can be either statically predefined by the network administrator or requested by the end station when needed. In this case, the end station requests the SVC connection using the SETUP message. The SETUP message contains the source and destination ATM address of the requested connection. The ingress node, which receives the call setup message, has to find out in its routing tables the path to the requested destination. IBM Nways 2230 implements the following addressing features:

- Different types of address support on same switch
- Tunneling and translating between different types of addresses
- Address route aggregation
- Calling party address insertion and replacement
- Address screening

2.8.9.2 Supported Address Formats

There are two different general types of address format that could be used to identify an ATM end station connected to the network. Each ATM user has to have at least one ATM unique address to be able to support switched virtual circuits (SVCs). IBM Nways 2230 supports the following address formats:

- ATM End System Address format (AESA)
- Native E.164 address format

The ATM Forum standard defines four AESA address formats. IBM Nways 2230 supports all types of AESA addresses with optional translation between different formats. For more details, refer to 2.8.9.6, “Address Translated and Tunneling” on page 67. The following formats are supported:

- Data Country Code (DCC)
- International Country Designator
- E.164
- Custom

All AESA addresses consist of 20 octets, which are divided into the following components:

Authority and Format Identifier (AFI) - This octet identifies the syntax of the rest of the address, as well as the authority which assigns the DCC, ICD or E.164 part of the AESA address. The following list of AFI is standardized:

- 0x39 for DCC
- 0x47 for ICD
- 0x45 for E.164

Initial Domain Identifier (IDI) - A code which identifies the subauthority that has allocated the address.

High-Order Domain-Specific Part (HO-DSP) - This field is defined by the authority, which is specified by AFI and IDI octets. An address should be constructed to enable efficient routing between different ATM networks.

End System Identifier (ESI) - Six octets identify the end system. This part is registered with the local ATM switch to enable the switch to differentiate between different connected end systems.

Selector (SEL) - This field is not used for identification of the end system, but could be used by the end system for some internal purposes.

Initial Domain Part (IDP) - Consists of the AFI and IDI fields which uniquely identify the authority responsible for allocating and assigning the domain specific part of the ATM AESA address.

Domain Specific Part (DSP) - Rest of the ATM AESA address, which consists of the HO-DSP, ESI and SEL fields.

The following figure shows the different types of the supported AESA address formats:

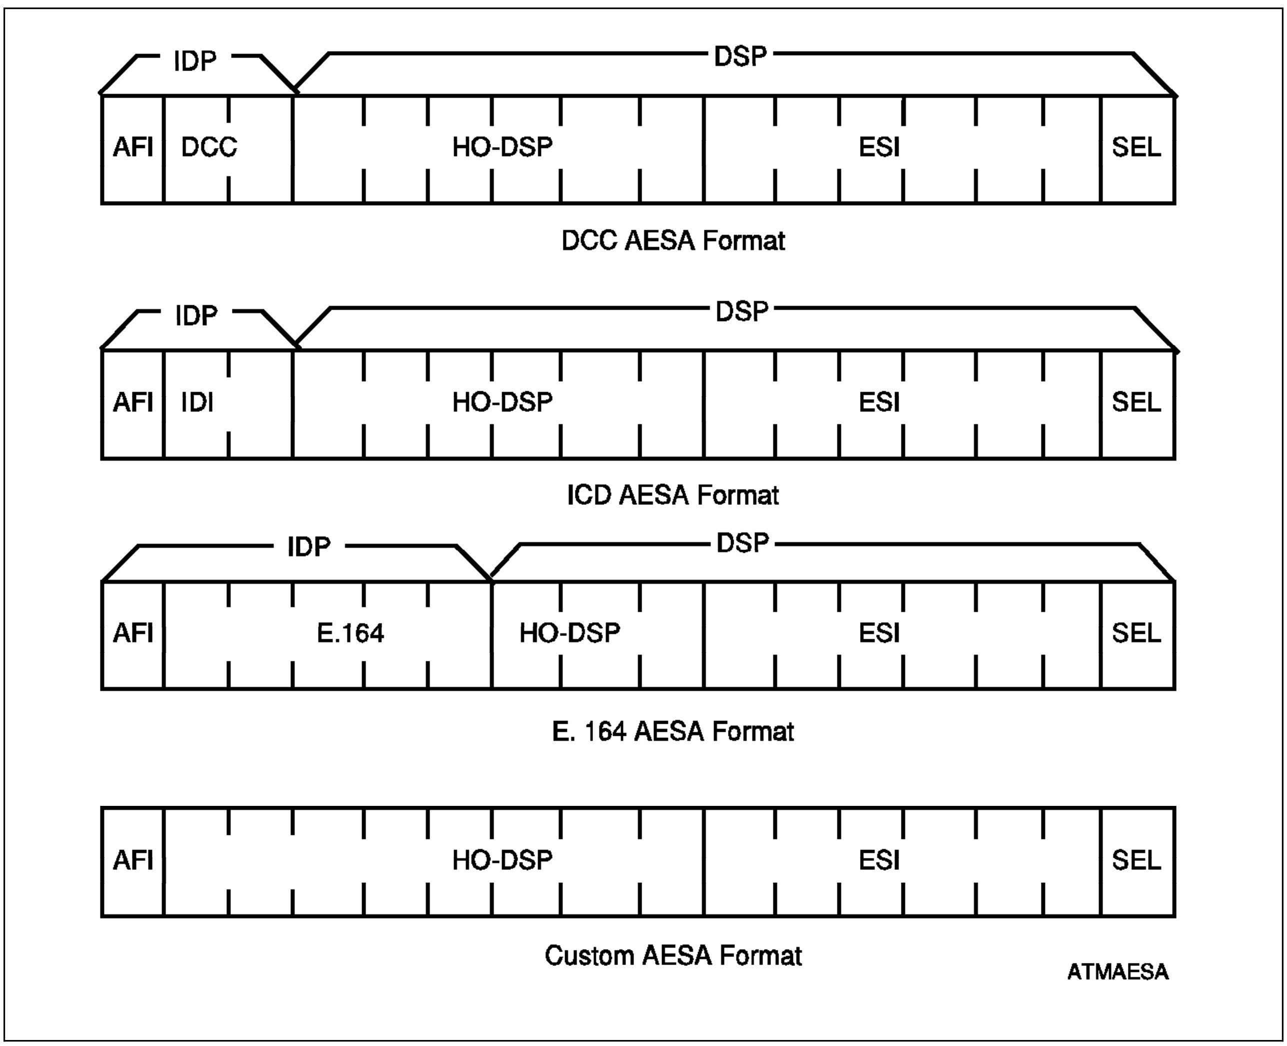

Figure 43. Address Formats

The last type of possible IBM Nways 2230 address format is the *Native E.164 Address Format*. Native E.164 addresses are represented as standard Integrated Service Digital Network (ISDN) numbers, the same as the one used for the telephone numbers. Native E.164 addresses consist of 1 to 15 BCD digits, which are placed in the 2 to 9 octets of the 20-octet AESA address. The rest of the AESA address octets are filled with zeros. The address is not divided into the different fields, such as AFI, IDI, HO-DPS, ESI and SEL.

IBM Nways 2230 implementation of the addressing and routing algorithm allows you to use more than one IDP/HO-DSP prefixes per switch. You are allowed to use different types of ATM addresses at the same time. IBM Nways 2230 routing algorithm enables you to divide the address into three portions: node prefix, port prefix and port address. This deep granularity of addressing space allows you to optimize the size of the routing tables, while providing great flexibility for the address allocations for ATM end stations.

Node prefix applies to all ports on the switch. It is used for the route aggregation and end system address construction during registration. You can configure multiple prefixes per node; however, you do not have to specify one if you have specified port prefixes or port addresses for all logical ports, which should support SVCs. You can define a node prefix to be a part of or a whole AESA or E.164 address. Node prefix does not have to be unique for a particular node. You can have several nodes with the same node prefix if you specify unique port prefixes for each node. In such a case, you have to increase granularity of the routing by propagating each port prefix, resulting in larger routing tables.

When the end station does not support the ILMI protocol to register its part of the ATM address, you can configure it manually. This feature enables SVC support for ATM end stations which do not implement the ILMI protocol.

If the logical port is configured for the UNI-DTE service, you can enter the ESI identifier, which is registered with the far end switch during the ILMI registration procedure. In this case, the IBM Nways 2230 registers the network prefix.

2.8.9.3 VNN Algorithm and E.164

2.8.3.1, “IBM Nways 2230 ATM Routing - Virtual Network Navigator” on page 57 discusses the IBM Nways 2230 algorithm in relation to the QOS routing. To be able to provide SVC services for the end users, the IBM Nways 2230 has to implement the routing algorithm to find the path between the source and destination ATM addresses. IBM Nways 2230 VNN routing algorithm distributes addressing information for all supported ATM address formats. IBM VNN implementation propagates the ports-associated addresses registered via ILMI or entered during configuration time through the network on a longest-match basis. In other words, VNN propagates only that part of the address (node prefix, port prefix or port address) which provides enough information for other nodes to make a routing decision. This approach enables you to build both a structured subnet ATM addressing scheme and an ad hoc address assignment. A subnetting structure should be preferred as the route aggregation results in the reduced size of routing tables in the network. The flat address space used in the ad hoc scenario could severely impact routing efficiency.

When using network-to-network connections, you can specify default route pointing to a selected logical port. When the called destination is unknown to the VNN algorithm, it routes the call to the logical port which is defined as default route. You are allowed to define more than one default route for your network; the one with the lower administrative cost is used.

Powerful hardware design with a dedicated processor for each I/O module greatly improves the SVC per second performance rating of the switch. Distributed design allows IBM Nways 2230 to build around 3000 SVCs/sec. For more information about the hardware implementation, refer to 3.3, "IBM Nways 2230 ATM Switch Models 600 and 650" on page 96.

2.8.9.4 Calling Address Screening

The call screening option can be used to determine if the local switch should process the call request or not. The decision is made at the ingress port of the network. This feature ensures that nonauthorized ATM workstations cannot build the circuit over the network. You can choose the screening using one of the following options:

Node Prefix Screening - Calling party address is checked against all configured node prefixes. If a match is found, the screening is successful and the node proceeds with the call request.

Port Prefix Screening - Screens the calling party address against all configured port prefixes.

Port Address Screening - Screens the calling party address against all configured port addresses.

There is no security implemented on the egress ports of the IBM Nways 2230 network.

2.8.9.5 Calling Address Insertion

IBM Nways 2230 Calling Address Insertion feature can be used when the user end station does not present its address in the call SETUP message. IBM Nways 2230 checks the SETUP message and, based on your configuration, uses one of the listed options:

Insert Calling Party Address - If the logical port received the call request, which does not contain the calling party information element, it inserts a configured address into the Calling Party Address field.

Replace Calling Party Address - If the call setup does not contain the calling party address, IBM Nways 2230 inserts the configured address. If there is a calling address, it overwrites it with the configured one. In such cases, call screening will be successful, no matter what the end stations put into the calling party information field.

Disabled - The logical port does not replace or insert the calling party element.

2.8.9.6 Address Translated and Tunneling

IBM Nways 2230 implements a rich set of features enabling you to define both local and remote gateways and ingress and egress address translated between private and public networks. These features increase the flexibility when interconnecting IBM Nways 2230 with other types of ATM networks.

Each SVC call SETUP message contains the calling party and called party information elements, which are identifying the source and destination port in the ATM network. In some situations, when the tunneling takes place, the calling party subaddress and called party subaddress could be used. Figure 44 on page 68 illustrates the call setup message format.

Calling Address Field	Called Address Field
Calling Subaddress Field	Called Subaddress Field

ATMSUP

Figure 44. ATM Setup Message Format

IBM Nways 2230 supports the following types of address translated and tunneling on either ingress or egress logical ports:

Tunneling - This option should be used when the call is routed through the network which is using different address domains. The following figure illustrates the tunneling functions:

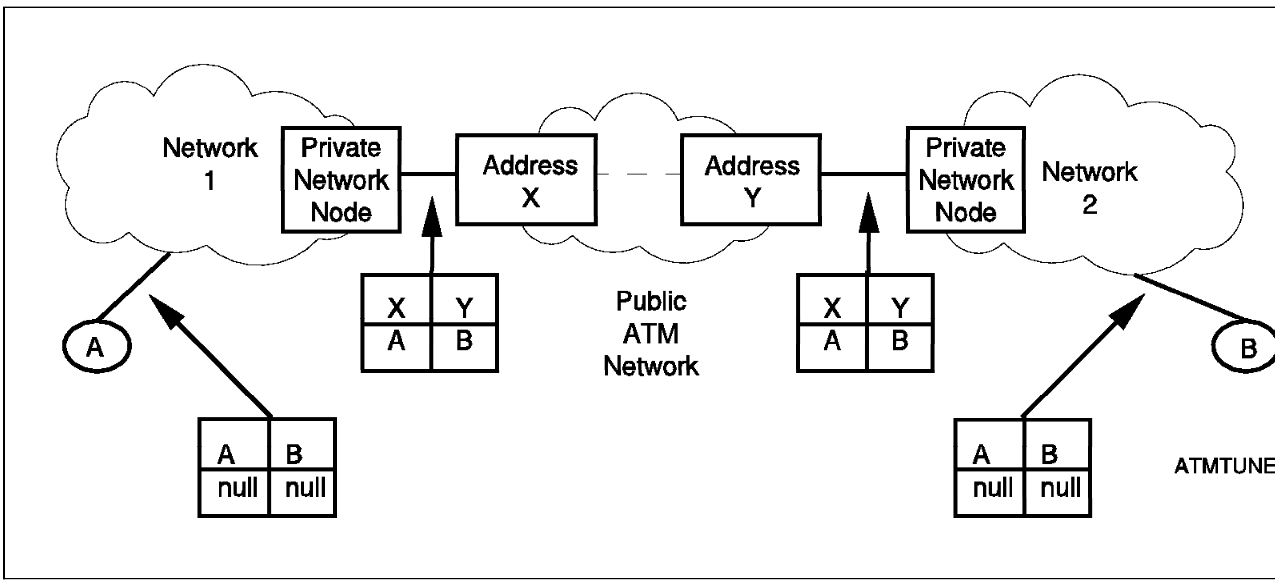

Figure 45. Address Tunneling in SVC Call Setup

The calling party address is substituted by the local gateway address and the called address is substituted by the remote gateway address on the egress logical port. The previous calling and called party addresses are carried in the subaddress fields. Remote gateway (ingress port) replaces the calling and called party address with the subaddress fields and completes the call setup to the destination end station.

E.164 Native to AESA Format Translated - You can use this option for translation of the E.164 native address to the AESA format. As the E.164 address is represented by up to 15 BCD digits, the rest of the AESA address is filled with zeros.

E.165 AESA to E.164 Format Translated - IBM Nways 2230 removes all the leading zeros as well as the trailing portion of the address starting with Fh.

Replace - When you select this option, the calling party address is replaced by the local gateway address, and the called address is replaced by the remote gateway address on the egress logical port. Figure 46 on page 69 illustrates usage of this option.

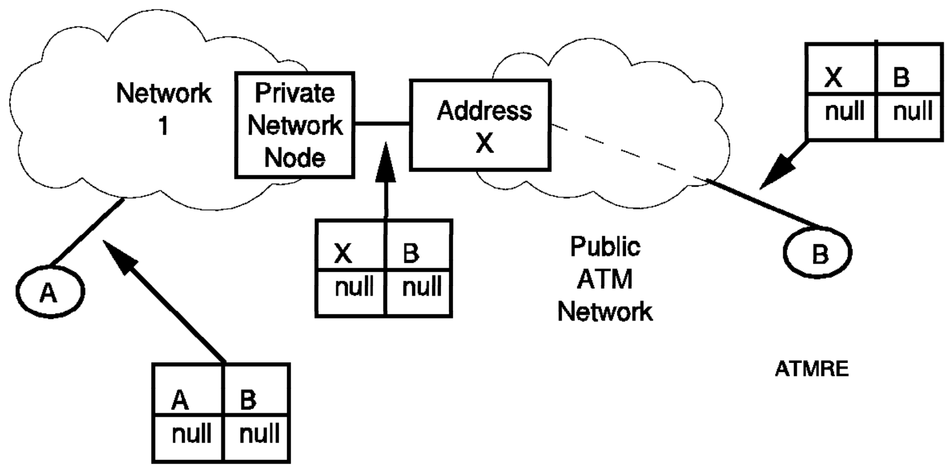

Figure 46. Replacement of ATM Addresses in the Call SETUP Message

2.8.10 IBM Nways 2230 Clock Synchronization

As soon as the IBM Nways 2230 starts to transport voice circuits or any other circuit emulation service, including the transport of a general data stream interconnecting time division multiplexors, synchronization becomes a very important issue. The IBM Nways 2230 can be configured to use one of the following clock sources:

- Use its stratum 3 oscillator
- Phase locked loop (PLL) to the clock signal derived from the I/O ports
- PLL to the external clock reference

The IBM Nways 2230 can derive the clock from the I/O module ports; you can configure two different ports to provide primary and secondary clock reference. If the primary clock reference fails, the switch starts to use the secondary clock reference; when the primary timing source recovers, IBM Nways 2230 automatically switches back to the primary reference. Two other external clock references can be connected to the switching module; for more information, refer to 3.3.2.2, “Switch Processor” on page 98. Primary and secondary clocks are distributed using the backplane to each I/O module. For each module, it is possible to specify if it synchronized from its local clock or if it should use the Primary System Clock or the Secondary System Clock.

IBM Nways 2230 implements both methods of circuit emulation synchronization:

Synchronous Residual Time Stamp (SRTS) - This method requires synchronization of both ingress and egress ATM switches to the same clock source. Data is sent together with the time stamps with the number of cycles of a reference clock period. These time stamps together with the reference clock are used to control an oscillator which drives the output link.

Adaptive Clock Method - When you are using the Adaptive Clock Method, the switch does not have to use the same clock reference for synchronization. Clock speed of the output is related to the number of bits in the output queue. When more bits are accommodated in the queue, clock speed is increased to empty it. If there is not enough data, the clock speed is decreased.

2.8.11 IBM Nways 2225 and 2230 Interconnection

IBM Nways WAN Broadband Switches hardware platform was designed for building large scalable multiservice networks. The IBM 2230 Switch should be used as a high-speed backbone ATM switch, while the IBM 2225 performs all the interworking and translated functions.

Current implementation of the IBM Nways 2225 and 2230 Switches software enables you to use the IBM Nways 2230 ATM network as a transport network for the IBM Nways 2225 Switch network. The following services could be used:

- ATM OPTimum trunk over IBM Nways 2230 network
- Frame relay or PPP/ATM service interworking
- Frame relay/ATM network interworking

IBM Nways 2225 Switch connects to the IBM Nways 2230 ATM network via an ATM UNI I/O module. As the IBM Nways 2225 Switch supports only an ATM DTE logical port, the IBM Nways 2230 has to be configured to perform ATM DCE functions.

Both networks run their own topology database. You have to define management gateways for both networks and all interworking circuits have to be configured for both networks separately. The following figure illustrates the IBM Nways 2225 and IBM Nways 2230 network.

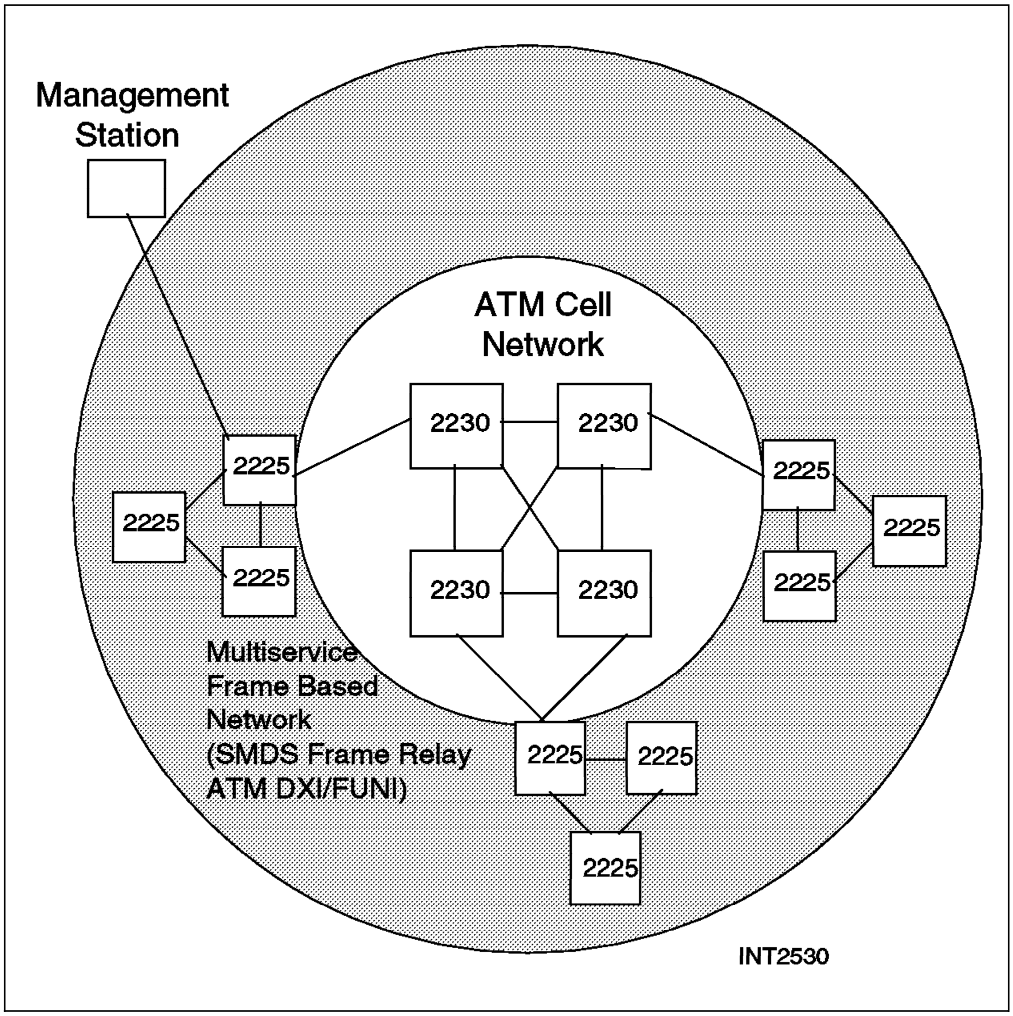

Figure 47. IBM Nways 2225 and 2230 Network

For future enhancements, refer to Appendix A, “Release 4.2 Preview for the Nways 2225” on page 127.

2.8.12 IBM Nways WAN Switches IP Addressing

The Nways 2219, 2225, 2230 WAN Switches use the IP addresses to identify switches inside the network. Each switch has assigned its unique IP address; this approach simplifies implementation of the management and allows you to implement efficient control and routing of the virtual circuits.

When you install the network, you allocate the class C IP network or subnetwork using the subnet mask 255.255.255.0 to the whole IBM Nways Switch network. You have to configure a static route pointing to the IBM Nways Switch, which will perform gateway functions for the management traffic. The IP host part of the switch address is automatically assigned by the network management station during the initial configuration. Nways 2219, 2225, 2230 WAN Switches use the the static mapping between the IP address and the node number. The host part of the IP address is identical to the node number. The network management station generates the IP addresses in the ascending order, starting with the IP address x.y.z.1. Using this IP addressing scheme limits the size of the network to 255 nodes. For future enhancements, refer to Appendix A, “Release 4.2 Preview for the Nways 2225” on page 127.

Chapter 3. Hardware Overview

This chapter provides detailed information on the IBM Nways 2219, 2225, and 2230 hardware architecture.

3.1 IBM Nways 2219 Frame Relay Switch

The IBM Nways 2219 Frame Relay Switch is a fully featured, high-performance, low-cost frame relay switch designed for building private, public carrier and hybrid frame relay networks. The IBM Nways 2219 is based on contemporary hardware and software technologies and industry standards to provide customers with cost-effective products and protect their investments. The modular design provides the flexibility of configuring node capacity according to a customer′s specific requirements. The IBM Nways 2219 provides PVC rate monitoring for congestion management and generates network usage statistics. These two capabilities are important for the management of wide area networks.

3.1.1 2219 Highlights

The IBM Nways 2219 is designed to take advantage of the latest technologies to achieve high reliability, cost effectiveness and future growth. The characteristics of the 2219 include:

- 6-slot modular frame relay switch
- High port density
- High throughput, low latency
- 5 I/O module slots
- Fully redundant fans and power supplies
- Hot swap, live insertion
- AC or DC power
- Flash memory
- RISC processing

3.1.2 IBM Nways 2219 Hardware Architecture

The following figure is a schematic view of the hardware architecture:

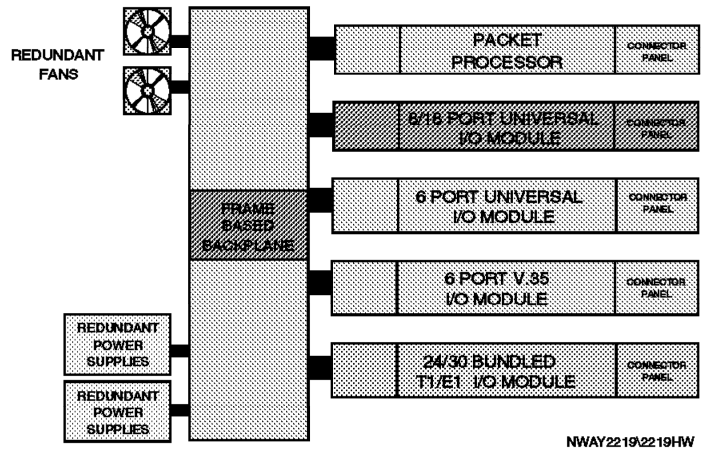

Figure 48. IBM Nways 2219 Model 250 Hardware Components

The various components are briefly described in the following text.

3.1.3 IBM Nways 2219 Components

The hardware architecture comprises three major components:

- One packet processor (PP)
- One or more input/output (I/O) modules
- One common enclosure and power/cooling system, with optional power supply and built-in fan system redundancy

3.1.3.1 Packet Processor

The packet processor is the circuit card with hardware and software operating at the frame or packet level. It performs processing, such as frame format validation, routing, queueing and protocol conversion. Nodes contain a single PP card.

The PP occupies one slot. It contains a high-speed RISC processor (Intel 960CA), 4 megabytes of instruction RAM for software, 4 megabytes of frame buffers to accommodate traffic bursts, 512 KB of flash memory for online software updates, and 128 KB of battery backed-up RAM for configuration storage.

The PP communicates with the I/O modules through an 800 Mbps backplane bus.

The PP has a throughput of at least 24000 pps and a full-duplex 12-Mbps I/O throughput.

The PP runs the network management functions of a node. The network management capacity includes Ethernet and serial interfaces to network management, nonvolatile memory and configuration storage, node status (fans, power) sense, fault-tolerance logic for operating two nodes onto one set of trunks and user links.

Note: The PP card is not hot-swappable.

3.1.3.2 2219 I/O Modules (IOM)

An I/O module manages the lowest levels of a node′s trunk or user interfaces. It performs physical, data link (frame) and multiplexing operations on external trunks and user links.

The I/O module will maintain its throughput with any number of ports, with no underruns or overruns, assuming that the total throughput does not exceed the maximum rating of 12 Mbps for the packet processor.

Hot swapping of an I/O module will not affect the operation of other I/O modules or the PP. The power load change of the I/O module hot swapping will not cause node problems.

Each I/O module has its own reset; in addition, each serial component on the I/O module has a reset provision. Each of these reset functions will allow module recovery to be confined to the module in question.

The PP contains the logic and the I/O contain the physical connections. IBM Nways 2219 does not use the front and rear card architecture.

The following table will give a summary of the supported I/O modules, followed by the technical details of each module:

Table 4. IBM Nways 2219 I/O Modules Summary

I/O module	Port Speeds	Port Capacity	Frame Relay
Channelized E1	1.920	1 30-bundle E1	yes
Channelized T1	1.536	1 24-bundle T1 (full or fractional)	yes
V.35	4.096 Mbps	6 V.35	yes
6-port Universal	4.096 Mbps	6 V.35, 6 X.21,	yes
8 or 18-port Universal	128 Kbps	8 V.24, 8 X.21, 18 V.24, or 18 X.21	yes

1-Port Channelized T1/E1: The 24/30 channelized T1/E1 module has built-in T1/E1 CSUs and delivers a D4 channel format T1 (or E1 equivalent) interface. This makes it easy and economical to interface to multiple customer sites over a single T1/E1 connection.

The module allows any n x DS0 channels the flexibility to act as either network, user port, or interswitch trunks. Users can map the DS0 channels to a maximum of 24/30 data links. Contiguous or noncontiguous n x DS0 channels compose each link.

- Physical interfaces: T1/E1 interface - 15-pin sub-D male connector; redundancy connector - 50-pin connector; E1 connector G.703 75 ohm balanced BNC connectors or 120 ohm; built-in CSU/DSU supporting zero encoding, jammed bit, B8ZS, timing, loop timing, internal timing.
- Status indicators: Single green LED lit for normal operation, blinking for initialization, off for module failure.

6-Port V.35: The 6-Port V.35 I/O module contains six V.35 ports, each of which is capable of data rates from 19.2 Kbps to 4.096 Mbps. Users can individually configure each of the V.35 ports on the module as DCE or DTE to provide any frame-based logical port functions. The individual connections can be made to a user device such as a router, bridge, or cluster controller. Individual connections can also be made to a network trunk via a DSU/CSU device.

For user devices, the connections can be frame relay or non-frame relay. If they are not frame relay, the IBM Nways 2219 will encapsulate or translate data links into frame relay format.

This module provides PVC rate monitoring for usage statistics and allows ports to be user or trunk interface.

- Physical interfaces: ISO 2573 V.35 35-pin female connector, an optional crossover to male with a cable.
- Timing options: internal clock, external clock or loop timing, independent transmit and receive path timing.
- Status indicators: Single green LED lit for normal operation, blinking for initialization, off for module failure.

6-Port Universal: The 6-Port universal I/O module is a base module for the IBM Nways 2219. It supports connections to a variety of synchronous interfaces. The module contains six universal ports that can operate at rates from 19.2 Kbps to 4.096 Mbps.

Users can individually configure each of the physical ports as DCE or DTE to provide any frame-based logical port functions.

This module allows customization of connector types through external I/O assembly; all 6 ports could be either V.35 DCE/DTE or X.21 DCE/DTE. If you want to connect a module to the DCE device the way a modem is, you have to use optional cross-over cable.

The following characteristics are included:

- Physical interfaces: CCITT V.35, 34-pin ISO 2593 (RS-422 drivers for the balanced signal, US only); X.21 15-pin D sub.
- Timing options: internal clock, external clock or loop timing, independent transmit and receive path timing.
- Status indicators: Single green LED lit for normal operation, blinking for initialization, and off for module failure.

8 and 18-Port Universal: This module is available in 8-port and 18-port configurations and supports X.21 or V.24 communications. Users can configure pairs of X.21 or V.24 interfaces, using one of the two external cable types as DCE or DTE to provide any frame-based logical port functions.

You can configure the 8-port or 18-port module with either all X.21 DCE female or V.24 DCE female interfaces. When connecting to a DCE device, optional cross-over cable is needed. The following figure shows the usage of the 8-port universal card.

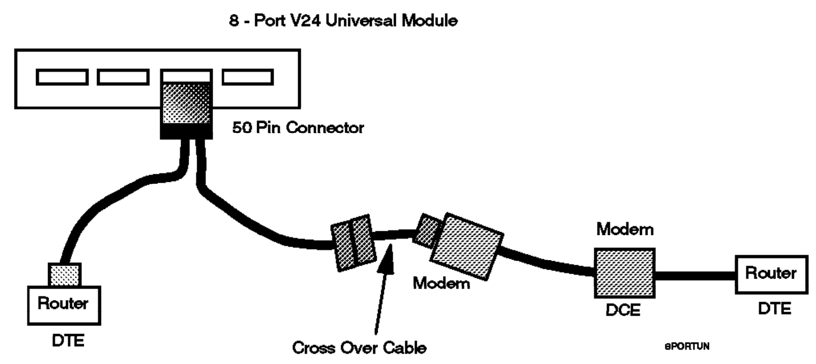

Figure 49. Connections to the 8-Port Universal Card

The individual connections can be made to a user device such as a router, bridge, or cluster controller. Individual connections can also be made to a network trunk via a DSU/CSU device.

For user devices, the connections can be frame relay or non-frame relay. If they are not frame relay, the IBM Nways 2219 will encapsulate or translate the data links into frame relay format.

This module provides increased concentration of the outer tier of a network. It allows aggregation of low-speed connections (2.4 Kbps to 128 Kbps, with 9.6 Kbps being the most common) into a switch that can provide local frame relay switching and channels traffic to backbone sites via T1/E1 circuits. The module is capable of saturating all 8 or 18 physical ports at 128 Kbps full-duplex speed. Each port is individually configurable at 2.4, 4.8, 9.6, 12, 16, 24, 32, 38.4, 48, 56, 64, 96, 112, or 128 Kbps.

This module is especially suited for concentration of multiple sites in a cost effective manner.

- Operational features provided: PVC rate monitoring for usage statistics; support for frame relay.
- Physical interfaces: ISO 2110 V.24 25-pin female connector; ISO 4903 X.21 15-pin female connector; optional cross-over to male with V.24 or X.21 cable.
- Clocking options: transmit and receive independently timed with internal clocking, external clock, and loop timing.
- Status indicators: Single green LED lit for normal operation, blinking for initialization, and off for module failure.

3.1.3.3 Physical Installation Data

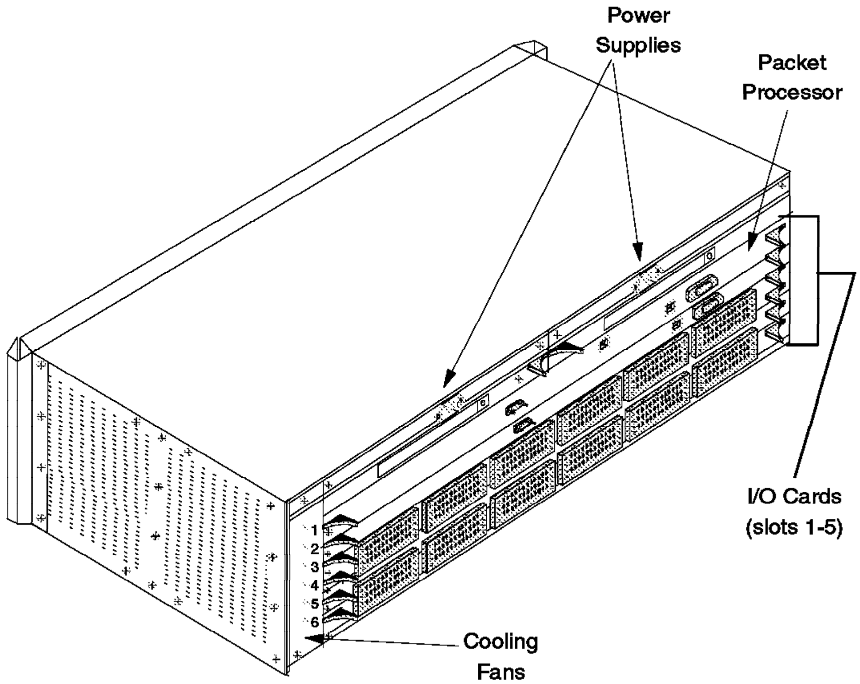

Figure 50. IBM Nways 2219 Rear View

Table 5. IBM Nways 2219 Physical Specifications

Specification	**Description**
Physical Characteristics	Basic switch includes a minimum of one power module, one fan module, and one packet processor. Capacity for five I/O modules.
Size	44.45 cm (17.5 in) wide x 22.225 cm (8.75 in) high x 28.575 cm (11.25 in) deep.
Weight	22.68 kgs (50 lbs) maximum
Thermal Dissipation	300 WATTS maximum, 512 BTU/hr.
Management Interfaces	Ethernet, RS-232, RJ-45.

Table 6. IBM Nways 2219 Electronic Specifications

Feature	**Specification**
120 VAC	3.0 amps maximum, 300 WATTS maximum, 50-60 Hz, single phase
240 VAC	1.0 amps maximum, 300 WATTS maximum, 50-60 Hz, single phase
-48 VDC	-39 to -76 VDC, 300 WATTS max

Table 7. IBM Nways 2219 Environmental Specifications	
Characteristic	**Requirement**
Ambient Operating Temperature	O to 49 degrees C
Relative Humidity	10% to 80% (noncondensing)
Operating Altitude	to 10000 feet
Ambient Storage Temperature	-40 to +65 degrees C, 95% relative humidity
Storage Altitude	-1000 to +30000 feet

3.2 IBM Nways 2225 MultiService Switch Models 400 and 450

The IBM Nways 2225 Multiservice Switch provides a flexible and cost-effective multiservice WAN platform for interworking among frame relay, SMDS ISDN and ATM. Interworking lets technologies at either end point of a connection to communicate seamlessly with one another or make use of new network transport services. Additionally, interworking delivers a clear migration path to ATM while taking advantage of today′s frame relay and SMDS services. By providing broadband packet services, the IBM Nways 2225 Multiservice Switch family gives users the flexibility of choosing the service that suits the particular requirements of their networked applications.

The IBM Nways 2225 Multiservice Switch provides fully redundant, high-performance, flexible broadband packet switches, utilizing a standards-based foundation suited for public and private networks.

The IBM Nways 2225 MultiService Switch family consists of two models: Model 400, which is an 8-slot unit, and Model 450, which is a 16-slot unit.

Note: There is no upgrade path from Model 400 to Model 450.

3.2.1 2225 Highlights

- Modular, scalable platform
- 6-14 I/O slots
- 1.2 Gbps backplane
- Fully redundant modules, fans and power supplies
- Hot swap, live insertion
- AC or DC power
- Flash memory
- RISC multiprocessing

3.2.2 IBM Nways 2225 Architecture

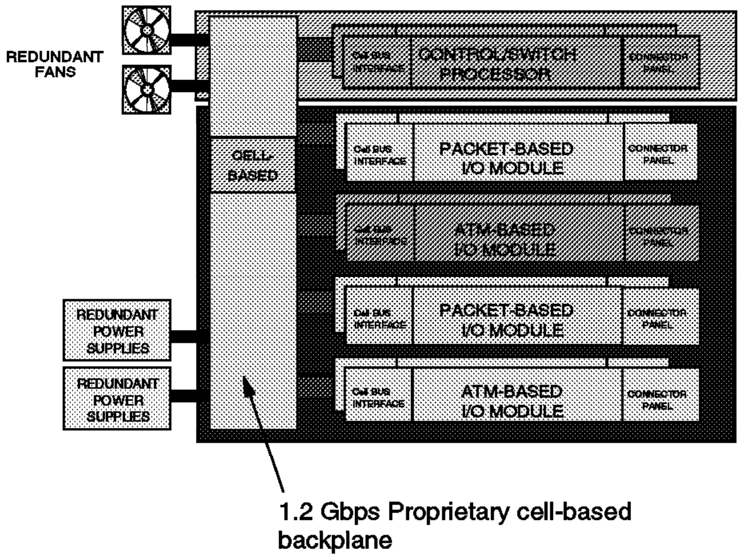

Figure 51. IBM Nways 2225 Hardware Components

The hardware architecture employs symmetrical RISC multiprocessing that consists of a control processor (CP) interacting with multiple input/output (I/O) modules. The CP provides network and system management and network routing functions, in support of the real-time switching functions provided by the multiple I/O modules.

The hardware platform is based on a 1.2 Gbps backplane bus.

The CP and I/O modules utilize the Intel i960 RISC processor to deliver scalable, high-performance packet switching needed for ATM and high bandwidth transmission environments, such as HSSI and E3/T3.

The following picture illustrates the overall task architecture, and the underlying processing performed by both the control processor and I/O modules (IOMs).

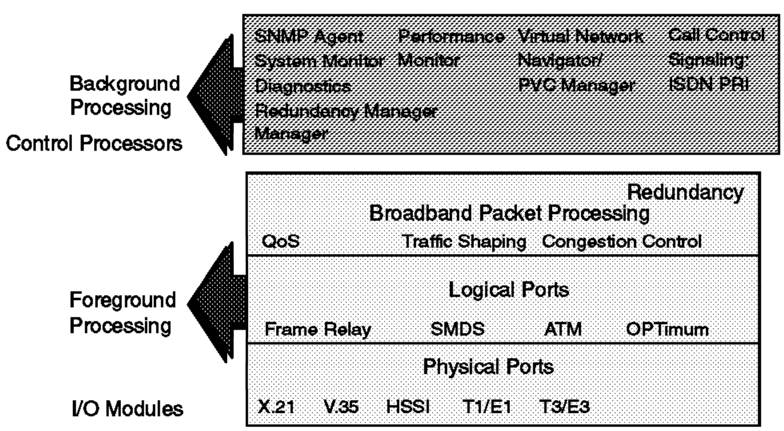

Figure 52. Processor Task Architecture

3.2.3 High Availability

The IBM Nways 2225 constitutes a highly reliable platform. The following features ensure a high degree of operational availability:

- Fully redundant CPs, I/O modules, power supplies and fan modules. All recoveries from failures are automatic.
- Live insertion and hot swap of all CP and I/O modules, power supplies and fan modules.
- Dynamic rerouting of PVCs using Open Shortest Path First (OSPF) algorithm in the event of network outage.
- Fault-tolerant PVC and ISDN backup trunk.

3.2.4 Redundancy

The IBM Nways 2225 Nways CP and IOP modules are optionally available in fully redundant pairs for high reliability networking requirements. Redundancy support is achieved by installing two identically configured CPs and/or IOP modules into neighboring slots in the front of the chassis while installing a single IOA (adapter) module interface into the back of the chassis. Redundant IOA modules occupy the space of the two neighboring slots in the back of the chassis to support the redundant configuration in the front of the chassis, presenting a single interface to the network.

When redundancy modules are installed and configured, the redundant partner continually polls the active module for status. If the redundant module detects a failure in the active module, the redundant module automatically takes over operation, thereby avoiding any service disruptions.

Redundancy and recovery comprise both hardware and software mechanisms that enable the switch to continue operation after certain types of failure. Each switch consists of multiple power supplies and cooling fans. The only restriction is the number of available slots.

Also, redundancy support is distributed in that CPs and IOP modules can be independently redundant. The changeover to a standby CP and/or IOP module is independent from other cards in the system.

The redundant IOP module slots are configured on the management station and sent to each IOP module. Also, hardware provides a status indicating whether an IOP module is connected to the special redundancy I/O panel, and the LEDs show which one is active.

The Redundancy Manager is a background task that runs on each CP and IOP module installed in the switch. It performs the following functions:

- Establishes an active and standby module for each pair of redundant modules
- Monitors the integrity of active modules, using keep alive messages over the cell bus; if a failure is detected in the active module, the standby module disables the active module and performs a warm boot to take over control
- Monitors the standby cards to ensure identical configurations are maintained between the active and standby card and, if a mismatch is detected, automatically updates the standby card configuration

3.2.4.1 IBM 2225 Power Limitations

The I/O modules are divided into two different types according to their consumption of the electrical power:

Type A - CP card, T1, E1, DSX and UIO cards

Type B - HSSI and ATM UNI cards

The Type A IOP cards use 40 WATTS and Type B cards use 45 WATTS. You can use any combination type of IOP cards, but you must not exceed the total power consumption of 700 WATTS.

3.2.4.2 IBM 2225 Front and Rear Card Concept

IBM Nways Switch uses the concept of front and rear cards. The front card provides all the logic and physical interface drivers and the rear card, which is sometimes referred to as the rear panel, provides only the wiring to the physical connectors. Physical cabling is connected to the passive elements. This approach allows you to change the front card with all the logic on it without rewiring all connections.

When you are using the redundant set of IOPs, special wiring for the rear panel is provided to enable connection to both front cards using one physical connector. The following figure illustrates the front and rear card concept:

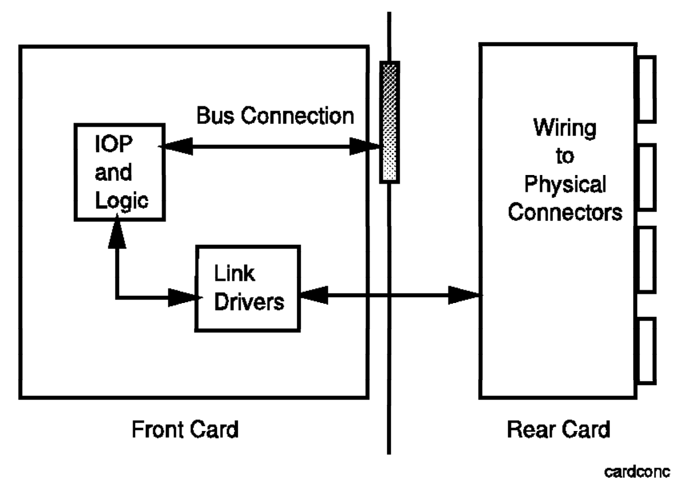

Figure 53. Front and Rear Card Concept

3.2.5 IBM Nways 2225 Components

This section discusses the components of the 2225.

3.2.5.1 Control Processor(s) (CPs)

Slot 1 is reserved for the main CP and slot 2 is reserved for the backup CP. These slots cannot be used for IOP modules. In addition, you cannot install a CP card in any other slot except slot 1 or slot 2.

There are two types of the control processor modules:

CP Basic - The CP basic is equipped with 4-8 MB, 2-8 MB of flash memory for online software updates, and 128 Kbps of battery backed-up RAM for configuration storage. It also implements special hardware for efficient allocation of processing power between packet processing and the background chores. High resolution timer hardware is implemented to allow precise implementation of time average queue length algorithms for congestion analysis.

CP Plus - The CP plus provides the same functions as the CP basic. However, CP plus contains more memory (16 MB of frame buffers) for use with high-speed I/O modules. The CP plus also supports SMDS ICI and billing functions.

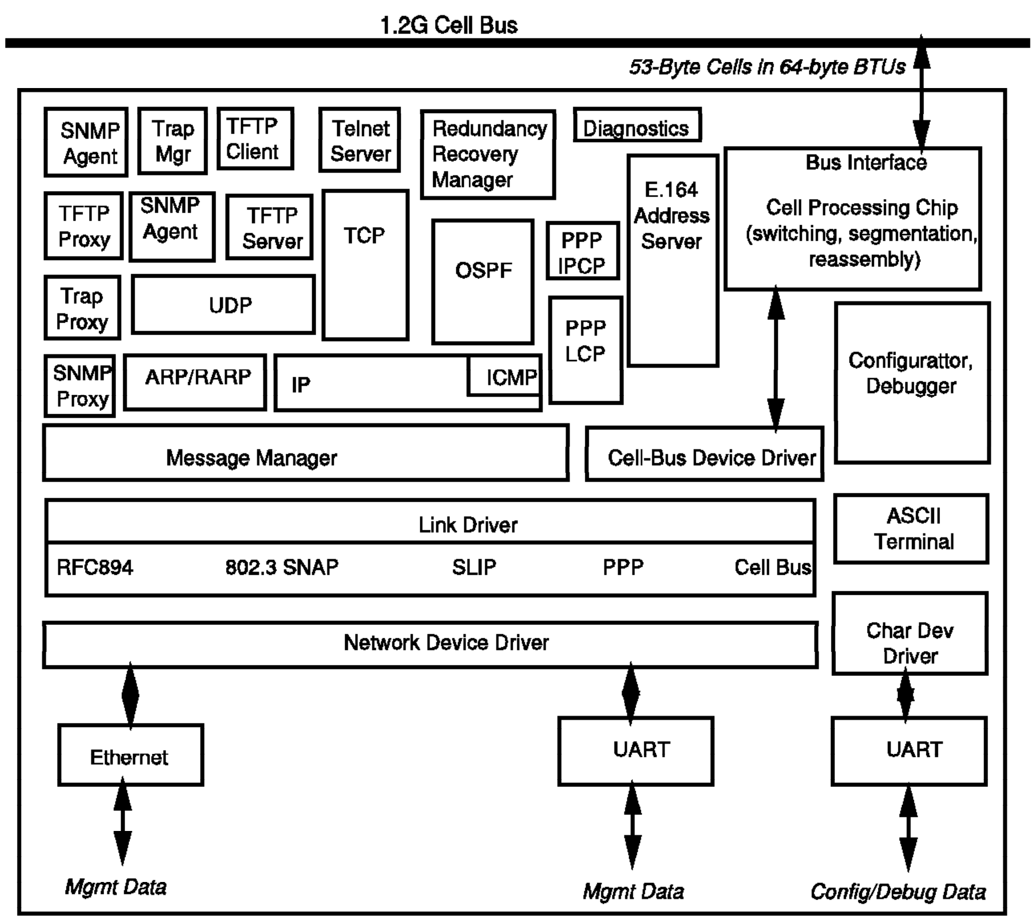

Figure 54. Control Processor Software

The software components residing in the CP are depicted in Figure 54. There are three processing layers:

1. Physical Layer

   In this layer, the CP supports communications with the following devices:

   - External communications devices via Ethernet and RS-232 ports. The Ethernet is for in-band network management, and this is where the Network Management Station will be attached to.

There are two RS-232 ports. One is reserved for out-of-band network management and console access via a dial-up modem, the other is for local console access only, for configuration and debugging purposes.

- The internal bus, shown at the top of Figure 54 on page 83, provides high-speed communications between the active CP and an IOP. A cell bus driver drives the cell processing/bus interface chip.

Each device driver consists of an ISR and subroutines to initialize the device and to transmit data to it.

2. Data Link Layer

   At this layer execute the appropriate link drivers which enable the CP to handle the protocols over Ethernet, UART and Bus Interface.

3. System and Network Services Layer

   One of the two UART ports on the CP is an ASCII character terminal under the control of the system console task, which interprets user commands to perform the following services: configuration initialization via script text file download; MIB command (GET/NEXT/SET) processor; debugger; install; show status; show statistics; reset; PING; telnet.

   The system management services consists of an SNMP agent, action routines, and trap manager. The SNMP agent on the CP represents the switch and acts as a peer to the Network Management Station. All software components which own the supported MIB objects need to register their GET, SET and GET-NEXT action routines with the SNMP agent as part of the initialization after the system is rebooted.

   When a new switch is booted for the first time, the SNMP agent is invoked by the configurator to process a configuration script file downloaded from the Network Management Station (NMS). The SNMP agent then invokes the registered SET action routines to initialize the configuration parameter values. These actions are located in the nonvolatile storage. They are retrieved by the software components for initialization purposes.

   When the SNMP agent receives an SNMP request during runtime, it dispatches the request to the registered GET/SET/GET-NEXT routine synchronously according to the object ID. If the object resides on the CP, a local procedure is invoked. Otherwise, the object resides on an IOP module, and the action routine invoked is actually the SNMP proxy client which sends the request to the real action routine residing in the IOP. The SNMP agent will generate an SNMP response when the SNMP proxy returns.

   When an alarm situation is detected by a CP software component (for example, fans failed), it invokes the trap manager, to send a trap to the NMS asynchronously. An IOP software component can send a trap message to the trap manager through the SNMP proxy.

   The network routing service uses OSPF task to find the best paths from the node to all configured destinations, according to the type of service (currently, guaranteed CIR throughput is the only type of service supported). It keeps the centralized routing table up to date for PVC/SVC establishment (VC path) and IP datagram routing (next hop). Once a PVC or SVC is established, the next-hop VC information is kept in the virtual circuit table for the critical loop to relay data frames. IP also forwards IP datagrams to a destination IP address along the shortest path.

   Though all ports are distributed on the IOPs, the OSPF task will construct its centralized data structures (for example, trunk interfaces and neighbors) during initialization and keep them up to date at runtime. Link drivers on the IOPs will notify the OSPF task when any of the ports change state. The OSPF task sends PDUs to and receives PDUs from trunk ports on the IOPs.

The TFTP client/server service goes into action when a file (for example, a new version of the system software) needs to be downloaded to a node; the NMS notifies the node by setting several MIB variables; the TFTP client on the CP then retrieves (GET) the file from the TFTP Server on the NMS. When a range of memory (for example, statistics log table) needs to be uploaded to the NMS, the TFTP client on the CP delivers (PUT) the memory image to the TFTP server on the NMS. When the backup CP needs to copy a range of memory (for example, nonvolatile storage) from the active CP, the TFTP client on the backup CP retrieves the file from the TFTP server on the active CP.

The SNMP proxy client registers SNMP GET/GET-NEXT/SET action routines on behalf of the objects on the IOPs. When an action routine is invoked, the proxy client sends the request to the proxy server on the target IOP. The response is then relayed back from the IOP to the proxy, then back to the SNMP agent. The SNMP proxy also delivers trap messages to the trap manager on behalf of the IOPs.

The IPCP and LCP tasks are responsible for establishing PPP connections for the NMS out-of-band access.

The CP Redundancy and Recovery Manager functionality has already been discussed.

The diagnostic service includes background processing for nondestructive memory tests, timer/counters, UART, etc. It runs continuously as the lowest priority background task. The results are logged into nonvolatile storage for monitoring and diagnostic purposes. Alarm traps will be sent to the NMS if errors are detected. The NMS can examine the results of the log of the background diagnostics.

3.2.5.2 2225 I/O Modules (IOMs)

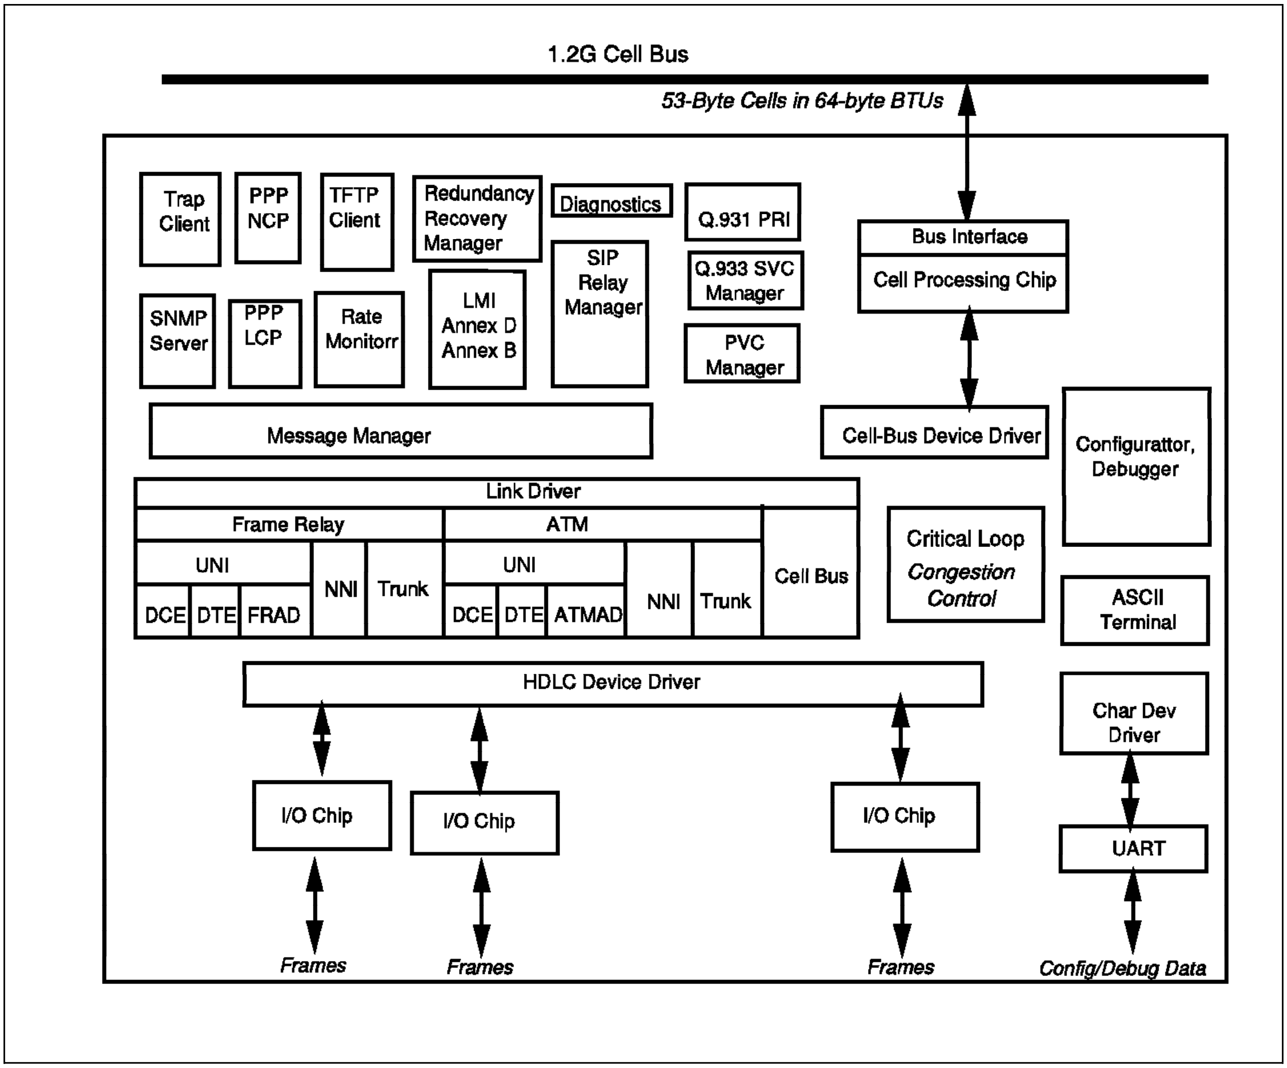

Figure 55. I/O Processor Software

The software components which are executed in the processors contained within the I/O modules are depicted in Figure 55.

On comparing this figure with Figure 54 on page 83, it is immediately apparent that the main differences reside at the physical and data link layer levels, which is understandable given the specific objectives of I/O modules.

All external communications media except RS-232 are designed to transfer in-band user data in HDLC variable-length frames in frame relay or fixed 53-byte cells in cell relay. RS-232 is for debugging use only.

At the data link layer resides the functional sublayers designed to provide data switching between the external communications ports and the internal bus: the frame relay sublayer, the ATM sublayer, the SMDS sublayer, and the BUS sublayer.

The frame relay sublayer uses the LAPF (the core HDLC) variable-sized frames specified in ANSI T1.617/618 and CCITT Q.922/933 as the user-to-network and network-to-network protocol format.

The ATM sublayer uses the standard 53-byte cells specified in ITU-T I.361 as the user-to-network, network-to-network and trunk protocol frame.

LMI maintains the operation of a UNI or an NNI, and it also reports status of virtual circuits asynchronously or via full status poll or response. For links where only PVCs are enabled, both LMI Revision 1 and ANSI Annex D are supported.

In addition to LMI, link drivers and the critical loop also perform the following management functions: collect traffic statistics for ports and circuits, collect QoS statistics for circuits, monitor operational status of ports and circuits, and send traps to the NMS if any change of status occurs; respond to the SNMP GET/GET NEXT/SET request from the NMS.

The background rate monitor task measures the user data arrival rate at the ingress node using the leaky bucket algorithm. Two counters are maintained by the rate monitor for each PVC on each user link. Bc_cnt is the number of committed bits allowed during the current time interval, and Be_cnt is the number of uncommitted bits allowed during the current time interval. The measurement time interval for each PVC is determined by the configured Bc/CIR and is measured by continuously running the timer, which is started when the switch is initialized. Bc_cnt and Be_cnt are initialized to the Bc and Be, respectively, of each user PVC. Each time a frame is received by the critical loop in an ingress node, the two counters are used to mark the frame to be red, amber or green. For each timer expiration, Bc_cnt and Be_cnt of each PVC are incremented by the burst and excess burst size allowed for a time interval, not to exceed the Bc and Be defined for the PVC.

The following table summarizes the main characteristics of I/O modules available in the IBM Nways 2225:

Table 8. IBM Nways 2225 I/O Modules Summary

IOP module	**Port Speeds**	**Port Capacity**	**ATM**	**SMDS**	**frame relay**
Universal IO	8 Mbps	8 V.35 or 8 X.21	yes*	yes	yes
Unchannelized T1/E1	1.544 or 2.048 Mbps	4 24-bundle T1 or 4 30-bundle E1	yes *	yes	yes
Channelized T1/E1	24 or 30 56/64 Kbps channels	4 24-bundle T1 or 4 30-bundle E1	yes*	yes	yes
HSSI	45 Mbps	2 HSSI	yes*	yes	yes
DSX-1	1.54 Mbps	10 24-bundle T1 or 10 30-bundle E1	yes*	yes	yes
Channelized DS3	56 Kbps through to 1.54 Mbps	Up to 28 T1 channels	no	no	yes
ATM UNI T3/E3	44.736 Kbps or 34.368 Kbps	1 T3 or E3 port	yes	no	no

* - ATM DXI

8-Port Universal IOP Module: This module supports connections to a variety of synchronous interface connections. Any port has the flexibility to act as either network (DCE), user port (DTE), interswitch trunk, or internetwork link (NNI). As a trunk connection, this module supports the ATM DXI, SMDS DXI, direct trunk, and OPTimum trunk interfaces. For user devices, the connections can be frame relay, UNI, or SMDS DTE. This module also supports connections from X.25, HDLC, or PPP and will assemble data links into the appropriate format for switching in a frame relay, SMDS, or ATM network.

The 8-port UIO module′s total speed is dependent on the packet size. However, 12 Mbps maximum is recommended for small sized packet transmissions.

The module consists of the front card, which contains all active circuits, and passive rear card with physical interfaces. The same front card can handle different types of rear cards. You can use one of the following types of rear cards: V.35 DCE, or X.21 DCE. You are not allowed to mix different types of ports using the same rear card.

IBM provides two different sets of the rear cards for redundant and nonredundant configuration. The redundant back card occupies two slots, providing eight physical ports, while the nonredundant version occupies only one slot, providing the same number of ports.

This module offers high network availability through a redundancy option. The redundancy option switches the line interface over to an optional redundant IOP module, operating in the same IBM Nways 2225 switch.

Operational Features

- PVC rate monitoring for usage statistics
- Allows ports to be user or trunk interface
- Cell and frame relay trunking
- Managed by SNMP
- Operates on 120v per switch (however, all cards have -48v)
- Supports any combination of frame relay DLCIs or ATM VCCs with an aggregate total of 4000
- Supports clocking per port up to 8 Mbps
- Supports up to 240 SMDS individual addresses

Interface Standards

- CCITT V.10
- CCITT V.11
- Network interfaces: V.35, and X.21

Physical Interfaces

- CCITT V.35 34-pin ISO 2593
- X.21, 15-pin D-sub
- Network interfaces: V.35, X.21

4-Port T1 IOP Module: This module contains four integral T1 CSU/DSUs and provides a D4 or ESF channel format T1 interface. This module comes in both channalized and unchannelized versions. This makes it economical and easy to interface to multiple sites over a single T1 connection. Any n x DS0 data link channel on this module has the flexibility to function either as network, user port, or interswitch trunk.

Users can map the DS0 channels on each of the T1 ports to a maximum of 24 HDLC data links per port. Contiguous or noncontiguous n x DS0 channels compose each HDLC data link. Users can configure each of the n x DS0 data link channels as DCE or DTE to provide a variety of logical port functions.

Operational Features

- PVC rate monitoring for usage statistics
- Up to 24 individual DSO data links over a single T1 connection (96 HDLC data links per IOP module)
- Support for contiguous or noncontiguous DS0 channels
- Ports can be network, user port, or interswitch trunks
- Contains four integral T1 CSU/DSUs supporting D4 Super Frame or Extended Super Frame (ESF)
- T1 and Fractional T1 interfaces
- Supports full T1 line and payload loopbacks, and DS0 level channel loopbacks
- Supports external, looped time, or internal clock input (external clock input is T1 via RJ-48 connector)
- Supports up to 1000 frame relay DLCIs
- Supports up to 240 SMDS individual addresses

Interface Standards

- AT&T Publications 62411 and 54016
- Bell TR-NPR-000054 and TR-TSY-000194
- CCITT G.703/4
- ANSI T1.403
- FCC Part 68
- Line Coding: AMI, B8ZS
- Framing: D4, ESF
- Network Interface: DS1

Physical Interfaces

- Four T1 ports
- 15-pin sub-D male connector, RJ-48 8-pin modular connector; both connector types supplied with integrated port autosence detection.
- Input clock source: BNC

4-Port E1 IOP Module: This module contains four integral E1 CSU/DSUs, and provides a CRC4 channel format E1 interface. This makes it economical and easy to interface to multiple sites over a single E1 connection. Any n x 64 Kbps data link channel on this module has the flexibility to function either as network, user port, or interswitch trunk.

Users can map the 64-Kbps channels on each of the E1 ports to a maximum of 30 HDLC data links per port. Contiguous or noncontiguous n x 64 Kbps channels compose each HDLC data link. Users can configure each of the n x 64 Kbps data link channels as DCE or DTE to provide a variety of logical port functions.

Operational Features

- PVC rate monitoring for usage statistics
- Up to 30 individual HDLC data links over a single E1 connection (120 HDLC data links per IOP module)
- Support for contiguous or noncontiguous DS0 channels
- Ports can be network, user port, or interswitch trunks
- Contains four integral E1 CSU/DSUs supporting CRC4 frame format
- E1 and Fractional E1 interfaces
- Supports full E1 line and payload loopbacks, and DS0 level channel loopbacks
- Supports external, looped time, or internal clock input (external clock input is E1 with 75-ohm BNC connector)
- Supports up to 1000 frame relay DLCIs
- Supports up to 240 SMDS individual addresses

Interface Standards

- CCITT G.703/4
- Line coding: HDB3
- Framing: CRC4
- Network Interface: E1

Physical Interfaces

- Four E1 ports, G.703 coaxial pair 75 ohm unbalanced, and G.703 symmetrical pair 120 ohm balanced; both connector types supplied on the same rear card.
- 15-pin sub-D male connector
- BNC connector
- Input clock source: BNC

2-Port HSSI IOP Module: This module provides connections at data rates of up to 45 Mbps. The two physical ports can each be configured at speeds that increment from 1.54 Mbps to 45 Mbps to support DTE or DCE connections. In addition, trunking of saturated T3/E3 circuits is supported via a HSSI to T3/E3 conversion device. The total throughput of the IOP module is 45 Mbps.

Users can individually configure each of the two physical ports as DCE or DTE ports to provide any frame relay logical port functions, as well as ATM DXI and SMDS DXI.

This module enables routers to connect to the network at LAN speeds; it also enables the switch to function as a high-speed trunk between switches or as a feeder link into a backbone central office class switch.

Operational Features

- Operates on -48v DC in a central office environment
- Allows flexibility with DTE, DCE and NNI frame relay interfaces
- Supports both ATM and SMDS DXI
- Adheres to all frame relay standards
- Supports non-frame relay services via direct FRAD and translated FRAD (PPP to RFC 1490)
- Hot swap and live insertion of all cards
- Managed by SNMP
- Supports any combination of frame relay DLCIs or ATM VCCs, with an aggregate total of 1000
- Supports up to 240 SMDS individual addresses

Interface Standards

- HSSI design specification Rev.2.11
- Line coding: NR2
- Framing: HDLC
- Network Interface: Data rate from 1.58 Mbps to 44.21 Mbps
- Programmable as DCE or DTE

Physical Interfaces

- Two HSSI ports
- Amplimite 50-pin receptacle

10-Port DSX-1 IOP Module: This module provides a Fractional T1 connection at a variable data rate (for example, 56 Kbps, 64 Kbps, 128 Kbps, or 1.54 Mbps) that is directly wired into the central office equipment without the use of CSU/DSUs. Users can individually configure each of the 10 DSX-1 ports as DCE or DTE to provide any frame relay logical port function, as well as ATM DXI and SMDS DXI.

The DSX-1 IOP module is intended for a switch located in a central office. The support for unchannelized T1 connections reduces unnecessary expense that network providers incur when using channelized T1 equipment to carry unchannelized data.

A typical configuration for a switch equipped with DSX-1 would be in a central office in a major metropolitan area that concentrates multiple remote T1 connections (via CSU or multiple interswitch trunks within a building). The central office switch, in turn, connects to another major metropolitan area via a HSSI-derived (over a CSU/DS3 conversion device) or ATM UNI 45-Mbps connection.

Operational Features

- Provides Fractional T1 connections
- Supports full T1 line and payload loopbacks
- Operates on -48v DC in a central office environment
- Allows flexibility with DTE, DCE and NNI frame relay interfaces
- Supports both ATM and SMDS DXI

- Adheres to all frame relay standards
- Provides non-frame relay services via direct FRAD and Translated FRAD (PPP to RFC 1490)
- Supports external (DSX-1), looped time, or internal clock input
- Supports any combination of frame relay DLCIs or ATM VCCs, with an aggregate total of 1000
- Supports up to 240 SMDS individual addresses

Interface Standards

- AT&T Publication 62411
- ANSI T1.403
- Line coding: AMI(jammed bit), B8ZS
- Framing: D4, ESF
- Network interface: DSX-1

Physical Interfaces

- Ten DSX-1 ports and external clock connection, RJ-48 8-pin modular connector
- Physical connector: RJ-48

1-Port Channelized DS3 I/O Module: This high-speed module enables IBM 2225 switches to provide 28 T1 connections at a variable data rate (56 Kbps to 1.54 Mbps). Each of the DS1 channels can individually be configured as DCE or DTE to provide any frame relay logical port function.

Operational Features

- Supports maximum throughput of 45 Mbps/port/module (256 byte frame size)
- Provides 28 T1 connections for frame-based traffic
- Supports fractional T1 connections
- Supports 938 maximum originating PVCs
- Supports 4000 PVCs
- Operates on -48 VDC in a central office environment
- Allows user or trunk interface port configuration
- Supports DTE, DCE and NNI frame relay interfaces
- Supports OPTimum trunking
- Adheres to all frame relay standards
- Supports external T1 clock source, loop timed, or internal clocking
- Live insertion capability
- Optional redundancy

Interface Standards

- AT&T Publication 62415
- ANSI T1.102
- Line coding: B8ZS
- Framing: D4, ESF (ITU-T)

- Application Mode: C-bit Parity, M13
- DS3 standards: ANSI T1.103, T1.107, T1.107a, T1.231, T1M1.3/91-003R3, Bellcore TR-NWT-000499, AT&T TR54016, RFC 1407
- DS1 standards: ANSI T1.403, T1.107, Bellcore TR-TSY-000312, Bellcore TR-NWT-000499, AT&T 62411, RFC 1406

Physical Interfaces

- The module has 1 DS3 port.
- Physical connectors: One RJ-48 8-pin modular connector (external clock; two 75 ohm BNC connectors (trx/rcv); one wire-wrap connector (external clock)

ATM UNI IOP Module: The most common configuration for a switch equipped with an ATM UNI IOP module is as a user-to-network interface between two switches. It provides high-performance trunk access to an ATM backbone network at T3 rates.

Operational Features

- Provides speeds of up to 45 Mbps (DS3/E3 speeds)
- Each module has an 8 Kbps port buffer for high-speed network communications
- Provides ATM access for multiple protocols (WAN, bridge and routing)
- Provides E3 and T3 interfaces
- Supports up to 4096 virtual circuits per module
- Provides a maximum CSPS-PDU size of 8192 bytes
- Provides interworking between frame and cell-based traffic

Interface Standards

- TR-TSY-000499
- ANSI T1.102
- CCITT G.703

Physical Interfaces

- One ATM-UNI port and external clock connection, COAX connector
- Physical connector: Coaxial cable

4-Port T1/E1 ISDN PRI IOP Module: This module supports an ISDN Primary Rate Interface. A typical application is for user devices to access the ISDN network via a Basic Rate Interface and, using the D-channel for signalling, establish a connection between one of their B-channels and a B channel on the 4-Port ISDN PRI IOP module. Once the connection is established, the user device communicates with remote users using a standard frame relay protocol over the B-channel. On a port-by-port basis, each port can be configure as an ISDN PRI or a standard channelized T1.

Operational Features

- Supports mixing the four ports as either ISDN PRI or Channelized T1/E1.
- Provides 23/30 B-channels for traffic and one D-channel for management.
- In addition to core ISDN functionality, B-channel polling and Caller ID Authentication are also supported.

- Supports the following Central Office switches: AT&T 4ESS, AT&T 5ESS, and Northern Telecom DMS-100.
- The switch type is dynamically configurable from the network management station.

Interface Standards

- AT&T Publications 62411 and 54016
- Bell TR-NPR-000054 and TR-TSY-000194
- ANSI T1.403
- FCC Part 68
- Line Coding: AMI (jammed bit), B8ZS
- Framing: D4, ESF
- Network interfaces: DS1

Physical Interfaces

- Four ISDN (T1) ports and external clock connection, RJ-48 8-pin modular connector
- Physical connector: RJ-48 (clock input), 15-pin Sub D male connector or RJ-48 8-pin modular connector (port connectors)

3.2.5.3 System Enclosure

The IBM Nways 2225 MultiServices Switch features one common enclosure and power/cooling system. The enclosure is the same for both models of the IBM Nways 2225 MultiServices Switch family, with the exception that slots 9-16 are sealed off and cannot be used in the Model 400.

3.2.5.4 Physical Installation Data

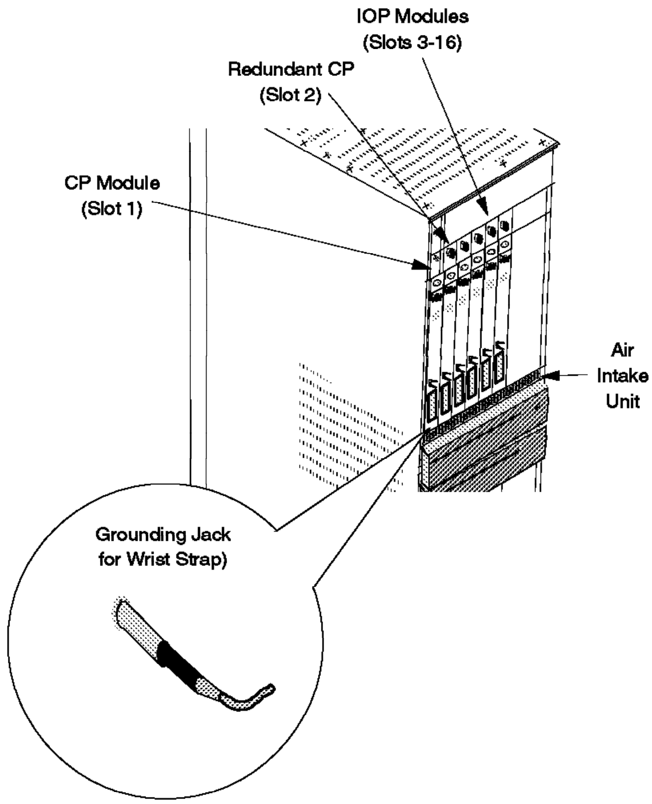

Figure 56. IBM Nways 2225 Front View

Table 9. IBM Nways 2225 Physical Specifications	
Specification	**Description**
Physical characteristics	Basic unit includes one power module, one cooling fan module, and one CP module. Model 400 has capacity for an additional CP and 6 IOP modules; Model 450 has capacity for an additional CP and 14 IOP modules. Both models support an additional power supply.
Size	48.26 cm (19.0 in) wide x 80.01 cm (31.5 in) high x 38.1 cm (15 in) deep
Weight	72.57 kgs (160 lbs) max
Thermal Dissipation	1300 WATTS maximum, 4433 BTU/hr AC; 1000 WATTS maximum, 3410 BTU/hr DC
Management Interfaces	Ethernet, RS-232

Table 10 (Page 1 of 2). IBM Nways 2225 Electronic/Electronic Specifications	
Feature	**Specification**
120 VAC applications	3.0 amps maximum, 300 WATTS maximum, 50-60 Hz, single phase

Table 10 (Page 2 of 2). IBM Nways 2225 Electronic/Electronic Specifications	
Feature	**Specification**
240 VAC applications	1.0 amps maximum, 300 WATTS maximum, 50-60 Hz, single phase
-48 VDC applications	-39 to -76 VDC, 300 WATTS max

Table 11. IBM Nways 2225 Environmental Specifications	
Characteristic	**Requirement**
Ambient Operating Temperature	O to 49 degrees C
Relative Humidity	10% to 80% (noncondensing)
Operating Altitude	to 10000 feet
Ambient Storage Temperature	-40 to +65 degrees C, 95% relative humidity
Storage Altitude	-1000 to +30000 feet

3.3 IBM Nways 2230 ATM Switch Models 600 and 650

The IBM Nways 2230 ATM Switch provides asynchronous transfer mode (ATM) services for both public carrier and private networks. This switch supports data, and circuit emulation traffic, allowing for the efficient sharing of network resources. The switch delivers high performance (internal switch port speed of 640 Mbps), and features dedicated signal processing, advanced distributed processing, and network topology management. In a typical application, the IBM Nways 2230 ATM switch is an ATM backbone switch, providing standards-based ATM WAN services.

There are two models of the Nways ATM Switch:

- **Model 600** offers a 4x4 matrix that runs at 640 Mbps per matrix thread, for a capacity of 2.5 Gbps nonblocking with 64 Kbps cell buffers in the switch fabric. The Model 600 has eight slots; however, only six can be used for I/O modules.
- **Model 650** offers an output-buffered nonblocking 8x8 matrix with each thread running at 640 Mbps for a total switching capacity of 5 Gbps, with 128 Kbps cell buffers in the switch fabric. The Model 650 is a 16-slot switch, but only 14 slots used for I/O modules. The new switching fabric will operate with all current I/O modules and will allow you to fully populate the switch with 4-port OC3 running in nonblocking configuration.

  IBM 2230 architecture enables easy upgrades. You can upgrade the Model 600 to Model 650 only by changing the central switching processor module. In case of a redundant configuration, you have to change both the main and backup SP.

3.3.1 2230 Highlights

- Line speeds range from T3 to OC3c
- PVC/PVC point-to-point and point-to-multipoint support
- ATM Virtual Channel Connection (VCC)
- Internetworking Services and Protocol translation via 2225
- Packet rate monitoring for network accounting and design

- Thousands of virtual circuits per physical port, 16KB virtual circuits per IOP module, and 240KB virtual circuits per switch
- Guaranteed hardware multicast and Quality of Service (QoS)
- Expansive cell buffers per line module for bursty data traffic
- Connection Admission Control algorithm
- Protocol translation features
- Usage-based billing capability
- Congestion management
- Optional redundant switch processor (SP) and power supplies for high-reliability networking
- Support for all four ATM classes of service: Constant Bit Rate (CBR, Variable Bit Rate - Real Time (VBR-RT), Variable Bit Rate - Non-Real Time (VBR-NRT), and Unspecified Bit Rate (UBR)
- Ten QoS classes implemented via hardware buffers: one CBR, four (VBR-RT), four (VBR-NRT), and one UBR
- SVC support for UNI 3.0, UNI 3.1, and IISP interfaces

3.3.2 Switch Hardware Architecture

This section describes the 2230 architecture, software and hardware in detail.

3.3.2.1 Fabric Architecture

The switch is internally an 8 x 8 switch, accommodating 14 I/O slots on the chassis. The ATM switch is inserted in either slot 8 or 9 in the chassis (switch redundancy is supported), and all of the I/O slots are radially wired to the ATM switch(es). This is how the I/O slots access the eight ports of the ATM switch:

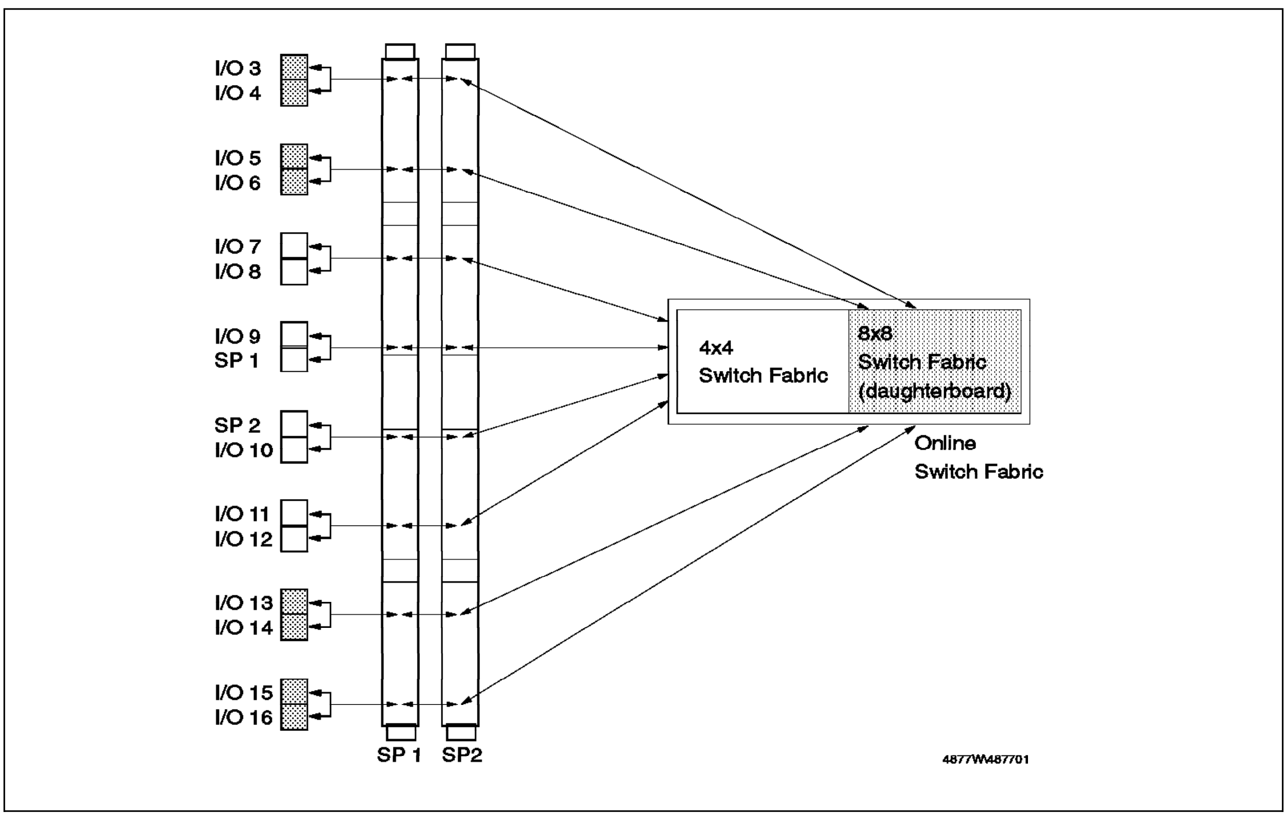

Figure 57. Switch Threads and I/O Ports

Shared threads are ideal for redundant pairs.

SP arbitrates thread access in order to ensure fair distribution of thread bandwidth between contending IOPs.

The speed of each thread is 640 Mbps; the switch can therefore take 4 x 155 Mbps sources, nonblocking, per thread. The CAC algorithm will limit the PCR to less than the thread bandwidth, total of all PCRs per one thread (two neighbor IOPs) cannot exceed the 640 Mbps. You cannot configure a blocking configuration for the CBR, VBR-RT and VBR-NRT traffic. However, the UBR cells will be dropped if the ingress ports of both neighbor modules will exceed the thread capacity. The switch processor will fairly distribute the thread bandwidth between the IOPs, which use the same thread.

At 155 Mbps, the switch is effectively a 32 x 32 nonblocking switch.

Since cell buffers are located on the individual IOPs as well as the switch fabric itself, the switch is considered a dual stage buffered switch. The size of the cell buffers located in the switch fabric is user configurable, whereas the cell buffers located on the IOPs are fixed size, depending upon the QoS.

3.3.2.2 Switch Processor

The IBM Nways ATM Switch employs proven symmetrical RISC multiprocessing, consisting of i960 CPUs on the switch processor(s) and every interface module, for enhanced overall performance and real-time functions.

The SP incorporates the silicon-based crosspoint matrix that constitutes the switching fabric of the switch. It also provides system-wide coordination functions for the platform, and participates in overall network wide functions, such as network management and operation, routing, signaling, and traffic management.

The SP redundancy feature allows either one of the two SPs to take over in the event of a failure. The active SP periodically sends out a sleeping pill message that contains the status of the node to the redundant module. If the redundant module does not see a sleeping pill message within a specified period, it will force the active module out of service and take over control of the node with a minimum of disruption; the hardware then prevents the out-of-service SP from performing any I/O activity.

The SP features a DB-9 connector on the front for module debug and manufacturing diagnostics, a DB-25 connector on the rear card for console access, running at 19.2 Kbps (modem can be ordered), and two rear-card integrated, hot-swappable, PCMCIA slots.

PCMCIA Socket A contains a Type I Ethernet module with 10BaseT interface, for connection to the NMS.

PCMCIA Socket B contains a Type III hard disk, with 170 MB of storage. The presence of a working hard disk is mandatory for the SP to function. The following data storage features are implemented:

128 KB Boot EPROM - The EPROM is in a socket and should never be changed; it holds the most primitive loading and test code.

1 MB Flash Memory - The flash memory is the nonvolatile, read-erase-write memory that holds the major program loading steps.

Parameter RAM (PRAM) - The PRAM holds trace and temporary states.

PCMCIA Type III hard disk - The hard disk is the central switch storage side. The disk drive holds the software code for each I/O module and for the SP itself. It also contains the network configuration, billing information and some trace information. Total storage capacity of the hard disk is at least 170 MB. PCDOS format is used.

The PCMCIA hard disk holds two copies of the switch node data. One copy is ACTIVE and the other one is BACKUP. Both copies contain application software and configuration data. You can copy the ACTIVE area to the BACKUP area and vice versa. Any fatal error in the disk drive causes the switch over to the standby switch processor. The following table illustrates the PCMCI hard disk files distribution:

Table 12. Switch Processor Operations

Data Type	Drive	Data	# of Area	Size of each Area
Software Image	C:	Sp images	2	2 MB
		IOM Type A images	2	2 MB
		IOM Type B images	2	2 MB
		IOM Type C images	2	2 MB
		IOM Type D images	2	2 MB
Config Data	C:	SP configurations	2	512 Kbps
		IOM Configurations	2 x 14	512 Kbps
Billing Info	D:	IOM Billing Info	14	6 MB

The switch has two system-wide timing references (primary and secondary). The timing reference can be: Internal Stratum 3 Clock Source; External Clock Input A (T1/E1); External Clock Input B (T1/E1); Derived from OC-3c line, DS-3 line, or DS-3 PLCP. The SP has two Phase Lock Loop (PLL) systems, which drive a node-wide reference. These two references are called primary and secondary clock reference and are distributed as a 2.048 Mhz signal pair to all I/O modules. Any I/O module can use either reference for its transmit timing on a port-by-port basis. The node timing signal is provided by the online SP.

3.3.2.3 Throughput

The IBM Nways ATM Switch fabric has 8x8 640 Mbps fabric streams. To the extent that the customer′s choice of interfaces (DS3/E3 to OC12) does not exceed the aggregate bandwidth of the switch, the switch is nonblocking. When the limits are exceeded, Call Admission Control negotiates traffic contracts guarantee QoS based on capacity.

Thus, when the system is operating under full load, the maximum sustained throughput is 5 Gbps.

The switch architecture supports fully loaded line speeds at DS-3, and OC-3c in a full-duplex mode, meaning that the lines can be fully saturated at the minimum intercell gap interval. The switch supports line speeds of DS-3 at 45 Mbps, and OC-3 at 155 Mbps.

There are no restrictions, within the fabric-size limits, to the concurrent operation of high-speed ports; for example, data handled by 112 T3 interfaces operating at 45 Mbps can be switched.

3.3.2.4 QuadPlane Architecture

The switch features four programmable parallel switch fabrics to provide an independent plane for each of the four classes of service defined by the ATM Forum (CBR, VBR-RT, VBR-NRT, UBR). Internally, it is a single stage of switching, with quad-plane buffering on the output ports. This is illustrated in the following figure:

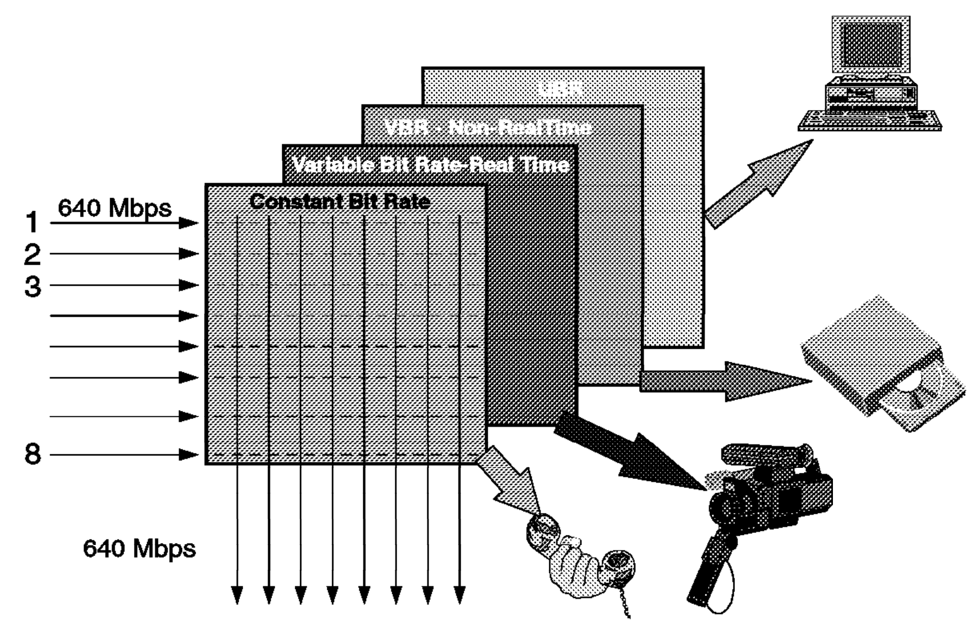

Figure 58. Quadplane Switch Architecture

Cells are switched into one of four parallel dedicated buffer structures that constitute the switch fabric. The buffers supporting variable bit rating are further subdivided with four individual threshold markers, creating a total of ten distinct classes of service.

The ATM switch microcode ensures that these ten classes are available and guarantees end-to-end service in the ATM network.

3.3.2.5 Traffic Management

The IBM Nways 2230 ATM Switch is optimized for transport of multimedia traffic. This switch uses a nonblocking, busless architecture in combination with a quadplane buffer fabric to ensure optimal traffic management. The quadplane architecture assigns a separate, software configurable buffer to the ATM Forum defined classes of service. The fabric can hold 128000 cells and is divided among the four previously mentioned planes.

Additionally, the I/O modules in the switch use a similar quadplane buffer to support QoS classes. I/O buffer capacities are 24000 cells (1.2 MB) per OC-3 and 8000 cells (424 Kbps) per DS3 port.

Unspecified Bit Rate (UBR) is implemented by a daughter card that attaches to the IOP, which performs value-added processing on a UBR class of service. The UBR option can handle up to 1000 VCs with an additional queue capacity of 64000 cells (on top of the 24000/8000 cell buffer on the IOP) provided on the daughter module. This processing can go up to 600 Mbps. The processing does not affect the other classes of service.

The CBR buffer plane is serviced with strict (that is, preemptive) priority over other planes to ensure that time sensitive traffic does not get stuck behind other traffic.

VBR-RT, VBR-NRT and UBR buffer planes are serviced in a *weighted* round-robin fashion. (Note: Users can select one of two servicing algorithms for each port on an IOP: either weighted round robin, which is the default, or fixed priority. The service algorithm is selected at the physical port level.) Bandwidth is divided into 32 slices, and slices are assigned to VBR-RT and VBR-NRT in proportion to the amount of established traffic. UBR is guaranteed at least one slice.

The weights are determined by the effective bandwidth which is computed based on Sustainable Cell Rate (SCR), Peak Cell Rate (PCR) and Maximum Burst Size (MBS) of each individual connection. Weighting based on the effective bandwidth calculated by the Connection Admission Control (CAC) algorithm ensures that the bandwidth assigned to each service class is adequate to guarantee the QoS contract specified in the call setup.

Trip points are also used in the buffers to allow for priority queueing as well as CLP=1 dismodule capabilities. One trip point is configured in both the CBR and the UBR buffers such that cells with CLP=1 are dismoduleed after the queue exceeds the threshold specified by the trip point. Additionally, both VBR buffers (RT and NRT) each have four separate trip points which process as follows:

Table 13. Subclass QoS Thresholds for VBR Traffic

Thresholds >	**1**	**2**	**3**	**4**	**Queue Full**
Priority 1 VC	pass	pass	pass	drop CLP=1	drop all
Priority 2 VC	pass	pass	drop CLP=1	drop all	drop all
Priority 3 VC	pass	drop CLP=1	drop all	drop all	drop all
Priority 4 VC	drop CLP=1	drop all	drop all	drop all	drop all

All the trip points are software configurable.

Per VC or VP, any connection supports up to three Generic Call Rate Algorithms (GCRA) per port. There is support of the leaky bucket configuration per connection. There is a tagging option for one GCRA per connection. For each connection, there is bidirectional policing, meaning that the policing is separate for each end of the connection. All combinations of ATMF 3.1.are supported. A full set of dismodule, marked, and pass counts per VC/VP is kept.

Connection Admission Control (CAC) is necessary to guarantee QoS; to ensure that finite network resources, coupled with statistical multiplexing environment, does not lead to unacceptable levels of cell losses and cell delays.

Basic CAC principle is offline table generation that maps an *effective bandwidth* for the parameters of a call or PVC provision command, plus a metric of link state at each buffer site in the network. The effective bandwidth is checked against the remaining bandwidth and, if the route is suitable, the call is connected. Oversubscription is also supported.

3.3.3 2230 IOPs

The IOPs are the components that terminate external equipment and adapt cells to the SP. For each line type, grade of reliability, and speed of interface, there is an I/O processor (IOP) and I/O adapter pair (IOA). There are many functions that are specific to each processor type. This introductory section covers the common attributes of all IOMs.

The switch offers 1:1 redundancy at the IOM (line module) level.

The top-front of each IOP has GOOD and FAIL LED indicators, which are green and red, respectively. At the bottom are redundancy led pairs that show which IOP of a pair is online, if redundancy is provisioned. For each type of ATM interface, there is a red and a yellow alarm LED: the red indicates that the port is in a loss of signal or framing state. During power up, the LEDs assume states that reveal the progress of the power up.

The front of any IOM has a UART connection to a DB-9. This is a female DB-9 conforming to the PC standard. It runs at 19200 bps async. This connector runs service terminal applications. In debug or manufacturing modes, the debug console has specific capabilities. The IOP connects to the backplane.

Only medium-reach lasers are currently available for the switch IOMs. Medium reach has a 15 to 17 km long reach, with up to 40 km to be available in a future release.

In its maximum configuration, the IBM Nways 2230 ATM Model 650 offers a 16-slot chassis, with two slots reserved for SPs; therefore, 14 slots are available for IOMs.

The table below provides the maximum configurations for each supported interface on the switch:

Table 14. Port Density Per Interface

Interface Type	Ports/IOM	Speed/Port(Mbps)	Max Ports/Switch
E3	8	34.368	112
DS3	8	44.706	112
OC3	4	155.52	56
STM-1	4	155.52	56

Table 15. Interfaces Supported at Physical Layer

Interface Type	Ports/IOM	Physical Interface
E3	8	BNC
DS3	8	BNC
OC3	4	Single or Multimode Fiber
STM-1	4	Single or Multimode Fiber

8-Port DS3 ATM UNI IOP Module: This module provides eight 44.736 Mbps interfaces. Each port can be configured individually as a UNI interface, Interim Inter-Switch Signalling Protocol (IISP) port, direct trunk, or OPTimum cell trunk.

The port buffers (8 KB each), combined with the 128 KB programmable cell buffers in the SP quadplane, give this module the flexibility, performance, and data integrity required for high-speed networking.

Operational Features

- Eight DS3 ATM connections at wire speed
- ATM UNI 3.0/3.1 and ATM IISP 3.1 (also known as P-NNI Phase 0), direct trunking, and OPTimum cell trunking
- All 4 ATM classes of service
- Programmable QoS scheduler
- UNI 3.0/3.1 signalling, with high signalling throughput via an i960 signal processor
- UNI 3.0/3.1 cell-bearing DS3 physical interfaces that support PLCP and direct (HEC) mapping modes of operation
- 16000 Multicast source connections (global to switch)
- 16000 Unicast connections (combined VCCs and VPCs) per module, with 2000 connections per port
- An 8 Kbps output cell buffer per port
- Path switching based on virtual path
- Extensive Operation and Management (OAM) cell processing (F4, F5 flows, and fault management)
- Wide range of MIB support (ILMI, AToM, and DS3)
- Ingress Usage Parameter Control (UPC) for cell policing based on UNI 3.0/3.1 GCRA algorithm
- Statistics collection on ATM and physical layers
- Internal, loop-timed, and system-timing, transmit-clocking options

Interface Standards

- ITU G.703
- ANSI t1.102

Other Standards Supported - DS3

- ANSI T1E1.1/94-002R1, T1.107, T1.107a, T1.403
- ATM Forum UNI 3.0/3.1
- Bellcore TR-NWT 001112 and TR-TSY-000499
- ITU G.804
- RFC 1407
- TR54014 (AT&T ACCUNET T45 and T45R)

Physical Interfaces

- Eight ATM UNI 3.0/3.1 cell-bearing DS3 ports supporting C-bit/M-framing, PLCP per TR-TSY-000773, and direct cell mapping per C.804

- BNC connector per ANSI T1.104

8-Port E3 ATM UNI IOP Module: This module provides eight 34.368 Mbps interfaces. Each port can be configured individually as a UNI interface, Interim Inter-Switch Signalling Protocol (IISP) port, direct trunk, or OPTimum cell trunk.

The port buffers (8 Kbps each), combined with the 128 Kbps programmable cell buffers in the SP quadplane, give this module the flexibility, performance, and data integrity required for high-speed networking.

Operational Features

- Eight E3 ATM connections at wire speed
- ATM UNI 3.0/3.1 and ATM IISP 3.1 (also known as P-NNI Phase 0), direct trunking, and OPTimum cell trunking
- All 4 ATM classes of service
- Programmable QoS scheduler
- UNI 3.0/3.1 signalling, with high signalling throughput via an i960 signal processor
- UNI 3.0/3.1 cell-bearing E3 physical interfaces that support PLCP and direct (HEC) mapping modes of operation
- 16000 Multicast source connections (global to switch)
- 16000 Unicast connections (combined VCCs and VPCs) per module, with 2000 connections per port
- An 8 Kbps output cell buffer per port
- Path switching based on virtual path only
- Extensive Operation and Management (OAM) cell processing (F4, F5 flows, and fault management)
- Wide range of MIB support (ILMI, AToM)
- Ingress Usage Parameter Control (UPC) for cell policing based on UNI 3.0/3.1 GCRA algorithm
- Statistics collection on ATM and physical layers
- Internal, loop-timed, and system-timing, transmit-clocking options

Interface Standards

- ITU G.703
- ITU G.705

Other Standards Supported - E3

- ITU G.751
- ITU G.832
- ITU G.804
- ATM Forum 94-0406R4
- RFC 1407

Physical Interfaces

- 8 ATM UNI 3.0/3.1 cell-bearing E3 ports supporting C-bit/M-framing, and direct cell mapping per C.804

- BNC connector per ANSI T1.104

4-Port ATM UNI OC3c IOP Module: This module provides four 155.52 Mbps interfaces. Each port can be configured individually as a UNI interface, Interim Inter-Switch Signalling Protocol (IISP) port, direct trunk, or OPTimum cell trunk.

The OC3 module comes with either 8 Kbps or 24 Kbps cell buffers per port; the port buffers and the 128 Kbps cell buffers on the SP are based on the switch′s quadplane architecture.

Operational Features

- Four OC3 ATM connections at 155.52 Mbps
- Either 8 Kbps or 24 Kbps cell buffers per port
- ATM UNI 3.0/3.1 and ATM IISP 3.1 (also known as P-NNI Phase 0), direct trunking, and OPTimum cell trunking
- All 4 ATM classes of service
- UNI 3.0/3.1 signalling, with high signalling throughput via state-of-the-art hardware switching
- 16000 virtual circuits (combined VCCs and VPCs) per module
- Extensive Operation and Management (OAM) cell processing (F4, F5 flows, and fault management)
- Wide range of MIB support (ILMI, AToM)
- RFC 1595 provisioning and monitoring
- Path switching based on virtual paths
- Ingress Usage Parameter Control (UPC) for cell policing based on UNI 3.0/3.1 GCRA algorithm
- Egress Encapsulated Forward Congestion Indicator (EFCI) marking
- Both single and multimode fiber optic transceivers

Interface Standards

- ANSI T1.105
- ANSI T1.106

Physical Interfaces

- 4 ATM UNI 3.0/3.1 cell-bearing OC3 155.52 Mbps ports

Physical Connectors

- Subscriber Connector (SC)

Signal Distance/Levels (Single-Mode Laser Optics)

- Up to 15 kms
- TX Power: -8dBm, -15dBm
- RX Sensitivity: -8dBm, -28dBm

Signal Distance/Levels (Multimode LED Optics)

- Up to 2 kms
- TX Power: -14dBm, -20dBm
- RX Sensitivity: -14dBm, -29dBm

Other OC3 Standards

- ATM Forum UNI 3.0/3.1
- ANSI T1M1.3/92-005R1
- Bellcore TR-NWT-001112
- Bellcore GR-253-CORE
- RFC SONET 1595

4-Port ATM UNI STM-1 IOP Module: This module provides four 155.52 Mbps interfaces. Each port can be configured individually as a UNI interface, Interim InterSwitch Signalling Protocol (IISP) port, direct trunk, or OPTimum cell trunk.

The STM-1 module comes with either 8 Kbps or 24 Kbps cell buffers per port; the port buffers and the 128 Kbps cell buffers on the SP are based on the switch's quadplane architecture.

Operational Features

- Four OC3 ATM connections at 155.52 Mbps
- Either 8 Kbps or 24 Kbps cell buffers per port
- ATM UNI 3.0/3.1 and ATM IISP 3.1 (also known as P-NNI Phase 0), direct trunking, and OPTimum cell trunking
- All 4 ATM classes of service
- UNI 3.0/3.1 signalling, with high signalling throughput via state-of-the-art hardware switching
- 16000 virtual circuits (combined VCCs and VPCs) per module
- Path switching based on virtual paths
- Up to 16000 point-to-multipoint source connections
- Extensive Operation and Management (OAM) cell processing (F4, F5 flows, and fault management)
- Wide range of MIB support (ILMI, AToM)
- RFC 1595 provisioning and monitoring
- Ingress Usage Parameter Control (UPC) for cell policing based on UNI 3.0/3.1 GCRA algorithm
- Egress Encapsulated Forward Congestion Indicator (EFCI) marking
- Both single and multimode fiber optic transceivers

Interface Standards

- ANSI T1.105
- ANSI T1.106

Physical Interfaces

- 4 ATM UNI 3.0/3.1 cell-bearing OC3 155.52 Mbps ports

Physical Connectors

- Subscriber Connector (SC)

Signal Distance/Levels (Single-Mode Laser Optics)

- Up to 15 kms

- TX Power: -8dBm, -15dBm
- RX Sensitivity: -8dBm, -28dBm

Signal Distance/Levels (Multimode LED Optics)

- Up to 2 kms
- TX Power: -14dBm, -20dBm
- RX Sensitivity: -14dBm, -29dBm

Other STM-1 Standards

- ATM Forum UNI 3.0/3.1
- ANSI T1M1.3/92-005R1
- Bellcore TR-NWT-001112
- Bellcore GR-253-CORE
- RFC SONET 1595
- ITU G.707
- ITU G.708
- ITU G.709
- ITU G.783

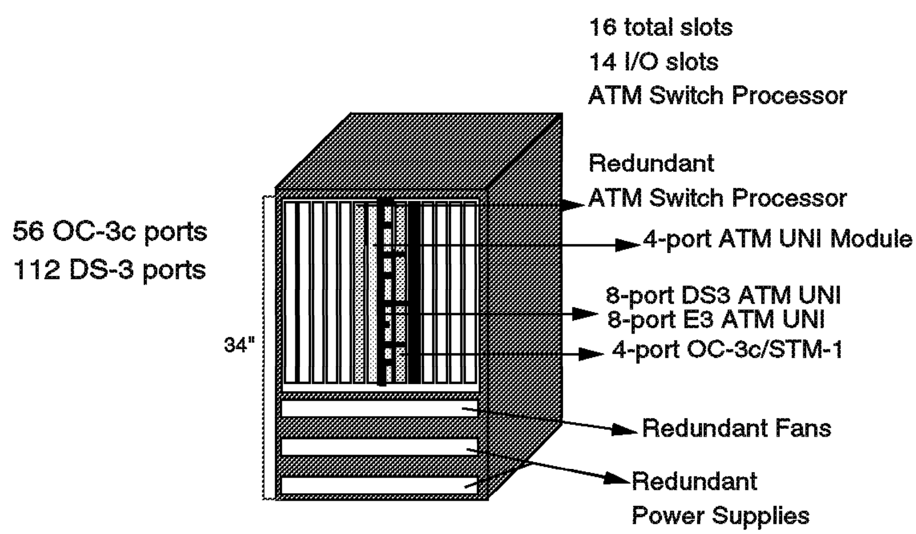

Figure 59. IBM Nways 2230 Model 650 Front View

3.3.4 Redundancy and Automatic Switchover

Redundancy is provided both at the common equipment level and at the I/O module level. Common equipment includes the processor/control module, the switching matrix and the power supplies and fans.

Redundancy of power supplies and fans is provided with instantaneous switchover. Additionally, when redundant power supplies and/or fans are installed, all such features work in load sharing mode. IBM Nways 2230 cooling system is separated into two distinct sections:

The I/O and CP modules are cooled by the fan tray. The cooling system is designed for fault tolerance; failure of a single fan will not cause an overheat condition. IBM provides an optional air filter, which may be added to the system. The fan tray can be serviced online and is accessible from the front panel of the enclosure.

Each power supply contains two fans, and optional air filters are accessible on the right side of the supply.

You can install up to two power supplies in the IBM Nways 2230 Switch. Redundant power supplies will operate in the load sharing mode; however, one power supply is fully capable of handling the load of the system. IBM Nways 2230 Switch can operate at both AC or DC line voltages. Like the power supplies, the line side of the system is redundant and independent. Failure of one side of the line will cause the shutdown of a single power supply. However, the other power supply will continue to run and is capable if carrying the full load.

The switch processor (SP) module supports redundancy through a hot-standby SP module to provide switchover in the case of fabric failure or any other failure on the SP primary module. A SP or I/O Processor (IOP) switchover may cause a momentary disruption of user traffic, but all established PVCs and SVCs will remain established.

The Redundancy Manager is a background task that runs on each SP module installed in the switch. It performs the following functions: establishes an active and redundant module for each pair of SPs; monitors the integrity of active modules by means of keep-alive messages over the cell bus; detects a failure in the active module and instructs the redundant module to disable the active module and perform a warm boot to become the active module; monitors the Parameter Random Access Memory (PRAM) in redundant modules to ensure identical configurations are maintained between the active and redundant module, and if a mismatch occurs, the active module automatically updates the PRAM in the redundant module.

If an active IOP fails, the redundant module reboots it and takes over as the active module. The changeover does not have any direct effect on other modules in the switch. They may, however, be indirectly affected by not being able to access information on the resetting module.

The redundancy scheme is not designed to provide a fault-tolerant system. If an SP fails, the changeover will cause loss of data, since data transfer flows through the SP. During the reboot, no new PVCs/SVCs can be established, no rerouting of existing PVCs/SVCs can occur, and no IP routing can occur through this node (that is, no network management traffic).

You can install IOPs in redundant pairs; however, not every IOPs can have a redundant one. Redundant modules have to share the same thread. The following redundant pairs are supported:

Slot 3 - Slot 4
Slot 5 - Slot 6
Slot 7 - Slot 8
Slot 11 - Slot 12
Slot 13 - Slot 14
Slot 15 - Slot 16

Table 16. IBM Nways 2230 Physical Specifications	
Specification	**Description**
Physical characteristics	Basic unit includes one power module, one cooling fan module, one SP module, and one SPA module mounted inside a chassis.
Size	48.26 cm (19.0 in) wide x 84.455 cm (33.25 in) high x 38.1 cm (15 in) deep
Weight	90.72 kgs (200 lbs) max
Thermal Dissipation	1400 WATTS maximum, 4433 BTU/hr AC; 1400 WATTS maximum, 4433 BTU/hr DC
Management Interfaces	Ethernet, RS-232

Table 17. IBM Nways 2230 Electronic Specifications	
Feature	**Specification**
90-132 VAC	12.0 amps maximum, 1400 WATTS maximum, 50-60 Hz, single phase
180-264 VAC	7.5 amps maximum, 1400 WATTS maximum, 50-60 Hz, single phase
-48 VDC	-48 to -76 VDC, 1400 WATTS maximum, 30 amps max

Table 18. IBM Nways 2230 Environmental Specifications	
Characteristic	**Requirement**
Ambient Operating Temperature	O to 50 degrees C
Relative Humidity	10% to 80% (noncondensing)
Operating Altitude	to 3050 meters (10000 feet)
Ambient Storage Temperature	-40 to +65 degrees C, 95% relative humidity
Storage Altitude	-305 to 9150 meters (-1000 to +30000 feet)

Chapter 4. Managing Networks Using Nways Wide Area Element Manager

This chapter provides detailed information on the Nways Wide Area Element Manager.

Generally, any type of wide area network requires a strong implementation of a management system. The management system should allow you to build a centralized or distributed network control environment, where all information such as configuration, performance, fault and billing information of the managed network should be concentrated. The IBM Nways Wide Area Element Manager fulfills these requirements.

The IBM Nways network management strategy is to provide an easy-to-use element management solution that can coexist with an enterprise manager and provide:

- A scalable architecture for future network growth
- A reliable network management environment
- Leverage other products for a total management solution

4.1 Network Management Standard Framework

The IBM Nways Wide Area Element Manager implementation fully complies with the industry management standards. The following management tasks are implemented:

- Configuration Management
- Performance Management including billing functions
- Fault Management and diagnostic procedures
- Security Management

In addition to these standard management tasks, the following are supported:

- Central code distribution
- Accounting and billing

The following sections briefly cover the implementation of these tasks in the IBM Nways Wide Area Element Manager.

4.1.1 Configuration Management

Initial network configuration or updates can be downloaded locally or remotely from the Nways Element Manager through the out-of-band method to battery-backed configuration memory in each switch.

The Nways Element Manager provides a logical network initialization procedure, from setting network wide parameters down to provisioning individual circuits in a stepwise manner. Operators use a series of pull-down menus provide theu user with a check lists for the proper configuration. Nways Element Manager provides default assignment of required parameters. It automatically prompts the operator with an error message to supply missing parameters, if necessary.

In addition, a template function allows the user to build a template to be used when programming similar parameters. These templates are saved in a file and can be retrieved through a menu.

The Nways Element Manager configuration support simplifies operator tasks wherever possible. For example, trunks and circuits can be defined with a single step even though these involve multiple network objects.

The Sybase relational database support provides summary reports of configuration information. Users can manipulate the information for troubleshooting, trending and historical analysis.

4.1.2 Fault Management

Nways Element Manager provides a suite of fault management facilities. Fault isolation is accomplished through user-initiated test sequences utilizing loopback tests at the individual port level. A variety of traps for alarm indications and statistics logging are available for all objects in the network of IBM 2219, 2225, and 2230 Nways wide area switches including trunks, physical and logical ports, and circuits.

4.1.2.1 Alarm Concept

Each Nways switch executes an SNMP agent, which provides an interface between the switch objects (that is, cards, logical ports, circuits, etc.) and the network management station. When an object of the switch changes its operational status or exceeds a predefined threshold, the switch SNMP agent automatically forwards the *alarm* to the network management station using standard SNMP TRAP messages. You can configure up to 16 Nways Element Manager stations or other SNMP management stations IP destination addresses for each switch in the network. Each switch can forward traps to the different management workstations. Transport delay for SNMP traps depends on many factors, such as network topology, network and management traffic load and the location of the NMS workstation and the switch which sent the TRAP. Alarms are not stored in the internal log of the IBM Nways switches; they are stored in the preconfigured NMS stations.

When the network administrator finds out that an object is in error inside the network, a SNMP poll and debug functions can be initiated to get more detailed information. The following figure illustrates the concept of SNMP traps and polling.

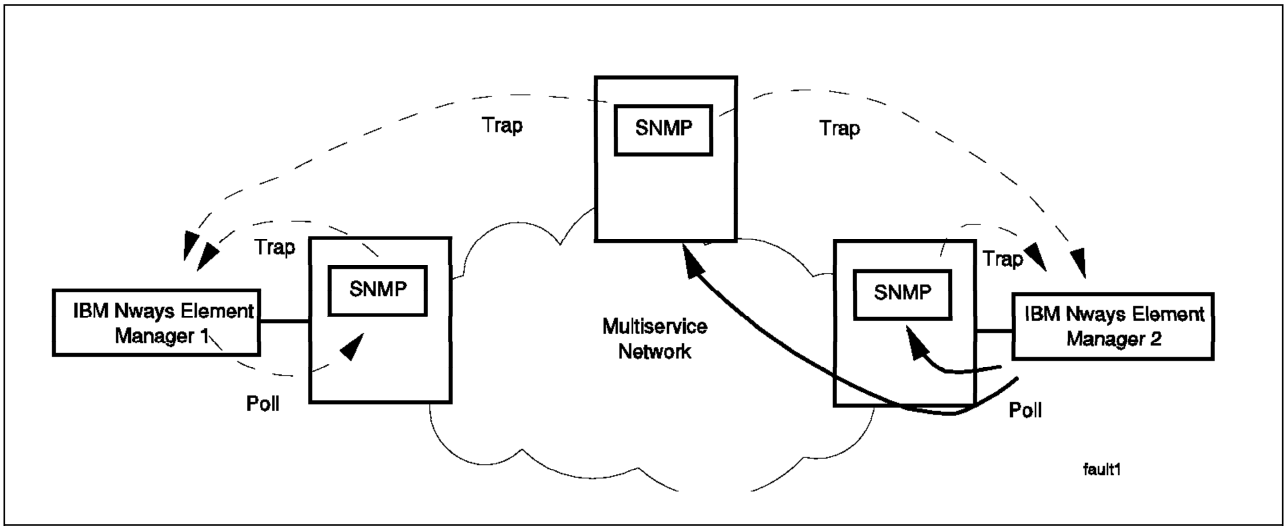

Figure 60. SNMP Trap and Poll Fault Management Concept

4.1.2.2 Alarm Processing

When the NMS workstation receives an SNMP trap, it processes it using the standard HP OpenView event application. Alarms are stored in the log file for future optional processing.

Dividing alarms into different categories enables you to look for a certain type of failure inside the network. Each alarm is stored with its source ID, the date, time stamp and its severity level. The alarm text describes the type of problem.

You can browse the alarm log using various types of filters. The topology map represents an alarm by changing the color of the affected object. Color change is propagated through the whole tree of views to enable quick identification of object(s) in error. Objects such as switch or card level LEDs are also graphically represented and can be observed for alarm conditions.

Additionally, the Nways Element Manager supports audible alarms to notify network operators about major failures in the network. You can specify which alarms are audible alarms.

4.1.2.3 OA&M Alarms Support

OA&M functions are supported only on the ATM cell-based interfaces. IBM Nways Switches support the OA&M (Operations administration and maintenance) cells. IBM implementation currently supports the following OA&M functions:

- Performance management flows are supported with congestion indications, traffic load indications, and utilization states.
- Generating and processing the VP and VC AIS OAM alarms cells.
- OA&M loopbacks.

OA&M alarms are used to signal a change of status of the circuit. IBM implementation supports OA&M AIS alarm for both virtual path and virtual circuit connections. You can either disable or enable processing of the OA&M cells per switch, per logical port or per virtual circuit.

The OA&M loopback feature enables you to verify the state of the circuits inside the IBM Nways Switch network. You can request loopback on any ATM VP or VC connection. Each request specifies the number of test cells which should be sent over the looped circuit. If you are running the loopback over the IBM Nways Switch network, you specify the number of hops which will be traversed by the test cells. Gradually increasing the number of hops enables you to test each segment of the connection. The IBM Nways manager provides you with statistical information: the number of received responses and the high, low and average response times.

4.1.2.4 IBM Nways Switches Diagnostics

Diagnostic functions are very important when operating any kind of WAN network. Remote diagnostics enable you to locate the failing part of the switch or isolate the failing connection inside the network. In addition to loopback diagnostics for frame relay PVC and ATM circuits, IBM Nways Switches support various diagnostic tests and physical port loopbacks. The following diagnostic tests are supported:

- IBM 2230 Foreground and Background diagnostics
- IBM 2219, IBM 2225 and IBM 2230 physical ports loopbacks

For more information about virtual circuits loopbacks, refer to 4.1.2.6, “IBM 2219, IBM 2225 and IBM 2230 Physical Ports Loopbacks” on page 114.

4.1.2.5 IBM 2230 Foreground and Background Diagnostics

IBM Nways 2230 Switch diagnostics help you to isolate hardware and software problems inside the network. Each I/O and CP module runs continuous *background diagnostics*. Background diagnostics monitor the switch components for potential failures or problems; they are executed automatically as a background process, in order not to interfere with switch operations. Background diagnostics send an alarm to the NMS workstation when one of the following problems occur:

- Corruption of different data structures
- Corruption of the code space

Alarms sent by the background diagnostic process are categorized into two groups: fatal errors and non-fatal errors; both types of alarms are automatically saved in the NMS workstation error log.

You can use foreground diagnostics either to test the problems that were indicated by the background diagnostics or to collect specific statistical data. You can run the foreground diagnostics to verify the status of the I/O module, physical/logical ports or to isolate the cause of a transmission error. There are two foreground diagnostic types available:

Internal test - Checks for the internal hardware or a physical port problem.

External test - Loops out the ports toward the connected equipment. This enables you to test the remote port′s ability to send and receive data.

The foreground diagnostic affects the user traffic on the tested I/O modules or ports.

4.1.2.6 IBM 2219, IBM 2225 and IBM 2230 Physical Ports Loopbacks

Various types of physical port loopbacks are available to test the integrity of the line. IBM Nways Switches provide you with statistics which help you identify errors on the physical connections.

4.1.3 Security Management

Nways Element Manager enables you to prevent unauthorized access to the network, providing protection of the NMS workstation from unauthorized access of the operating system. The following features are supported to protect your network:

- Three access levels to the NMS with different rights protected by passwords
- Auditing trail functions

4.1.3.1 Security Password Protection

Nways Element Manager supports secure access to the network through password protection. You can define the following three levels of passwords:

Administrator Password - User with this password is allowed to create new operators. You should define only one administrator per network.

Operator Password - Operator password allows you to manage and configure all features inside the network.

Provisioning Password - These users can configure only logical ports, trunks and virtual circuits and monitor the network.

None - If you do not have any password, you can open different views of the network; however, you cannot perform any operation.

4.1.3.2 Audit Trail

The Audit Trail function keeps a record of all changes which have occurred in the network. You can retrieve this information from the database. Each record has its date and time stamp and operator name, if applicable. The Audit Trail function logs the following network activities:

- Invalid log in.
- Log in or log off.
- Add, modify, or delete any object in the network, such as a switch, logical port or a circuit.
- Reboot of a switch or module.
- Download switch software, initialization of the script file or PRAM SYNC file.
- Standby module takes over in redundant pair.
- Add, delete or modify a Network Management Station path or entry.
- User session timeout.

Note: The Audit Trail does not contain any information about the bit error rate on the network trunks, status changes of circuits or trunks, etc. This type of information is available using the fault management functions of IBM Nways Switches.

4.1.4 Debugger Diagnostics

For more in-depth diagnostics capabilities, not possible from the NMS, you can log on to the IBM Nways Switch either using the Telnet application or via an ASCII terminal. The IBM Nways Switch allows you to issue many types of commands to examine the current status of the system. Using the debugger, you can dump memory locations, set and examine the parameters of the switch objects, etc. Debugger commands are out of the scope of this document.

4.1.5 Centralized Code Control

IBM Nways Wide Area Element Manager supports centralized code distribution. Each CP module has an image of operational code. If you want to upgrade the operational code to the new release, you overwrite the NVRAM with the new image of the code.

Using the IBM Nways Wide Area Element Manager, several SNMP variables are set to inform the CP to download the new version of the code. The code is downloaded using the TFTP protocol to the PP, CP, or SP which holds all operational codes for itself and all I/O modules. In case of a redundant configuration, the active module is downloaded first, then a backup module synchronizes its code image with a primary one. The IBM Nways 2219 and 2225 Switches store the operation code images in the processor module memory; the IBM Nways 2230 switch uses a PCMCIA hard disk.

You can use the IBM Nways Wide Area Element Manager to find out which release of code is running in a particular switch.

4.2 Flexible Network Management Access

Nways Element Manager manages the network of the IBM 2219, 2225, and 2230 Nways wide area switches in several ways, allowing the operator flexibility on how and from where the network is managed. Nways Element Manager supports in-band management via an Ethernet connection to a local switch or by locating the Nways Element Manager and the switch on the same IP network.

Alternatively, Nways Element Manager can access any of the IBM 2219, 2225, and 2230 Nways Wide Area Switches in a network out-of-band via a dial modem. Nways Element Manager supports both in-band and out-of-band network connections to ensure that the network nodes are always accessible.

4.2.1.1 Network Management Access to the IBM Nways 2219 and 2225 Switches

The IBM Nways Element Manager uses the IP protocol to transport the SNMP messages between the management station and the switches. To be able to manage your switches, you have to have IP connectivity to one of the nodes inside your network. This node provides the *gateway* between your network application and the managed nodes. IBM implementation supports multiple gateways and multiple management centers. This feature enables the creation of the backup management center. For more information, refer to 4.4.2, "Backup Workstation" on page 124.

IBM Nways 2219 and 2225 have implemented the IP protocol stack to be able to communicate with the workstation. The IP protocol stack is executed by the central processor in the IBM Nways 2225 and by the packet processor in the IBM Nways 2219 switch. These processor cards have the following types of IP interfaces:

Direct Ethernet Access - Your Management workstation is directly connected to the IBM Nways Switch. IBM Nways 2219 PP and IBM 2225 CP cards use the Ethernet AUI interface to connect to the Ethernet network. *Direct* means that no router exists between the management station and the gateway switch.

Indirect Ethernet Access - This configuration is very similar to the Direct Ethernet Access setup. The only difference is that between the management workstation and the gateway switch is some type of routed IP network. In this case, you define the IP static routes on both sides, the management workstation and the gateway switch. The following figure shows the access to the IBM Nways Switch network using the Ethernet connection to the processor cards.

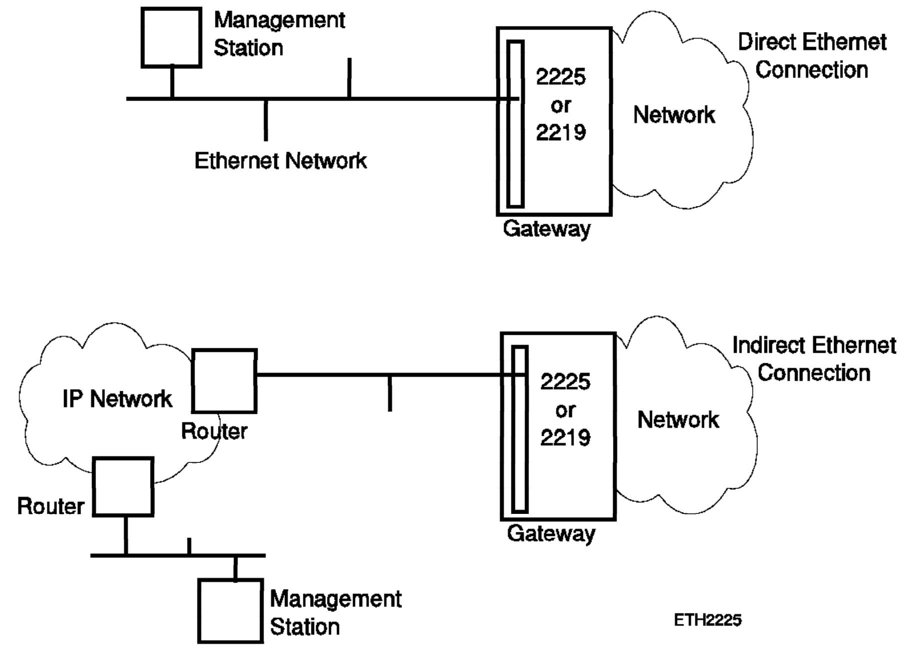

Figure 61. Connection Between the Network Management Station and Switches

SLIP Connection - Each processor card has one asynchronous serial interface, which can be used for SLIP connection. You can connect the network management station (NMS) either directly to the gateway switch using the null modem cable, or you can use dial-in asynchronous modems for remote node operation. SLIP connection should be used as a backup connection to a node, which is not able to communicate over the IBM Nways network. The following figure illustrates connection to the IBM Nways Switch using the SLIP protocol.

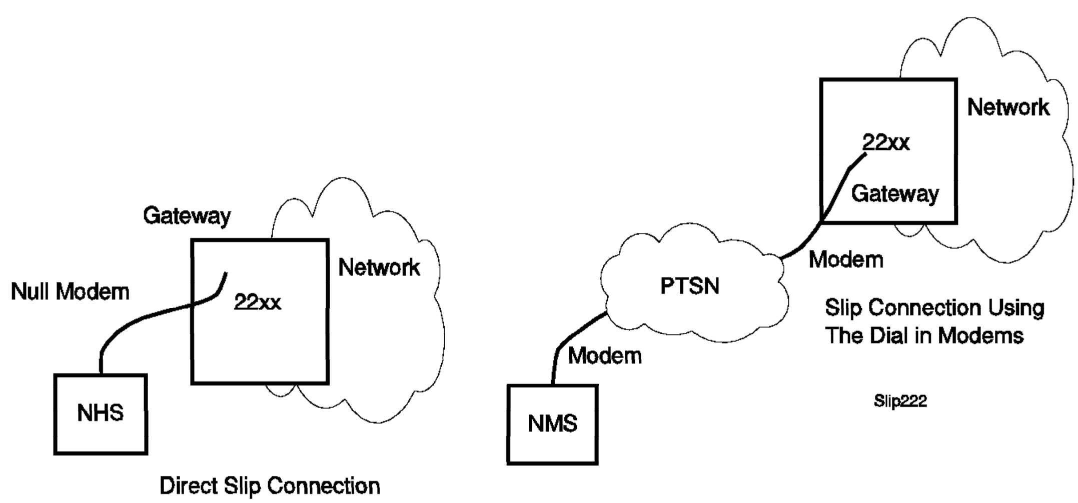

Figure 62. SLIP Connection to IBM 22xx Nways Switches

Management DLCI - CP and PP implement the frame relay encapsulation of the IP protocol according to the standard RFC 1490. You can reach any gateway node using frame relay and an external router with the frame relay interface. In this case, you define a DLCI between the router and the gateway node to carry all the management traffic. The router is connected to one of the frame relay UNI-DCE ports of the IBM Nways 2219 or 2225 Switch. Note that you are not able to configure the management DLCI for frame relay interworking; in other words, you cannot use an ATM-connected router to manage the IBM Nways Switch network. The following figure illustrates the connection of the management station using the frame relay management DLCI.

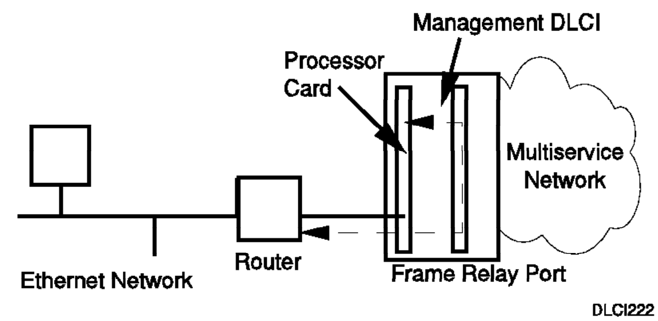

Figure 63. Connection of NMS Using the Management DLCI

Gateway node IP connection allows you to manage all IBM Nways 2219 and IBM 2225 Switches in the network. Each trunk has to dedicate 5% of the bandwidth for management traffic. Management traffic could be completely separated from the user data network by configuring dedicated management trunks.

4.2.1.2 Network Management Access to the IBM Nways 2230 Switch

IBM Nways 2230 Switch provides you with similar access options as the IBM Nways 2219 and 2225 Switches. You can use Ethernet PCMCIA 10BaseT adapter for LAN connection to the gateway switch. Both options, direct and indirect connection, are supported. Outbound access using the serial asynchronous port on the SP card is supported as well.

IBM Nways 2230 Switch supports direct access using the ATM network. A management workstation could be connected either directly to the gateway switch using the ATM network interface card or indirectly via the router with the ATM interface. IBM Nways 2230 Switch implements the Classical IP according to standard RFC 1577. The following figure illustrates the possible connection over the ATM network.

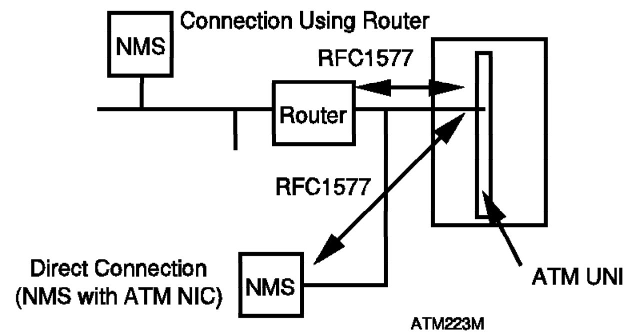

Figure 64. Management Access to the IBM 2230 Using an ATM Network

4.2.1.3 Management of the Hybrid Network

Today, implementation of the IBM Nways Manager allows you to manage all switches using one central management station. However, you have to define two different gateways, one for the IBM Nways 2230 Switch network and one for the IBM 2219 and IBM 2225 Nways switch network. No management flows exist between the IBM Nways 2225 and IBM 2230 switches. IBM Nways Wide Area Element Manager provides you with two different maps.

4.2.2 Accounting and Billing Management

Accounting and billing features are essential for public network operators; however,they could be useful in private networks as well. Accounting information in private networks will help the operator understand the traffic distribution between different departments in the organization and the distribution of the traffic inside the network. This type of information could be used for future network planning. Operators of the public network will use this information to charge network users.

4.2.2.1 Bulk Statistics

IBM switches implement local collection of statistical information, which can be transferred to the NMS workstation and stored in the bulk statistic database. You can specify which type of statistics should be collected by the switch. IBM switch collects requested statistical information and keeps it in memory. You can transfer these statistics as a raw statistics file in IBM Nways proprietary binary file. IBM Nways Wide Area Element Manager processes these files and stores the calculated values in the bulk statistics database. The following picture illustrates the bulk statistics collection mechanism.

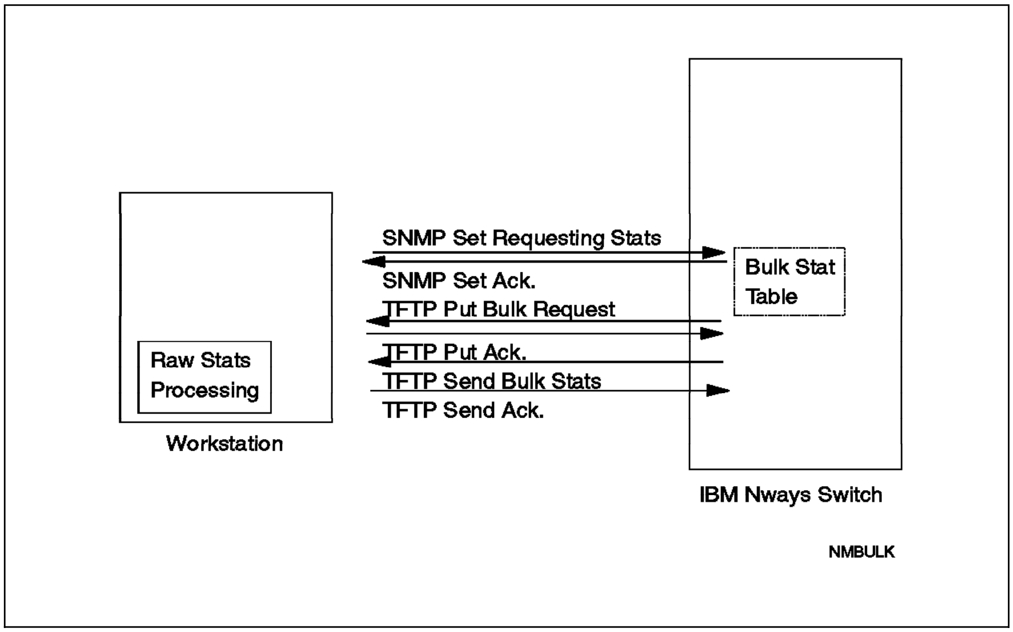

Figure 65. Bulk Statistics Collection

The bulk statistics collection is based on a distributed and self-contained approach. You can use parallel collections inside the IBM Nways Switch network; it is also possible to run many bulk statistics applications simultaneously.

Note that the bulk statistics use a different approach than the performance management, which collects the data using the SNMP polling. Current bulk statistics implementation provides you with gross statistics measured over 15-minute periods.

Release 4.1 bulk statistics provides you with a raw bulk statistics file, which is transferred to a workstation. The following picture illustrates the measured values.

PC timestamp

STDX timestamp

trunk count

ifIndex	ifOperStatus	ifSpeed	ifInOctets
ifInUcastPkts	ifInNUcastPkts	ifInDiscards	ifInErrors
ifInUnknownProtos	ifOutOctets	ifOutcastPkts	ifOutNUcastPkts
ifOutDiscards	ifOutErrors		
		ifIndex	ifOperStatus
ifSpeed	ifInOctets	ifInUcastPkts	ifInNUcastPkts
ifInDiscards	ifInErrors	ifInUnknownProbs	ifOutOctets
ifOutUcastPkts	ifOutNUcastPkts	ifOutDiscards	ifOutErrors

circuit count

cktSrcIfIndex	cktSrcDlci	cktVcState	cktPriority
cktOde	cktCir	cktBc	cktBe
cktInFrames	cktInDEFrames	cktInODEFrames	cktInFECNFrames
cktInBECNFrames	cktInDiscards	cktInOctets	cktInDEOctets
cktInODEOctets	cktOutFrames	cktOutDEFrames	cktOutODEFrames
cktOutFECNFrames	cktOutBECNFrames	cktOutOctets	cktOutDEOctets
cktOutODEOctets	cktOutLostFrames	cktOutLostDEFrames	cktOutLostODEFrames
cktOutLostOctets	cktOutLostDEOctets	cktOutLostODEOctets	cktRtMinDelay
cktRtMaxDelay	cktRtAvgDelay		
		cktSrcIfIndex	cktSrcDlci
cktVcState	cktPriority	cktOde	cktCir
cktBc	cktBe	cktInFrames	cktInDEFrames
cktInODEFrames	cktInFECNFrames	cktInBECNFrames	cktInDiscards
cktInOctets	cktInDEOctets	cktInODEOctets	cktOutFrames
cktOutDEFrames	cktOutODEFrames	cktOutFECNFrames	cktOutBECNFrames
cktOutLostDEFrames	cktOutLostODEFrames	cktOutLostOctets	cktOutLostDEOctets
cktOutLostODEOctets	cktRtMinDelay	cktRtMaxDelay	cktRtAvgDelay

NMBULK1

Figure 66. Bulk Statistics Measure Values

IBM has implemented a Raw Data Translator stand-alone application, which processes the bulk statistics retrieved from the switch. This application provides you with a format suitable for importing to spreadsheet or database. IBM Nways Wide Area Element Manager executes the Raw Data Translator each day at midnight to produce a comma-delimited ASCII stream and to import the data into the Sybase database. Bulk statistics are collected at 15-minute intervals. This is configurable.

IBM Nways Wide Area Element Manager billing system provides a network operator with standard accounting formats which comply with the automatic message accounting (AMA) formats defined by Bellcore Automatic Message Accounting Format (BAF) requirements standards. IBM has currently implemented the billing only for SMDS services; however, billing applications are under development.

IBM Nways Wide Area Element Manager distributed architecture includes distributed accounting servers. These servers interact with IBM Nways Switches to collect the accounting information for billing record generation. This distributed approach is sufficiently flexible and extensible to fit into any service provider billing system.

IBM Nways Wide Area Element Manager billing system supports or will support the following services:

SMDS Billing - based on the standards presented in Bellcore *TR-TSV-000775 Usage Measurement Generic Requirements in Support of Billing of Switched Multimegabit Data Service*.

ATM Billing - The first version of the billing system supports accounting and usage measurement for intranetwork ATM SVCs according to requirements derived from Bellcore *GR-1100-CORE*, as well as input from the ATM network service provider industry.

Future versions of the billing system will provide accounting and usage measurement functionality for internetwork circuits, frame relay to ATM interworking circuits, and PVC services on both the IBM Nways 2225 Switches and IBM Nways 2230 Switches.

Frame Relay Billing - IBM Nways 2225 Switches support collection of frame relay statistics for billing purposes. IBM Nways Wide Area Element Manager provides you with the following information in the billing records for frame relay:

- Source and destination DLCI address
- Number of successfully delivered frames
- Total byte count of the successfully delivered frame relay frames
- Start and end date and time stamps of measurement period

4.3 Nways Wide Area Element Manager

IBM Nways Wide Area Element Manager will provide support for frame relay, ISDN, SMDS and ATM network. You should not operate any of the IBM Nways switches without the network management station. IBM Nways Wide Area Element Manager does not only manage switches, but it provides all necessary microcode updates as well.

IBM has implemented the distributed version of the IBM Nways Wide Area Element Manager to enhance performance and provide operators with a scalable management solution. Release 2 of the IBM Nways Wide Area Element Manager is based on the fully distributed multiserver architecture.

IBM Nways Wide Area Element Manager Release 2 provides network operators with a published application programming interface (API) to enable them to customize the network management according to their needs. The *distributed architecture* enables you to distribute management functions across several workstations. This approach increases flexibility and scalability when building large IBM Nways Switch networks. Operators can distribute management functions to fit their needs. You can implement the management of the network based on one UNIX workstation and later on, when the network grows, distribute management functions over more workstations.

4.4 IBM Nways Wide Area Element Manager Requirements

This section provides a summary of IBM Nways Wide Area Element Manager hardware and software requirements.

4.4.1 Hardware Requirements

IBM recommends that a workstation be dedicated to supporting the Nways Element Manager. It is possible, however, to add Nways Element Manager to an existing HP OpenView platform. The following text details how to calculate the amount of memory and the swapfile and hard drive size required.

The following hardware specification is required for a single workstation architecture where HP OpenView, Nways Element Manager, and Sybase reside on the same platform. This solution is recommended for up to 10 users.

- 17-inch or 20-inch monitor
- 1/4-inch tape drive
- CD-ROM drive
- 3.5-inch diskette drive

For 1-10 switches and up to three users:

- Sun SPARC 5 Model 110
- 1.5 GB hard drive
- 96 MB RAM

For 11-30 switches and up to 10 users:

- Sun SPARC 20 Model 71
- 2.1 GB hard drive
- 128 MB RAM

For 30+ switches and up to 10 users:

- Sun SPARC 20 Model 151 or Model HS11
- 2.1 GB hard drive
- 256 MB RAM

Two workstations (HP Server and Sybase server): The following hardware specification is required if the number of remote users is greater than 10. This assumes two workstations: one an HP Server where HP OpenView and Nways Element Manager reside and a second called a Sybase server where HP OpenView, Sybase and Nways Element Manager reside. In addition, if the customer foresees deploying multiple NMS workstations, it is recommended that Sybase reside on a separate platform.

Note: The CNM Proxy Agent support can be run on the Sybase server; it cannot be run on the HP Server.

Common requirements for the HP Server include the following:

- 17-inch or 20-inch monitor
- 2.1 GB hard drive
- 1/4-inch tape drive
- CD-ROM drive
- 3.5-inch diskette drive

For up to 25 users:

- Sun SPARC 20 Model 151 or Model HS11
- 128 MB RAM

For 25-50 users:

- Sun SPARC 20 Model 151 or Model HS11 or SPARC 1000E
- 256 MB RAM

Common requirements for Sybase Server

- 17-inch or 20-inch monitor (if 17-inch use S17)
- 2.05 GB hard drive
- 256 MB RAM
- 1/4-inch tape drive
- CD-ROM drive

X-Terminal remote stations: It is recommended that X-Terminals which are to be used to display HP OpenView have 4 MB of standard user memory and 3 MB of video memory.

Modem attachment: A V.32 modem (internal or external) is required. The customer must also provide any necessary modem cables.

Ethernet attachment: If local Ethernet connectivity to the Network Management Station is desired, then an Ethernet adapter and cabling is also required.

4.4.2 Backup Workstation

It is also recommended that the customer purchase a workstation for backup that has the same configuration as the online system. The database can be backed up so that if there is a hardware failure on the workstation, the system can be quickly reloaded. This backup system can not be online at the same time as the primary system.

4.4.3 Software Requirements

The Nways Element Manager requires:

- Sun Solaris 2.4 Operating Environment and patches (Sun will advise on patches required depending on the platform.)
- Sun Motif Windows Manager 1.2.4
- Sybase Open Server Version 4.9.2 (Note: If using two workstations as specified in 4.4.1, "Hardware Requirements" on page 123, then the HP Server does not require this Sybase support, the Sybase server does.)
- HP OpenView Network Node Manager Version 3.31 (includes SNMP Management Platform software)

X-Terminal remote stations require X11 Release 5.

The maximum number of hops in the network is limited to 16. This includes hops for any routers between the Network Management station and the gateway in one of the IBM 2219, 2225, and 2230 Nways wide area switches.

The number of X-Terminals remotely accessing Nways Element Manager is limited to 50.

Network Management traffic through the gateway should be kept at 5% or lower utilization of the processor.

For remote workstations running X-Windows, it is recommended that Solaris and Motif be installed so that less processing power is required when accessing the server.

4.4.4 Planning Information

Nways Element Manager (5765-A16) is distributed in one package with the following materials:

- One QIC (1/4-inch) tape with the product image
- International Program License Agreement
- International Program License Pointer Sheet
- Proof of Entitlement

4.4.5 Security, Audibility, and Control

Nways Element Manager products use the security and audibility features of the Sun SPARC hardware and Solaris and HP OpenView software.

User management is responsible for evaluation, selection, and implementation of security features, administrative procedures, and appropriate controls in application systems and communication facilities.

Appendix A. Release 4.2 Preview for the Nways 2225

IBM plans to implement several major enhancements to the Nways 2225 switch code. These enhancements will be implemented mainly by an IBM Nways switch code upgrade; however, new I/O modules may be available as well.

Note: At the time this redbook was written there was no information on Release 2.0 for the Nways 2230 switch. These features discussed in this section are unannounced. Do not commit to customers until formally announced. Please refer to your announcement letters. The following software enhancements are included in the Release 4.2 for the 2225:

- OSPF and rerouting algorithm enhancements
- Internal IP address scheme enhancement
- Frame relay enhancements
- IBM Nways 2225 and 2230 switches interconnection enhancements
- SMDS enhancements
- ISDN enhancements
- Network Management enhancements

The following new I/O modules are included with the Release 4.2 of code:

- ATM UNI cell switching module
- ATM Interworking unit I/O module

A.1 OSPF and Rerouting Algorithm Enhancements

IBM has implemented OSPF and rerouting algorithm enhancements with the following major goals:

- Improved rerouting policy in the IBM Nways switch network
- Improved performance of the OSPF algorithm
- Add new metric for frame relay to improve efficient handling of delay sensitive traffic such as voice
- Enable the exchange of routing information between IBM Nways 2225 and IBM Nways 2230 switches

A.1.1 Distributed OSPF Implementation

Current implementation of OSPF uses only the central processor module of the IBM Nways 2219 and 2225 switches. The IBM Nways 2225 switch I/O modules are equipped with powerful Intel 960 processors. This architecture enables the distribution of routing intelligence to each I/O module in the IBM Nways 2225 switch.

Each I/O module in intermediate nodes will be able to take a routing decision for management traffic, reducing the load on the central processor. Under this distributed architecture, each I/O module will have its own network topology map, and will be able to make routing decisions for new circuits without consulting a centralized resource. OSPF link state advertisement (LSA) is aggregated and propagated only by the central processor. This approach limits the network overhead traffic.

A.1.2 OSPF Delay Metric

IBM will add new OSPF routing metrics, which will provide new QoS parameters for frame relay traffic. Currently, Release 4.1 uses the following metrics:

- Trunk virtual bandwidth
- Trunk administrative cost

When you configure a delay sensitive frame relay PVC for voice traffic, IBM Nways switch routing algorithm chooses the best path over the network based on virtual bandwidth availability and administration costs. After establishing a circuit over such a path, the priority queues are used to minimize delay. This approach is not optimal from a delay standpoint; the path with the higher amount of virtual bandwidth could introduce higher delay than other paths.

The new OSPF metric adds a delay metric, which consists of both propagation and transmission delays. Delay is automatically measured by the trunk protocol during the trunk initialization, so it cannot include delays introduced by frame queuing and trunk congestion. The value of this metric is symmetric.

The OSPF algorithm uses one or more metrics for each trunk. This implementation enables better routing decisions for different types of traffic.

Release 4.2 will enable you to choose one of the following metrics, which will be used for routing a particular virtual circuit:

- Administrative Cost
- Propagation Delay

The VC manager uses the CAC algorithm to satisfy bandwidth requirement for a virtual circuit. However, the routing metric used to find out the optimal path will be either administrative cost or propagation delay. This new approach will change the load balancing algorithm. The number of reroutes to achieve optimal distribution will be minimized as the available bandwidth metric changes are not considered. The circuits will be rerouted according to the following algorithm assuming the routing priority vector equal for all VCs in the network.

- If the present path does not meet the QoS and there is a better path, then reroute.
- If the present path meets the QoS and there is a better path, then reroute.
- Otherwise, do not reroute.

Note: When using Release 4.2 software, you will not be able to configure the reroute rules such a Negative-to-Positive, Unrestricted and Negative-to-Anything.

A.1.3 Classes of Routing Priority

IBM Nways switches, running operating software prior to Release 4.2, reroute all virtual circuits without any distinction. In case of trunk failure, all circuits are rerouted over the backup path, even if there is no sufficient bandwidth to satisfy requested QoS for the rerouted virtual circuits. The virtual circuits with mission critical traffic will be affected by the virtual circuits carrying less important data.

IBM will implement a priority rerouting algorithm providing distinction between virtual circuits according to their importance. A PVC with higher configured priority will select the more optimal path which satisfies its QoS. A higher priority VC can force lower priority traffic away from optimal path to suboptimal paths. This approach will ensure that higher priority VCs will always be given preference over lower priority VCs. Such preference will result in providing better QoS to higher priority VCs.

The VCs are divided into different groups according to their priority. IBM 2225 switch will try to satisfy requested QoS for groups of VCs with higher priority. The priorities are specified by a pair of integer numbers:

- Bandwidth priority - can be specified in range from 0 to 15 (0 = highest)
- Preempting priority - can be specified in range from 0 to 15

The combination of these two priorities reflects the importance of the the VC. If you configure for all VCs the priority vector (0,0), the network operates in a manner identical to prior software releases. A network which contains a mixture of priority levels performs additional processing, to ensure that higher priority VCs always get the lower cost path through the network. The full bandwidth is guaranteed whenever possible, even at the cost of lower priority VCs.

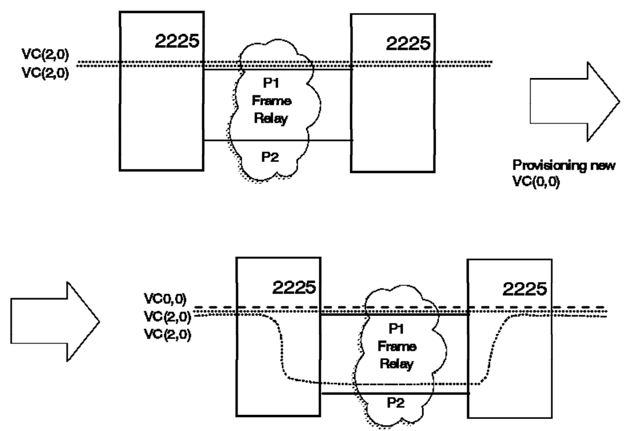

Figure 67. Priority Routing in the IBM Nways 2225 Network

Consider the example illustrated in Figure 67. There are two paths P1 and P2 between the IBM Nways 2225 switches. The cost of path P1 is equal to 100 and the cost of the path P2 is equal to 200. Assuming that the virtual bandwidth of path P1 is entirely consumed by the VCs with routing priority vector (2,0), no other VCs would be allowed over path P1. The new VCs would be normally forced to use the higher cost path P2, which may involve higher delays and more hops. Suppose a user is provisioning a VC with routing priority vector (0,0) between the switch A and B. The system will detect that path P1 is entirely populated by the virtual circuits with priority routing vector (2,0). The switch running the Release 4.2 software will force enough (2,0) VCs from the P1 path to ensure that the P1 path has enough bandwidth to satisfy QoS of the VC with the priority routing vector (0,0). As a result, some (2,0) priority VCs will be forced onto the P2 path. Note that this algorithm is not limited to the described example. When the network topology is more complex, IBM Nways switch priority rerouting algorithm consider multiple paths over multiple hops.

A.1.3.1 Priority Routing Rules

IBM Nways switch will implement the following rules to preserve routing priority inside the network:

New circuits provisioning rules - A new provisioned VC with higher routing priority will select a cost optimal path which satisfies its QoS. The VCs with lower bandwidth priority on the selected path will be ignored. Higher bandwidth priority VCs will force the lower bandwidth priority VCs from their current path until the available bandwidth is enough to accommodate their QoS and is positive. VCs forced from their path due to the bandwidth priority are chosen in the following order:

- The virtual circuits with lowest preempting priority are chosen first.
- The group of virtual circuits of the same preempting priority are forced from their path according to their configured CIR (CIR reflecting equivalent bandwidth). The VCs with higher CIR are chosen first.
- Within a group of VCs with same CIR and preempting priority, VC identifier is used.

Trunk Failure Recovery rule - IBM Nways switches will try to reroute the VCs in response to trunk failure. The IBM Nways switches will still allow a trunk to reach a negative virtual bandwidth. However, you should specify which VCs will be kept in the network, using the routing priorities. The following priority rules are used:

- A higher priority VC will choose an optimal path without taking into account the bandwidth consumed by lower bandwidth priority VCs. These could be forced onto non-optimal paths based on the rules used for new virtual circuit provisioning.
- Lower bandwidth priority VCs will not be allowed to traverse a trunk, where at least one higher bandwidth priority VC is present, and the virtual bandwidth would become negative.
- The virtual trunk bandwidth could become negative if there are no more VCs to force off the trunk.

Balance Rerouting

IBM Nways switches periodically test the route for each VC. The new virtual circuits provisioning rules are used for load balancing.

A.2 Frame Relay Enhancements

Release 4.2 will add several new functions, enabling utilization of network bandwidth more efficiently. Congestion management enhancements will provide improved services to well-behaved users of the network. The following improvements are planned:

- Asymmetric CIR
- Application or Customer specific routes
- Closed Loop Congestion Management
- Configurable Congestion Thresholds

A.2.1 Asymmetric CIR

The previous releases of IBM Nways 2225 switch code used the same CIR in both circuit directions. This implementation is far from optimal when dealing with applications with asymmetric bandwidth demand. Release 4.2 will enable you to specify different CIRs for each direction of the VC. This feature will give you higher degree of control over the network traffic, and allow you to use the network bandwidth more efficiently.

A.2.2 Application and Customer Specific Routes (ASR/CSR)

ASR/CSR feature will enable you to ensure that only selected PVCs will utilize specific trunks through the network. Application specific routes are defined so that only certain virtual circuits are allowed to use reserved trunks.

A.2.3 Closed Loop Congestion Management

Congestion Management implementation prior to the Release 4.2 has used the BECN and FECN bits to avoid the congestion in the network. The IBM Nways switch sets the FECN and BECN bits for all frames traversing the congested trunk. User devices, which receive these frames, are supposed to reduce the rate at which they inject data into network. Unfortunately, many devices ignore FECN and BECN bits; these devices continue to overload the network, affecting the traffic of well-behaved devices, which reduce their transmission rates. This leads to the unfair bandwidth distribution, with ill-behaved devices monopolizing the network and causing possible frame discards at congested resources.

IBM will implement a close loop congestion management algorithm, which will reduce the rate of excess data at the ingress port during the congestion in the network. The reduction in excess data will be done proportionally to the preconfigured Be value of a PVC. In this case, PVC traffic is reduced more fairly. The close loop congestion management algorithm was designed with the following major objectives:

- Provide a fair reduction of the excess bandwidth for the end-user devices based on the configured Be.
- Reduce the excess traffic from ill-behaved end devices.
- Guarantee data transfer for all end devices operating within the CIR and Bc without discard due to congestion.

Implementation of the closed loop congestion algorithm includes two parts:

- **Congestion OSPF link state advertisement** - The OSPF agent monitors each *N* second the congestion state of all switch trunks. If one or more links become congested, OSPF sends a link state advertisement to all other nodes in the network. Conversely, if the trunk changes its status from congested to noncongested for *Nc* seconds, OSPF sends a LSA indicating new state of the trunk. The OSPF algorithm updates its routing table according to congested status of the network trunks. The *N* and *Nc* timers will be configurable per trunk.
- **Reducing the PVCs rate at the ingress ports** - The IBM Nways 2225 switch examines all ingress PVCs. If it finds a PVC routed through one or more congested trunks, then the ingress port reduces its excess traffic flow into the network. The amount of the reduction is based on the worst congestion state of all trunks in the path.

A.2.4 Configurable Congestion Threshold

Release 4.1 used a hard coded congestion threshold for congestion determination. Release 4.2 will enable you to configure the following threshold parameters

- Mild congestion threshold
- Severe congestion threshold
- Absolute congestion threshold

This feature provides you with higher flexibility in determining which conditions will cause traffic discarding.

Based on the card type, the threshold cannot exceed a maximum value. You have to configure the threshold to meet the following criteria: Mild<Severe<Absolute. You can configure the threshold per logical port; the following table illustrates the maximum value for different type of I/O modules:

Table 19. Maximum Threshold Values

Card	Maximum value allowed per Lport	
	56 Byte buffers	**Bytes**
8 Port UIO	5450	305200
10 Port DSX	4668	261408
4 Port Chan T1/T1_PRI	225	12600
4 Port Chan E1/E1_PRI	174	9744
4 Port UnChan T1	5408	302848
4 Port UnChan E1	5408	302848
2 Port HSSI	23632	1323392
1 Port ATM UNI	60799	3404744

A.3 SMDS Enhancements

IBM will improve the IBM Nways 2225 switch SMDS services. The following functions are planned for new code release:

- Delivering of the group address datagrams
- Traffic management
- SMDS billing
- SMDS addressing

A.3.1 Delivering of SMDS Group Address Datagrams

Prior to Release 4.2 all group addressed datagrams were sent to each switch in the same OSPF domain, whether or not it was providing SMDS services.

New implementation will send the frame only to these switches, which are providing SMDS services to the end users, and who are members of the group address. Throughput of the network is improved by eliminating unnecessary group address traffic.

A.3.2 Configurable Traffic Management

Previous releases transport all SMDS datagrams as frames with fixed priority=3 and color=green. New implementation will enable you to choose a priority and color of datagrams received on a particular SMDS logical port.

A.3.3 Addressing Improvements

Country Code Address Prefix - Support for and International E.164 address prefix of up to 4 digits will be added. This feature allows extension of IBM Nways 2225 switch SMDS networks across international borders.

16 Digit Address Mask - The SMDS address mask is now configurable in the range from 6 to 16 digits of E.164 address. This enhancement applies to both individual addresses and group addresses.

A.4 IBM Nways 2225 and 2230 Interconnection Improvement

The Release 4.2 code for IBM Nways 2225 switches will implement improvements to the Release 4.1 IBM Nways 2225 and 2230 interconnection scenario. IBM has focused on the single routing OSPF domain and the ability to use cell trunks between the IBM Nways 2225 and 2230 switches. To achieve this goal, the following function was added to the IBM Broadband Switch platform:

Single OSPF domain - IBM Nways 2225 and 2230 switches will use a compatible OSPF routing algorithm and be able to exchange the topology information. Both switches are in a single OSPF routing domain. The following routing policies are implemented:

- Frame traffic can be established or rerouted over frame and cell-based trunks
- Cell traffic can be rerouted only over cell-based trunks.

Single PVC Establishment - Single configuration action will configure the PVC traversing both IBM Nways 2225 and 2230 network. You do not have to configure separate PVCs for IBM Nways 2225 and 2230 network.

ATM Direct Cell Trunk - The IBM Nways 2225 ATM UNI module is upgraded to support direct cell trunk logical port configuration. This feature integrates IBM Nways 2225 and IBM Nways 2230 via using a common trunking protocol in both networks. ATM Direct Cell Trunk supports the routing exchange protocol, which is necessary to integrate both networks into one OSPF routing domain.

Single Map for management - With the support for new trunks, all IBM Nways 2225 and IBM 2230 switches will be manageable from a single map.

A.5 Release 4.2 Class B Addressing

2.8.12, "IBM Nways WAN Switches IP Addressing" on page 71 briefly describes the IP addressing in Release 4.1, which requires all switches to use network mask 255.255.255.0, which limits the size of the network up to 255 nodes. Release 4.2 supports the subnet mask 255.255.0.0, if a class B IP network is to be used. Static mapping between the host part of the address and the node ID allows the use of 16-bit node IDs. The Release 4.2 will support a much larger number of switches in a single network.

Multiple Class C addresses are also supported.

A.6 Hardware Enhancements

IBM has developed new hardware I/O modules providing higher flexibility when interconnecting the IBM Nways 2225 to the IBM Nways 2230 switch and adding the ATM switching functions to the IBM Nways 2225 switch platform. The following modules are supported:

- ATM UNI Cell Switching I/O module
- ATM Interworking Unit OC3 I/O module

A.6.1 ATM UNI Cell Switching I/O Module

The ATM cell switching module is a one port DS3 card supporting the current frame-to-cell capability of the ATM DS3 I/O module, including:

- Frame relay/ATM network interworking
- Frame relay/ATM service interworking
- ATM OPTimum trunking
- ATM direct cell trunk

In addition to these functions, the ATM cell switching module will support pure cell switching. Both ATM UNI DTE and DCE logical port types are supported. Using this new module, you can connect the ATM end device such a router directly to the IBM Nways 2225 switch.

I/O module support the CBR and VBR class of service using the segmentation and reassembly via AAL-1 and AAL-5. IBM Nways 2225 switch can perform the ATM cell switching in addition to interworking functions.

A.6.2 ATM Interworking Unit I/O Module

The ATM Interworking Unit I/O Module supports one optical fiber (OC3c/STM1) physical port at speeds up to 155 Mbps. This I/O module supports the same functions as the current ATM UNI DS3 module, including:

- Frame relay/ATM network interworking
- Frame relay/ATM service interworking
- ATM OPTimum trunking
- ATM direct cell trunk

The I/O module supports CBR and VBR classes of service using the segmentation and reassembly via AAL-1 and AAL-5. This module will be mostly used for the high speed trunk connectivity to the IBM Nways 2225 switch, in cases where the IBM Nways 2225 and 2230 switches are collocated.

A.7 IBM Nways Wide Area Element Manager Enhancements

The IBM Nways Wide Area Element Manager will add new features, improving fault resiliency. The following features should be added:

- Common IBM Nways Wide Area Element Manager for IBM Nways 2225 and IBM Nways 2230 switches.
- Automatic database replication
- Bulk statistics improvement

A.7.1 Common IBM Nways Wide Area Element Manager for the IBM Nways WAN Switches

The IBM Nways Wide Area Element Manager integrates the Nways 2219, 2225, and 2230 switches into a single topology map. New PVCs traversing mixed network can be configured via single action. For more information refer to A.4, "IBM Nways 2225 and 2230 Interconnection Improvement" on page 133.

A.8 ISDN Enhancements

The following features will be added to the ISDN access support:

- ISDN user authentication
- EuroISDN and Australia ISDN versions
- Multilink PPP

A.8.1 ISDN Authentication

Previous ISDN implementation uses only the calling address verification for user authentication. Any end-user device using the same calling number was able to connect to the IBM Nways multiservice switch network. Release 4.2 allows ISDN dialup user authentication via password protection. IBM Nways 2225 switches implements the following authentication services:

- Password Authentication Protocol (PAP)
- Challenge Handshake Authentication Protocol (CHAP)
- Remote Authentication Dial-In User Service (RADIUS)

The authentication feature assumes that the end-user ISDN device uses PPP protocol to communicate with the IBM Nways multiservice switch network. During the PPP negotiation between the end-user device and switch, an authentication method (PAP or CHAP) is negotiated. The end-user device identifies itself to the Nways 2225 switch using PAP or CHAP protocol. Somewhere in the network is a RADIUS Authentication server which is used to authenticate a user upon establishment of the connection. When a user logs into the network, information about the user is sent to the RADIUS Authentication server. If the authentication fails, the connection is terminated. The following picture illustrates usage of the RADIUS server.

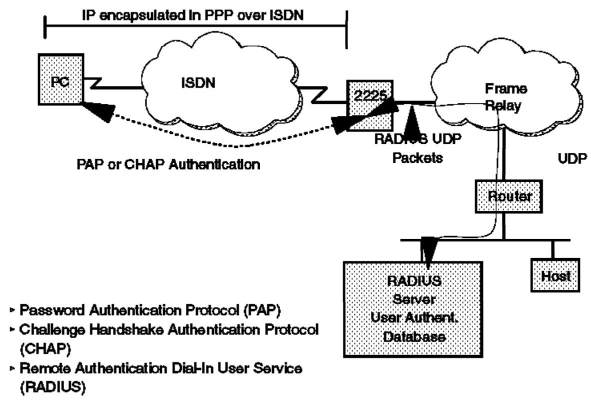

Figure 68. Usage of the RADIUS Authentication Server

A.8.1.1 PAP Authentication

When the PAP authentication method is used, the user sends to the IBM Nways 2225 switch its ID and password. The switch in turn sends this information to the RADIUS Authentication server using a RADIUS UDP packet. The RADIUS server searches its user database for end-user authentication. Based on the search result, the RADIUS server sends either *fail* or *pass* message to the switch, which either terminates or processes the connection.

A.8.1.2 CHAP Authentication

The challenge-handshake authentication protocol implements a higher level of security then the PAP authentication method. The PAP sends a user ID and password unencrypted over a ISDN network. The CHAP uses an encryption, so the password itself is protected.

When the user tries to logon to the network, an ingress switch generates a string of random characters, which is sent to the end-user device. The user hashes the string using the a password and sends the hash value back to the switch. IBM Nways 2225 switch then sends both the original and the hashed value via RADIUS UDP packet to the RADIUS Authentication server. The RADIUS Authentication server uses its copy of user password to hash the original string. If the result matches with the hash string value received from the end-user device, the user passes authentication, otherwise the connection is terminated.

Note: RADIUS UDP packets are encrypted using a secret hash key, in both PAP and CHAP.

A.8.2 Multilink PPP

Multilink PPP allows you to combine multiple PPP links into one logical data pipe. This feature provides you with greater throughput for PPP connections by combining the bandwidth of several physical lines. The multilink operation coordinates multiple independent links between two end devices and IBM Nways 2225 switch, to provide a virtual link with greater bandwidth.

The datagrams between the two multilink connected systems are split between the different physical lines. The multilink PPP was originally designed to exploit multiple ISDN channels, but this feature is applicable to any situation in which multiple PPP links connect two systems such as routers. The datagrams are fragmented and sent over multiple physical link. The receiving system reassembles the multilink PPP packets using the sequencing information in the multilink packet headers.

Assuming that there are no PPP connections between the end device and the IBM Nways 2225 switch, the multilink PPP bundles are created as follows:

1. The end-user system tries to establish first connection with IBM Nways 2225 switch. During the PPP negotiations the end system indicates that it supports multilink PPP and specifies an endpoint discriminator, which will be used in determining the appropriate multilink PPP bundle ID. The ISDN I/O module also supports multilink PPP.
2. After the initial PPP negotiation, if the authentication is enabled for current logical port, PAP or CHAP RADIUS authentication is performed. If the authentication fails, the connection is terminated.
3. The ISDN I/O module determines the multilink PPP bundle ID. If the authentication methods are disabled, the switch uses the endpoint discriminator as a multilink bundle ID. If an authentication method is used, the multilink bundle ID is a combination of both endpoint discriminator and the authentication end-user device ID.
4. If the the end-user device establishes a second call to the same switch, the same procedure is used to calculate the multilink PPP bundle ID. If the ID matches the one calculated for the firt call, the PPP link is combined with the first one.

A.8.3 ISDN Standards Enhancements

IBM will add ISDN support for the E1 interface card on the ISDN PRI module. This feature enables you to use the ISDN in Europe and Australia. The following types of the ISDN signaling was added:

- EuroISDN
- Australian signalling standard

Appendix B. Frame Relay Supported Standards

Frame Relay	Compliance
FRF.1 User-to-Network Implementation Agreement (UNI)	
ANSI T1.403	COMPLIANT
CCITT Recommendation V.35	COMPLIANT
CCITT Recommendation G.703 (2048 Kbps)	COMPLIANT
CCITT Recommendation G.704 (2048 Kbps)	COMPLIANT
CCITT Recommendation X.21	COMPLIANT
Data Transfer	COMPLIANT
Permanent Virtual Connection (PVC) Procedures	COMPLIANT
FRF.1.1 User-to-Network Implementation Agreement (UNI)	
ANSI T1.403 - 1989 Carrier to Customer DS1 Metallic Interface	COMPLIANT
CCITT Recommendation V.35	COMPLIANT
CCITT Recommendation G.703 (2048 Kbps)	COMPLIANT
CCITT Recommendation G.704 (2048 Kbps)	COMPLIANT
CCITT Recommendation X.21	COMPLIANT
ANSI 530-A-1992	COMPLIANT
Electrical Characteristics	COMPLIANT
High Speed Serial Interface (ANSI/TIA/EIA/-612-1993)	COMPLIANT
DS3 Interface (44.736 Mbps)	COMPLIANT
E3 Interface (34.368 Mbps)	COMPLIANT
High Speed Interface (HSSI)	COMPLIANT
CCITT V.36/V.37 (1988) Interface (2 to 10 Mbps)	COMPLIANT
Data Transfer	COMPLIANT
Congestion Control Procedures	COMPLIANT
Permanent Virtual Connection (PVC) Procedures	COMPLIANT
Switched Virtual Connection (SVC) Procedures	Planned
FRF.2 Network-to-Network Implementation Agreement (NNI)	
DS1 Interface (1544 Kbps)	COMPLIANT
CCITT Recommendation G.703 (2048 Kbps)	COMPLIANT
Data Transfer	COMPLIANT
Congestion Management	COMPLIANT
Permanent Virtual Connection (PVC) Procedures	COMPLIANT
Network Performance Parameters	COMPLIANT
FRF.2.1 Network-to-Network Implementation Agreement (NNI)	
DS1 Interface (1544 Kbps)	COMPLIANT
CCITT Recommendation G.730 (2048 Kbps)	COMPLIANT
CCITT Recommendation G.704 (2048 Kbps)	COMPLIANT
DS3 Interface (44.736 Kbps)	COMPLIANT
E3 Interface (34.368 Kbps)	COMPLIANT

Frame Relay	**Compliance**
Transmission over ATM	COMPLIANT
High Speed Serial Interface (HSSI)	COMPLIANT
CCITT V.36/V.37	COMPLIANT
Data Transfer	COMPLIANT
Congestion Management	COMPLIANT
Permanent Virtual Connection (PVC) Procedures	COMPLIANT
Network Performance Parameters	COMPLIANT
Application of Event Driven Procedures at the NNI	COMPLIANT
Country Code and National Network Identifier	COMPLIANT
FRF.7 Frame Relay PVC Multicast Service and Protocol Description Implementation Agreement	
One-Way Multicast Service	COMPLIANT
Two-Way Multicast Service	COMPLIANT
N-Way Multicast Service	Not Planned
General	COMPLIANT
Bandwidth Management	COMPLIANT
Congestion Management	COMPLIANT
Procedures	COMPLIANT
Interworking Requirements	COMPLIANT
FRF.8 Frame Relay/ATM PVC Service Interworking	

Appendix C. Special Notices

This publication is intended to help networking specialists, customers, and IBM business partners who plan to install and implement IBM Nways WAN switch products. The information in this publication is not intended as the specification of any programming interfaces that are provided by 2219, 2225, and 2230 product name(s). See the PUBLICATIONS section of the IBM Programming Announcement, 196-067 for more information about what publications are considered to be product documentation.

References in this publication to IBM products, programs or services do not imply that IBM intends to make these available in all countries in which IBM operates. Any reference to an IBM product, program, or service is not intended to state or imply that only IBM′s product, program, or service may be used. Any functionally equivalent program that does not infringe any of IBM′s intellectual property rights may be used instead of the IBM product, program or service.

Information in this book was developed in conjunction with use of the equipment specified, and is limited in application to those specific hardware and software products and levels.

IBM may have patents or pending patent applications covering subject matter in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to the IBM Director of Licensing, IBM Corporation, 500 Columbus Avenue, Thornwood, NY 10594 USA.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact IBM Corporation, Dept. 600A, Mail Drop 1329, Somers, NY 10589 USA.

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The information contained in this document has not been submitted to any formal IBM test and is distributed AS IS. The information about non-IBM (″vendor″) products in this manual has been supplied by the vendor and IBM assumes no responsibility for its accuracy or completeness. The use of this information or the implementation of any of these techniques is a customer responsibility and depends on the customer′s ability to evaluate and integrate them into the customer′s operational environment. While each item may have been reviewed by IBM for accuracy in a specific situation, there is no guarantee that the same or similar results will be obtained elsewhere. Customers attempting to adapt these techniques to their own environments do so at their own risk.

The following terms are trademarks of the International Business Machines Corporation in the United States and/or other countries:

AIX	IBM
Nways	RISC System/6000

The following terms are trademarks of other companies:

C-bus is a trademark of Corollary, Inc.

PC Direct is a trademark of Ziff Communications Company and is used by IBM Corporation under license.

UNIX is a registered trademark in the United States and other countries licensed exclusively through X/Open Company Limited.

Microsoft, Windows, and the Windows 95 logo are trademarks or registered trademarks of Microsoft Corporation.

Java and HotJava are trademarks of Sun Microsystems Inc.

CascadeView is a trademark of Cascade Communications Inc.

Virtual Network Nagavator is a trademark of Cascade Communications Inc.

OPTimum is a trademark of Cascade Communications Inc.

HP OpenView is a trademark of Hewlett Packard.

Sybase is a trademark of Sybase, Inc.

Solaris is a trademark of Sun MicroSystems.

Other trademarks are trademarks of their respective companies.

Appendix D. Related Publications

The publications listed in this section are considered particularly suitable for a more detailed discussion of the topics covered in this redbook.

D.1 International Technical Support Organization Publications

For information on ordering these ITSO publications see "How To Get ITSO Redbooks" on page 145.

- *High-Speed Networking Technology: An Introductory Survey*, GG24-3816
- *IBM 2220 Nways BroadBand Switch: Concepts and Products*, SG24-4307
- *Asychronous Transfer Mode (ATM) Tutorial*, SG24-4625
- *Networking BroadBand Services (NBBS) Architectural Tutorial,* GG24-4486
- *Frame Relay Guide*, GG24-4463

D.2 Redbooks on CD-ROMs

Redbooks are also available on CD-ROMs. **Order a subscription** and receive updates 2-4 times a year at significant savings.

CD-ROM Title	**Subscription Number**	**Collection Kit Number**
System/390 Redbooks Collection	SBOF-7201	SK2T-2177
Networking and Systems Management Redbooks Collection	SBOF-7370	SK2T-6022
Transaction Processing and Data Management Redbook	SBOF-7240	SK2T-8038
AS/400 Redbooks Collection	SBOF-7270	SK2T-2849
RISC System/6000 Redbooks Collection (HTML, BkMgr)	SBOF-7230	SK2T-8040
RISC System/6000 Redbooks Collection (PostScript)	SBOF-7205	SK2T-8041
Application Development Redbooks Collection	SBOF-7290	SK2T-8037
Personal Systems Redbooks Collection (available soon)	SBOF-7250	SK2T-8042

How To Get ITSO Redbooks

This section explains how both customers and IBM employees can find out about ITSO redbooks, CD-ROMs, workshops, and residencies. A form for ordering books and CD-ROMs is also provided.

This information was current at the time of publication, but is continually subject to change. The latest information may be found at URL http://www.redbooks.ibm.com.

How IBM Employees Can Get ITSO Redbooks

Employees may request ITSO deliverables (redbooks, BookManager BOOKs, and CD-ROMs) and information about redbooks, workshops, and residencies in the following ways:

- **PUBORDER** — to order hardcopies in United States
- **GOPHER link to the Internet** - type GOPHER.WTSCPOK.ITSO.IBM.COM
- **Tools disks**

  To get LIST3820s of redbooks, type one of the following commands:

  ```
  TOOLS SENDTO EHONE4 TOOLS2 REDPRINT GET SG24xxxx PACKAGE
  TOOLS SENDTO CANVM2 TOOLS REDPRINT GET SG24xxxx PACKAGE (Canadian users only)
  ```

  To get lists of redbooks:

  ```
  TOOLS SENDTO WTSCPOK TOOLS REDBOOKS GET REDBOOKS CATALOG
  TOOLS SENDTO USDIST MKTTOOLS MKTTOOLS GET ITSOCAT TXT
  TOOLS SENDTO USDIST MKTTOOLS MKTTOOLS GET LISTSERV PACKAGE
  ```

  To register for information on workshops, residencies, and redbooks:

  ```
  TOOLS SENDTO WTSCPOK TOOLS ZDISK GET ITSOREGI 1996
  ```

  For a list of product area specialists in the ITSO:

  ```
  TOOLS SENDTO WTSCPOK TOOLS ZDISK GET ORGCARD PACKAGE
  ```

- **Redbooks Home Page on the World Wide Web**

  http://w3.itso.ibm.com/redbooks

- **IBM Direct Publications Catalog on the World Wide Web**

  http://www.elink.ibmlink.ibm.com/pbl/pbl

  IBM employees may obtain LIST3820s of redbooks from this page.

- **ITSO4USA category on INEWS**
- **Online** — send orders to: USIB6FPL at IBMMAIL or DKIBMBSH at IBMMAIL
- **Internet Listserver**

  With an Internet E-mail address, anyone can subscribe to an IBM Announcement Listserver. To initiate the service, send an E-mail note to announce@webster.ibmlink.ibm.com with the keyword subscribe in the body of the note (leave the subject line blank). A category form and detailed instructions will be sent to you.

How Customers Can Get ITSO Redbooks

Customers may request ITSO deliverables (redbooks, BookManager BOOKs, and CD-ROMs) and information about redbooks, workshops, and residencies in the following ways:

- **Online Orders** (Do not send credit card information over the Internet) — send orders to:

	IBMMAIL	Internet
In United States:	usib6fpl at ibmmail	usib6fpl@ibmmail.com
In Canada:	caibmbkz at ibmmail	lmannix@vnet.ibm.com
Outside North America:	dkibmbsh at ibmmail	bookshop@dk.ibm.com

- **Telephone orders**

United States (toll free)	1-800-879-2755
Canada (toll free)	1-800-IBM-4YOU
Outside North America	(long distance charges apply)
(+45) 4810-1320 - Danish	(+45) 4810-1020 - German
(+45) 4810-1420 - Dutch	(+45) 4810-1620 - Italian
(+45) 4810-1540 - English	(+45) 4810-1270 - Norwegian
(+45) 4810-1670 - Finnish	(+45) 4810-1120 - Spanish
(+45) 4810-1220 - French	(+45) 4810-1170 - Swedish

- **Mail Orders** — send orders to:

IBM Publications	IBM Publications	IBM Direct Services
Publications Customer Support	144-4th Avenue, S.W.	Sortemosevej 21
P.O. Box 29570	Calgary, Alberta T2P 3N5	DK-3450 Allerød
Raleigh, NC 27626-0570	Canada	Denmark
USA		

- **Fax** — send orders to:

United States (toll free)	1-800-445-9269
Canada	1-403-267-4455
Outside North America	(+45) 48 14 2207 (long distance charge)

- **1-800-IBM-4FAX (United States)** or **(+1) 415 855 43 29 (Outside USA)** — ask for:

  Index # 4421 Abstracts of new redbooks
  Index # 4422 IBM redbooks
  Index # 4420 Redbooks for last six months

- **Direct Services** - send note to softwareshop@vnet.ibm.com

- **On the World Wide Web**

Redbooks Home Page	http://www.redbooks.ibm.com
IBM Direct Publications Catalog	http://www.elink.ibmlink.ibm.com/pbl/pbl

- **Internet Listserver**

  With an Internet E-mail address, anyone can subscribe to an IBM Announcement Listserver. To initiate the service, send an E-mail note to announce@webster.ibmlink.ibm.com with the keyword subscribe in the body of the note (leave the subject line blank).

IBM Redbook Order Form

Please send me the following:

Title	Order Number	Quantity

- **Please put me on the mailing list for updated versions of the IBM Redbook Catalog.**

First name Last name

Company

Address

City Postal code Country

Telephone number Telefax number VAT number

- Invoice to customer number

- Credit card number

Credit card expiration date Card issued to Signature

We accept American Express, Diners, Eurocard, Master Card, and Visa. Payment by credit card not available in all countries. Signature mandatory for credit card payment.

DO NOT SEND CREDIT CARD INFORMATION OVER THE INTERNET.

List of Abbreviations

ABR	Available Bit Rate
ATM	Asychronous Transfer Mode
CBR	Constant Bit Rate
CDV	Cell Delay Variation
CLR	Cell Loss Ratio
CP	Control Processor
CPE	Customer Premise Equipment
CSU	Channel Service Unit
DCE	Data Communications Equipment
DLCI	Data Link Connection Identifier
DSU	Data Service Unit
DTE	Data Terminal Equipment
DXI	Data Exchange Interface
FRAD	Frame Relay Access Device
HDLC	High-Level Data Link Control
HSSI	High-Speed Serial Interface
IOA	I/O Adapter
IOP	I/O Processor
ILMI	Interim Link Management Interface
IP	Internet Protocol
ISDN	Integrated Digital Network Exchange
ITSO	International Technical Support Organization
LAP	Link Access Protocol
LMI	Link Management Interface
MIB	Management Information Base
NMS	Network Management Station
OA&M	Operations Administration & Management
OSPF	Open Shortest Path First
PP	Packet Processor
PPP	Point-to-Point Protocol
PRAM	Parameter Random Access Memory
PRI	Primary Rate Interface
PVC	Private Virtual Circuit
QOS	Quality of Service
SLIP	Serial Line over Internet Protocol
SMDS	Switched Multimegabit Data Service
SNMP	Simple Network Management Protocol
SP	Switch Processor
SVC	Switched Virtual Circuit
UNI	User-to-Network Interface

Index

Numerics

A

B

C

D

F

G

H

I

L

M

N

P

R

S

T

U

V

Printed in U.S.A.

SG24-4777-00