Networking Services Technology Overview

Ascend Communications, Inc.

Product Code: 80001
Revision 03
August 1996

Contents

2 Frame Relay Services

List of Figures

About This Guide

This section describes the purpose, audience, and organization of the *Networking Services Technology Overview.* It also includes a list of related Cascade documentation.

Purpose

The *Networking Services Technology Overview* provides technical information on Cascade's networking services and on the technologies used by these services. This guide is intended to serve as a companion to both the *CascadeView/UX Network Configuration Guide* and the *Cascade 500 Network Administrator's Guide*, as well as a stand-alone reference.

The *Networking Services Technology Overview* is not an exhaustive overview of the various technologies. The primary goal of this guide is to help Cascade network administrators understand and properly provision network services. This guide emphasizes the aspects of the technologies that are particularly relevant to Cascade products.

Audience

The *Networking Services Technology Overview* is intended for anyone who wants to understand the networking services supported on the Cascade WAN switching platforms. It also provides background information for the network administrator or operator who is responsible for configuring a Cascade switch and provisioning services.

Readers should be familiar with basic networking concepts and terminology.

Document Conventions

In this document, important terms are ***italicized*** in bold typeface when accompanied by a definition or discussion of the term.

An arrow and a shaded box, as shown below, means reader take note. These notes point out special information about Cascade's implementation of the networking service or technology.

Sample text.

Organization

This guide is organized into five chapters as follows:

- Chapter 1, “Introduction,” introduces the reader to Cascade’s networking services and provides an overview of Cascade’s current products, their components, and key features.
- Chapter 2, “Frame Relay Services,” discusses the background and basic concepts of Frame Relay, as well as Cascade-specific details for understanding and provisioning Frame Relay services on the Cascade switch.
- Chapter 3, “SMDS Services,” discusses SMDS concepts with special attention to low-speed SMDS architecture. This chapter also includes information on how the Cascade switch functions in the role of SMDS Access Server/Switch and for provisioning SMDS services.
- Chapter 4, “ATM Services,” discusses the underlying concepts of ATM. This chapter also defines how the Cascade 500 ATM switch implements pure ATM cell switching.
- Chapter 5, “ISDN Services,” discusses the basic concepts of ISDN. This chapter also discusses Cascade’s ISDN fault-tolerant services for supporting mission-critical network applications.

The Glossary defines many of the technical terms used in this guide.

The Index provides a cross-reference listing of topics and entries for information contained in this guide.

Related Cascade Documentation

Following is a list of Cascade publications containing related information:

- *Cascade 500 Hardware Installation Guide*
  (Product Code: 80011)
- *Cascade 500 Network Administrator's Guide*
  (Product Code: 80012)
- *CascadeView/UX Network Management Station Installation Guide*
  (Product Code: 80014)
- *CascadeView/UX Network Configuration Guide*
  (Product Code: 80017)
- *CascadeView/UX Diagnostic and Troubleshooting Guide*
  (Product Code: 80018)
- *Cascade B-STDX 8000/9000 Hardware Installation Guide*
  (Product Code: 80005)
- *Cascade STDX 6000 Hardware Installation Guide*
  (Product Code: 80006)
- *STDX and B-STDX MIB Definitions*
  (Product Code: 80015)

For specific release-related information about the Cascade switch software, refer to its accompanying product release notes.

1

Introduction

Welcome to the Cascade Multiservice WAN Platforms

The Cascade STDX and B-STDX family of switches, along with the Cascade 500 ATM switch, provide a flexible, cost-effective, multiservice Wide Area Network (WAN) platform for internetworking among Frame Relay, Switched Multimegabit Data Service (SMDS), and Asynchronous Transfer Mode (ATM) networks. By providing all three broadband packet services, Cascade gives you the flexibility of choosing the service that suits your particular network applications.

All of Cascade's switches accommodate multiple IO modules to support numerous interface specifications, data rates, and protocols. Together, these platforms comprise an extensible family of switches that provides a standards-based foundation to meet the needs of a public carrier or a privately managed WAN.

STDX 6000

The STDX 6000 is a 6-slot, modular, Frame Relay switch. It provides a Frame Relay infrastructure for either public carrier or privately managed WAN-based networks.

B-STDX 8000/9000

The B-STDX 8000 is an 8-slot, modular, multiservice platform. It is designed to provide cost-effective solutions for low-density to medium-density sites within public network carriers and enterprise end-user networks.

The B-STDX 9000 is a 16-slot, high capacity, modular, multiservice platform for internetworking among broadband networks. The B-STDX 9000 is designed to accommodate large networks with high-bandwidth requirements.

Cascade 500

The Cascade 500 is a 16-slot, modular ATM switch. It is designed to play a key role at the center of public network carriers' ATM, Frame Relay, and SMDS service networks, as well as in large enterprise end-user networks.

For specific information about the type of Cascade switch you have and its hardware architecture, refer to its accompanying hardware installation guide. In this guide, references to ***Cascade switch*** *apply to all models unless specified otherwise.*

Networking Services

As a multiservice WAN platform, Cascade switches support Frame Relay, SMDS, and ATM networking services. In addition, Cascade provides ISDN fault-tolerant services for enhancing network reliability to support mission-critical applications.

Frame Relay Services

Cascade IO modules can act as Data Communications Equipment (DCE) or Data Terminal Equipment (DTE) devices at transmission rates of 2.4 Kbps to 45 Mbps, providing the Permanent Virtual Circuit (PVC) functions are configured as follows:

- Frame Relay switch interface (UNI-DCE)
- Frame Relay feeder interface (UNI-DTE)

- Frame Relay Network-to-Network Interface (NNI)
- Leased line trunk
- OPTimum inter-switch trunk over Frame Relay

If the incoming traffic is not in Frame Relay format, the Cascade switch can either encapsulate framed HDLC data into Frame Relay format (FRAD) for protocols such as X.25/HDLC and SNA/SDLC, or translate the traffic in point-to-point protocols, such as IP and IPX, into Frame Relay data.

For more information about Frame Relay, refer to Chapter 2, “Frame Relay Services.”

Switched Multimegabit Data Services (SMDS)

Cascade provides Switched Multimegabit Data Services (SMDS) defined by Bellcore and the SMDS Interest Group for SMDS DXI (Data Exchange Interface), SNI (Subscriber-Network Interface), and SSI (Server-Switch Interface) at speeds ranging from 56 Kbps to 45 Mbps.

Cascade SMDS services enable the Cascade switch to perform actual SMDS switching by routing SMDS data packets over Cascade switch trunks using the OSPF routing protocol. You can connect Cascade B-STDX 8000/9000 switches to an SMDS Distributed Queue Dual Bus through an SMDS Data Service Unit (DSU).

For more information about the SMDS Access Server and switching system services, refer to Chapter 3, “SMDS Services.”

Asynchronous Transfer Mode (ATM) Services

The Cascade B-STDX 8000/9000 switch provides ATM DXI DTE and DCE and ATM UNI Version 3.1 DTE and DCE interfaces. The Cascade 500 ATM switch provides a DS3/E3 ATM UNI interface, as well as an OC3c/STM-1 and OC12c/STM-4 ATM UNI interface.

In addition, the Cascade family of switches provides the following:

- ATM capability and standards compliance according to ATM UNI specifications to build reliable ATM wide area networks.
- Existing network services, such as Frame Relay and SMDS voice or video circuits, to make use of the ATM network.

- Conversion of Frame Relay to ATM using the established standards for ATM adaptation using AAL5.
- Circuit emulation using the established standards for ATM adaptation using AAL1.

For more information about ATM implementation on the Cascade 500 ATM switch, refer to Chapter 4, "ATM Services."

ISDN Fault-Tolerant Services

Fault tolerance is a method for providing redundancy in hardware systems to protect against downtime should one of the redundant systems or components fail. For the Cascade B-STDX 8000/9000, fault tolerance is provided via redundancy IO modules and power supplies. Fault tolerance is also provided by the capability for the hot-swapping replacement of parts while the switch is running.

ISDN fault-tolerant services are software functions that offer an additional measure of reliability for supporting the uninterrupted flow of critical traffic between source and destination. Some of these functions are integral to the networking service, such as dynamic rerouting, while others are optional features.

To increase reliability beyond that provided by the hardware and networking service, Cascade provides additional software-based services designed to support mission-critical applications. This added reliability is critical to financial institutions where any interruption in service could be devastating to their business.

The first of these ISDN fault-tolerant networking services is ***ISDN Trunk Backup***. ISDN Trunk Backup refers to the capability of the switch to set up one or more backup trunks to replace a primary trunk, and then reroute all PVCs from the primary trunk to the backup or other available trunks.

The PVCs remain rerouted on the backup trunks until trunk restoration occurs. During trunk restoration, the PVCs are rerouted from the backup trunk(s) and the backup trunk(s) are cleared.

Cascade provides two additional fault tolerant services:

- Fault Tolerant PVC
- Access Failure Recovery

Fault-Tolerant PVC allows a set of backup ports on the B-STDX 8000/9000 switch to restore connections from a failed data center to the backup data center. When enabled, Fault-Tolerant PVC reroutes all affected Frame Relay PVCs to the backup set of ports.

Access Failure Recovery provides failure recovery service through the ISDN Primary Rate (PRI) IO Module for the B-STDX 8000/9000 switch. If access to the switch fails, an ISDN dial-backup call is automatically placed to the Cascade switch and a new connection is established.

For more information about ISDN fault-tolerant services, refer to Chapter 5, “ISDN Fault-Tolerant Services.”

Service Network Backbones

Current data service demand dictates the need for an ATM backbone network configuration because ATM is the only technology that can carry aggregated customer traffic at speeds of 100 Mbits and above. In addition, ATM provides the migration to new data service applications such as video conferencing and usage-based billing.

Figure 1-1 shows a network of Cascade B-STDX 9000 switches connected to a Cascade 500-based ATM backbone. The Frame Relay-to-ATM internetworking capabilities of the B-STDX 9000 switches bring Frame Relay service traffic onto the ATM backbone, while the Cascade 500 switches scale the backbone to meet the demand.

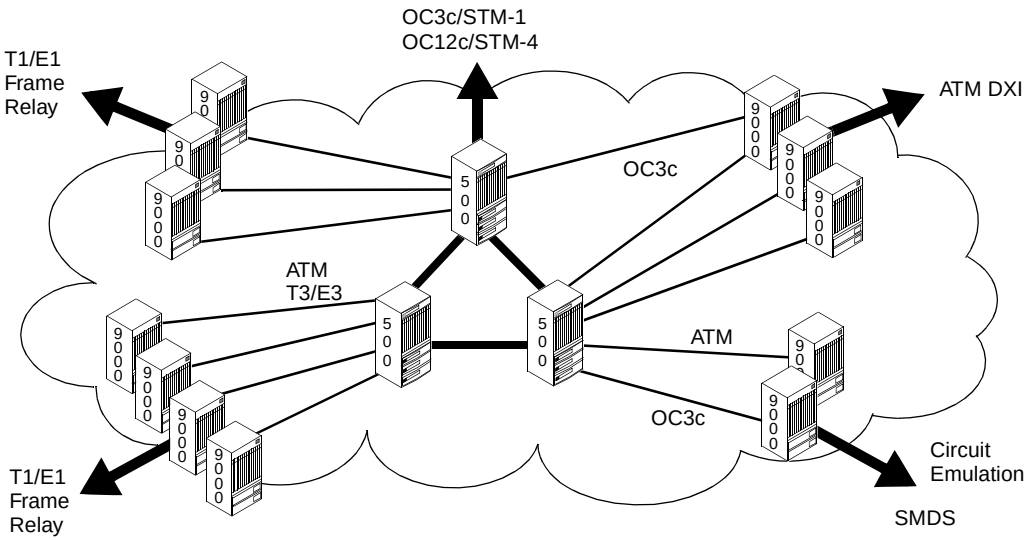

Figure 1-1. Cascade B-STDX With Cascade 500 Backbone

With the growth of Frame Relay, IP internet services require a scalable, high-bandwidth infrastructure for interconnecting multiservice access switches. The Cascade 500 provides the scalable, service-ready implementation that is needed for real ATM backbone deployment.

As Figure 1-2 shows, once ATM service deployment has accelerated, additional Cascade 500 switches can be deployed for native ATM services at T3/E3 speeds.

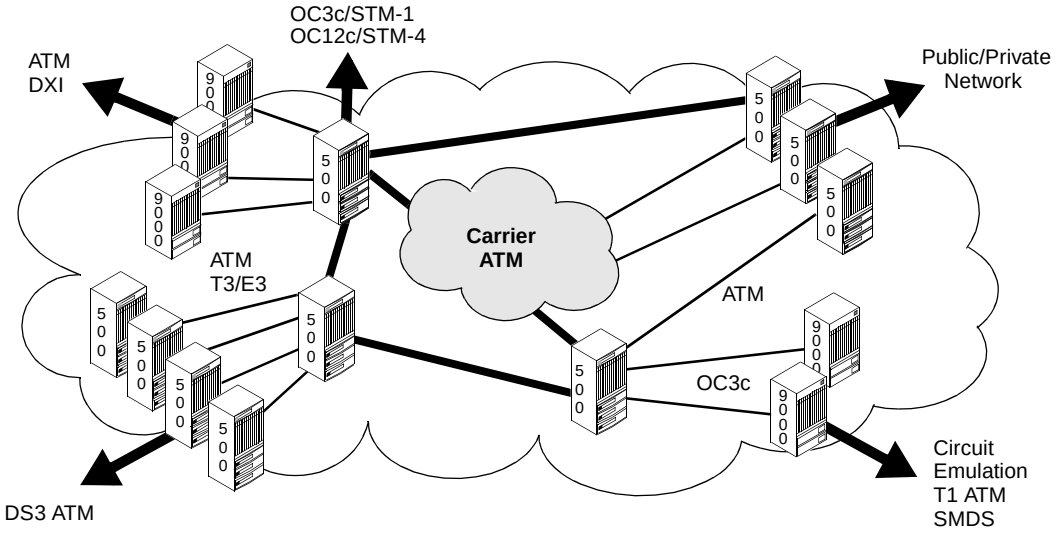

Figure 1-2. Cascade Network With Native ATM T3/E3 Services

Enterprise Network Backbones

As enterprise network providers become more like service providers, they must deliver non-stop data service for integrating remote offices within the campus information systems infrastructure. Enterprise networks are migrating to greater use of high-bandwidth ATM services, while simultaneously deploying ATM in campus backbones.

As Figure 1-3 shows, by using the Cascade 500 as the collapsed campus backbone, network planners can integrate their ATM LANs and WANs.

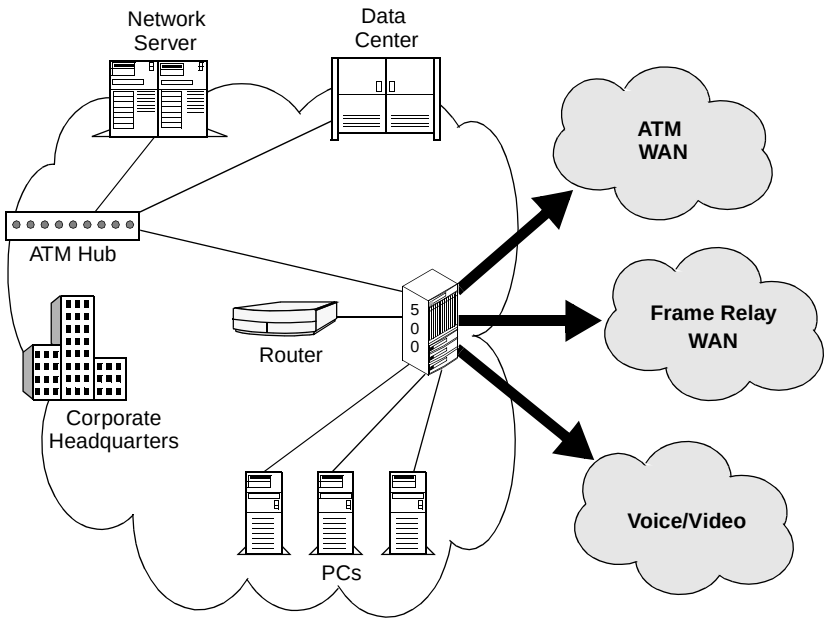

Figure 1-3. Cascade 500 as a Campus Backbone

Configuration Flexibility

The Cascade family of switches offers a high degree of configuration flexibility. The software component of the switch provides for flexibility in the design and implementation of the network, and reliability and ease in managing the network.

Because Cascade uses a symmetrical architecture, all ports are software-defined as one of the following:

- Trunks
- User-to-Network Interface (UNI) links
- Network-to-Network Interface (NNI) links

You can use any port for any of the software configurable interfaces, regardless of the type of service. The Cascade symmetrical architecture, referred to as ***Logical Port Architecture***, lets you mix user network connections and inter-switch trunk connections on any single IO module, allowing for high capacity in a compact platform and economical use of IO modules.

Open Packet (OPTimum) Trunking

OPTimum trunking is a software component that allows public data networks based on Frame Relay, SMDS, or ATM to be used as trunk connections between Cascade switches.

The OPTimum trunk requires either a User-to-Network Data Terminal Equipment (UNI-DTE) feeder or a Frame Relay NNI to be configured on the same physical port. This configuration allows for link management exchange, as shown in Figure 1-4.

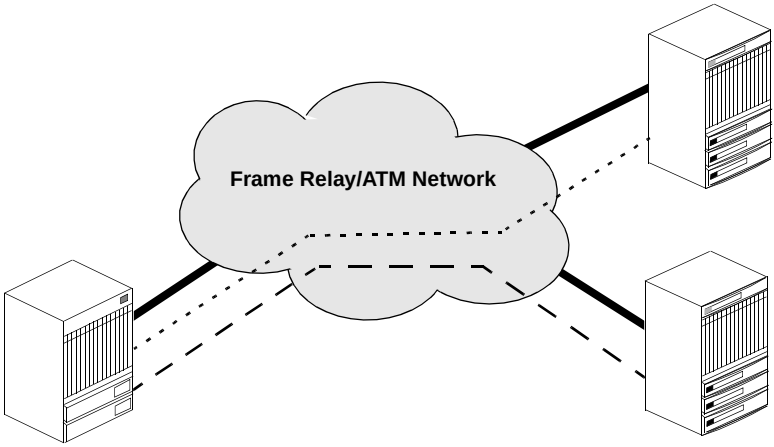

Figure 1-4. Cascade OPTimum Trunking

Comprehensive Management

Cascade's software platform provides sophisticated, standards-based capabilities for managing and controlling wide area networks, while integrating easily with existing network management systems. ***CascadeView***, the network management application, is built on Hewlett-Packard OpenView.

The Cascade switch has a native Simple Network Management Protocol (SNMP) agent supporting the Cascade Enterprise Management Information Base (MIB), standard MIB-II, and other MIBs required by the Frame Relay and ATM Forums.

The network management capabilities include the following features:

- Open Shortest Path First (OSPF) virtual circuit routing
- Rate monitoring
- Fault management
- Configuration management
- Usage statistics generation and collection

Managing the Cascade switch is simplified by the features described in the following sections.

HP OpenView Platforms

HP OpenView is a multivendor platform and already supports many communications devices that allow a management station to manage Cascade switches and other vendors' equipment. OpenView software lets you create a graphical network map and use pull-down menus to monitor, diagnose, and control objects on the network.

Both HP OpenView UNIX and HP OpenView Windows are adopted as the standard Cascade management platforms. HP OpenView UNIX runs on Sun Sparc 10/20 workstations. HP OpenView Windows runs on Intel-based 486 PCs.

Standards-based CascadeView

The Cascade switch has a native SNMP agent, and therefore can be managed from any SNMP management system supporting the standard MIBs and the Cascade Enterprise MIB extensions.

The CascadeView Network Management Station (NMS) communicates with the Cascade switches through the Internet Protocol (IP) using in-band or out-of-band management connections. CascadeView supports in-band management over the following connections:

- Ethernet
- Frame Relay in-band management Data Link Connection Identifier (DLCI)
- SMDS in-band management

Alternatively, CascadeView can access any Cascade switch in a network out-of-band over a dial-up modem.

With the ***Multi-NMS option***, multiple management stations can be configured in a server environment, allowing network management stations in different locations to have either read access to, or scheduled write access control of the network database.

For more information about the Multi-NMS option, refer to the *CascadeView/UX Network Management Station Installation Guide*.

Configuration Management

CascadeView provides a logical, step-by-step, network map creation procedure from setting network-wide parameters to provisioning individual Permanent Virtual Circuits (PVCs). Information in menus provides a series of choices to further simplify the configuration process.

CascadeView also provides summary reports of configuration information to document the current network configuration.

Performance Monitoring and Fault Management

The Cascade switch monitors the inbound and outbound traffic for each PVC and each port. Performance statistics display real-time information in a graphical format for a selected set of network objects.

CascadeView enables you to graph either the throughput on an entire switch or the throughput of individual physical and logical ports, circuits, or trunks. The statistical information varies according to your specific selection criteria.

Figure 1-5 illustrates an example of a trunk throughput selection with Statistics Type set to “Throughput (bps).” Extensive fault isolation is accomplished through user-initiated test sequences using loopback tests at the individual port level. (This display is only available on the CascadeView Windows platform.)

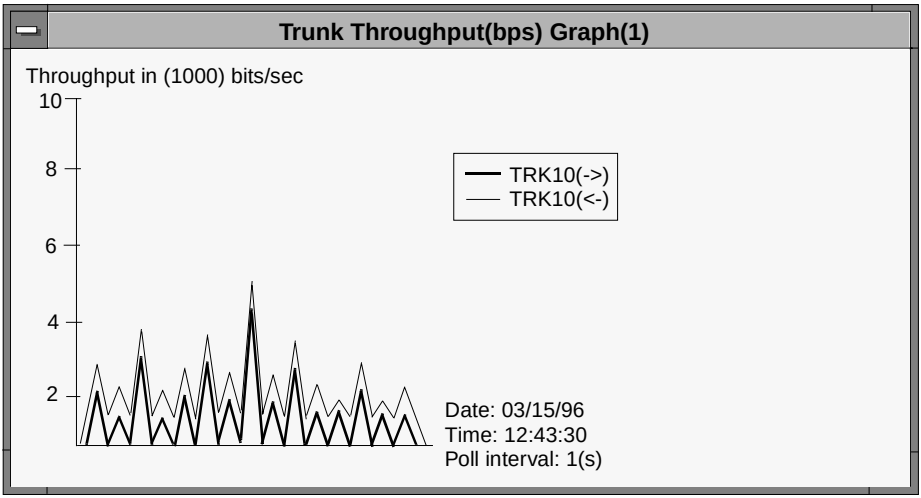

Figure 1-5. Trunk Throughput Graph in CascadeView Windows

CascadeView's background diagnostic capability has report generation for background diagnostic results for periodic analysis. A variety of traps for alarm conditions or statistics logging is available for all objects in the Cascade network, including switches, trunks, physical ports, logical ports, and PVCs.

For more information about CascadeView's statistics and diagnostic capabilities, refer to the *CascadeView/UX Diagnostic and Troubleshooting Guide*.

Security Management

CascadeView supports secure access to the Cascade switch network through two levels of password protection. One level is intended for the network administrator and the other level for the network operator. The network administrator can modify the network operator password to determine who has read-only and read-write control over the network.

2

Frame Relay Services

Background

Frame Relay is a packet switching technology that is simple, fast, and efficient. Frame Relay is particularly well suited for bursty LAN traffic. LAN traffic presents high transmission speeds and high bandwidth requirements for the WAN. Moreover, while individual bursts of LAN traffic outstrip the capabilities of remote links, their overall utilization is less than 20 percent.

What is needed is a mechanism to provide higher speeds without buying unnecessary bandwidth. Frame Relay was designed to meet these requirements, and is a viable solution for LAN applications that require high-density connectivity to remote sites.

To meet the demands of today's higher bandwidth requirements, for example voice, video, and imaging applications, and the low tolerance for network downtime, Frame Relay is changing from a technology for basic LAN interconnection to a more advanced application.

Technology Fundamentals

The following sections outline the basic concepts for understanding and provisioning Frame Relay services.

What Is Frame Relay?

Frame Relay is a multiplexed data networking technology supporting connectivity among user equipment, such as routers, and between user equipment and carriers' Frame Relay network equipment. The Frame Relay protocol supports data transmission over a connection-oriented path.

In Frame Relay, data is divided into variable-length frames, all of which include addressing information. These frames are then forwarded to a Frame Relay network, which then tries to deliver the frames to their specified destination over an assigned virtual connection.

High Throughput and Low Delay

High throughput and low delay are important characteristics of a Frame Relay network. Unlike X.25 networks designed for unreliable analog lines, Frame Relay takes advantage of the more reliable data communication lines, such as fiber optic. The Frame Relay protocols are confined to Layers 1 and 2 of the OSI protocol stack. Frame Relay does not offer error correction services, but instead relies on the upper-level protocols at the end stations to provide reliability. Without this additional overhead, Frame Relay is able to run at very high speeds (up to 45 Mbps).

Frame Relay benefits from the statistical gains of packet switching and makes efficient use of bandwidth by its ability to multiplex individual packets over one physical interface onto the WAN.

Virtual Circuits

Frame Relay uses multiple, uni-directional ***virtual circuits*** for bi-directional end-to-end connectivity. These virtual circuits are logical connections to various endpoints throughout the WAN. The endpoints, or users, of a Frame Relay network are often multiprotocol routers, as shown in Figure 2-1. Many virtual circuits can traverse a single physical interface.

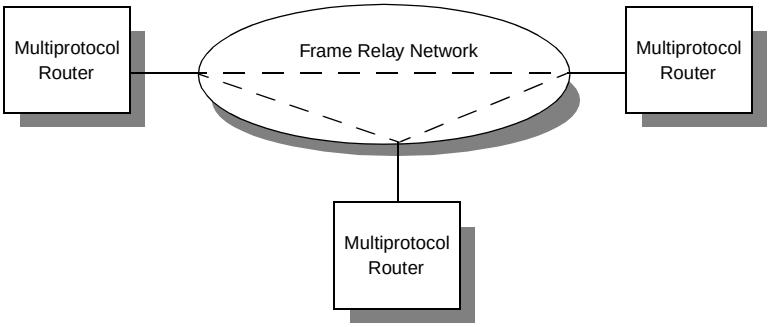

Figure 2-1. Virtual Circuits

Frame Relay uses ***Permanent Virtual Circuits (PVCs)*** and ***Switched Virtual Circuits (SVCs)***. PVCs, which are most prevalent today, are similar to dedicated private lines and are available by subscription from a service provider. SVCs allow dial on-demand services to be available over Frame Relay networks.

Data Link Connection Identifiers

PVCs are identified by a 10-bit address called a ***Data Link Connection Identifier (DLCI)***. DLCIs typically have local significance only (but not necessarily), and are used to uniquely identify the logical endpoints of a virtual circuit.

In Figure 2-2, a PVC between Los Angeles and New York is identified by the set of DLCIs that define the endpoints. New York is assigned DLCI 80, Los Angeles is assigned DLCI 111 and DLCI 120, and Atlanta is assigned DLCI 40. When New York sends frames to Los Angeles, it uses DLCI 80. When the frames arrive in Los Angeles, the DLCI is changed to 111.

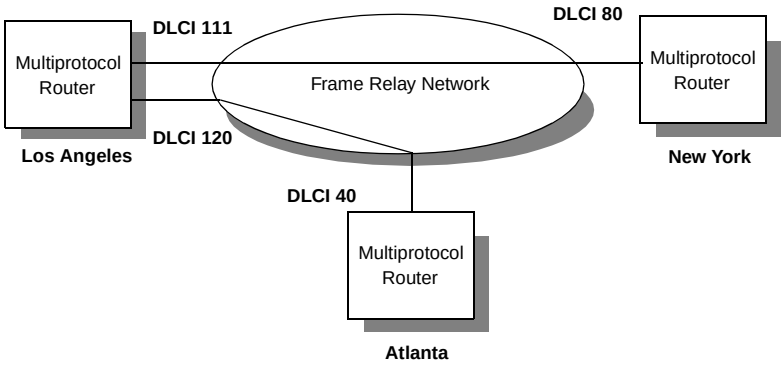

Figure 2-2. Data Link Connection Identifiers

DLCIs identify specific end devices and are typically configured to have "local significance" only. This way, DLCIs can be reused at non-overlapping endpoints. From this addressing viewpoint, the Frame Relay network appears as a LAN.

From a configuration viewpoint, the DLCIs that map to the logical ports at each end of a virtual circuit cannot be reused. Consequently, each DLCI assigned to a particular logical port must be unique. This rule is automatically enforced by the CascadeView software.

Frame Relay standards support a theoretical limit of 975 DLCIs over one physical interface when using “local significance.” With “global significance,” 975 DLCIs can be applied to the entire network. Consequently, the entire network is limited to 975 DLCIs since each DLCI must be uniquely identified.

Interface Definitions

The interface that enables the Customer Premise Equipment (CPE) Frame Relay device access to a Frame Relay network is called the ***User-to-Network Interface (UNI)***. Unlike ISDN and X.25, the UNI implementation agreement preceded vendors’ implementations of the standard. Therefore, implementation of the UNI standard is identical throughout the industry.

Frame Relay circuits or packets can traverse across multiple Frame Relay networks. This enables two Frame Relay networks to connect to each other. The interface between adjacent Frame Relay networks is called the ***Network-to-Network Interface (NNI).***

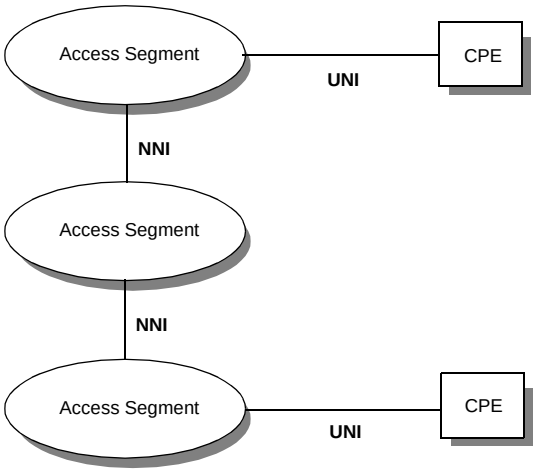

Figure 2-3. User-to-Network /Network-to-Network Interfaces

Frame Relay service providers can offer ***Access Services*** and ***Transit Services***. Access Services provide a direct interface to the end user. Transit Services provide PVC connections to other Access Services.

A carrier can provide either Access or Transit Services for a given PVC, or in the case of a Long Distance Interexchange Carrier (IXC), both Access and Transit Services for a given PVC. A carrier can also offer Access and Transit Services simultaneously for different PVCs.

PVCs that span multiple Frame Relay networks are called ***Multi-network PVCs***. Multi-network PVCs are interconnected across networks using the NNIs, as shown in Figure 2-4.

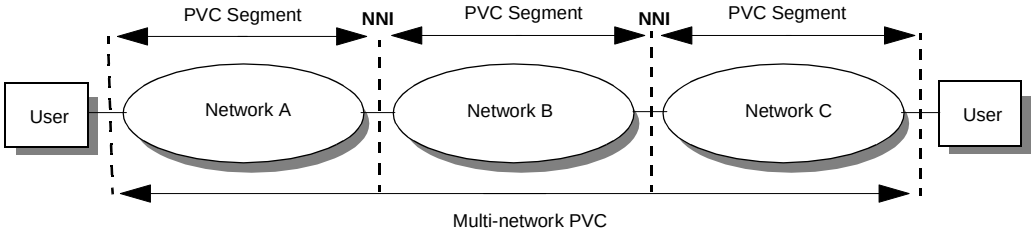

Figure 2-4. Multi-network PVC

Link Management Interface

The link management interface provides the user with dynamic notification of the addition and deletion of PVCs. It also monitors each network connection through a periodic heartbeat "keep alive" polling process. The ***heartbeat polling*** process is an exchange of sequence numbers between the network and the user device to ensure that both are operational. This polling process provides information about the user device's physical connection to the network through a Link Integrity Verification Report, as well as a Full Status Report about the status of all PVCs.

The user device (DTE) initiates the link management polling process by sending a status enquiry.

There are three types of link management protocols that can be used in a Frame Relay network:

- Local Management Interface (LMI)
- ANSI Annex D
- CCITT Annex A

In the early days of Frame Relay, the LMI specification was created by an industry group consisting of StrataCom, Digital Equipment Corp., and Northern Telecom (and assisted by Cisco). This specification was developed to create a mechanism for managing the early deployments of Frame Relay networks.

Annex D and Annex A protocols were subsequently drafted to be the "official" management mechanism for Frame Relay. Many vendors today support all three link management protocols. LMI uses ***DLCI 1023*** for signaling purposes. The ANSI and CCITT protocols use ***DLCI 0***.

Cascade supports the LMI, ANSI, and CCITT protocols for link management. Additionally, Cascade allows the switch to be configured for autodetection, and to have the link management protocol disabled. When configured for autodetection, the switch automatically senses the type of link management protocol in use and responds appropriately.

Rate Monitoring and Enforcement

Rate monitoring is a set of rules that describes traffic flow. The sender has a mechanism to ensure that the transmission of its guaranteed packets function in a certain way. The network knows what kind of traffic to expect and monitors the behavior of the traffic.

Guaranteed packets must be delivered according to some time constraint and with high reliability. ***Best effort packets*** are delivered to the best of the network's ability after meeting the requirements for delivering the packets.

The standard rate monitoring definitions for Frame Relay are as follows:

Committed Information Rate (CIR) – The rate at which the network agrees to transfer data under normal conditions. CIR is defined for each PVC.

Committed Burst Size (Bc) – The maximum number of bits, during time interval T, the network agrees to accept traffic under normal conditions. Bc is defined for each PVC.

Excess Burst Size (Be) – The maximum number of uncommitted bits, during time interval T, the network agrees to accept traffic above the committed burst size Bc. Be is defined for each PVC.

Discard Eligible Bit (DE) – A bit in the Frame Relay header used to indicate a frame is eligible for discard by a congested node.

Rate enforcement is implemented on a per DLCI basis on user links at the ingress of the Cascade 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.

Committed Rate Measurement Interval (T) – Time interval (T) is the length of time data is received before a rate enforcement determination is made. T is calculated according to the following formula:

$$T = \frac{Bc}{CIR}$$

A typical T is one second. It is determined from the values given for Bc and CIR. For example:

$$T = \frac{56KBits}{56Kbps} = 1 \text{ sec}$$

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)
- Over the excess burst size

Cascade uses the colors green, amber, and red to describe and categorize packet frames for rate monitoring and enforcement.

Green Frames

If the number of bits received during the current time interval, including the current frame, is less than Bc, the frame is designated as a ***green frame***. Green frames are never discarded by the network, except under extreme congestion conditions.

Amber Frames

If the number of bits received during the current time interval, including the current frame, is greater than Bc but less than Be, the frame is designated as an ***amber frame***. Amber frames are forwarded with the DE bit set, and are eligible for discard if they pass through a congested node.

Red Frames

If the number of bits received during the current time interval, including the current frame, is greater than Be, the frame is designated as a ***red frame***. Red frames are forwarded with the DE bit set when Cascade's Graceful Discard feature is enabled. When the Graceful Discard feature is disabled, red frames are discarded.

Congested nodes that must discard packets use the color designations to determine which frames to discard. Red frames are discarded first, followed by amber and then green.

Figure 2-5 illustrates how rate monitoring and enforcement determine bandwidth allocation.

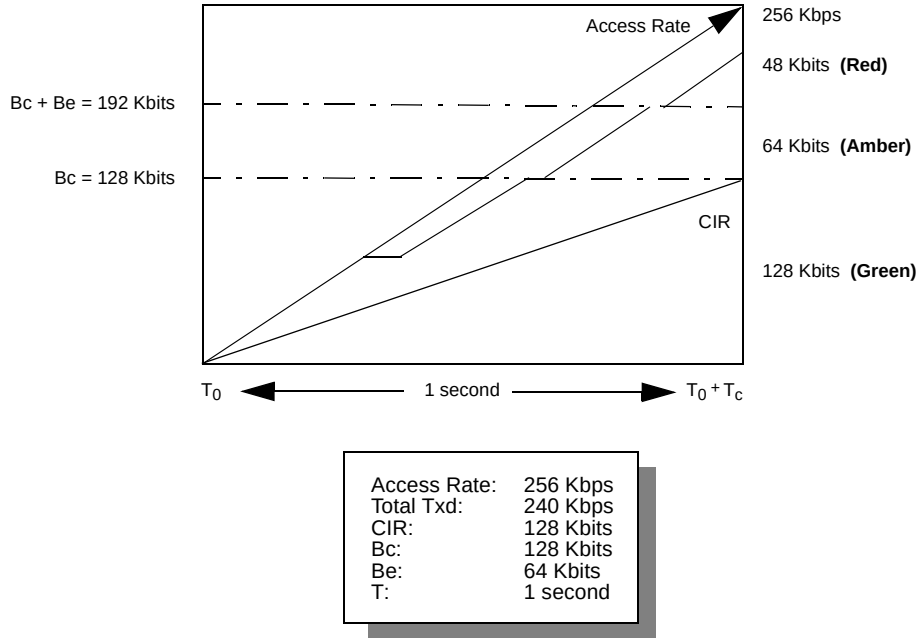

Figure 2-5. Determining Bandwidth Allocation

Congestion Avoidance

In the header of all frames, there is a field for the Forward Explicit Congestion Notification Bit ***(FECN)*** and the Backward Explicit Congestion Notification Bit ***(BECN)***. The Frame Relay network uses the FECN and BECN bits to notify the edge nodes of congestion in the network.

FECN informs the destination device of congestion for destination-controlled protocol suites. BECN informs the source device of congestion for source-controlled protocol suites. Both the FECN and BECN bits are set by the network, not the user. Therefore, there is no obligation for endpoint nodes to take any regard of the bits.

FECN and BECN bits work by indicating the existence of network congestion to the higher-layer protocols, so that they can reduce the rate of information flow into the network until the congestion clears. However, many Frame Relay CPE devices cannot use this information. The FECN and BECN bits are often ignored or simply counted by internetworking devices, such as routers, to provide an indicator of congestion in the network.

Frame Format

Frame Relay uses a simple framing structure. In addition to the starting and ending flags, there is an address field that is normally two octets, an information field that is normally less than 4096 octets, and a two-octet CRC for error detection. This simple framing structure allows Frame Relay to operate at higher speeds with low latency.

Figure 2-6 illustrates the simple framing structure and format for Frame Relay.

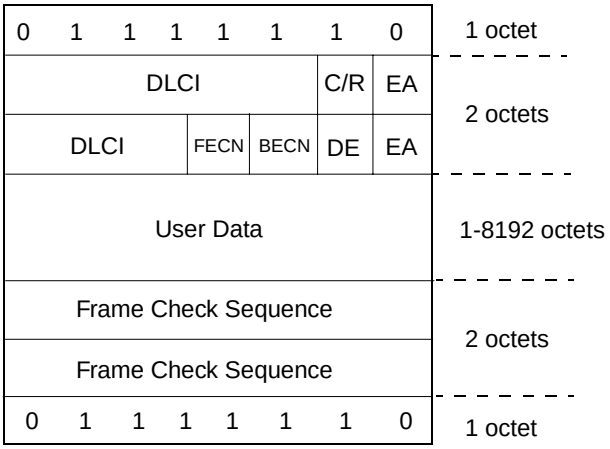

Although Cascade supports a maximum length of 8192 octets, a maximum length of 4096 octets is recommended. The 2-octet CRC cannot guarantee detection of errors for user data fields that are longer than 4096 octets.

Figure 2-6. Frame Relay Frame Structure

The framing structure consists of the following format:

Flags: Interframe delimiters, hex 0x7E (HDLC).

DLCI: 10-bit address that identifies a PVC.

C/R: Identifies control frames as commands or responses.

EA: Extended Addressing (currently not used).

DE: Discard Eligibility identifies frames that are not guaranteed delivery.

FECN: Used by the network for congestion notification.

BECN: Used by the network for congestion notification.

User Data: The Cascade network does not interpret or modify the contents of user data except for the virtual circuit dedicated to the management protocol.

LAN/WAN Interconnect Using Frame Relay

Numerous remote sites can be interconnected using a Frame Relay network. Packets can be either bridged or routed. If routed packets are traversing the network using Frame Relay as the transport protocol, then routing control packets must also traverse the network (unless a static route was configured). Either way, the bridged or routed packet is data within the Frame Relay cloud. The network does not interpret or modify the contents of the user data.

If multiple protocols are being sent over the Frame Relay network, the receiving device must be able to determine what type of protocol is encapsulated in the received frame. Special fields are inserted into the frame header to act as protocol identifiers.

There are different mechanisms for protocol identification. The original space allocated in the Network Level Protocol ID (NLPID) field was insufficient to handle the number of protocols requiring Frame Relay support. As a result, Frame Relay borrows two different mechanisms from the LAN environment for protocol identification:

- The LLC SNAP header
- The Ethernet MAC type field

Cascade Frame Relay Switching

The remainder of this chapter describes in more detail how Cascade implements specific aspects of Frame Relay services in its
STDX 6000 and B-STDX 8000/9000 switches.

Congestion Management

As data travels through the Cascade network and is queued for transmit, the state of each transmit queue is checked for pending congestion. A time average algorithm, Average Queue Length (AQL), is executed each time a frame is queued for transmit. The AQL is calculated and compared against a precalculated threshold.

The threshold is calculated so that when the AQL is less than or equal to the threshold, maximum throughput and minimum delay are achieved. For channelized cards, the default values assigned to thresholds are dependent on the number of DS0s/TS0s you assign to each logical port. A DS0/TS0 is a 64 Kbps channel used in T1/E1 transmission. There are 24 DS0 channels available in a T1 line and 31 TS0 channels available in an E1 line.

There are three high-water thresholds used in determining congestion:

1. Mildly congested
2. Severely congested
3. Absolutely congested

If the AQL exceeds the mildly congested AQL threshold, but is less than the severely congested AQL threshold, the state of the link is considered ***mildly congested***.

When the link is mildly congested:

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 to another link.

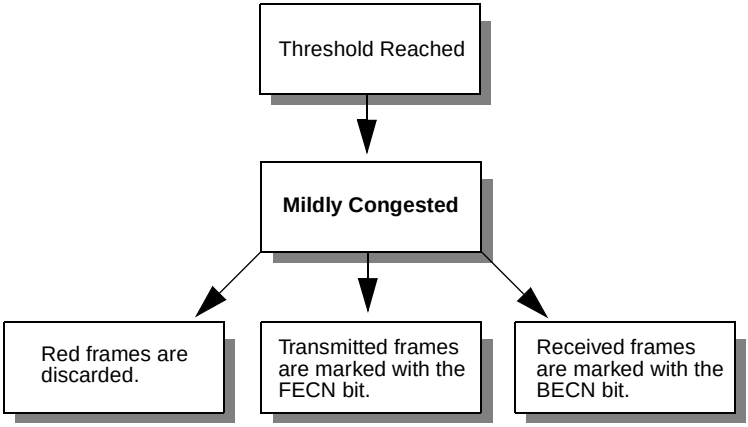

Figure 2-7. Mildly Congested State

If the AQL exceeds the severely congested AQL threshold, but is less than the absolutely congested AQL threshold, the state of the link is considered ***severely congested***. In this state, the continued forwarding of amber and red frames would compromise the delivery of green frames.

When the link is severely congested:

1. All incoming red frames are discarded.
2. All incoming amber frames are discarded.
3. FECN and BECN bits are set.

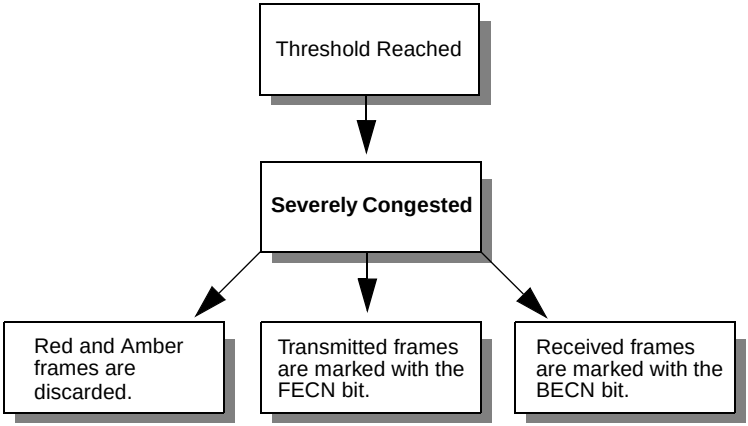

Figure 2-8. Severely Congested State

If the AQL exceeds the absolutely congested AQL threshold, the state of the link is considered ***absolutely congested***. Consequently, there is no room on the queue for any packets, regardless of the type. When the link isabsolutely congested:

1. All incoming frames are discarded.
2. FECN and BECN bits are set.

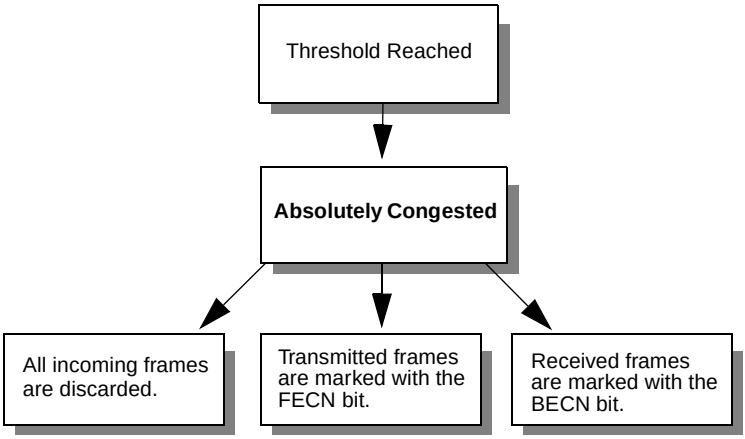

Figure 2-9. Absolutely Congested State

Rate Enforcement

Cascade switches maintain two counters for each DLCI on each user link:

- **Bc_cnt** is the number of committed bits allowed during the current time interval.
- **Be_cnt** is the number of uncommitted bits allowed during the current time interval.

A time interval is one second. It is measured by a continuously running one second timer started when the Cascade switch is initialized.

Bc_cnt and Be_cnt are initialized to the Bc and Be, respectively, and set for each user DLCI. In addition, two adjustment variables, Bc_adj and Be_adj, are defined for each user DLCI and set to Bc(1/T) and Be(1/T), respectively. These values represent the number of committed and uncommitted bits allowed during any one second interval. Bc_cnt and Be_cnt are also used when re-adjusting during each timer expiration.

Each time a frame is received by an ingress node, the following algorithm (Figure 2-10) is used to update the counters. Frames received with the DE bit already set are always counted against Be.

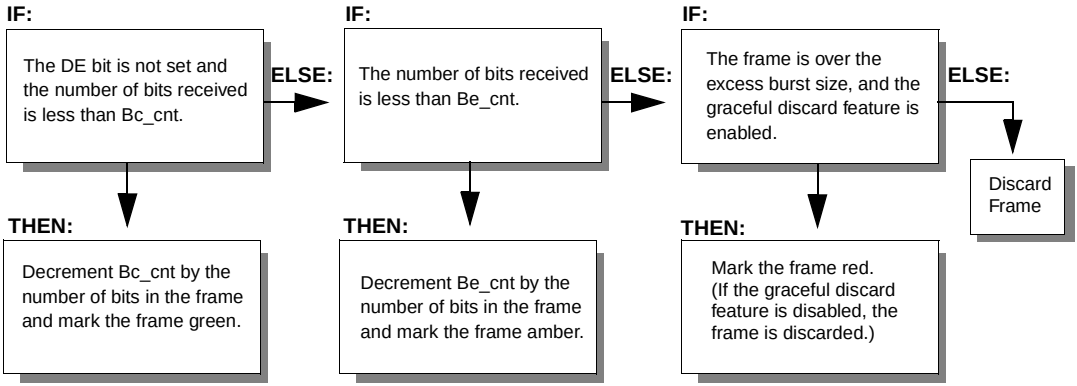

Figure 2-10. Algorithm for Updating Counters

For each timer expiration, the Bc_cnt and Be_cnt of each user DLCI are updated as follows:

- Bc_cnt = min (Bc, Bc_cnt+Bc_adj)
- Be_cnt = min (Be, Be_cnt+Be_adj)

OSPF Routing

Open Shortest Path First (OSPF) is the Internet standard for Interior Gateway Protocol (IGP). Network nodes use this routing protocol to determine the shortest path over which to send traffic.

OSPF is a link state routing protocol that determines the shortest paths throughout the network. OSPF creates very little additional overhead on the network and can reroute traffic quickly around failed links. OSPF is also scalable, and therefore accommodates very large networks. For details on OSPF, refer to *Internet RFC 1131*.

Upon initial activation of virtual circuits, Cascade switches use OSPF to determine the best paths through the network. As frames are received by the Cascade switch, the Frame Relay header is stripped and a Cascade trunk header is prepended.

Keep in mind that OSPF is only used for route calculation within the Cascade switching network for virtual circuit setup. During startup, each Cascade switch broadcasts its local state, including usable interfaces and reachable neighbors throughout the network.

Within a Cascade switching network, switches use OSPF to form adjacencies with their directly connected neighbors. Adjacencies can be thought of as highly developed conversations between nodes, and are formed by passing HELLO packets. Once adjacencies are formed, Cascade switches monitor the state of their directly attached links.

A Cascade switch sends out Link State Advertisements (LSAs) to other Cascade switches on the network whenever the state changes, for example, link up and link down.

In addition, a Cascade switch sends out ***Link State Updates***, which consist of the state information of all links. Network nodes collect Link State Updates and place them in a ***Link State Database***. Each Cascade switch runs the Djkstra algorithm on the Link State Database, which results in a ***tree of shortest paths*** throughout the network.

Figure 2-11 illustrates a Cascade network based on OSPF routing.

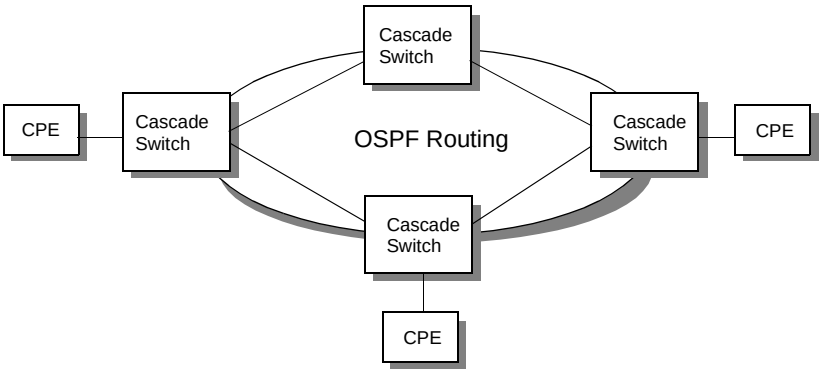

Figure 2-11. OSPF Routing

The shortest path to a destination is determined by adding link state costs. Each link has an assigned metric. The cost of a path is determined by adding the costs of each link along the path. The route with the lowest path cost that meets the QoS requirements for the virtual circuit is chosen first.

All Cascade switches run exactly the same algorithm in parallel. From the synchronized topological database, each Cascade switch constructs a tree of shortest paths with itself as the root. The best route to each destination can be derived from this shortest-path tree. External devices, such as the NMS connected to an Ethernet or Serial line, appear on the tree as leaves.

OSPF Metrics

OSPF does not require any additional configuration in Cascade switches. On PVC activation, OSPF transparently calculates the ***Shortest Path First (SPF)*** route and maps PVCs to the appropriate one. PVCs are mapped to the shortest path according to bandwidth availability.

As PVCs are mapped to a route through the network, bandwidth availability decreases. This causes the OSPF link-state metrics for bandwidth usage for that route to increase. When new trunks are added to the network or down trunks are reactivated, OSPF recalculates the shortest path tree. A load-balancing algorithm reroutes PVCs that are not using the shortest route.

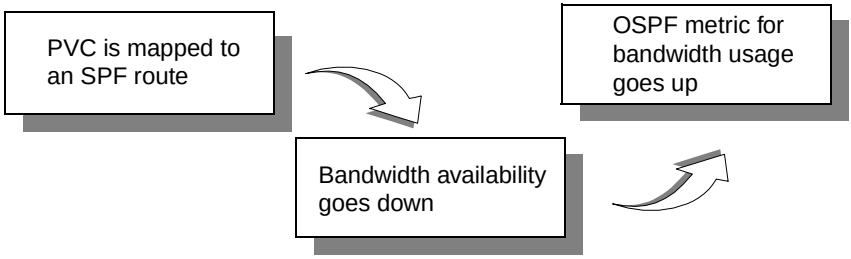

Figure 2-12. Dynamic Load Balancing

The ***OSPF Trunk Administrative Cost Metric*** is a function of OSPF that allows you to control the specific path a virtual circuit takes through the network. You can select a shorter hop path regardless of the available bandwidth on other, perhaps longer paths.

In general, the OSPF Trunk Administrative Cost Metric enables you to do the following:

- Assign an administrative cost to each trunk in the network.
- Configure a trunk cost value between 1 and 65534 (inclusive), with a default value of 100.

To ensure that all Quality of Service (QoS) guarantees are met, the available bandwidth of each trunk and any QoS requirements are considered prior to optimizing administrative cost for path selection.

The OSPF Trunk Administrative Cost Metric function is particularly useful when establishing a virtual circuit and rerouting around a trunk or switch failure. For example, if there are multiple paths with sufficient bandwidth to meet Committed Information Rate (CIR) requirements to the destination, the path with the lowest administrative cost is chosen. On initial establishment, if there is not enough bandwidth to satisfy CIR requirements, the circuit is not established. On reroute, the circuit is re-established even if all available paths have insufficient bandwidth.

Cascade switches support a Define Path function that allows you to manually define a PVC path, thereby bypassing OSPF to make PVC routing decisions. Manually defined paths provide additional control over routing, such as when conducting performance bench tests for route comparisons or when forcing critical traffic over a known reliable path. Additionally, an Alternate Path option allows OSPF to reroute this PVC should a trunk fail in the manually defined path. Both functions are configured through CascadeView.

The following list summarizes the key OSPF routing functions:

CIR Guaranteed Bandwidth ToS (Type of Service) Routing – If there are multiple paths that have sufficient bandwidth to meet CIR requirements to the destination, the path with the lowest administrative cost is chosen. If multiple paths exist with sufficient available bandwidth and the same administrative cost, the path with the largest available bandwidth is chosen. If there is no path to the destination that has sufficient available bandwidth, the circuit is re-established using the path with the lowest administrative cost.

Route Recovery – When a node or trunk is down, new shortest-path trees for those affected PVCs are recalculated immediately at the ingress nodes due to fast convergence of the link state updates. The PVCs are then rerouted to the new route. Recovery time is typically under four seconds. PVC rerouting is reported to the NMS.

Load Balancing – This algorithm distributes new PVCs equally over all routes with the same administrative cost. You can tune the load balancing algorithm to load balance from negative trunks to positive trunks, positive trunks to positive trunks, or disable the load balancing algorithm.

Efficient Network Throughput – Network throughput is maximized as a result of enforcing the CIR-based routing in the same time as allowing excess burst (Be/T) to be delivered using the uncommitted bandwidth. Continuous route re-evaluation and load balancing also improves the network throughput. To optimize this capability, PVCs carrying constant bulk traffic can be assigned larger CIR to guarantee the bandwidth required.

Support for Manual (pre-defined) Routes – In case of failures on the defined path, you can pre-define a PVC path. You can also reroute PVCs to an OSPF-defined path, which dynamically reroutes the PVC if that pre-defined path fails.

Virtual Bandwidth

End-user applications determine the bandwidth requirements for DLCIs. You can over-subscribe bandwidth in the Frame Relay network to take advantage of the statistical nature of packet-switching. ***Virtual bandwidth*** is calculated to allow for over-subscription. The key factors for defining virtual bandwidth are as follows:

- Bandwidth can be over-subscribed.
- 5% of bandwidth is reserved for control traffic.
- Trunk Oversubscription Factor (K) is configured to calculate Virtual Bandwidth:

$$0.95\% \text{ (configured bandwidth)} \times \frac{K\ (\%)}{100} = \text{Virtual Bandwidth}$$

Figure 2-13. Determining Virtual Bandwidth

The routing for PVCs is done dynamically by OSPF. On initial activation of a PVC, OSPF reserves bandwidth on the trunks along the shortest path that meets QoS requirements. In addition, if no path exists with a sufficient amount of available bandwidth, the PVC will not be established. The amount of reserved bandwidth is deducted from the reserved bandwidth pool.

The formula used to determine virtual bandwidth is only used for allocating the initial path for the PVC. If the trunk fails, requiring all active PVCs to use an alternate path, the bandwidth reservation scheme is no longer enforced. Consequently upon failure, active PVCs are rerouted even if it results in a negative bandwidth condition.

Configuring the trunk utilization factor at a higher percentage allows more available virtual bandwidth to be defined over the trunk. A value of 200% in the K factor effectively doubles the available virtual bandwidth. Cascade reserves 5% for management traffic.

You set the trunk utilization factor when you configure the trunk in CascadeView. Typically, the trunk utilization factor is set according to a projection of how much of the CIR will actually be used (i.e., accessed simultaneously by network devices) at a given time.

Logical Configurations

Physical port and logical port configurations are independent of each other. The physical port configuration, such as a T1 connection, supplies clock source, timing, and signal-strength settings. The logical port configuration determines the protocol interaction between the Cascade switch and one of the following:

- User equipment
- Another Cascade switch
- Another network

Therefore, the Cascade switch can play a variety of roles within the Frame Relay network. That role determines how the Cascade switch interacts with the attached equipment. This behavior is configured on a per port basis in the logical port configuration.

Cascade supports the following logical port configurations for Frame Relay:

- Frame Relay Switch (UNI-DCE)
- Frame Relay Feeder (UNI-DTE)
- Frame Relay NNI
- Frame Relay Direct FRAD
- Frame Relay PPP Translation FRAD

- Frame Relay OPTimum trunk
- Frame Relay Direct Line trunk

Each of these logical port configurations has special functions that enable particular network tasks to be performed. These configurations are described in the following sections.

Frame Relay Switch

The Frame Relay User Network Interface DCE (UNI-DCE) performs Frame Relay DCE functions for link management purposes, and communicates with a Frame Relay DTE on the other end of the link.

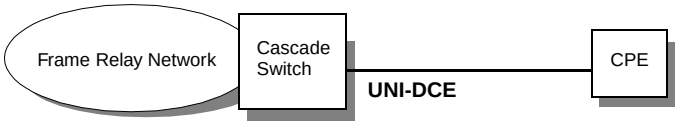

Figure 2-14. Frame Relay UNI-DCE Logical Port Definition

Frame Relay Feeder

The Frame Relay User Network Interface DTE (UNI-DTE) performs Frame Relay DTE functions specified for link management. In this configuration, the Cascade switch acts as the DTE to connect to a Frame Relay DCE network switch.

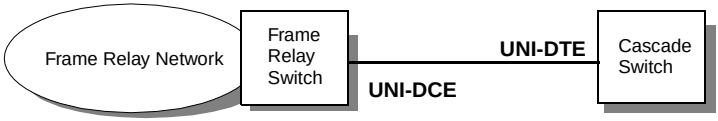

Figure 2-15. Frame Relay UNI-DTE Logical Port Definition

Frame Relay NNI

The Frame Relay Network-to-Network Interface (NNI) logical configuration enables two Frame Relay networks to connect using a standard protocol. The Frame Relay networks can be using the same or dissimilar switches.

The NNI port performs both Frame Relay DTE and DCE functions specified for link management.

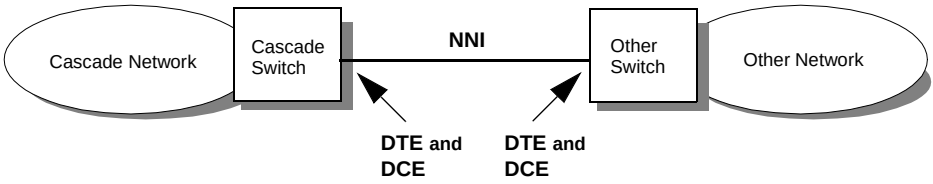

Figure 2-16. Frame Relay NNI Logical Port Definition

The NNI logical port configuration is used for connecting two networks together for bi-directional messaging. Once established, the NNI connection enables both networks to share status about the state of the PVCs in each network. However, not all public Frame Relay products provide NNI service.

Direct FRAD

The Frame Relay Direct FRAD (Frame Relay Assembler/ Disassembler) logical configuration enables the port to perform Frame Relay encapsulation/de-encapsulation for HDLC/SDLC-based protocols. The FRAD function encapsulates HDLC/SDLC traffic entering the network into a Cascade frame and then de-encapsulates it upon exiting the network. This function is restricted to one point-to-point PVC.

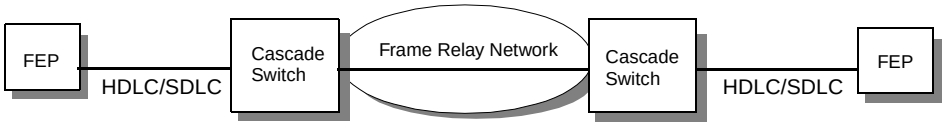

Figure 2-17. Frame Relay Direct FRAD Configuration

PPP Translation FRAD

The Frame Relay PPP (Point-to-Point protocol) Translation FRAD logical configuration enables the DTE device configured for PPP to communicate with another DTE device on the network configured for Frame Relay using RFC 1490 multiprotocol encapsulation. This configuration allows for a single circuit to be established between the two devices.

Cascade implements this feature by stripping the PPP header, translating the frame payload from PPP encapsulation to RFC 1490 encapsulation, and then applying an appropriate Frame Relay header. From the ingress Cascade switch through the network to the destination CPE device is a Frame Relay PVC. Frame Relay is completely transparent to the PPP devices.

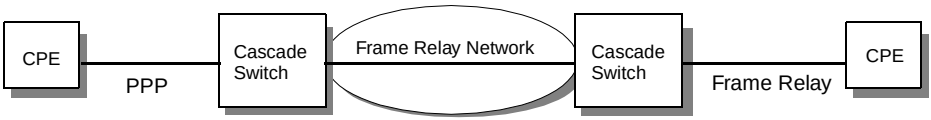

Figure 2-18. Frame Relay PPP Translation FRAD Configuration

Frame Relay OPTimum Trunk

The Frame Relay OPTimum Trunk logical configuration provides a trunk through a switched Public Data Network (PDN) between Cascade switches. This feature is referred to as Cascade Open Packet Transport (OPTimum) trunking.

The Frame Relay trunk requires either a UNI-DTE feeder or a Frame Relay NNI to be configured on the same physical port to allow for the Link Management exchange between the two connections. This tunneling configuration maintains the Cascade Frame Relay header.

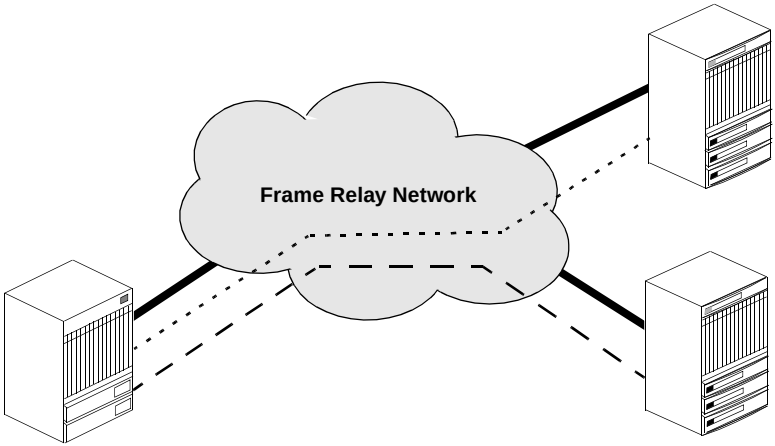

Figure 2-19. Frame Relay OPTimum Trunk Logical Port Configuration

Direct Line Trunk

The Direct Line trunk logical configuration allows for a direct trunk connection to another Cascade switch. The trunk connection is used to carry traffic destined for other switches in the network using Cascade's trunk protocols. This configuration takes advantage of Cascade's value-added features, such as manually defined paths, and green, amber, and red rate monitoring designators for frames.

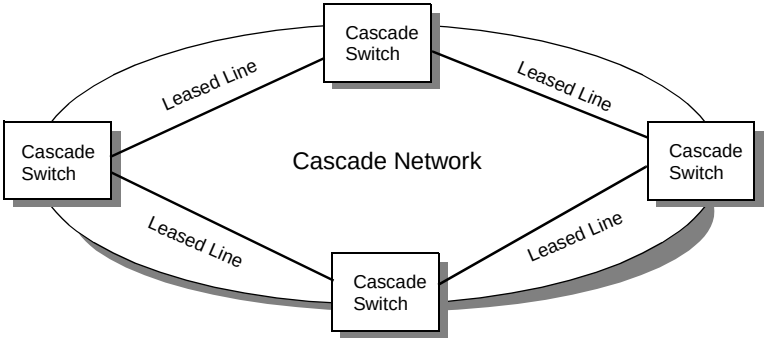

Figure 2-20. Direct Line Trunk Logical Port Configuration

Frame Relay Multicast

Frame Relay multicasting enables a device to forward a frame on a particular DLCI into the Frame Relay network. The Frame Relay network then broadcasts the frame to a predefined list of destinations. The network manager preconfigures the list of destinations.

Cascade's implementation of Frame Relay multicast provides a point-to-multipoint frame delivery service. This service is connection-oriented. To send multicast data, the network manager must first create individual PVCs from the site that sends the broadcasts to the sites that are to receive the broadcasts. After the PVCs are defined, the network manager must establish a ***Multicast Group***.

A Multicast Group consists of a multicast DLCI with a list of member PVC DLCIs participating in a multicast communication. The Multicast Group is a logical entity providing multicast service to all members.

Multicast is uni-directional. The incoming frame is multicast to multiple destinations, as shown in Figure 2-21.

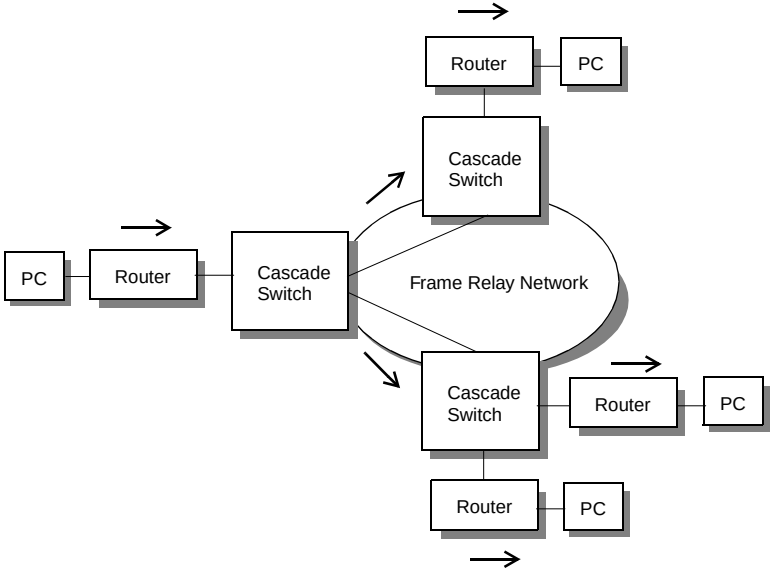

Figure 2-21. Frame Relay Multicast

Each Cascade B-STDX 8000/9000 supports 32 Multicast Groups. Although the number of members in a Multicast Group is theoretically unlimited, there are performance considerations that restrict the number of members.

Frame Relay In-Band Network Management

Frame Relay in-band network management enables the NMS operator to remotely manage the Cascade network over Frame Relay to support Simple Network Management Protocol (SNMP) packet transports within IP on User Datagram Protocol (UDP) envelopes.

Cascade supports Frame Relay in-band network management through the configuration of a Management DLCI. A ***Management DLCI*** is used when the NMS is connected to a LAN on which a router with a Frame Relay connection to the gateway switch resides.

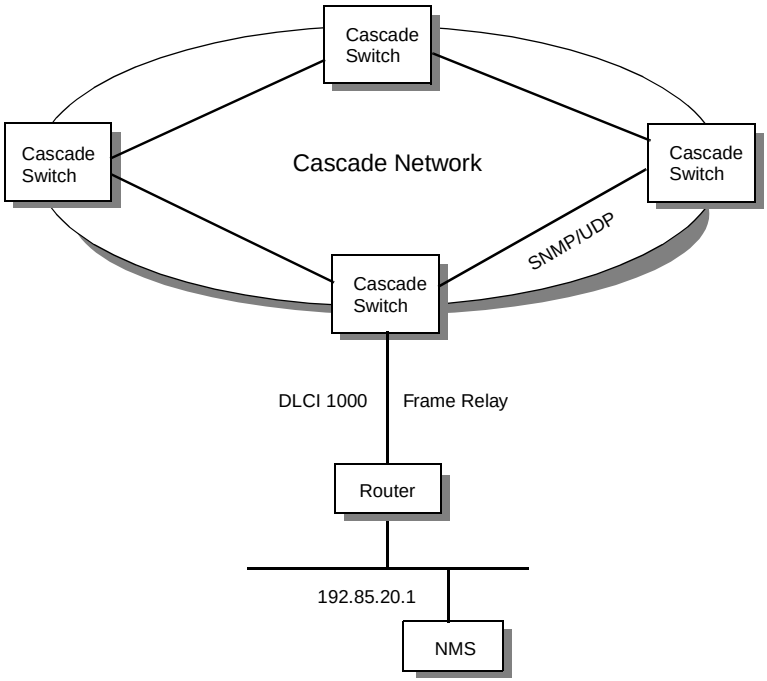

Figure 2-22. Frame Relay In-Band Network Management

The router attached to the Cascade switch must have a route configured for a reserved DLCI that points to the IP address for the gateway switch.

Other network management configurations include SMDS In-Band Management described in Chapter 3, and Ethernet, Indirect Ethernet, and SLIP, which are described in the *Cascade B-STDX 8000/9000 Hardware Installation Guide*.

Frame Relay Management MIB

To provide customers with the ability to monitor the characteristics of the Frame Relay network to which they are connected, Cascade supports the Frame Relay Forum standard Customer Network Management (CNM) MIB. The CNM MIB defines all of the available parameters regarding the characteristics of a Frame Relay network.

To implement this feature, CascadeView includes both a CNM MIB and a ***proxy agent***. The Cascade proxy agent enables users to access Cascade switch configuration and operational status information. The CNM proxy agent is based on SNMP. It supports NNI and UNI-DCE logical ports.

CNM supports the following Frame Relay Forum standards and MIBs:

- RFC 1604 Frame Relay (Service MIB)
- RFC 1232 (DSI MIB)
- RFC 1213 MIB II (System and Interfaces Group)

Figure 2-23 shows an example of a Cascade CNM implementation.

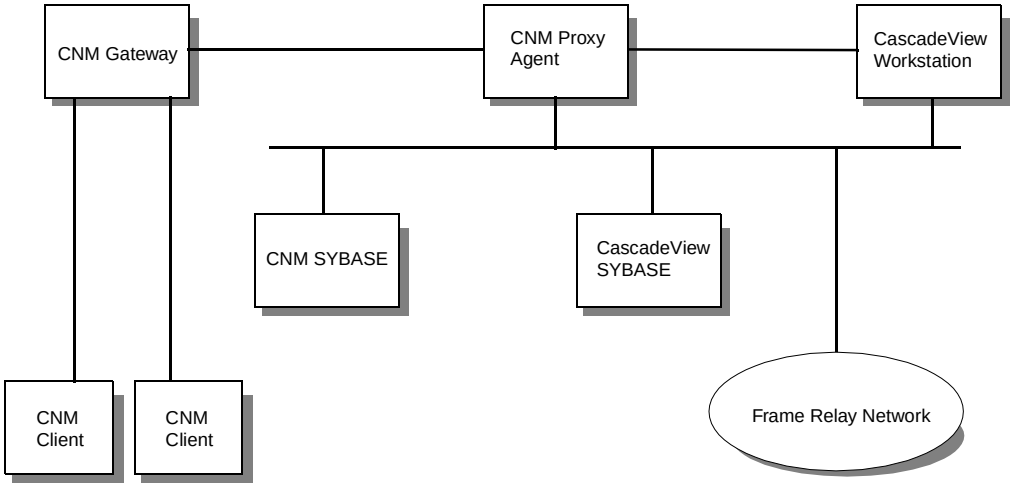

Figure 2-23. Customer Network Management (CNM) MIB

For more information on Cascade's CNM MIB, refer to the *STDX and B-STDX MIB Definitions*.

3

SMDS Services

Background

Switched Multimegabit Data Service (SMDS) is a public data switching service that provides LAN-like features and performance across wide geographic areas. SMDS was developed by Bellcore and is offered as a service by Local Exchange Carriers (LECs) in many metropolitan areas.

SMDS uses the same fixed-size cell relay technology as Asynchronous Transfer Mode (ATM). However, SMDS is a ***connectionless*** service that provides a flexible any-to-any communication capability.

Technology Fundamentals

The following sections outline the basic concepts for understanding SMDS services.

What Is SMDS?

SMDS is a high-speed data service that currently offers access at rates up to Digital Signal Level 3 (DS3). SMDS has many applications, including LAN interconnections and high-speed remote database access. The cost of SMDS is typically a flat rate, defined by an access class.

In the original SMDS architecture, Customer Premise Equipment (CPE) connects to the SMDS Switching System (SS) through a Distributed Queue Dual Bus (DQDB) ***Subscriber-Network Interface (SNI)***, as shown in Figure 3-1. The CPE uses the ***SMDS Interface Protocol (SIP)*** to communicate with the network supporting SMDS across the SNI.

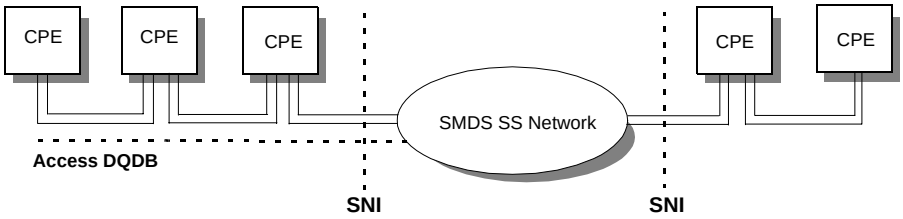

Figure 3-1. DQDB-based SMDS Architecture

The SNI is the interface between the network supporting SMDS and the subscriber-owned equipment. The CPE attaches to an access facility that links it to an SMDS Switching System over a dedicated path. Only data that originates from or is destined to the CPE is transported across the SNI.

The SMDS Switching System validates the source address associated with each data unit in an SMDS address assigned to the SNI from which the data was sent. The source address and destination address screening features validate the addresses used.

Connectionless Protocol

Each SMDS data unit is transmitted independently, and contains addresses that identify both the sender and receiver. The service provider assigns SMDS addresses that identify the SNI from which the data unit was sent and/or the SNI for which it was destined.

There are two types of addresses:

Individual addresses — An individual address is uniquely assigned to a single SNI. Each SNI can have a maximum of 16 addresses.

Group addresses — A group address, which is analogous to a LAN multicast address, allows an SMDS data unit to be delivered to multiple SNIs.

For more information on individual and group addresses, refer to “Individual and Group Addressing” on Page 3-15.

The ITU-T E.164 standard defines the SMDS address format. In the United States, individual and group SMDS addresses consist of the prefix 1, followed by a 10-digit number. In other countries, SMDS addresses begin with the appropriate country code. An Address Type field preceding the address identifies whether the address is an individual or a group address.

The SMDS Interface Protocol (SIP) defines how the CPE communicates with the network supporting SMDS across the SNI. SIP is a connectionless protocol based on the DQDB protocol defined in the IEEE 802.6 standard. It consists of three protocol levels that control the customer’s access to the network.

Level 3 — Contains the user data, source E.164 address, destination E.164 address, and other header information.

Level 2 — A 53-byte cell that contains the segmented L3_PDU and segmented header information.

Level 1 — The physical layer that contains the physical level interface.

A connectionless service such as SMDS routes each data unit independently to the destination. There are no connections to establish because each data unit is independent of the previous or subsequent one.

Compare a connectionless service to the process of mailing letters to a particular address through the postal service. The letters travel by the most efficient method to reach the destination. At the destination, the letters are given to a local mail carrier and delivered. No one knows the path the letters took to reach the destination. In comparison, a connection-oriented service, such as Frame Relay and ATM, is analogous to the process of establishing a telephone connection; the call creates an end-to-end route over which the information is sent.

Low-speed SMDS Architecture

Low-speed SMDS architecture enables those sites having low traffic volume to economically access the SMDS network. Low-speed SMDS enables customer premise equipment to access SMDS at a lower speed (n x DS0) using non-DQDB-based interfaces.

Low-speed SMDS differs from DQDB-based SMDS both in access speed and in the data transport mechanism used. Instead of using DQDB Level 2 and Level 1 interfaces, low-speed SMDS uses industry-standard HDLC and synchronous links as its Level 2 and Level 1 interfaces.

This SMDS architecture introduces a new element called the ***SMDS Access Server***.

The SMDS Access Server gives carriers the ability to combine high performance T3 connections with their SMDS switching systems, while offering SMDS as a service to users at speeds from 56 Kbps or 64 Kbps to T1/E1.

The SMDS Access Server architecture uses the ***Data Exchange Interface Protocol/Subscriber-Network Interface (DXI-SNI) Protocol***. The DXI-SNI allows a CPE to exchange Level 3 PDUs with the SMDS Access Server using DXI (HDLC-based) at Level 2. This frame-based protocol provides DS0 SMDS access at 56 Kbps and 64 Kbps.

Figure 3-2 illustrates the low-speed SMDS Access Server architecture using DXI-SNI.

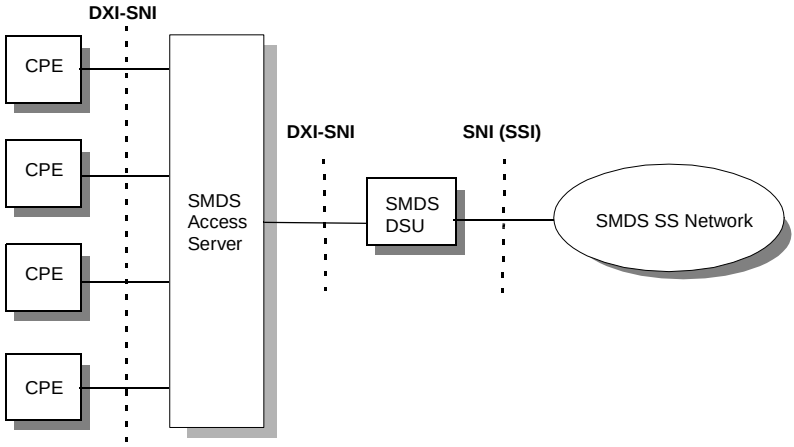

Figure 3-2. Low-speed SMDS Access Server Architecture

Cascade's SMDS Access Server and Switching Services

The remainder of this chapter describes in more detail how Cascade implements specific aspects of SMDS Access Server functionality on the B-STDX 8000/9000 WAN platform.

SMDS Access Server With DXI Switching

Cascade's DXI switching capability for SMDS enables the routing of SMDS packets between Cascade switches, effectively creating Cascade SMDS routing domains. This functionality, combined with support for low-speed SMDS access, is called ***SMDS Access Server/Switching***.

The Cascade SMDS Access Server/Switch configuration incorporates the following key functions:

Network Management — Provides support for the NMS to remotely manage the Cascade network using SMDS services.

System — Provides support for the individual address, group address, source address validation, address screening, and heart beat poll.

Operational — Provides support for the memory administration, maintenance, network traffic management, network data collection, and customer network management.

Performance and Quality of Service — Provides support for SMDS OPTimum Path and the Cascade SMDS Management Information Base (MIB).

The CascadeView/UX network management system manages the SMDS Access Server/Switch using SNMP commands from the Network Management Station (NMS). The Cascade SMDS Access Server/Switch has provision-driven functions to complete SMDS orders and provision the appropriate network resources to support new SMDS customers.

For more information about configuring SMDS services, refer to the *CascadeView/UX Network Configuration Guide*.

OSPF Routing

The OSPF routing protocol creates, maintains, and distributes the routing tables for adjacent Cascade switches in the SMDS network. OSPF treats SMDS addresses as IP destinations.

Using OSPF as the routing mechanism, the SMDS Access Server/Switch can route SMDS packets from one Cascade switch to another over the SMDS network or directly over high-speed trunk lines. The Cascade switch can also access a SMDS Switching System supporting high-speed, DQDB-based cell switching over a SMDS Data Service Unit (DSU).

The interface between the SMDS Access Server and the external DSU is the ***Data Exchange Interface (DXI)***. The interface between the DSU and the Switching System is the ***SMDS-to-Access Server Interface (SSI)***. In this scenario, the Cascade switch exchanges Level 3 data units with the DSU. The DSU then formats SIP Level 2 and Level 1 cells to access the SMDS Switching System.

Figure 3-3 illustrates the following:

- The Cascade SMDS Access Server with DXI-based Level 3 switching
- SMDS cell-based (Level 2 and Level 1) switching

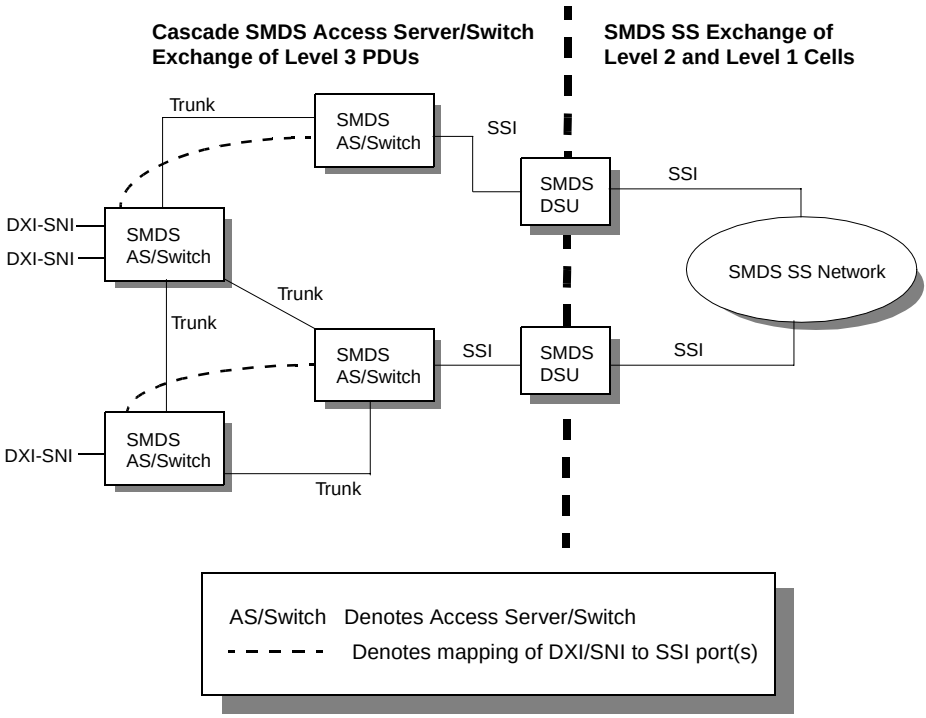

Figure 3-3. Cascade SMDS Switching

As shown in Figure 3-3, the SMDS Access Server/Switch can route SMDS data packets between switches in the OSPF routing domain, and multiplex data from DXI-SNI into the SMDS-to-Access Server Interface (SSI).

SMDS Access Server/Switch Configurations

The SMDS Data Service Unit (DSU) connects to the SMDS Switching System (SS) through the SMDS-to-Access Server Interface (SSI), thereby providing the access server functionality to forward data to the SMDS Switching System, or to provide Cascade's OPTimum trunking capability.

Cascade's OPTimum trunking allows Cascade switches to be connected across SMDS, Frame Relay, and ATM networks.

You can configure the same physical port with one or more logical trunk ports for SMDS OPTimum trunking. This capability allows a single physical port to be shared among several trunks. The Cascade switch also performs local switching to minimize the unnecessary use of switching and transmission resources in the network.

Figure 3-4 through Figure 3-6 show three configurations for Cascade's implementation of the SMDS Access Server with DXI switching architecture.

Figure 3-4 shows each Cascade switch connected through SSIs and trunks. SMDS data packets are sent from one switch to another through trunks. SMDS data units are sent through SSIs only when their destinations are not served by a Cascade switch.

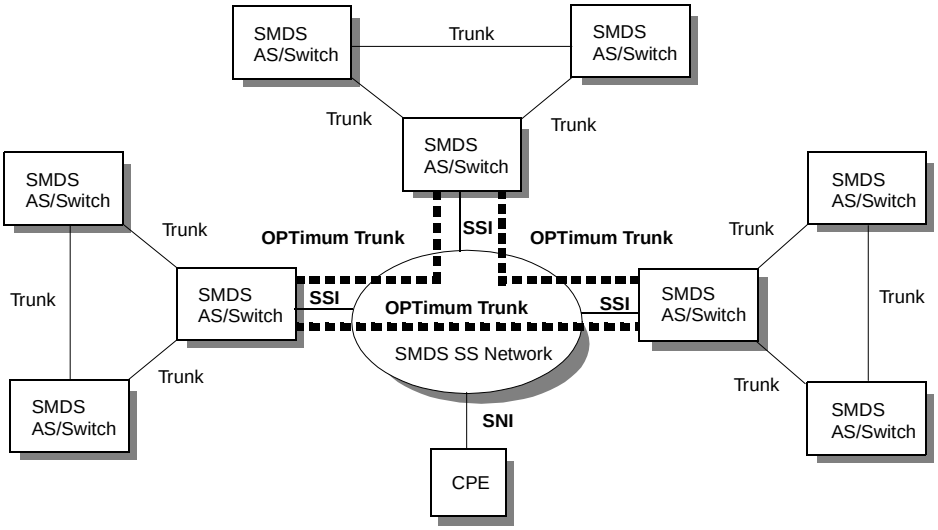

Figure 3-4. Cascade SMDS Access Server Networks Connected Through SSIs and SMDS OPTimum Trunk

Figure 3-5 shows several small Cascade SMDS Access Server networks connected to the SMDS SS network. Each small SMDS Access Server network consists of several Cascade switches connected with trunks.

SMDS data packets are sent from one switch to another in the same routing domain through trunks. However, SMDS data packets are sent from one domain to another through SSIs. Since there are no trunks connecting these small SMDS Access Server networks, SNMP/IP packet routing relies on routing and proxy Address Resolution Protocol (ARP). These addressing entities are explained later in the "SMDS Addressing" section.

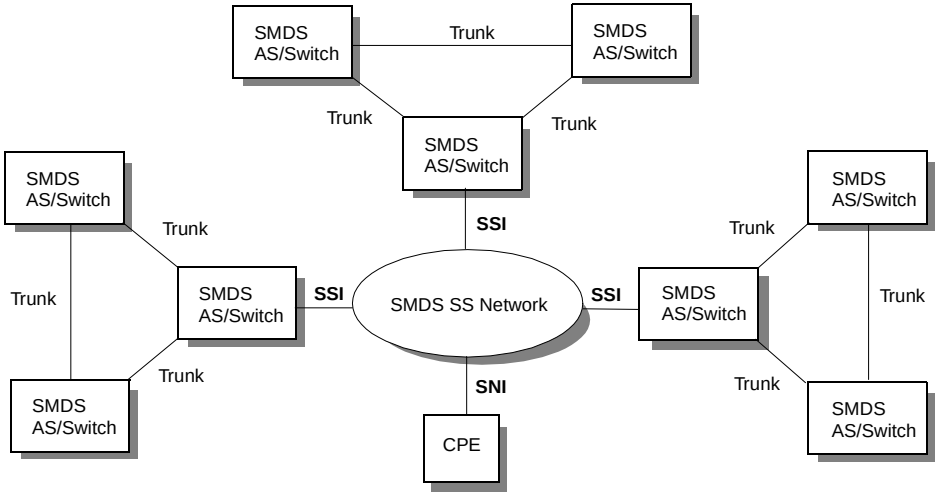

Figure 3-5. Cascade SMDS Access Server Networks Connected Through SSIs Only

Figure 3-6 shows all of the Cascade switches connected together with direct trunks.

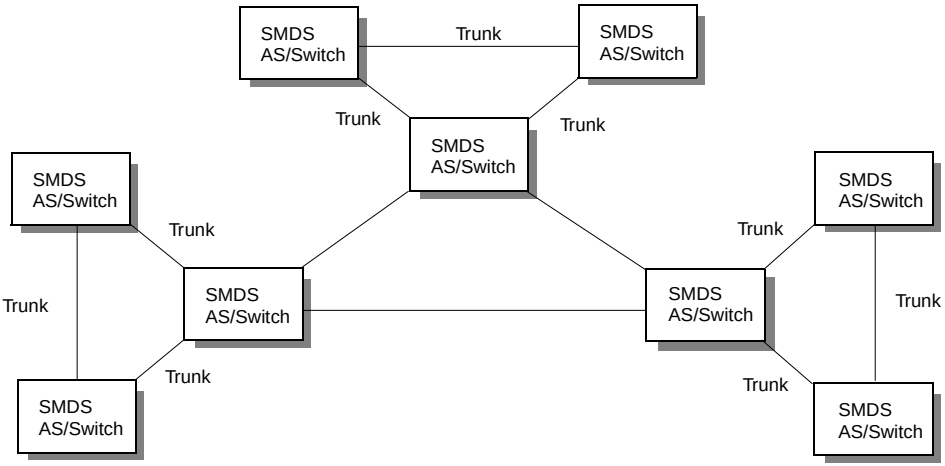

Figure 3-6. Cascade SMDS Access Server/Switch Networks Connected Through Direct Trunks

SMDS Addressing

This section describes how Cascade implements its SMDS addressing architecture for the SMDS Access Server/Switch.

Area IDs and Subscriber Numbers

In the Cascade SMDS Access Server with DXI switching architecture, an SMDS address is divided into two parts:

- Area number (Area ID)
- Subscriber number

The Area ID can start with any digit. The length can be up to eight digits (four bytes long for Binary Coded Decimal encoding).

The Cascade switch software automatically defines the Area IDs by applying a mask to the individual address. Figure 3-7 shows this addressing scheme.

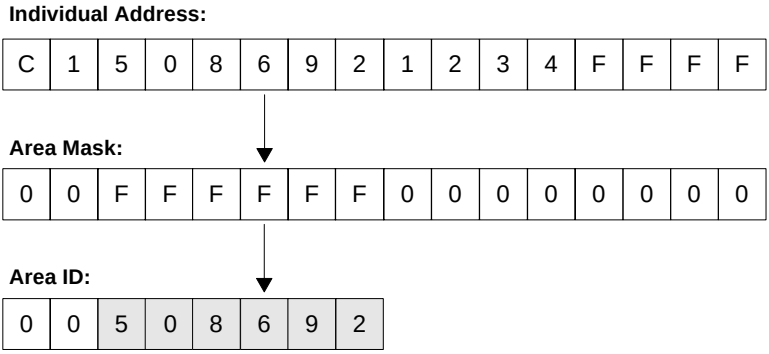

Figure 3-7. SMDS Access Server/Switch Addressing

When a new SMDS address is advertised on the network, the Cascade switch software applies the SMDS Address Mask to the individual address, and searches the database for the resultant Area ID. If the Area ID is not found, the new Area ID is added to the database and mapped to the switch's IP address.

Using CascadeView, you can set the SMDS Address Mask to mask from one to eight digits of the SMDS address. The SMDS Address Mask setting is applied to all SMDS addresses in the Cascade network.

Figure 3-8 shows how Area IDs and subscriber numbers are assigned in the SMDS network.

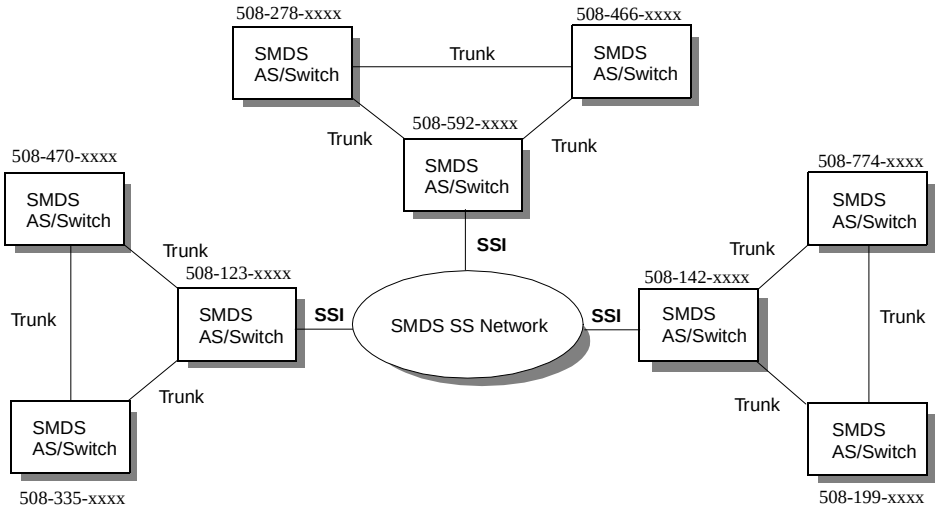

Figure 3-8. SMDS Area IDs and Subscriber Numbers

In Figure 3-8, the Area IDs are defined by the first six digits of the SMDS individual address (for example, 508335). The “xxxx” portion of the SMDS address represents a range of subscriber numbers. Area IDs are mapped internally to IP addresses and used by OSPF to establish routing tables.

Individual and Group Addressing

SMDS uses two types of addressing to transport Level 3 PDUs:

Individual Addressing — Individual addressing enables a CPE to send data units to a single address.

Group Addressing — Group addressing enables a CPE to send the same data unit to multiple addresses.

When a CPE sends a group addressed data unit, the network supporting SMDS delivers copies of the data unit to a set of DXI-SNIs identified by the individual destination addresses defined in the group.

In CascadeView, SMDS source and destination address screening features enable you to restrict incoming and outgoing data to particular CPE destinations, as needed. Figure 3-9 illustrates the association between individual addresses, Area IDs, and IP addresses.

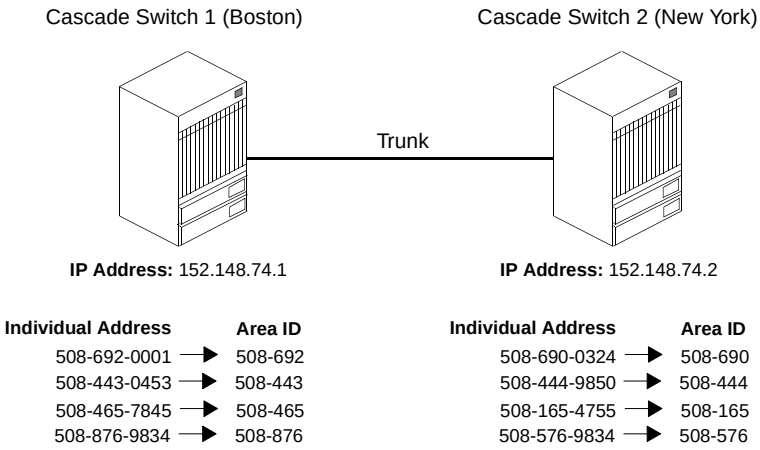

Figure 3-9. SMDS Access Server/Switch Addressing

Routing and Virtual Paths

To route SMDS packets in the OSPF domain, the Cascade switch establishes a Virtual Circuit (VC) between the source and the destination switch. This routing method shares the advantages of the Cascade Frame Relay PVC support, including

- In-sequence packet routing
- Low overhead (five bytes long)
- High performance

A Virtual Path is established between each pair of Cascade switches (with SMDS service) in the OSPF routing domain.

Figure 3-10 illustrates the Virtual Path structure in the Cascade switch. The Virtual Paths are node-to-node connections not port-to-port connections. Keep in mind that you can have only one Virtual Path between a pair of switches.

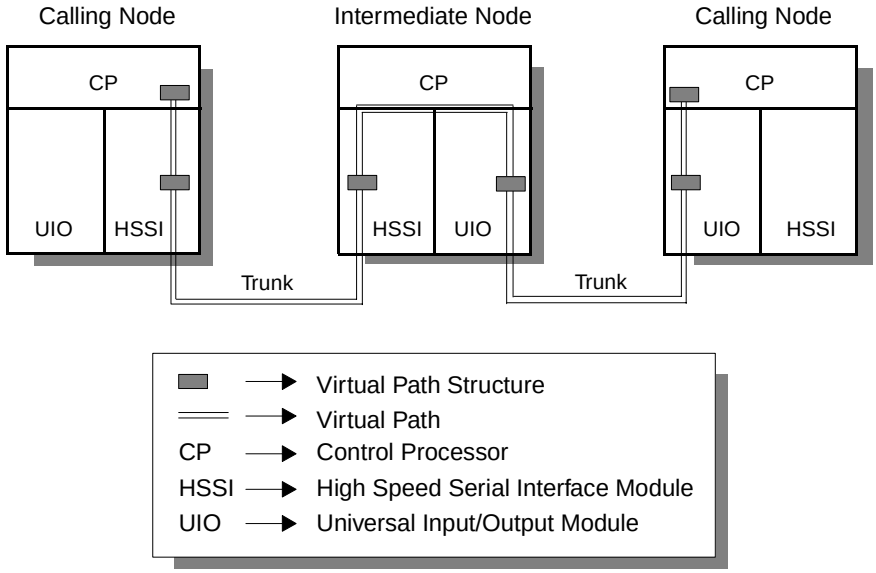

Figure 3-10. Virtual Path Structure

Figure 3-11 illustrates the concept of Virtual Paths used to route SMDS packets between pairs of Cascade SMDS Access Server/Switches in the same OSPF routing domain.

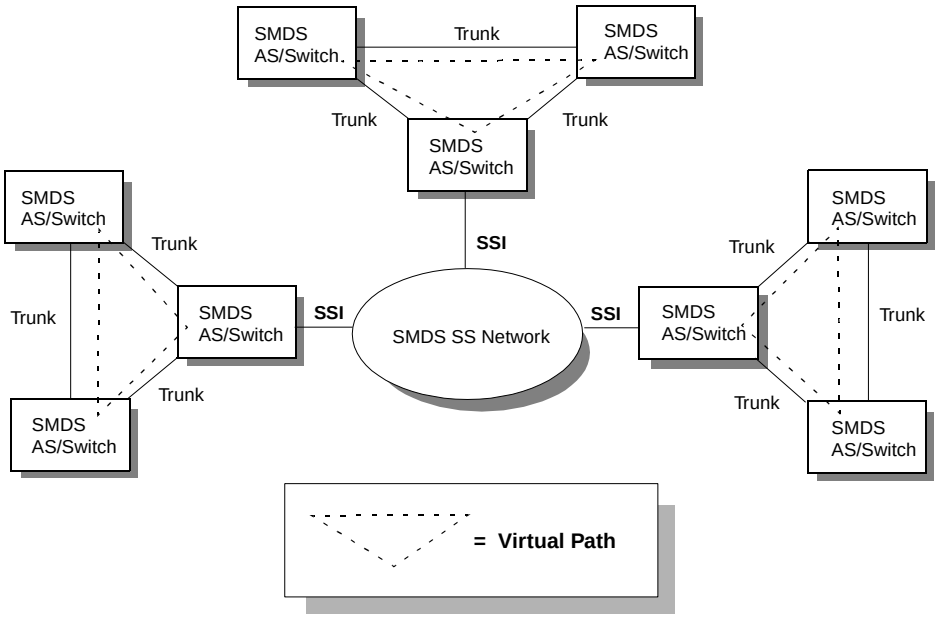

Figure 3-11. Virtual Paths Between Pairs of Cascade SMDS Access Server/Switches

The Virtual Path information for routing SMDS packets through a Cascade network is constructed during the OSPF protocol message exchange between the switches at initialization and/or when a node becomes active.

Routing Individual Addressed Packets

The SMDS Switching System routes SMDS individual addressed packets according to the following rules:

1. If the Area ID of the destination address is in the Cascade switch, route the packet to another DXI-SNI locally in this switch.

2. If the Area ID of the destination address is not in this Cascade switch, but is in the Cascade OSPF routing domain, route the packet following the virtual path to the destination node. The destination node then routes the packet to the destination DXI-SNI according to the entire SMDS address.

3. If the Area ID of the destination address is not in the Cascade OSPF routing domain, and the input DXI-SNI port maps to an SSI, route the packet locally to the SSI (if the DXI-SNI and the SSI are in the same switch), or route the packet to the node to which the SSI is attached (if the DXI-SNI and the SSI are not in the same switch). In the latter case, the destination node then routes the packet to the appropriate SSI.

Figure 3-12 illustrates these rules.

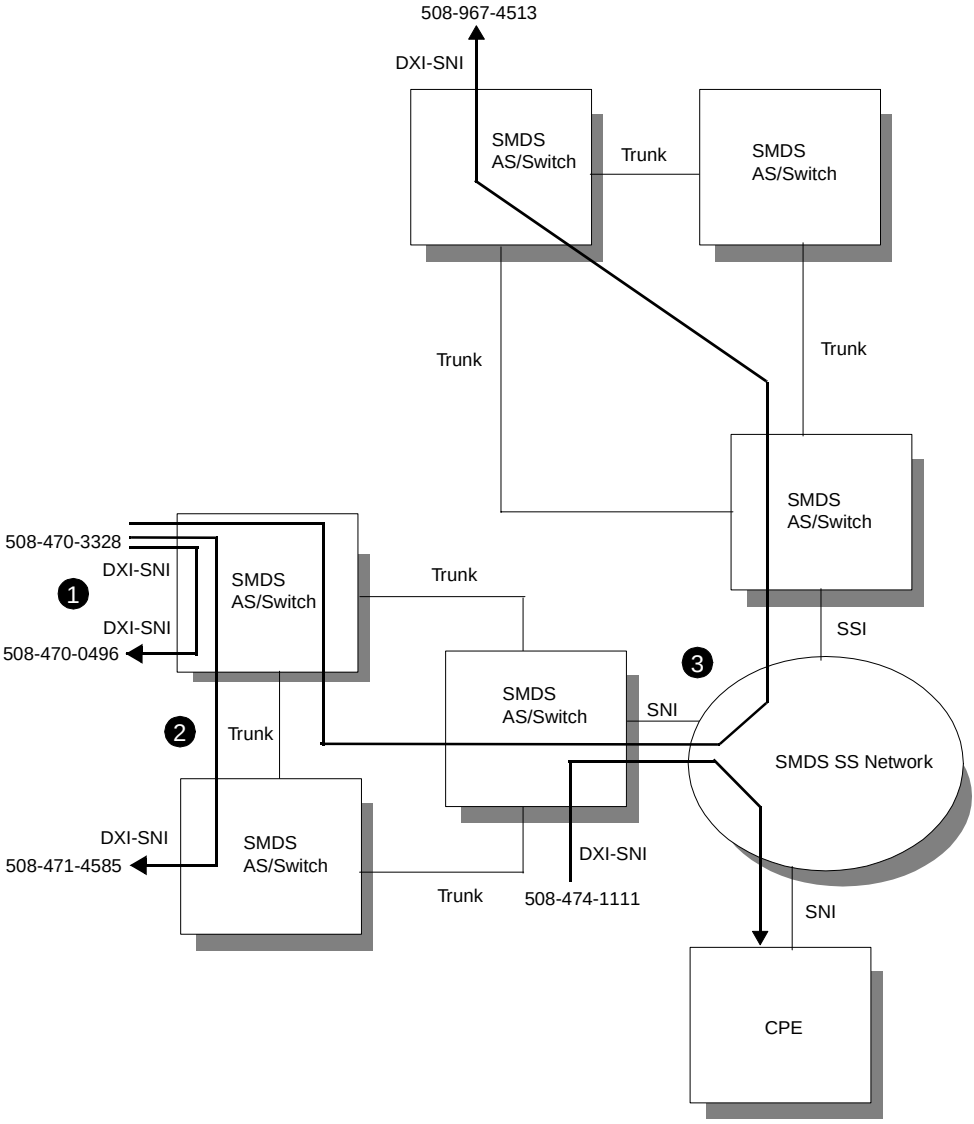

Figure 3-12. Individual Addressed Packet Routing

SMDS address screening validates the source address of each data unit to ensure that sender of an SMDS data unit does not indicate a false source address. In addition, t SMDS Access Server/Switch validates the following:

- The source address specified by the sending CPE is an individual address
- The source address specified by the sending CPE is legitimately assigned to the DXI-SNI from which it was sent

Routing Group Addressed Packets

When the SMDS Access Server/Switch receives a group addressed packet from a DXI-SNI, it forwards the packet using the following process:

1. Routes the packet to each DXI-SNI whose address is in the same group, in the sa switch, and maps to the same SSI.
2. Blindly sends a copy of the packet to each switch with SMDS service in the Casc OSPF routing domain.
3. For each switch that receives a copy of the packet, sends a copy to each of the switch's DXI-SNIs whose address is in the same group and maps to the same SS
4. Any switch that connects to the SMDS SS network sends a copy to the network through the SSI.
5. If a packet is received from the SSI, it is not re-sent over an OPTimum trunk configured on the receiving switch.

Proxy ARP

Proxy ARP service is used to manage the Cascade SMDS Access Server/Switches throughout the SMDS network. In Figure 3-13, the NMS is attached to the SMDS Switching System through a router.

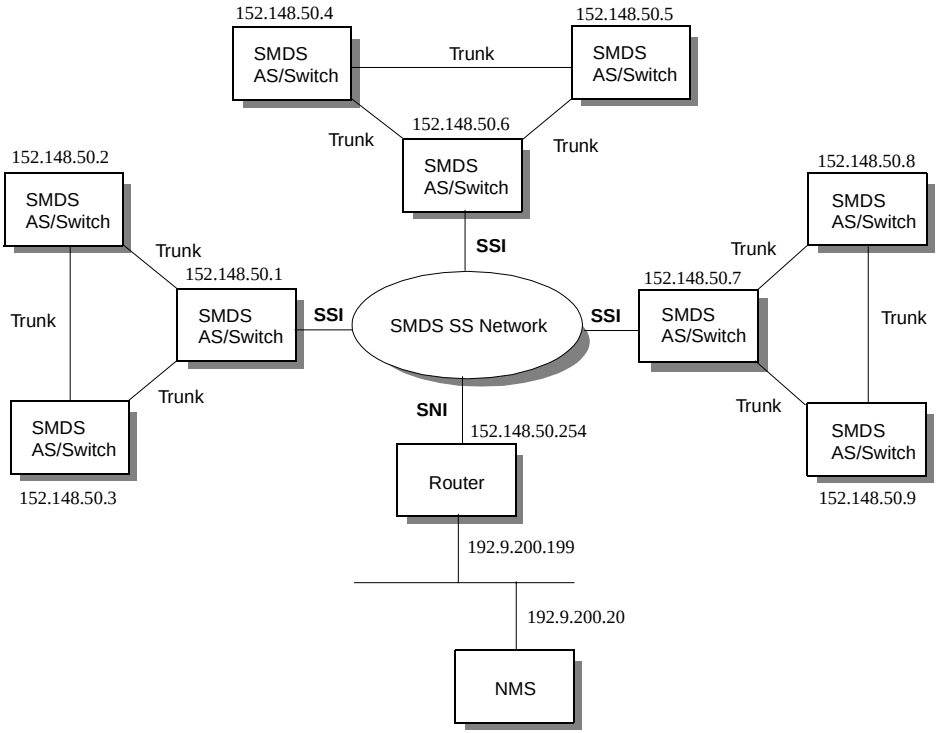

Figure 3-13. Cascade Switches Support Proxy ARP

The NMS manages the network traffic using SMDS In-Band network management. To be managed from this NMS, all SMDS Access Servers/Switches must be in the same IP subnet.

With an attached SSI, the SMDS Switching System network provides proxy ARP service for SMDS Access Server/Switches that are not attached to the SMDS Switching System network in the routing domain.

SMDS In-Band Network Management

SMDS In-Band network management enables the NMS operator to remotely manage the Cascade network using SMDS services to transport the SNMP/UDP/IP protocol packets.

Cascade supports two configurations for SMDS In-Band management of the Cascade network:

- In the first configuration, the Cascade switch uses the SSI-DTE (SMDS-to-Access Server Interface/Data Terminal Equipment) to communicate with the router/NMS through the SMDS Switching System.
- In the second configuration, the Cascade switch uses the DXI-SNI (Data Exchange Interface Protocol/Subscriber Network Interface) to communicate with the router/NMS.

In-Band Management Using SSI-DTE

Using SSI-DTE, the SMDS Access Server/Switch functions like a CPE that is directly attached to the SMDS Switching System network. Figure 3-14 illustrates a sample SSI-DTE configuration.

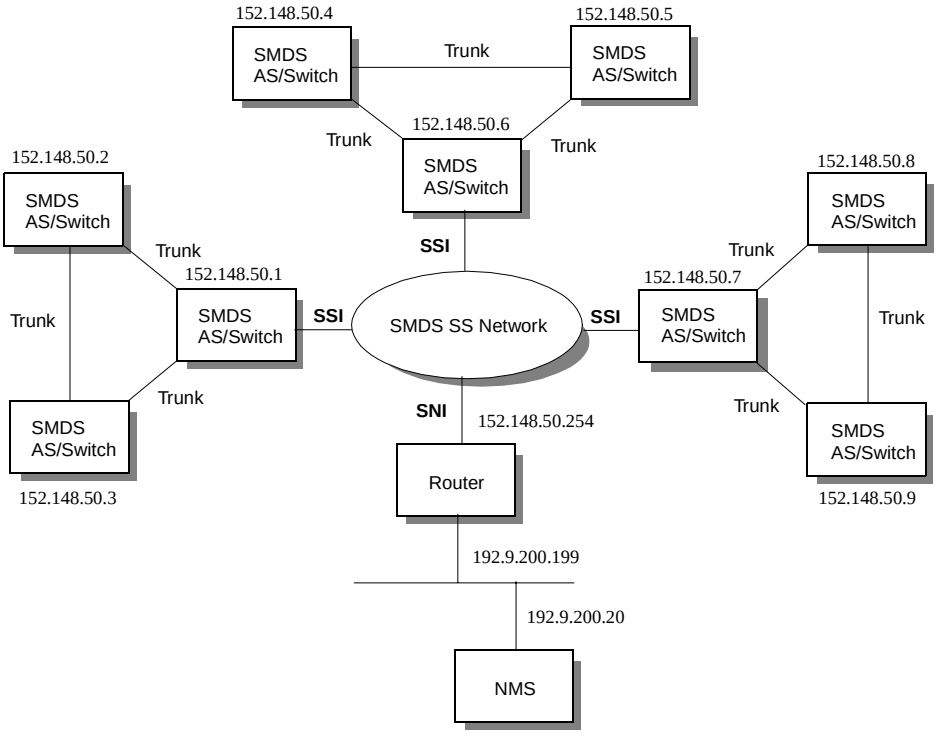

Figure 3-14. SSI-DTE In-Band Management

Using SSI-DTE, the SMDS Access Server/Switch does the following:

- Assigns an SMDS individual address on the SSI-DTE for the Cascade switch. This address is also used for the OPTimum Trunk source address.
- Assigns an SMDS group address for the multicast ARP requests.
- Assigns an IP address for the SSI-DTE.
- Supports ARP over SMDS.
- Specifies the SSI-DTE as the interface and the router as the next hop to reach the NMS.

In-Band Management Using DXI-SNI

Using DXI-SNI, the SMDS Access Server/Switch communicates to the router/NMS using a DXI-SNI point-to-point connection.
Figure 3-15 illustrates a sample DXI-SNI configuration.

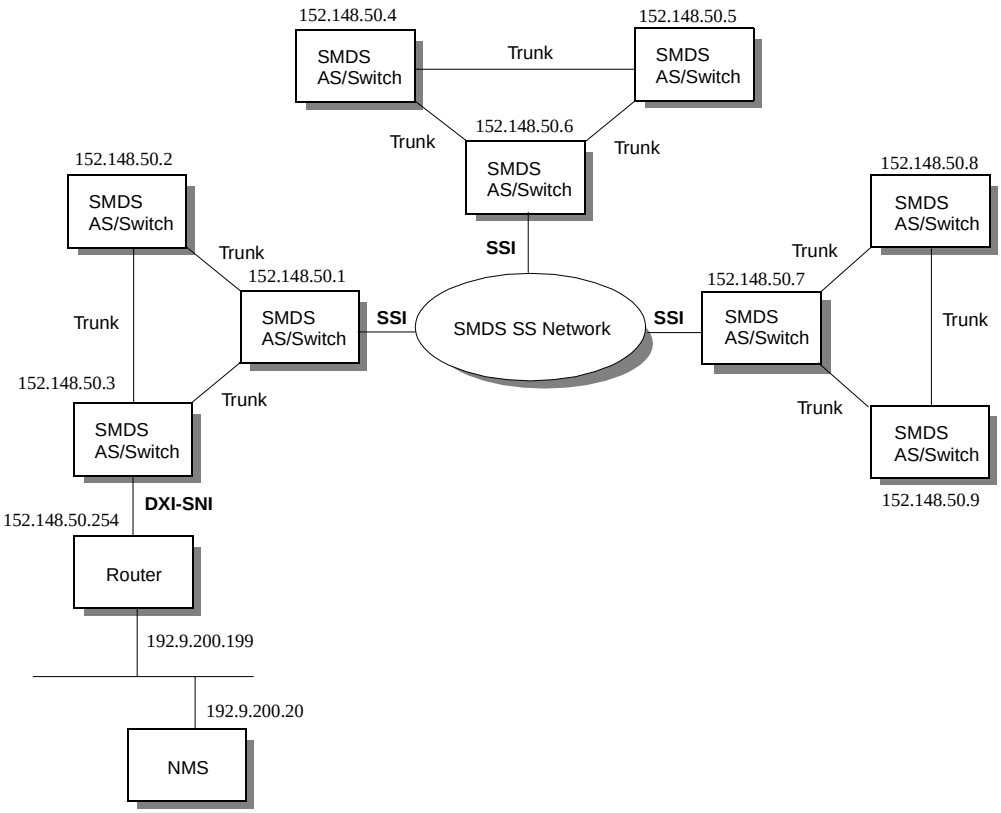

Figure 3-15. DXI-SNI In-Band Management

Using DXI-SNI, the SMDS Access Server/Switch does the following:

- Assigns an SMDS Individual address on the DXI-SNI for the Cascade switch
- Assigns an SMDS Group address for the multicast ARP requests
- Assigns an IP address for the DXI-SNI
- Supports ARP over SMDS
- Specifies the DXI-SNI as the interface and the router as the next hop to reach the NMS

4

ATM Services

Background

Asynchronous Transfer Mode, or ATM, is the latest networking technology intended to support a wide range of advanced high-speed data applications. ATM is designed to remove the barriers between local and wide-area networks by providing seamless interconnection for LAN interworking. ATM eliminates these barriers by removing data transmission speed as an issue and providing flexible bandwidth-on-demand.

In the past, companies built large networks to accommodate a specific kind of data transmission. These included voice networks, data networks, and television networks. Many times this led to both a duplication of effort and tremendous cost outlays. Because many of these networks were built for peak load conditions, the average utilization was typically very low, leading to excessive costs.

Therefore, it became important for companies to find ways of using a single network infrastructure and assigning bandwidth on an as-needed basis. ATM lets both private corporations and public service providers build unchannelized networks to make more efficient use of the underlying bandwidth on the network.

Currently, many users require data services above DS1, but not yet DS3. By offering scalable rates from 1.5 Mbps to 155 Mbps or higher, ATM services can make the wide-area network transparent for applications. Unlike Frame Relay or other data services, ATM can easily accommodate delay sensitive traffic such as voice and video.

Technology Fundamentals

The following sections describe some of the basic ATM concepts. This discussion is not intended to be a comprehensive overview on the emerging ATM networking technology. For additional information, learning materials, and the latest developments in the ATM standard, consult the ATM Forum specifications.

What Is ATM?

ATM is based on connections, not channels. The term “asynchronous” refers to the way in which ATM achieves its unchannelized bandwidth allocation. ATM only sends data associated with a connection when there is actual data to send. This is in contrast to channelized or Time Division Multiplexed (TDM) networks, where even if a channel is idle, a special bit pattern must be sent in every time slot representing a channel.

Unlike X.25 or Frame Relay, ATM uses very short, fixed-length packets called ***cells***. The ATM cell is 53 bytes long, consisting of a five-byte header containing an address, and a fixed 48-byte information field or ***payload***. In contrast, Frame Relay uses a two-byte header and a variable-length information field, as shown in Figure 4-1.

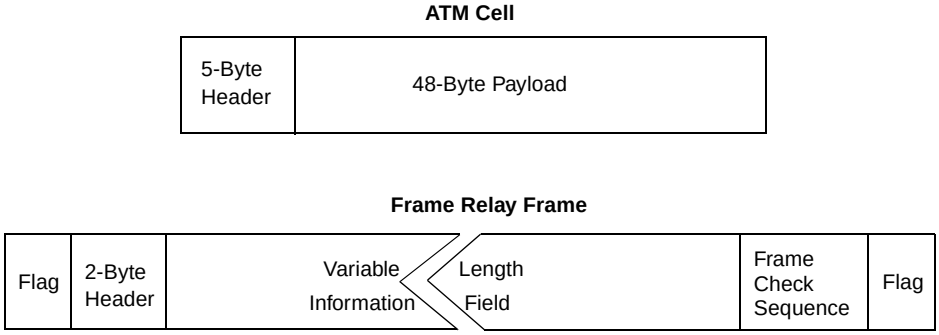

Figure 4-1. Cells vs. Frames

For more information on ATM cells, refer to "ATM Cell Structure" on Page 4-12.

ATM functionality corresponds to the Physical Layer and a portion of the Data Link Layer of the Open Systems Interconnection (OSI) Reference Model.

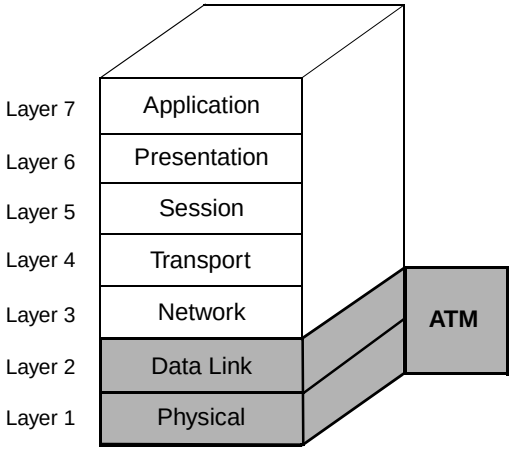

Figure 4-2. ATM and the OSI Reference Model

The ATM protocol functionality must be implemented in both user equipment such as hubs and routers, and network elements such as switches. The network does not know anything about the end-to-end application that is running over an ATM connection after the connection has been established.

How ATM Works

ATM is application transparent. The ATM cell size is a compromise between the long frames generated by data communications applications and the short repetitive needs of voice traffic. As such, ATM allows a free mix of data, voice, and video within the same application.

ATM supports four service classes to handle the various data types on a network. This ensures optimal network usage and guaranteed end-to-end delivery.

The four ATM service classes include the following:

Constant Bit Rate (CBR) — Handles digital information, such as video and digitized voice, that must be represented by a continuous stream of bits.

Variable Bit Rate-Real Time (VBR-RT) — Handles the packaging of special delay-sensitive applications, such as packet video, that requires low cell delay variation between endpoints.

Variable Bit Rate Non-Real Time (VBR-NRT) — Handles packaging for the transfer of long, bursty data streams over a pre-established ATM connection. This service is also used for short, bursty LAN traffic.

Available Bit Rate/Unspecified Bit Rate (ABR/UBR) — Handles LAN traffic.

There are currently three different types of ATM Adaptation Layers (AALs) to handle the different types of ATM traffic. For more information on the ATM Adaptation Layer (AAL), refer to
"ATM Adaptation Layer (AAL)" on Page 4-10.

ATM takes information such as voice, video, and data from multiple sources and multiplexes it into a cell stream, as shown in Figure 4-3.

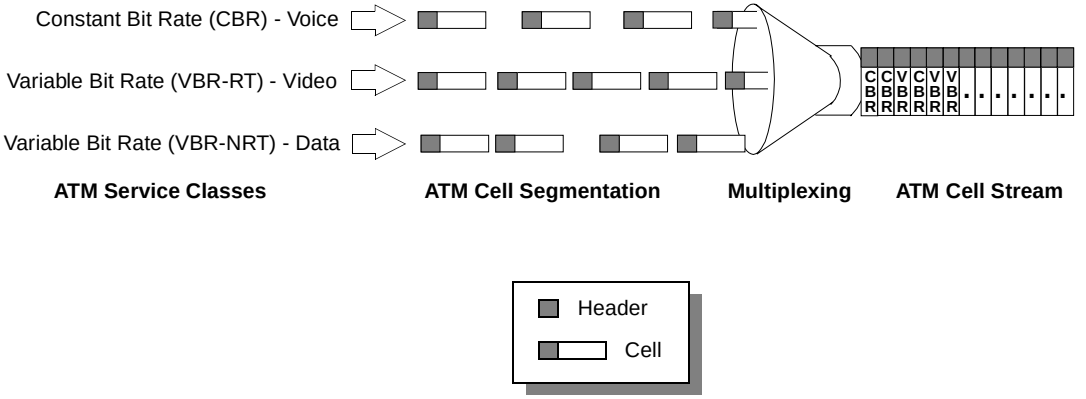

Figure 4-3. ATM Multiplexing

Multiplexing defines the means by which several streams of data share a common physical transmission medium. Switching takes the instances of a physical transmission medium containing multiplexed information streams and rearranges the information streams between input and output.

Therefore, information from one physical link in a specific multiplex position is switched to another output physical link. Conventional LANs, such as Ethernet and FDDI, use shared media where only one node can transmit at a time.

Bandwidth Efficiency

As a cell relay technology, ATM differs from Time Division Multiplexing, or TDM. TDM is a synchronous mode of data communication that digital telephone networks have long used.

TDM networks move information in fixed-length 8-bit "time slots." Because TDM networks are designed to carry voice, 8,000 time slots per second (the rate of sampling used for the digitalization of voice) are allocated to each connection. This concept is what creates the basic 64-Kbps single digital channel.

Moving information around in time slots works only if the time slots are synchronized. Information arriving on a node can expect to find its allocated time slot waiting to carry the information out. If input and output time slots are skewed from each other, the output slot might "depart" before it was filled. This form of transmission is sometimes called "synchronous transfer" because network clocks are synchronized to assure time slots match throughout the network.

Figure 4-4 represents a sample TDM data stream. Time slot "User A" occurs right after the Frame Bit (F/B) and every third time slot thereafter. To locate User A data, start with the time slot after the Framing Bit, and count every third time slot thereafter. This is the synchronous nature of TDM.

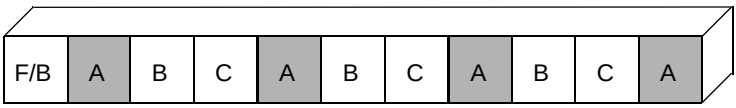

Figure 4-4. Time Division Multiplexing

Unlike TDM, ATM attempts to allocate bandwidth more efficiently by giving users access to the entire communications channel on-demand. If the channel is in use, a user may have to wait to gain access. However, because cells are small and fixed-length compared to frames, the delay is minimal and can be controlled.

Using ATM multiplexing, for example, User A can transfer data using the entire bandwidth of the channel. User B inserts messages occasionally, as needed, and User C takes up almost no capacity.

Figure 4-5 represents a sample ATM data stream. Users can access the channel on a somewhat random basis. Unlike TDM, a user's traffic occurs asynchronously with respect to any framing information in the channel, hence the asynchronous nature of ATM.

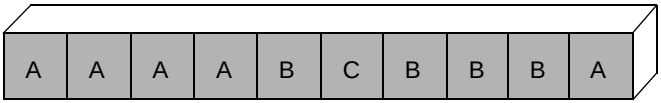

Figure 4-5. ATM Multiplexing

The ATM Standard

The ATM standard is defined as part of the Broadband Integrated Services Digital Network (B-ISDN) standard. Architecturally, ATM comprises the bottom three layers of the B-ISDN protocol stack, as shown in Figure 4-6.

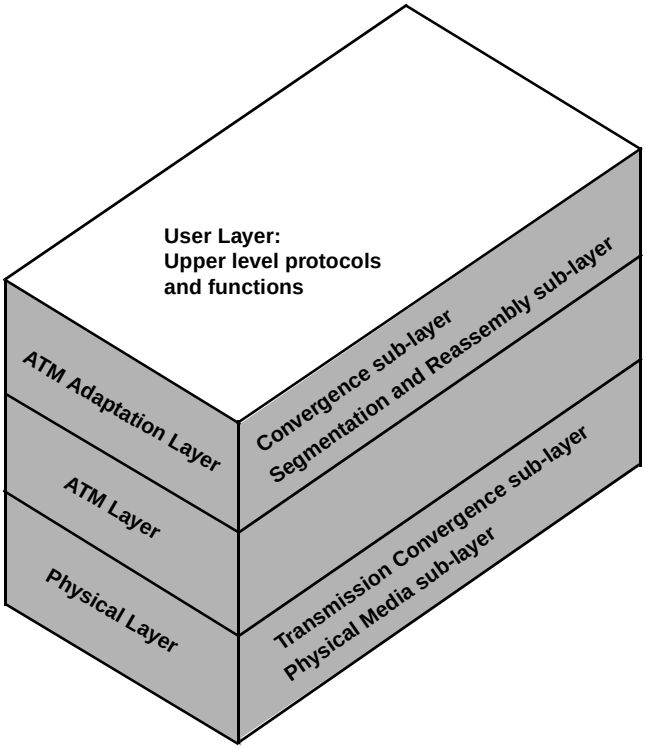

Figure 4-6. B-ISDN Reference Model for the ATM Standard

The three layers have the following functions:

Physical Layer — Defines the electrical or physical interface, line speeds, and other physical characteristics. The Physical Media sub-layer provides the bit transmission capability, including bit alignment, line coding, and electrical/optical conversion. (This corresponds to the Physical Layer of the OSI Reference Model shown in Figure 4-2.)

In addition to the generation and recovery of transmission frames, current ATM physical media includes the following:

Synchronous Optical Network (SONET) — A Bellcore specification that is currently being used in worldwide public data networks. It is a synchronous optical network-based User-Network Interface (UNI), either public or private, operating at speeds from 51 Mbps to 2 Gbps over single-mode optical fiber.

OC3c SONET — A SONET-based fiber-optic UNI, either public or private, operating at 155.52 Mbps over single-mode and multimode optical fiber.

Digital Signal Level 1 (DS1) — A public or private UNI operating at 1.544 Mbps over coaxial cable.

Digital Signal Level 3 (DS3) — A public or private UNI operating at 44.736 Mbps over coaxial cable.

The Transmission Convergence sub-layer is part of the Physical Layer and does the following:

- Generates and recovers the transmission frame
- Performs cell framing and recovery
- Performs bit timing
- Generates header error-control sequence

ATM Layer — Defines the cell format and provides for the transport of fixed-length cells between the endpoints of a virtual connection. The ATM Layer also provides multiplexing functions to establish multiple connections across a single UNI.

The ATM Layer performs the following functions:

- In the transmit direction, multiplexes ATM cells from individual virtual channels into one cell stream.
- At the receiving end, demultiplexes arriving cell streams into individual cell flows appropriate to the virtual channel.

ATM Adaptation Layer (AAL) — Defines the process of converting information from the upper layers into ATM cells. This layer deserves special attention because of the important role it plays in handling the different types of ATM traffic.

ATM Adaptation Layer (AAL)

The ATM Adaptation Layer (AAL) runs from the end system and is transparent to the ATM network. The AAL is organized into two logical sublayers:

Convergence sub-layer — Adapts the services provided by the ATM Layer to the requirements of the upper layers.

Segmentation and Reassembly sub-layer — Segments upper layer information into a size suitable for the payload of an ATM cell, and reassembles the contents of the ATM cell payload into upper layer information.

There are currently three different types of AALs to handle different types of ATM traffic.

Type 1: Constant Bit Rate (CBR) Services — Type 1 is an isochronous, constant bit-rate service for audio and video applications.

Type 3/4: Variable Bit Rate (VBR) Data Transfer Services — Type 3/4 resulted from the merging of Type 3 (connection-oriented) and Type 4 (connectionless) VBR Data Transfer services. Type 3 handles packaging for the transfer of long, bursty data streams over a pre-established ATM connection. Type 4 manages the packaging of short, bursty data streams such as LAN traffic, but is encumbered by additional overhead.

Type 5: Simple and Efficient Adaptation Layer (SEAL) — Type 5 is an extension of the Type 3 AAL. It simplifies the Segmentation and Reassembly sub-layer portion of the Adaptation Layer to pack all 48 bytes of the ATM cell payload with data. Type 5 makes ATM look like high-speed Frame Relay. It also assumes that only one message is crossing the UNI at a time. That is, multiple end-users at one location cannot interleave messages on the same virtual circuit, but must queue them for sequential transmission. Type 5 is mainly used for data applications.

The International Telecommunications Union-Telecommunications Sub-section (ITU-T) has defined certain service classifications based on how bits are transmitted, the required bandwidth, and the types of connections required. These include the following:

Class A (CBR) — Connection-oriented, constant bit rate data with a timing relationship between source and destination, for example 64 Kbps digital voice.

Class B (VBR-RT) — Connection-oriented, variable bit-rate video and audio with a timing relationship between source and destination.

Class C (VBR-NRT) — Connection-oriented, variable bit-rate with no timing relationship between source and destination, for example Frame Relay traffic.

Class D (UBR/ABR) — Connectionless, variable bit-rate data with no timing relationship between source and destination, for example SMDS traffic.

Figure 4-7 summarizes the AAL Service Classes.

	Class A	Class B	Class C	Class D
Timing	Required		Not Required	
Bit Rate	Constant	Variable		
Connection Mode	Connection-Oriented			Connectionless
Applications	Voice, Video	Compressed Voice, Video	Frame Relay	SMDS

Figure 4-7. AAL Service Classes

ATM Cell Structure

As shown in Figure 4-8, the ATM cell consists of two parts: a 5-byte header and a 48-byte information field, or payload. For networking purposes, only the header is significant.

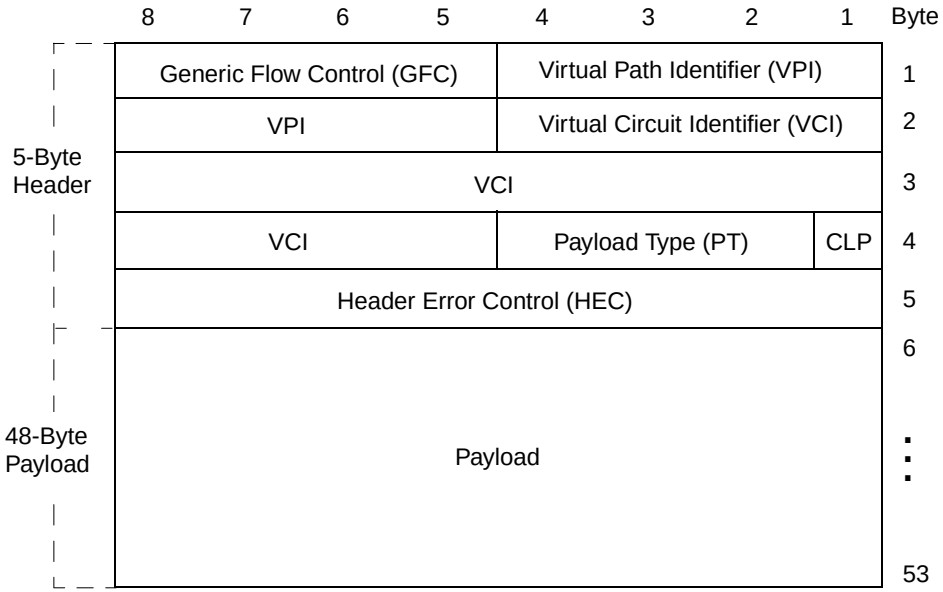

Figure 4-8. ATM Cell Structure

The ATM cell contains the following fields:

Generic Flow Control (GFC) — Controls the flow of traffic across the UNI and into the network. Currently, this field is not used.

Virtual Path Identifier (VPI) and Virtual Circuit Identifier (VCI) — Provides addressing identifiers used to route cell traffic. Because of their routing significance, the VPI and VCI addressing identifiers are described further in the next section.

Payload Type (PT) — Indicates the type of information carried by the cell. This designation is used by the network or terminating equipment to provide various traffic handling mechanisms.

Cell Loss Priority (CLP) — Indicates the cell Discard Eligible (DE) status, depending on current network congestion conditions.

Header Error Control (HEC) — Provides protection against misdelivery of cells due to addressing errors.

Payload — Follows the HEC field and contains 48 bytes of user data.

ATM Connections: VPIs and VCIs

ATM cells are sent between two points over a shared facility composed of ***Virtual Channels (VCs)***. These connections can be established on-demand (as a switched service), or pre-provisioned (such as Frame Relay PVCs).

A Virtual Channel is used to describe a connection between two communicating ATM entities associated by:

- One identifier value called the Virtual Path Identifier (VPI)
- One unique identifier value called the Virtual Channel Identifier (VCI)

A ***Virtual Path (VP)*** is used to describe multiple, bi-directional transport of ATM cells belonging to virtual channels that are associated by a common VPI. A virtual channel can consist of the following:

- A group of ATM links
- User equipment to a central office switch link
- Switch-to-switch link
- Switch-to-user equipment link

All communications proceed along this same VC, which preserves cell sequence and provides a certain quality of service. Figure 4-9 illustrates the concept of VCs and VPs.

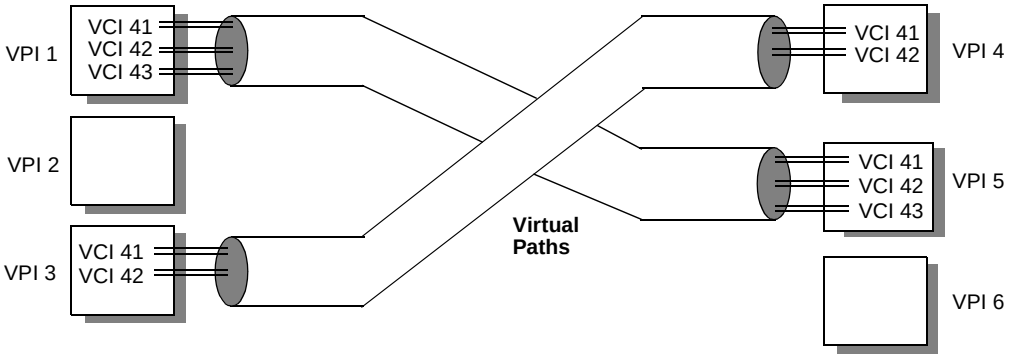

Figure 4-9. Virtual Channels and Virtual Paths

As previously mentioned, the cell header contains both a VPI and VCI, allowing a cell to be given a unique VC identifier and be associated with a particular virtual path. As shown in Figure 4-9, the same VCIs can be used within different VPs.

The VPI and VCI are used only for establishing connections between two ATM entities, not the end-to-end connection.

The ATM switch initially looks at the VPI. If there is a VP configured for that VPI, the ATM switch looks no further and switches cells based on the VPI. However, if there is no VP configured for that VPI, the switch looks at the VCI and switches cells based on the combined VPI/VCI.

VPIs and VCIs, like DLCIs for Frame Relay, have local significance at the switch. As the message is forwarded, the switch may change the VPI, and in the case of an isolated VC, may also change the VCI.

Cascade's Implementation of ATM Services

Cascade supports pure ATM cell switching on its Cascade 500 platform, as well as interworking with other established technologies, such as Frame Relay and SMDS, on its B-STDX platform. The remainder of this chapter discusses how the Cascade 500 ATM switch implements pure ATM cell switching.

Physical and Logical Port Architecture

A logical port refers to the type of service that is mapped to a physical port on the Cascade 500. Usually, there is a one-to-one correspondence between a logical port and a physical port. In the case of a channelized interface or an OPTimum trunk, there is a one-to-one mapping at the channel or virtual circuit/virtual path level. This enables multiple logical ports to be mapped to a single physical port.

The Cascade 500 supports the following physical port types:

- 8-port DS1/E1 ATM UNI
- 8-port DS3/E3 ATM UNI
- 4-port OC3c/STM-1 ATM UNI
- Single port OC12c/STM-4 ATM UNI

Some of the physical port parameters include the following:

- Port administration status
- Transmit clock source
- Bandwidth
- Cell payload scramble
- MIB interface number

For a complete list of physical port parameters, refer to the *Cascade 500 Network Administrator's Guide*.

The Cascade 500 supports the following logical port types:

ATM UNI DCE — Configures the logical port to communicate with an ATM CPE over ATM PVCs and SVCs.

ATM UNI DTE — Configures the logical port to communicate with an ATM switch over ATM PVCs and SVCs. The Cascade 500 acts as an access concentrator feeding multiple Frame Relay and/or ATM PVCs and SVCs to the ATM network via this logical port. This type of connection is generally used to connect a Cascade private network to a public network, and as a feeder for Cascade OPTimum trunks.

ATM IISP DCE and ATM IISP DTE — Enables you to configure Interim Inter-switch Signaling Protocol (IISP) DCE or DTE ports that connect a Cascade 500 to another vendor's ATM switch, or a Cascade 500 to another Cascade 500. These services statically route SVCs through a mixed-vendor switch network.

ATM Direct Trunk — Enables you to make a direct trunk connection between two Cascade 500 switches.

ATM OPTimum Cell Trunk — Enables a logical port to communicate with a peer Cascade 500 over an ATM Permanent Virtual Path (PVP). Before you can configure this type of logical port, you must first define an ATM UNI DTE logical port on the physical port.

ATM Cell Switching

Wide-area ATM switches must meet the demands of heavy traffic periods and effectively utilize network bandwidth while minimizing traffic congestion. In addition, they must provide guaranteed delivery for all types of traffic, including voice, video, and multimedia applications.

The Cascade 500 supports the following ATM service classes, each with scalable Quality of Service (QoS) levels:

- Constant Bit Rate (CBR)
- Variable Bit Rate-Real Time (VBR-RT)
- Variable Bit Rate Non-Real Time (VBR-NRT)
- Available Bit Rate/Unspecified Bit Rate (ABR/UBR)

For a complete description of the four ATM service classes, refer to "ATM Adaptation Layer (AAL)" on Page 4-10.

In general, the Cascade 500 implements scalable and configurable ATM service classes by implementing four buffer planes, each associated with one of the four ATM service classes. Internally, the Cascade 500 performs dual-stage buffering, with buffering taking place on both the switch module and on each line card.

Cells coming in and out of the buffer planes are processed according to a scheduler algorithm, where each buffer plane is serviced according to its ATM service class priority and the actual traffic entering the Cascade 500. This algorithm ensures that all ATM service classes are attended to at some point.

To guarantee levels of performance for the various types of traffic, each ATM service class has varying QoS requirements. For example, video-based services have stringent delay variation and cell loss ratio objectives, while other applications have different QoS requirements.

The following sections describe how the Cascade 500 implements scalable and configurable QoS levels for the ATM service classes.

Cascade 500 QoS Parameters

On the Cascade 500, when configuring QoS parameters for ATM service classes, you can select the following on a per-port basis:

Dynamic — Enables the bandwidth allocation to change dynamically according to bandwidth demands. Dynamic bandwidth allocation pools the remaining bandwidth for a connection. This includes bandwidth that has not already been allocated to a specific queue or assigned to a connection.

Fixed — Specifies the percentage of bandwidth you want to reserve for that ATM service class. For example, if you set CBR or VBR traffic to Fixed, you are specifying the maximum bandwidth to reserve for this type of traffic. If the network requests a circuit that exceeds the fixed value, the circuit cannot be created.

If specific ATM service classes are set to Dynamic, any remaining bandwidth percentage is allocated to those ATM service classes as needed. For example, if CBR is fixed at 30%, UBR is fixed at 25%, and the two VBR classes are set to Dynamic, the remaining 45% of bandwidth will be dynamically allocated between the two VBR service classes.

You can also select one of the following routing metrics for each ATM service class. Routing optimizes network resources by avoiding congested paths and finding less congested paths, thereby helping to guarantee fast data delivery.

Cell Delay Variation — This routing metric measures the average variation in delay between one cell and the next, expressed in fractions of a second. When emulating a circuit, cell delay variation measurements allow the network to determine if cells are arriving too fast or too slow.

End-to-End Delay — This parameter measures the time it takes a cell to get from one end of a connection to the other.

Admin Cost — This parameter measures the administrative cost associated with the logical port. The administrative cost is specified by the network administrator, allowing for manual routing.

Cascade 500 ATM Traffic Descriptors

When you create either a Permanent Virtual Circuit (PVC) or a point-to-multipoint circuit on the Cascade 500, you can select one of several traffic descriptors. Traffic descriptors specify which traffic parameters are used for traffic control. Traffic descriptors also determine the number and type of cells that are admitted into a congested queue, and whether or not high-priority cells are tagged as low-priority cells when traffic exceeds the traffic parameter thresholds. Traffic descriptors are as follows:

Peak Cell Rate (PCR) — Peak Cell Rate is the maximum allowed cell transmission rate (expressed in cells per second). It defines the shortest time period between cells, and provides the highest guarantee that all network performance objectives (based on cell loss ratio) are met.

Sustainable Cell Rate (SCR) — Sustainable Cell Rate is the maximum average cell transmission rate allowed over a given period of time on a given circuit. It allows the network to allocate sufficient resources to guarantee that all network performance objectives are met. This parameter applies only to VBR traffic; it does not apply to CBR or UBR traffic.

Maximum Burst Size (MBS) — Maximum Burst Size is the maximum number of cells that can be received at the Peak Cell Rate.
If the burst is larger than anticipated, the additional cells are either tagged or dropped. This parameter applies only to VBR traffic; it does not apply to CBR or UBR traffic.

CLP=0 — Specifies the high-priority cell stream (cells whose Cell Loss Priority bit is set to 0).

CLP=1 — Specifies the low-priority cell stream (cells whose Cell Loss Priority bit is set to 1).

CLP=0+1 — Specifies the aggregate cell stream (cells whose Cell Loss Priority bit is either 0 or 1).

Tagging — Tagging refers to the method of identifying a high-priority cell (CLP=0) as a low-priority cell (CLP=1), as opposed to simply dropping the cells from the cell stream when the CLP=0 cell stream is non-conforming.

Best Effort — Sets a "Best Effort" bit in the cell header. The network attempts to deliver traffic that exceeds the limits of the traffic contract. However, there are no guarantees that traffic will be delivered.

For information on traffic descriptor combinations, refer to the *Cascade 500 Network Administrator's Guide*.

Delivering ATM QoS Guarantees

To guarantee that a specified amount of data is delivered, traffic contracts must be configured for each of the ATM service classes. While data above the traffic contract can still be delivered if there are network resources available, data that exceeds the traffic contract can be delayed or lost.

To make sure incoming traffic does not exceed its traffic contract, the Cascade 500 uses a traffic policing process or ***Usage Parameter Control (UPC)***.

If a cell exceeds the traffic contract, the Cascade 500 does one of the following:

- Delays the traffic until the congestion goes away and there is available bandwidth to deliver the traffic
- Tags the cell by setting the CLP bit to 1
- Discards the cell

To regulate VBR-RT and VBR-NRT traffic at the entry point of the network, the Cascade 500 implements a version of the open-loop congestion control mechanism, known as the ***Leaky Bucket Algorithm***.

Conceptually, a leaky bucket traffic shaping scheme works as shown in Figure 4-10. When data is to be sent, the sending host places the data flow's cells into the bucket. Cells drain out of the bottom of the bucket and are sent onto the network. The rate is enforced by a regulator, the SCR, at the bottom of the bucket.

The bucket's size, the MBS, limits how much data can build up waiting for entry onto the network. If the flow, the PCR, presents more data than the bucket can store, the excess data eventually begins to spill over the top of the bucket. Consequently, cells are delayed or discarded.

The primary effect of the Leaky Bucket Algorithm is to force a burst source of data into a flow of equally spaced cells.

Peak Cell Rate (PCR) = The rate at which the bucket can fill.
Sustainable Cell Rate (SCR) = The rate at which the bucket can drain.
Maximum Burst Size (MBS) = The size of the bucket.

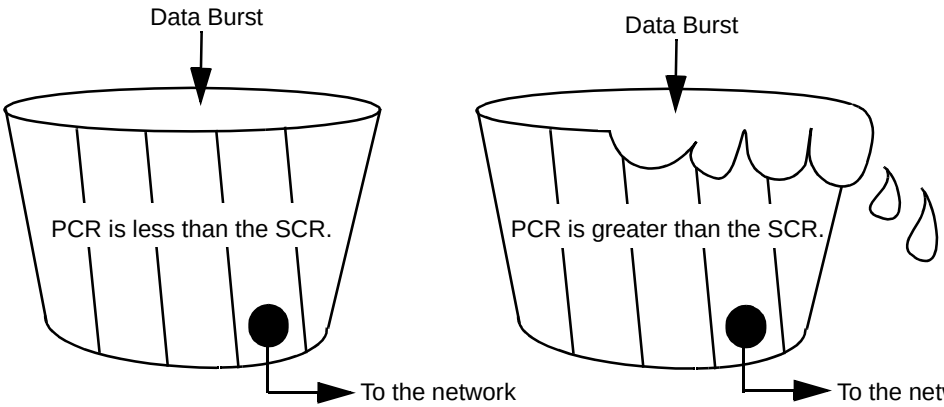

Figure 4-10. Leaky Bucket Algorithm

For CBR traffic, there is only a PCR for each cell stream. If the PCR is exceeded, the buffers and queues continue to fill until there is no more room, and cells are delayed or discarded.

Closely related to traffic management is network congestion control. The Cascade 500 uses a ***Connection Admission Control (CAC)*** algorithm to enable you to control circuit creation on physical ports based on QoS objectives. The following section briefly describes CAC.

Connection Admission Control (CAC)

The Cascade Connection Admission Control (CAC) algorithm performs connection admission control for all ATM service classes. The CAC enables you to control circuit creation on physical ports based on QoS objectives.

For example, the CAC computes an equivalent bandwidth value for a circuit based on the circuit's traffic parameters and requested Cell Delay Variation (CDV) and Cell Loss Ratio (CLR). The circuit's effective bandwidth, somewhere between the circuit's SCR and PCR, indicates how tightly packed the circuits can be on trunks while still guaranteeing CDV and CLR objectives.

The Cascade 500 finds a path for the circuit that has sufficient bandwidth available (i.e., whose available bandwidth on each trunk segment is greater than or equal to the circuit's effective bandwidth).

The Cascade 500 supports two CAC algorithms:

- The default Cascade CAC algorithm
- The customized CAC algorithm

Default Cascade CAC

The default Cascade CAC algorithm is based on a queuing analysis study of the output ports. Using both fluid flow and Gaussian queuing analysis techniques, a conservative closed-form expression for effective bandwidth usage is obtained. This expression takes into consideration the following:

- The traffic characteristics of the circuit to be configured, including Peak Cell Rate (PCR), Sustainable Cell Rate (SCR), and/or Maximum Burst Size (MBS)
- The current load on the physical port
- The buffer sizes
- The desired cell loss ratio objectives

The expression produces a conservative estimate of the required bandwidth to guarantee the desired cell loss ratio and provide good port utilization. If the circuit will adversely impact guaranteed delivery, the circuit is not created.

Customized CAC

The customized CAC algorithm allows network providers to customize the allocation of bandwidth based on a predefined rule. The effective bandwidth of a VBR-NRT or VBR-RT connection with some SCR value is computed as:

Bweff = SCR f1 f2

where: **f1** and **f2** are scale factors specified from CascadeView.

These factors depend on the physical port type and the SCR range of values, respectively. No ATM service class guarantees are provided when the customized CAC algorithm is used.

Over-subscription and CAC

When defining a logical port, you can do the following:

- Specify a particular bandwidth (capacity) for that logical port
- Configure upper limits on the portions of the bandwidth that can be used by different ATM service classes
- Supply over-subscription factors for each ATM service class (except CBR)

The available bandwidth for a given class on a logical port is calculated as follows:

(over-subscription factor **x** upper limit % of the QoS**)** **x** logical port bandwidth

When configuring a circuit, the inflated (over-subscribed) available bandwidth is compared against the effective bandwidth of the circuit to allow for over-subscription. For example, you could configure a logical port to have a capacity of 40 Mbps and specify the following bandwidth usage:

- CBR traffic 20%

- VBR-RT traffic 30%
- VBR-NRT traffic 40%
- UBR traffic 10%

You can also specify that VBR-NRT traffic can be over-subscribed 200%, while other ATM service classes cannot. In this case, the initially available bandwidth is as follows:

- 8 Mbps for CBR traffic (.2 **x** 40 = 8)
- 12 Mbps for VBR-RT traffic (.3 **x** 40 = 12)
- 32 Mbps for VBR-NRT traffic (.4 **x** 40 **x** 200% = 32)
- 4 Mbps for UBR traffic (.1 **x** 40 = 4)

In this example, if your network tries to send 32 Mbps of VBR-NRT traffic over the port, it is sent if there is little or no other traffic from the other ATM service classes. However, if the other ATM service classes are fully utilizing their reserved bandwidth, some of the traffic is dropped or delayed.

When over-subscription is used, the QoS for the ATM service class is not guaranteed because you are effectively overriding the CAC algorithm.

The CAC is integrated with Cascade's routing algorithm so that the route calculation meets the ATM service class guarantees. Cascade uses a variation of the Open Shortest Path First (OSPF) standard for end-to-end routing called the ***Virtual Network Navigator (VNN)***. The following section briefly describes the VNN.

Virtual Network Navigator (VNN)

In general, the VNN optimizes network resources by avoiding congested paths and finding less congested paths, thereby helping to guarantee fast data delivery.

The VNN establishes a network path based on a virtual circuit's ATM service class requirements, including

- Bandwidth
- End-to-end delay
- Cell Delay Variation (CDV)
- Cell Loss Ratio (CLR)

The VNN uses a "Best Route" calculation algorithm to optimize network resources. In addition, the VNN incorporates policy enhancements that allow you to:

- Select specific trunks
- Create virtual private networks
- Dedicate trunks to control traffic

Figure 4-11 shows the configuration of VNN trunk links on a Cascade network. Note that the VNN links are both physical links and logical links traversing alternative services.

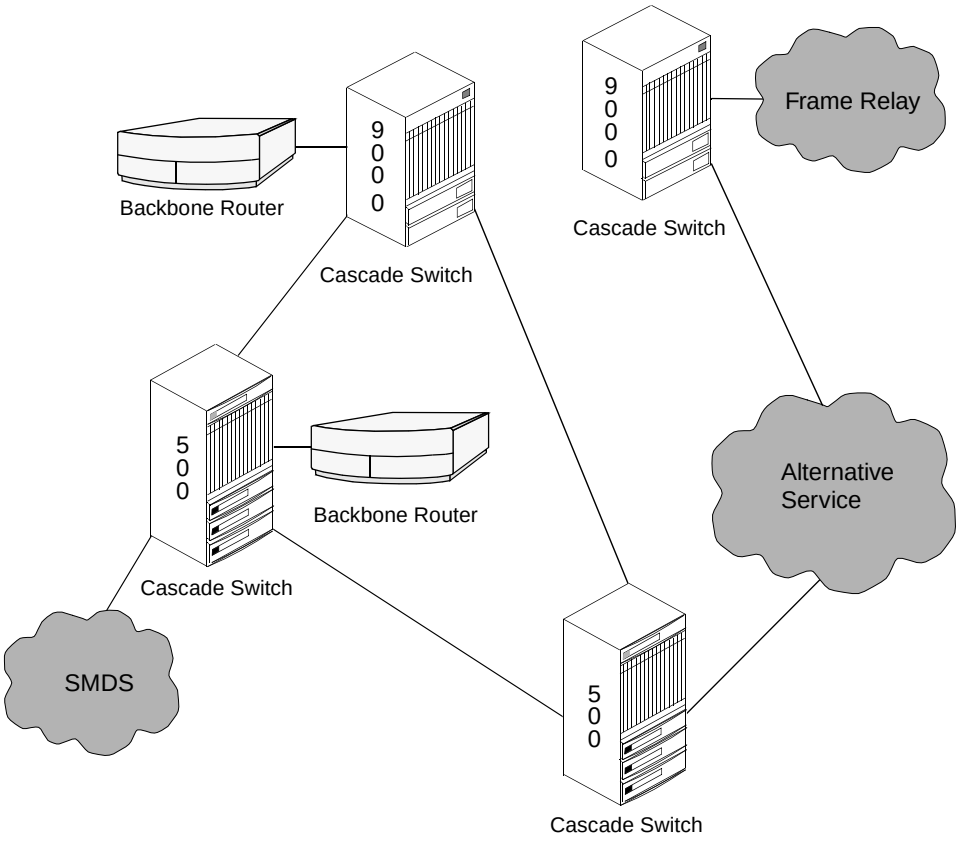

Figure 4-11. Virtual Network Navigator (VNN) Configuration

The VNN has three major components:

- A topology database capturing the physical topology, ATM service classes, and QoS levels on each ATM link
- A topology database distribution algorithm (OSPF based)
- A "Best Route" calculation algorithm for optimizing network resources

These components operate in conjunction with the ***Virtual Circuit Manager (VC Manager)*** services resident on every Cascade switch in a network. VC Manager is the switching intelligence in each node that builds the input port-to-output port VC mappings required for the basic operation of any switching device. VC Manager also performs the VC acceptance decision that compares the bandwidth resources required by an additional connection against the actual available resources on a link, thereby helping to prevent circuits from being created in congested areas of the network.

VNN Metrics

Each VNN link in a Cascade network is characterized by the following metrics:

Available bandwidth for each ATM service class — Indicates the available bandwidth on the trunk for each ATM service class traffic, expressed in Kbps. The value can be from 1 Kbps to 16 Gbps. The values of this metric are asymmetric (that is, they can be different in each link direction).

Delay — Indicates the static delay of the trunk, consisting of both propagation and transmission delay. Delay is measured by the VNN protocol when the link initially comes up. The delay measurement does not include queuing delays; therefore, it does not reflect trunk congestion. The trunk protocol dynamically measures the delay imposed by the trunk to a granularity of one millisecond.

Trunk capability — Indicates static trunk capabilities (e.g., packet or cell).

Version identifier — Indicates the version of routing software currently executing in the switch. This allows the gradual introduction of new routing functionality.

When a new route needs to be created, these four metrics are analyzed, and the VNN picks the route that best guarantees data delivery for the type of data being delivered.

VNN Routing Options

The following policy-based routing options are available in a VNN trunk-link configuration:

Administrative path control — You can assign costs to individual trunks. Usually, the administrative cost reflects the trunk's distance or tariff. Those trunks with higher costs are avoided. When there are two paths that satisfy a connection's ATM service class requirements, the path with the lowest total administrative cost is chosen. The larger the value given to a trunk, the less likely it will be used for virtual circuits. There is also an option to set the administrative cost of each trunk to its measured delay. Selecting this option provides least-delay routing.

Management trunks — You can configure trunks to be used exclusively for management traffic. These trunks do not carry any user data or routing information that is not required for network management purposes. This option is used by carrier customers to satisfy certain regulatory requirements prohibiting subscriber data transport across Local Access and Transport Area (LATA) boundaries. It is mainly used to enable the management of multiple networks from a single Network Management Station (NMS), even when user traffic between the networks is prohibited by government regulations.

Virtual private network trunk — You can dedicate trunks to specific end-users, thereby creating multiple virtual private networks from a single network of Cascade switches. Binding between a set of private trunks and customer ports indicates whether the trunks have been dedicated to a particular subscriber's virtual private network. If so, only those ports associated with that subscriber are allowed to establish VCs over the trunk.

Virtual path capability — You can specify whether or not the trunk supports the routing of ATM virtual paths.

Switched Services

ATM cell relay service supports both on-demand Switched Virtual Connections (SVCs) and non-switched Permanent Virtual Connections (PVCs). PVCs are typically used by institutions between fixed corporate locations.

In contrast to PVCs, SVCs provide flexibility for establishing dynamic connections between locations that are not fixed. SVCs optimize network performance because circuits only exist while needed. In addition, SVCs guarantee data delivery because they can be created dynamically, provided the network bandwidth is configured to accept the SVC creation parameters.

SVCs are established and dismantled using a ***signaling protocol***. Signaling is the act of transferring service-related information between the user and the network and among network elements. Signaling takes place between the user and the network over the User-Network Interface (UNI) and between network elements over the Network Node Interface (NNI).

The signaling protocols in UNI 3.0/3.1 support four ATM service classes:

- Constant Bit Rate (CBR)
- Variable Bit Rate-Real Time (VBR-RT)

- Variable Bit Rate Non-Real Time (VBR-NRT)
- Available Bit Rate/Unspecified Bit Rate (ABR/UBR)

Later versions of the UNI Specification are expected to allow discreet signaling of individual QoS parameters.

Signaling Protocol Stack

The signaling protocol stack is implemented in the user's equipment and in the network peer. It consists of the following layers, in descending order, as defined by the Open Systems Interconnection (OSI) Reference Model:

Layers 7-4	Call Control
Layer 3	Q.2931
Layer 2a	Signaling ATM Adaptation Layer (Service Specific Convergence Function)
Layer 2b	Signaling ATM Adaptation Layer (Service Specific Connection Oriented Protocol)
Layer 2c	ATM Adaptation Layer (Type 5)
Layer 2d	ATM Layer
Layer 1	DS3 framing and SONET framing

For detailed information about the Signaling Protocol Stack, refer to the *ATM User-Network Interface (UNI) Specification*, Version 3.2.

Call Control Functions

For on-demand SVCs, the cell relay service allows you to provide information regarding the type of service you require. For example, call control for point-to-point calls involves the following procedures:

- Establishing a call at the originating interface
- Establishing a call at the destination interface
- Clearing a call
- Handling error conditions

The signaling is done during the call establishment phase, except for multipoint calls where parties can be added or dropped as needed. The call establishment phase is followed by an information transfer phase and then a call release phase.

Transporting Signaling Messages

Signaling messages must be reliably transported between network peers. One important mechanism for ensuring reliability is the ***Signaling ATM Adaptation Layer (SAAL)*** function.

The SAAL resides between the ATM Layer and Q.2931 in the user's equipment. The purpose of the SAAL is to provide reliable transport of Q.2931 messages between peer Q.2931 entities, such as an ATM switch and host, over the ATM Layer.

The SAAL is subdivided into the following two parts:

Common Part (CP) — Represents the functionality common to all users requiring a connection-oriented, variable bit-rate information transfer. It provides uninsured information transfer and a mechanism for detecting corruption of information carried in the SAAL frames.

Service-Specific Part (SSP) — Represents the protocol and procedures associated with the signaling needs of the UNI. It provides data recovery.

For detailed information about the SAAL, refer to the *ATM User-Network Interface (UNI) Specification,* Version 3.1.

SETUP Message Capabilities

The SETUP message is one of the more important signaling messages. It enables an ATM device to select the desired bandwidth and QoS levels using appropriate information elements to establish a connection.

The SETUP message is sent by the calling user to the network and by the network to the called user to initiate call establishment. Some of the key information elements of the SETUP message include the following:

- AAL parameters
- ATM user cell rate
- Broadband bearer capability
- Called party number
- Calling party number
- Connection identifier
- QoS class

Defining SVC Addresses on the Cascade 500

The Cascade 500 includes an SVC and routing processor on every line card. Consequently, the Cascade 500 achieves very high SVC setup rate performance.

There are two broad types of addresses used to configure SVC parameters on the Cascade 500:

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

These are described in the following sections.

ATM End System Address (AESA) Formats

There are four AESA formats supported:

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

All AESA address formats consist of 20 octets (40 hex digits). Each of these address formats contains the following components:

Authority and Format Identifier (AFI) — The AFI part of the AESA address identifies the authority that allocates the DCC, ICD, or E.164 part of the AESA address, as well as the syntax of the rest of the address. The following are valid AFIs:

- 0x39 for DCC
- 0x47 for ICD
- 0x45 for E.164
- A user-specific code for custom prefixes/addresses

Initial Domain Identifier (IDI) — A hex code that identifies the sub-authority that has allocated the address. The format depends on the address type:

DCC, ICD — Consists of 2 octets (4 hex digits)

E.164 — Consists of 8 octets in binary coded decimal (BCD) format (1-15 hex digits, plus a trailing Fh; also, if less than 15 digits are entered, leading zeros are required to fill the 8 octets)

High-Order Domain-Specific Part (HO-DSP) — The authority specified in the AFI/IDI octets determines the format of this field. It describes the hierarchy of the addressing authority, and conveys topological significance. It should be constructed to facilitate routing through interconnected ATM subnetworks. The general format for each address type is as follows:

DCC, ICD — Consists of 10 octets (20 hex digits)

E.164 — Consists of 4 octets (8 hex digits)

Custom — Consists of 12 octets (24 hex digits)

End System Identifier (ESI) — A 6-octet (12 hex digit) field that identifies the end system. This is typically an IEEE Media Access Control (MAC).

Selector (SEL) — A 1-octet (2 hex digit) field that is not used for ATM routing, but may be used by the end system.

Initial Domain Part (IDP) — Consists of the AFI and IDI fields.
It uniquely identifies the administrative authority responsible for allocating and assigning the Domain Specific Part.

Domain-Specific Part — Consists of the HO-DSP, EDI, and SEL fields.

Native E.164 Address Format

Native E.164 addresses are the standard Integrated Services Digital Network (ISDN) numbers, including telephone numbers. Native E.164 addresses consist of 1-15 ASCII digits. For example, standard 10-digit United States telephone numbers, such as 508-555-1234, are native E.164 addresses.

Unlike AESA address formats, native E.164 addresses are not broken down into an AFI, HO-DSP, ESI, and SEL portion. When a native E.164 address is translated to E.164 AESA format, the native E.164 address is stored in octets 2-9 of the 20-octet AESA address, while the HO-DSP, ESI, and SEL portions are filled with zeros.

Conversely, when an E.164 AESA address is translated to native E.164 address format, the AFI, HO-DSP, ESI, and SEL portions, as well as any leading zeros in the 8-octet AESA E.164 address, are stripped off to produce the native E.164 address.

About Address Registration

Address information in a switch is used both for determining the proper route for calls and for calling party screening. When used for route determination, the switch advertises an appropriate subset of its configured node prefixes, port prefixes, and port addresses to all other switches in the network. When used for calling party screening, the switch uses the configured node prefixes, port prefixes, and/or port addresses to determine whether or not a call should be accepted by the network.

To perform these two functions at a UNI, both the user and the network need to know the ATM addresses that are valid at the UNI. Address registration provides a mechanism for address information to be dynamically exchanged between the user and the network, enabling them to determine the valid ATM addresses that are in effect at a UNI.

Address registration applies only to UNI ports on which ***Interim Local Management Interface (ILMI)*** is enabled. Any ILMI-eligible node or port prefix will be transferred from all ILMI-enabled private UNI-DCE ports and all ILMI-enabled public end-system UNI-DCE ports to their peer DTE devices. Node prefixes are not exchanged from "public switch" UNI-DCE ports. Only port prefixes are exchanged from these ports.

ILMI-eligible prefixes include the following:

- All native E.164 node prefixes
- All 13-octet (104-bit) AESA node prefixes
- All native E.164 port prefixes
- All 13-octet (104-bit) AESA port prefixes

The network side of the UNI provides the network prefix that consists of the IDP and HO-DSP portions. The user side of the UNI provides the remaining portion of the address, including the IEEE MAC address (the ESI portion), and the SEL portion of an ATM address; this forms the user part of the address. This is shown in Figure 4-12.

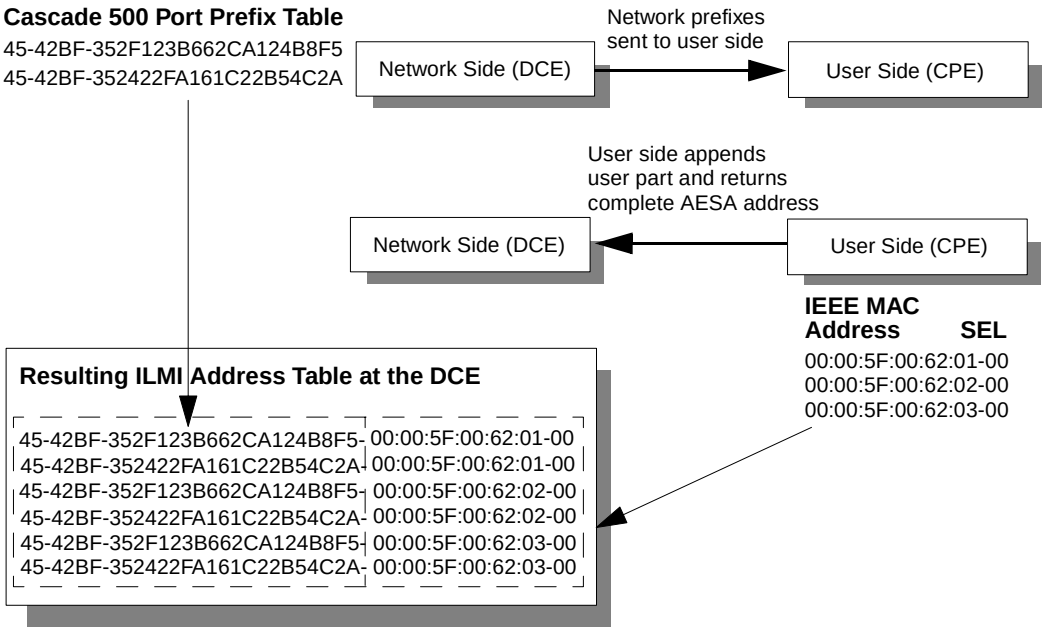

Figure 4-12. Address Registration

About Route Determination

The node prefixes, port prefixes, and port addresses that are configured on network nodes are used to determine the route for a given SVC. The route is determined by a "best match" hierarchy, starting from the left-most digit of the called party address.

For example, if you have three nodes configured with the following combination of native E.164 address information, each call is routed for a certain called party address.

	Node 1	Node 2	Node 3
Node Prefixes	508 6	None	508 603
Port Prefixes	508551 508552 508553	5085 508553 6035	508554 508555
Port Addresses	5085511111 5085511112 5085511113 5085555555 5085555556	None	None

Called Party Address	Node	Routing Determination
5085511234	1	Port prefix 508551 on Node 1 is a longer match than port prefix 5085 on Node 2 and node prefix 508 on Node 3.
5085555555	1	This calling party address exactly matches a port address defined on Node 1. This is a longer match than port prefix 5085 on Node 2 and port prefix 508555 on Node 3.

Called Party Address	Node	Routing Determination
5085555557	3	Port prefix 508555 on Node 3 is a longer match than port prefix 50855 on Node 2 and node prefix 508 on Node 1.
5085561111	2	Port prefix 5085 on Node 2 is a longer match than node prefix 508 on Node 1 and node prefix 508 on Node 3.
6175551111	1	Node prefix 6 on Node 1 is the only match.
6035551111	2	Port prefix 6035 on Node 2 is a longer match than node prefix 6 on Node 1 and node prefix 603 on Node 3.
6038558888	3	Node prefix 603 on Node 3 is a longer match than node prefix 6 on Node 1. There is no matching prefix or address on Node 2.
5085531111	1 or 2	Since the longest match occurs on both Nodes 1 and 2, the Admin Cost value assigned to port prefix 5085 on each node determines where the call is routed. The call is routed to the node with the lowest Admin Cost value for port prefix 5085.

Called Party Address	Node	Routing Determination
5145551234	None	The call is not routed to any of these nodes because there are no matching node prefixes, port prefixes, or port addresses. If, however, you set up a default route on a port being used for network-to-network connections, all non-matching calls are routed to that port.

About Configuring Node Prefixes

Node prefixes apply to all ports on the Cascade 500 and are used for routing aggregation and address registration. You can configure multiple node prefixes. However, you do not need to configure any node prefixes if you have port prefixes or port addresses defined on the node.

At the very least, a node prefix consists of the two Authority and Format Identifier (AFI) digits of the AESA address, or at least one digit of the 1-15 digit native E.164 address. You can define the node prefix to be part or all of the AESA or E.164 address. For example, for E.164 addresses that begin with 508555, you can configure the node prefix as 5 (at a minimum), 50, 508, 5085, and so on. The level of granularity you need to define depends on your network.

Node prefixes do not have to be unique to a particular node. For example, you can define node prefix 508 on multiple nodes. However, if you do so, you need to define port prefixes or port addresses to provide more granularity for routing determination.

About Address Translation

Address translation allows an SVC to be created across several networks that are using different addressing schemes. Calling party and called party addresses are stored as information elements in the SETUP message that is sent to initiate call setup. In some situations, calling party and called party sub-addresses are also stored as information elements in the SETUP message.

The following factors determine how address translation occurs:

- Whether or not local and/or remote gateway addresses are defined on the egress port
- The type of translation (tunnel or replace) selected as the egress address translation mode
- The numbering plan of the signaled calling and called addresses

Calling party and called party processing are independent. In the case of a native E.164 called party or calling party address, the related sub-address field is always set to null, since the sub-address field cannot carry native E.164 addresses.

Using ingress address translation, the calling party sub-address (if it is not null) overwrites the calling party address at the ingress port, and the called party sub-address (if it is not null) overwrites the called party address.

For more information on address translation, refer to the *Cascade 500 Network Administrator's Guide*.

The following examples illustrate the state of the calling party/called party address and sub-address elements of the SETUP message at various points along the connection.

- Egress tunneling is enabled on Network 1's egress port.
- Ingress tunneling is enabled on Network 2's ingress port.
- Local Gateway address X is configured to a prefix on Network 1's egress port, and the prefix corresponds to B.
- Remote Gateway address Y is configured to a prefix on Network 1's egress port, and the prefix corresponds to B.

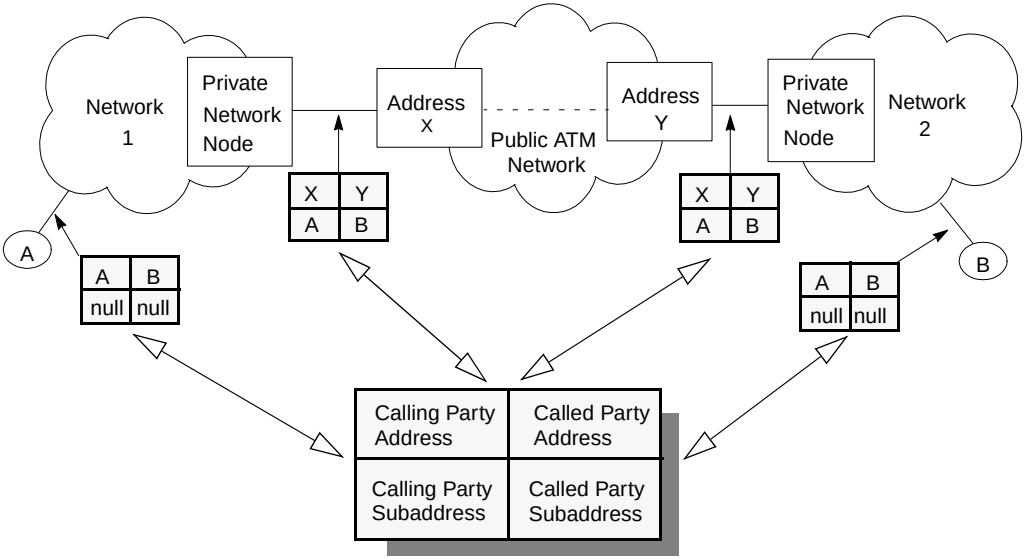

- Egress tunneling is enabled on Network 1's egress port.
- Ingress tunneling is enabled on Network 2's ingress port.
- No local Gateway address is defined on Network 1's egress port.
- Remote Gateway address Y is configured to a prefix on Network 1's egress port, and the prefix corresponds to B.

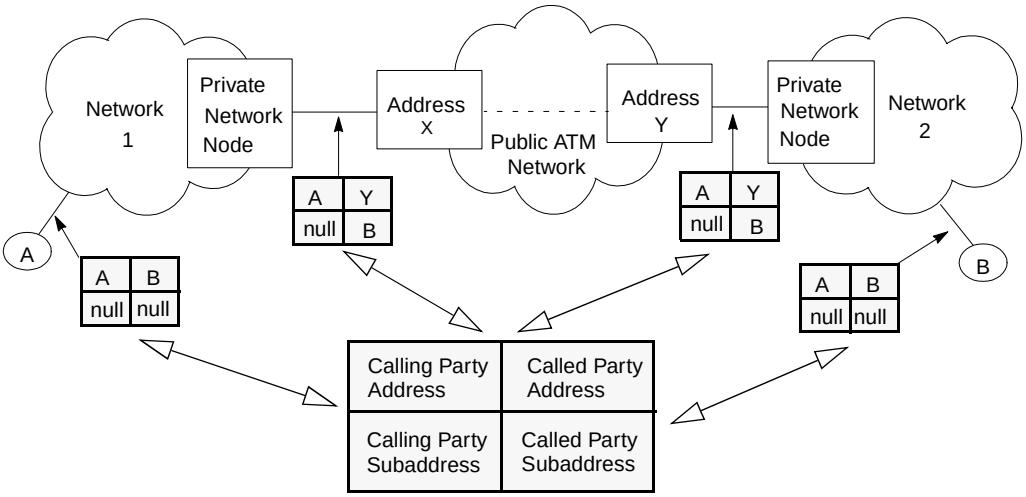

5

ISDN Services

Background

Integrated Services Digital Network, ISDN, is a digital network architecture that has been in use by various telephone companies for almost 30 years.

ISDN was initially deployed during the 1960's in an effort to increase the speed and quality of existing communications. Today, ISDN is being brought into homes and offices to replace conventional analog telephone service and redefine local and business communications throughout the world.

Technology Fundamentals

ISDN communications provide many types of benefits and services. The following sections describe ISDN and the various benefits and services it provides.

What Is ISDN?

ISDN is a set of network protocols that enables a wide range of digital communication services. Some of the benefits derived from ISDN include

- Improved Remote Access
- Faster Call Setup
- Improved Quality of Service
- Device Integration
- Dynamic ISDN Connections
- Bandwidth-on-demand

Improved Remote Access

With an ISDN, dial-in digital communications is achieved at a rate that surpasses conventional analog service.

Analog communication is often times limited to 56 Kbps. An ISDN exceeds 56 Kbps with an ISDN Basic Rate Interface (BRI). Using one of two B channels on the BRI, dial-in communications is increased from 56 to 64 Kbps. Using both B channels, an effective bit rate of 128 Kbps is achieved. Figure 5-1 illustrates how devices can send their signals across an ISDN BRI.

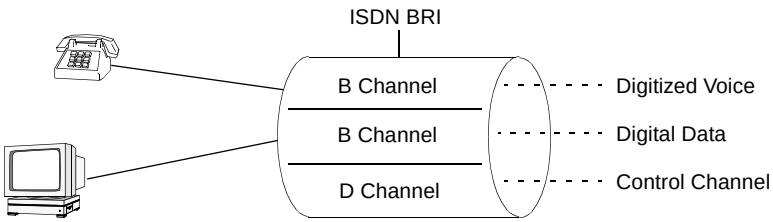

Figure 5-1. ISDN and Device Integration

Faster Call Setup

An ISDN connection is set up much faster than a typical analog connection since there is no modem involved and no negotiation process.

Improved Quality of Service

Analog communications are known to be error prone due to the nature of the analog signal and the method used to transfer the signal across the copper wire. With an ISDN, the signal produced is more reliable resulting in a higher Quality of Service.

Device Integration

One of the most important ISDN benefits is device integration. Using the same copper wire that links most homes and businesses to the existing telephone network, ISDN provides the integration of digital signals from various devices over a single network interface. Communication that once required numerous wire pairs can now take place over one wire pair passing digital signals through logical B channels.

These devices are connected to an ISDN line using an ISDN Network Termination One (NT-1) device. The ISDN NT-1 device connects up to eight ISDN devices to a single ISDN line.

For more information on ISDN NT-1 devices, refer to "ISDN NT-1 Devices" on Page 5-7.

Dynamic ISDN Connections

Unlike leased lines, ISDN connections are established on an as-needed basis. A connection is activated only when it is needed, so charges are incurred only for the time the connection is up, plus the call setup charges.

One very common use of ISDN is as a backup to a leased line. An ISDN connection is kept in reserve in case the leased line goes down. When the ISDN connection is no longer needed, it is terminated. Initiating and terminating an ISDN connection is done automatically. Figure 5-2 illustrates ISDN and leased-line backup.

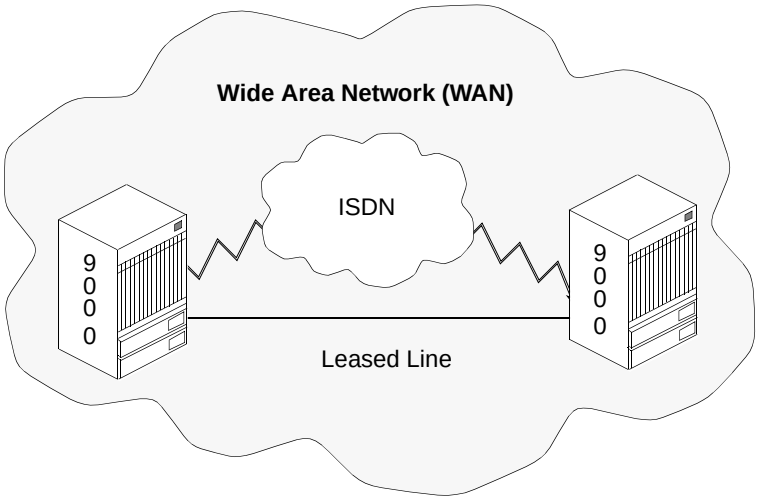

Figure 5-2. ISDN and Leased-Line Backup

Bandwidth-on-demand

Since an ISDN can automatically initiate and terminate calls, it is a cost effective way to add capacity to a leased line that at times reaches full capacity. When the leased line is saturated, an ISDN connection is established. When the connection is no longer needed, it is terminated.

How ISDN Works

Using ISDN, a conventional analog signal is replaced with a 64 Kbps digital signal. This digital signal is transported over an ISDN connection using B and D channels. B channels carry digital voice and/or data. D channels initiate, receive, and control calls.

ISDN is available on two types of interfaces, a Basic Rate Interface (BRI) or a Primary Rate Interface (PRI).

Basic Rate Interface (BRI)

Basic rate is an interface which is most often used in a home or small office. A single BRI services two telephones that can each call a different destination simultaneously. Both calls are sent over the same copper wire.

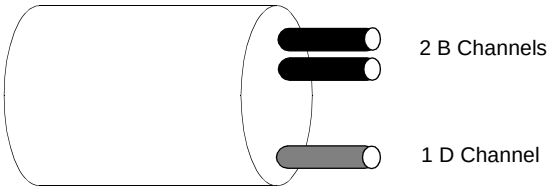

Figure 5-3. Basic Rate Interface (BRI)

A BRI consists of two B channels operating at 64 Kbps and one D channel operating at 16 Kbps. The total user data rate is 144 Kbps.

A BRI provides better voice quality, higher data rates, lower error rates, faster call setup times, and much more flexibility when compared to the conventional telephone network.

Primary Rate Interface (PRI)

Primary rate is an interface that is most often used in large commercial sites. Primary rate connections are often used as high-speed trunks for transferring large files and other continuous data streams. They can also be subdivided with a multiplexer to provide channels for multiple devices.

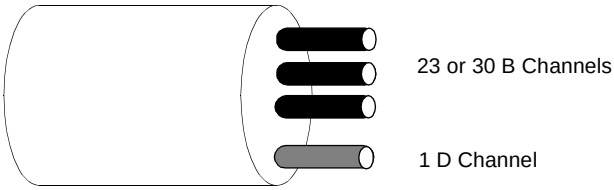

Figure 5-4. Primary Rate Interface (PRI)

A PRI can have one of two configurations. The North American PRI configuration consists of 23 B channels and one
D channel that operate at 64 Kbps. The total data rate is
1.536 Mbps.

The European PRI configuration consists of 30 B channels and one D channel that operate at 64 Kbps. The total bit rate is
1.984 Mbps.

The D channel provides signaling to initiate and terminate calls. This signaling adheres to the OSI model and operates in the Physical, Data-Link, and Network protocol layers. These protocols are used to define message types that are sent between the customer and the local loop to set up and maintain services.

ISDN NT-1 Devices

The ISDN Network Termination One (NT-1) is a device used in North America to connect the customer's data or telephone equipment to the local telephone loop. Originally, the ISDN NT-1 device was designed as a piece of equipment that would be owned and maintained by the service provider. With the breakup of AT&T, the Regional Bell Operating Companies (RBOCs) are no longer allowed to own equipment that resides on customer premises. In the United States, the ISDN NT-1 device is purchased, installed, and maintained by the customer.

The ISDN NT-1 device connects terminal equipment and terminal adapter equipment to the local telephone loop. Terminal equipment includes ISDN-compatible telephones and computers. Terminal adapters are devices used to connect non-ISDN-compatible equipment. Figure 5-5 shows how an ISDN NT-1 device connects this type of equipment.

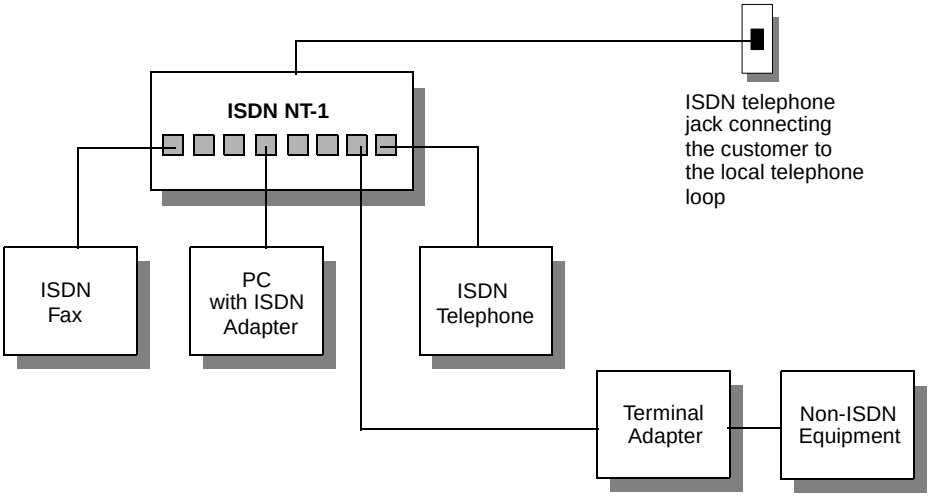

Figure 5-5. ISDN NT-1 Device

In the United States, an ISDN NT-1 device connects to the customer's equipment on one side, and the telephone company's wire on the other side of the device. On the customer's side, up to eight devices can be connected and supported by the ISDN NT-1 device.

ISDN Telephone Services

There are a number of additional telephone services that are provided with an ISDN, including the following:

Multiple Telephone Numbers — This service provides a way for a single interface to have multiple telephone numbers. This service is often used in homes or small offices.

Caller Identification — This service provides a way for the telephone number of the incoming call to be displayed.

Caller Identification Restriction — This service provides a way for the caller to prevent the telephone number from being displayed to the called party.

Cascade's Implementation of ISDN Services

Cascade provides ISDN support through its HyperPATH remote access software and its B-STDX 8000/9000 hardware platforms.

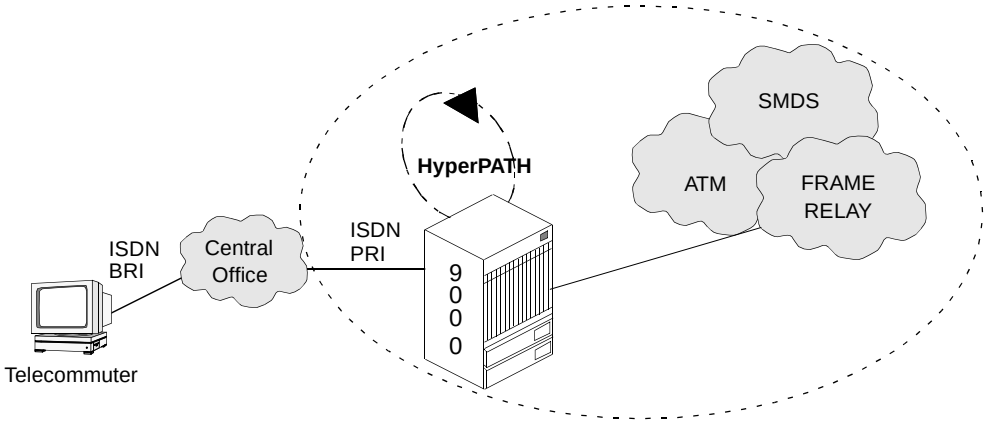

Figure 5-6. ISDN and Remote Access

The HyperPATH remote access software supports ISDN dial-in communications over a BRI connection. The incoming ISDN call connects to a PRI logical port on a B-STDX 8000/9000 switch and is then transferred to its destination using HyperPATH's internetworking services. The ISDN internetworking support includes PPP-to-Frame Relay and ATM services.

The B-STDX 8000/9000 hardware platforms provide ISDN fault-tolerant services in addition to the HyperPATH-related support. SMDS and native Frame Relay logical ports are also supported. The remainder of this chapter discusses Cascade's ISDN fault-tolerant services. For information on HyperPATH, refer to the *HyperPATH User's Guide*.

ISDN Fault-Tolerant Services

Fault tolerance is a method for providing redundancy in hardware systems to protect against downtime should one of the systems or components go down. For the Cascade switch, fault tolerance is supported via redundant IO modules and power supplies. Fault tolerance is also provided through "hot swapping," a practice which enables the replacement of parts while the Cascade switch is running.

ISDN fault-tolerant services are software functions that offer an additional measure of reliability for supporting the uninterrupted flow of critical traffic between source and destination. Some of these functions are integral to the networking service, such as dynamic rerouting, while others are optional features.

To increase reliability beyond the measure provided by the hardware and networking service, Cascade introduces a new set of software-based services designed to meet the reliability requirements of mission-critical applications. This added reliability is critical to financial institutions such as banks and brokerages where any interruption in service could be devastating to the business. These software-based services include the following:

Fault-tolerant PVC — Allows a set of backup ports on the Cascade B-STDX 8000/9000 to restore connections from a failed data center to the backup data center. When enabled, fault-tolerant PVC automatically reroutes all affected Frame Relay circuits to the backup set of ports.

Access Failure Recovery — Provides failure recovery service through the ISDN Primary Rate (PRI) IO Module for the B-STDX 8000/9000 switch. If access to the switch fails, an ISDN dial-backup call is automatically placed to the switch and a new connection is established.

ISDN Trunk Backup — Allows the Cascade switch to set up one or more backup trunks to replace a primary trunk, and reroute all PVCs from the primary trunk to the backup or other available trunks. The PVCs remain rerouted on the backup trunks until trunk restoration occurs. During trunk restoration, the PVCs are rerouted from the backup trunk(s) and the backup trunk(s) are cleared.

Cascade ISDN Trunk Backup Service

ISDN Trunk Backup can be initiated automatically upon failure of a Cascade switch trunk line or at a scheduled time as specified by the network operator. The ISDN Trunk Backup service can also be initiated manually at any time by the network operator. Restoration of the original (primary) trunk can be initiated in the same manner.

The hardware component of the ISDN Trunk Backup function is simply a UIO card with a V.35 or X.21 interface. All signal pins required by an ISDN terminal adapter (ISDN-TA) are connected to the UIO card. The Cascade switch initiates the ISDN trunk backup over ISDN circuits via the assertion of control signals on each port. The ISDN-TA initiates the ISDN call procedure based on the status of these control signals.

ISDN trunk backup is distinct from the ISDN PRI support available in the Cascade switch. The ISDN PRI support enables the Cascade switch to receive ISDN call requests and establish ISDN connections on the ISDN PRI card; whereas, ISDN trunk backup is performed via connections on the UIO card and is simply a vehicle for two ISDN-TAs to establish a connection across an ISDN network.

ISDN trunk backup performed via ISDN is transparent to the Cascade switch. The ISDN connection is established when the ISDN-TA recognizes the modem control-signal transitions from the UIO card on the connected port due to a Cascade switch link-up.

For information on ISDN call setup, refer to "ISDN Call Setup Procedure" on Page 5-17.

Figure 5-7 shows an example of how ISDN trunks can be used to back up a primary trunk.

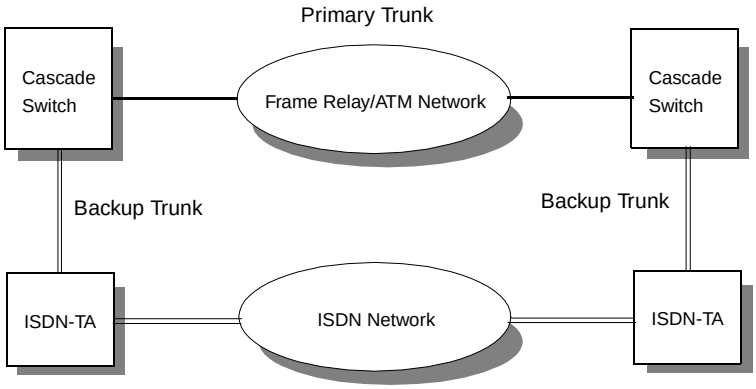

Figure 5-7. Trunk Backup Using ISDN Trunks

As shown in Figure 5-7, the primary trunk connecting the two Cascade switches can be of any type supported by the Cascade switch. The ISDN Trunk Backup function can be used to back up primary trunks regardless of whether they are a Frame Relay network trunk or a simple leased line. Because the ISDN backup trunk is relatively low speed, ATM Virtual Channels/Circuits (VCCs) are not supported. ATM VCCs can be rerouted on alternate ATM trunks.

The rerouting of PVCs by the Cascade switch is governed by OSPF and a load-balancing algorithm. When PVCs from a trunk are rerouted or restored, their paths are determined by these algorithms.

In the case of a primary-trunk failure for instance, all the PVCs may not be rerouted onto the backup trunk. Some PVCs may be routed via alternate paths. In the case of the restoration of a primary trunk, all the PVCs that were originally routed on the primary trunk may not be restored to the primary trunk, but may be routed via alternate paths.

Guidelines for ISDN Trunk Backup

A backup trunk can have a total bandwidth that is less than the primary trunk. For example, when ISDN trunks are used as backup circuits, the backup trunk has a bandwidth range from 64 Kbps to 1.5 Mbps. However, this is not necessarily a problem. To provide additional backup bandwidth, multiple backup trunks can be used to back up a single primary trunk.

When a backup trunk is activated, all of the PVCs from the associated primary trunk are rerouted according to the following guidelines:

- OSPF rules
- Limitation of backup trunk line speed

Rerouting PVCs Using OSPF

The rerouting of PVCs occurs according to Cascade's OSPF rules, which are applied as follows:

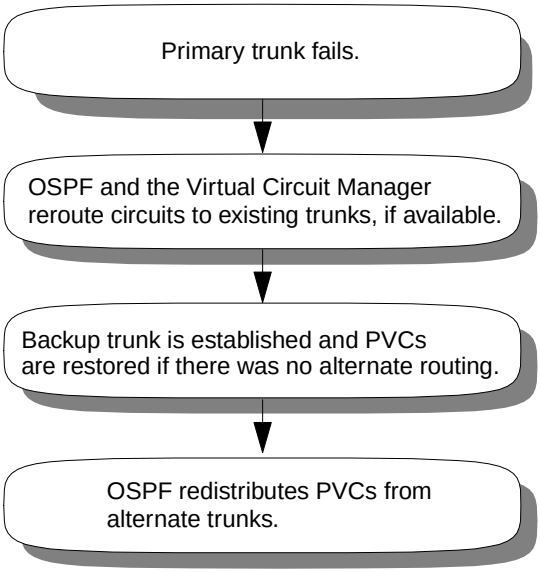

The rerouting takes place without regard to network congestion.
In the case when the backup-trunk bandwidth is less than the primary-trunk bandwidth, congestion can occur because the sum of the rerouted PVC's Committed Information Rates (CIRs) can be greater than the bandwidth of the backup trunk.

When congestion occurs, frames are transmitted based on the PVC's priority (1, 2, or 3). If a state of severe congestion is reached, frames are discarded in the reverse order of priority.

Backup Trunk Line Speed

In addition to OSPF rules, the highest transmission rate of one PVC is limited by the circuit speed of the backup trunk. For example, if a primary trunk contains a PVC with a CIR of 512 Kbps, and an ISDN trunk with a line speed of 64 Kbps is being used as a backup trunk, the highest transmission speed is limited to 64 Kbps for PVCs on that trunk.

To provide sufficient backup bandwidth, the network operator can define up to eight 64 Kbps backup trunks to back up a single primary trunk.

When a backup is initiated on a primary trunk, the trunk backup function in the Cascade switches on both sides of the trunk generate a link-up for the backup port. However, for certain backup trunk types, call setup direction (i.e., which of the two Cascade switches connected by a trunk initiates the call setup) is important. Therefore, the network operator must specify to the NMS during trunk configuration, which node or switch is to initiate the call setup.

When defining trunks, any given backup trunk can only back up a single primary trunk, and at least one backup trunk is needed for every primary trunk requiring backup.

ISDN Backup Trunk Operation

The Cascade switch attempts to establish the ISDN backup trunk in response to the conditions described previously. This section describes how the Cascade switch responds to those conditions. For details on configuring trunk backups, refer to the *CascadeView/UX Network Configuration Guide*.

Backup Upon Trunk Failure

The first type of automatic trunk backup takes place after a trunk failure. A ***trunk failure*** occurs when the Cascade switch determines that a trunk is no longer available. The processing to make this determination begins with the receipt of a ***link-down*** from the link level for either a direct trunk line or for a network trunk such as Frame Relay or ATM.

To avoid rerouting in the case where a trunk only goes down momentarily, the Cascade switch monitors the trunk for a period of time after the link-down indication is received. If a link-up or circuit-up indication is not received during this period, the backup process is initiated. This period of time is called the ***trunk failure threshold*** and is set from the Network Management Station (NMS). If a link-up or circuit-up indication is received during this period, backup processing is not initiated.

Once the trunk failure threshold has been met, the Cascade switch determines if one or more backup trunks have been configured for the failed primary trunk. If a backup trunk(s) exists and the network operator has previously enabled "backup on trunk failure" for the failed trunk, the Cascade switch initiates a backup trunk setup.

If "backup on trunk failure" is enabled, but no backup trunk exists, the Cascade switch sends an SNMP trap to the NMS indicating that trunk backup has failed. Otherwise, backup processing is not initiated, and no traps are sent to the NMS.

From a configuration role only, the NMS informs the two Cascade switches that define the trunk connection that "backup on trunk failure" has been either enabled or disabled for a trunk. If "backup on trunk failure" is enabled, each Cascade switch receives both the trunk failure threshold and the trunk restoration threshold from the NMS.

Scheduled Backup

The second type of automatic trunk backup takes place at a time specified in advance by the network operator. The operator can schedule the initiation and termination of primary-trunk backups via the NMS. The NMS informs the two Cascade switches that define the trunk connection, via an SNMP SET command, that a scheduled backup must be initiated for a trunk.

Upon receipt of this SET command, the Cascade switch determines if a backup trunk(s) has been configured for the primary trunk. If a backup trunk(s) exists, the Cascade switch initiates a backup trunk setup. Otherwise, backup processing is not initiated and the Cascade switch sends an SNMP trap to the NMS indicating that scheduled trunk backup has failed.

Operator Initiated Backup

Manual trunk backup occurs as a result of network operator commands. The operator can initiate or terminate primary-trunk backups at any time via the NMS. Operator commands always have the highest priority over all types of manual backup initiation and termination.

Through an SNMP SET command, the NMS informs the two Cascade switches that define the trunk connection that a manual backup must be initiated for a trunk.

Upon receipt of the backup notification SET command, the Cascade switch determines if a backup trunk(s) has been configured for the primary trunk. If a backup trunk(s) exists, the Cascade switch initiates a backup trunk setup. Otherwise, backup processing is not initiated, and the Cascade switch sends an SNMP trap command to the NMS indicating that manual trunk backup has failed.

ISDN Call Setup Procedure

For ISDN trunk backup, establishment of the backup trunk consists of the ***ISDN call setup procedure***. Associated with each ISDN backup trunk is a logical port configured as DTE. To start the backup trunk, the Cascade switch invokes a link-up routine for its backup port. The link-up routine generates the following sequence of modem control-signal transitions from the UIO card to the ISDN-TA:

- For the Cascade switch port configured as DTE - raise DTR and RTS
- For the ISDN-TA port - detect ER and RS up

When setting up an ISDN circuit, the call setup direction is important. When defining a backup trunk on the NMS, the network operator must specify which node initiates the ISDN call setup. The NMS informs the Cascade switch of this as part of the backup-trunk configuration data that it sends to the switch.

To indicate the completion of the call setup for each backup trunk, the source and destination ISDN-TAs generate the following sequence of modem-control signal transitions:

- For the ISDN-TA port - raise CS, DR, and CD
- For the Cascade switch port configured as DTE - detect CTS, DSR, and DCD up

If the completion indication for setting up the backup trunk is returned by the ISDN-TA, the Cascade switch then invokes a link-down routine to bring down the primary trunk. This prevents the trunk from being used as a route for any PVCs.

Once the backup trunk or trunks have been successfully set up, each PVC carried by the primary trunk is rerouted by the OSPF routing algorithms. OSPF determines on which trunk to route each PVC. The chosen trunk may be an existing primary trunk or one of the backup trunks defined for the primary trunk.

In the event that no completion indication is received within five seconds of invoking the link-up routine, the Cascade switch enters into a ***Call Setup Retry Cycle***. The Cascade switch retries the call setup after a period of time called the ***Retry Interval***. Each cycle consists of a defined number of retries, called the ***Retry Number***, separated in time by the retry interval. The Cascade switch then waits for a period of time called the ***Cycle Interval*** before initiating another retry cycle. Figure 5-8 illustrates this process.

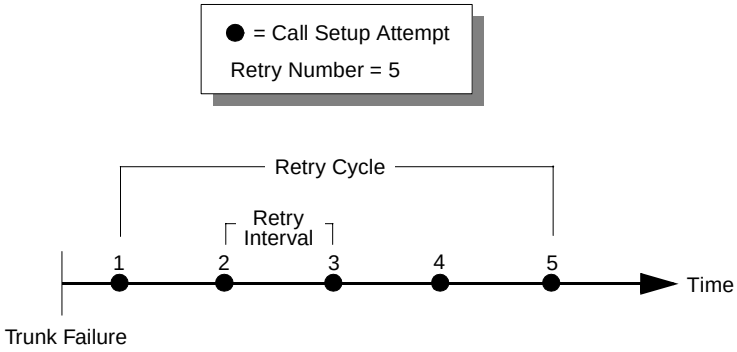

Figure 5-8. Call Setup Retry Cycle

The network operator specifies the retry interval, retry number, and cycle interval when defining the primary trunk. The NMS informs the Cascade switch of these parameters and all other backup trunk configuration data that it sends to the switch via SNMP SET commands.

If the call setup does not succeed after three retry cycles, the Cascade switch sends a backup failure SNMP trap to the NMS, indicating that the backup trunk was not established.

The retry cycle continues indefinitely until one of the conditions for backup trunk termination is met. At that point, the Cascade switch ceases its efforts to establish the backup trunk.

Primary-Trunk Restoration

Primary-trunk restoration consists of rerouting the PVCs that were carried on the backup trunk and releasing the circuit used for the backup trunk. All rerouting is done by OSPF and all PVCs may not necessarily be rerouted on their original primary trunk.

Primary-trunk restoration can only occur when one of the following conditions for backup-trunk termination is met:

- The primary trunk is available and the backup was caused by a trunk failure.
- The primary trunk is available and the backup is scheduled for termination.
- A network operator entered a termination command.

The Cascade switch determines that the primary trunk is available in the following manner. A link-up message from a failed trunk indicates to the Cascade switch that the trunk is back on-line. To avoid losses in service if restoration is attempted on a trunk that is cycling between the up and down states, a trunk is only considered available by the Cascade switch if it remains in the link-up state for a minimum amount of time. This period of time is called the ***Trunk Restoration Threshold*** and is set from the NMS.

Manual trunk restoration has highest priority and takes place without regard for the reason that the backup was initiated, nor the current state of the primary trunk. If OSPF cannot find an available trunk, PVCs routed on the backup trunk are suspended.

ISDN Call Release Procedure

For ISDN trunk backup, release of the backup trunk occurs by the ISDN call release procedure. For each ISDN backup trunk that must be released, the Cascade switch invokes a link-down routine for the backup port associated with the backup trunk. The link-down routine generates the following sequence of modem control signal transitions from the UIO card to the ISDN-TA:

- For the Cascade switch port configured as DTE - lower DTR and RTS
- For the ISDN-TA port - detect ER and RS down

The ISDN-TA responds by generating this sequence of signal transitions to the UIO card of the Cascade switch:

- For the ISDN-TA port - lower CS, DR, and CD
- For the Cascade switch port configured as DTE - detect CTS, DSR, and DCD down

The ISDN-TA then releases the ISDN circuit. Once the ISDN circuit is released, the Cascade switch sends a trunk status SNMP trap to the NMS indicating that the backup trunk is no longer available and a second SNMP trap indicating that the backup trunk has been released, to allow its status to be changed to "defined."

A

Glossary

absolute congestion

In Frame Relay, a congested condition in the network that occurs when the queue length reaches a third threshold and there is no more room on the queue for any packets, regardless of the type of packet.

access rate

The data rate of the user access channel. The speed of the access channel determines how quickly (maximum rate) the end user may inject data into the network. See also *bandwidth.*

address

The logical location or identifier of a network node, terminal, pc, peripheral device, or location in memory where information is stored. See also *network address.*

address mask

A bit combination used to describe which portion of an SMDS address refers to the network (or subnet) and which part refers to the host. Sometimes referred to as mask. See also *subnet mask.*

alternate path

An optional automatic feature of OSPF (Open Shortest Path First) that reroutes the PVC should a trunk fail within a manually defined path.

amber frames

Cascade's own class of packet frames used to identify packets as they travel through the Frame Relay network. The network forwards amber frames with the Discard Eligible bit set; therefore the packet is eligible for discard if it passes through a congested node.

Annex D

A synchronous polling scheme used for the link management of a Frame Relay channel, where the user polls the network to obtain status information on the PVCs configured on the channel. Annex D exchanges this information using DLCI 0.

area ID

One of two portions of the SMDS address. Can start at any digit and the length can be up to eight digits (4 bytes long for BCD encoding).

Asynchronous Transfer Mode (ATM)

A method used for transmitting voice, video, and data over high-speed LAN and WAN networks. See also *cell relay.*

ATM Service Interworking Feeder

A service that enables Frame Relay network traffic to be fed into an ATM network, enabling a Frame Relay end-user to communicate with an ATM end-user.

ATM/DXI trunk

See *OPTimum PVC Trunk.*

ATM/DXI trunk interface

An ATM circuit used as a trunk between two Frame Relay networks that are built with Cascade switches.

backbone

The part of a network that carries the bulk of the network traffic, e.g. over Ethernet cabling, fiber-optic cabling.

Backward Explicit Congestion Notification (BECN)

A bit in the Frame Relay header that indicates the frame has passed through a congested node from traffic traveling in the opposite direction.

bandwidth

The transmission capacity of a computer or a communications channel.

bandwidth-on-demand

A WAN feature that enables users to dial up additional bandwidth as their applications demand.

Bc

See *Committed Burst Size.*

Be

See *Excess Burst.*

BECN

See *Backward Explicit Congestion Notification.*

best-effort packets

Packets delivered to the best of the network's ability, after the requirements for delivering the guaranteed packets are met. See also *guaranteed packets*.

broadband network

A type of network that allows for the transmitting of large amounts of information, including voice, data, and video over long distances using the same cable.

broadcast

A message that is sent to all users currently logged into the network.

burst mode

A method of data transmission in which information is collected and then sent in a single high-speed transmission, rather than one character at a time.

byte

A series of consecutive binary digits that are operated upon as a unit (for example, an eight-bit byte).

CascadeView/UX

The UNIX-based graphical user interface used to configure and monitor a Cascade network.

CBR

See *Constant Bit Rate.*

cell

Any fixed-length data packet. For example, ATM uses fixed-length, 53-byte cells. See also *cell relay.*

Cell Loss Priority

A field in the ATM cell header that indicates the eligibility of the cell for discard by the network under congested conditions.

cell relay

A form of packet transmission that uses a fixed-length, 53-byte cell over a packet-switched network; also known as Asynchronous Transfer Mode (ATM).

cell switching

An operational feature of cellular networks that enables callers to move from one location to another without losing the call connection. The cellular system is designed to switch calls to a new cell with no noticeable drop in the conversation. Cell switching is sometimes called “handing off.” While not noticeable in voice communications, the approximate 300 milliseconds this switching takes can be a problem in data transmission.

channel

Any connecting path that carries information from a sending device to a receiving device. May refer to a physical medium (e.g., coaxial cable) or a specific frequency within a larger channel.

channel bank

Equipment that converts multiple voice signals to time division multiplexed (TDM) signals for transmission over a T1 or E1 line.

Channel Service Unit

A device that functions as a certified safe electrical circuit, acting as a buffer between the customer’s equipment and a public carrier’s WAN.

CIR

See *Committed Information Rate.*

circuit

A communications channel or path between two devices.

circuit switching

A temporary communications connection that is established as needed between a sending node and a receiving node.

CLP

See *Cell Loss Priority.*

Committed Burst Size

The maximum amount of data, in bits, that the network agrees to transfer under normal conditions, during a time interval Tc. Committed Burst Size is defined for each PVC.

Committed Information Rate

The rate at which the network agrees to transfer information under normal conditions. The rate is averaged over a minimum increment of time, Tc.

Committed Rate Measurement Interval

The time interval during which the user is allowed to send only Bc committed amount of data and Be excess amount of data. In general, the duration of Tc is proportional to the burstiness of the traffic. Tc is computed from CIR and Bc as Tc=Bc/CIR.

communications protocol

A standard way of communicating between computers, or computers and terminals; also a hardware interface standard, such as RS-232C for communication between DTE and DCE devices.

congestion

The point at which devices in the network are operating at their highest utilization. Congestion is handled by employing a congestion avoidance mechanism. See also *mild, severe,* and *absolute congestion*.

Constant Bit Rate

A Quality of Service class defined by the ATM Forum for ATM networks. CBR is used for connections that depend on precise clocking to ensure undistorted delivery of bits.

Control Processor

A module that makes up the hardware architecture of a B-STDX 8000/9000 switch. A CP provides network and system management and routing functions in support of the real-time switching functions provided by the multiple, IO Processor modules (IOPs).

CP

See *Control Processor.*

CRC

See *Cyclic Redundancy Check.*

CRC error

A condition that occurs when the CRC in a frame does not agree with the CRC frame received from the network.

CSMA/CD

See *Carrier Sense Multiple Access Collision Detect.*

CSU

See *Channel Service Unit.*

Cyclic Redundancy Check

A calculation method used to check the accuracy of digital transmission over a communications link.

D4-format

In T1 transmission, 24 channels per T1 line, where channels are assigned sequentially.

data bits

In asynchronous transmission, the bits that actually contain the data being sent. Also called “payload” in some transmission methods.

Data Bus (DB) connector

A cable connector used to connect devices to parallel or serial ports. The number following DB indicates the number of pins in the connector (e.g., DB-25 connectors have 25 pins).

Data Carrier Detect

A hardware signal, defined by the RS-232-C standard that indicates that the device is on-line and ready for transmission.

Data Communications Equipment (DCE)

Any device that connects a computer or terminal to a communications channel or public network.

Data Exchange Interface (DXI)

A specification, described in RFC 1483, that defines how a network device can be used to convert data for interworking between different network services (i.e., Frame Relay-to-ATM).

Data Link Connection Identifier

A 10-bit address that identifies PVCs. See also *Locally/ Globally Significant DLCIs*.

data-link layer

The second of seven layers of the ISO/OSI model for computer-to-computer communications. This layer ensures data flow and timing from one node to another by synchronizing blocks of data and controlling the flow of data.

data packet

One unit of information transmitted as a discrete entity from one network node to another. In packet-switched networks, a data packet is a transmission unit of a fixed maximum length that contains a header, a set of data, and error control information.

Data Service Unit

A device that connects DTE to digital communications lines. A DSU formats the data for transmission on the public carrier WAN, and ensures that the carrier's requirements for data formats are met.

Data Set Ready

A hardware signal defined by the RS-232-C standard that indicates that the device is ready to operate.

Data Terminal Equipment (DTE)

Any device, such as a terminal or computer, that is connected to a communications device, channel, or public network.

DCE

See *Data Communications Equipment.*

D-Channel

The data channel in ISDN used for control signals and customer data. In Primary Rate Interface (PRI) ISDN, the D-Channel operates at 64 Kbps.

DE

See *Discard Eligible.*

dedicated line

A communications circuit used for one specific purpose, and not used by or shared between other users.

dedicated server

A computer on the network that functions only as a server performing specific network tasks.

define path

A function that allows a manual path to be defined for the PVC, thereby bypassing the OSPF (Open Shortest Path First) algorithm to make PVC routing decisions.

delay

In communications, a pause in activity, representing the time that a message must wait for transmission-related resources to become available.

Digital Signal (Digital Service)

A classification of digital circuits. The DS defines the level of common carrier digital transmission service. DS-0 = 64 Kbps (Fractional T1), DS-1 = 1.544 Mbps (T1), DS-2 = 6.312 Mbps (T2), DS-3 = 44.736 Mbps (T3), and DS-4 =274-176 Mbps (T4).

Discard Eligible (DE)

A bit in the Frame Relay header used to indicate that a frame is eligible for discard by a congested node.

DLCI

See *Data Link Connection Identifier.*

domain

A network community of users sharing the same database information.

DS

See *Digital Signal.*

DS0

A 64-Kbps channel used in T1 transmission. There are 24 DS0 channels in a T1 line.

DS1

A standard digital transmission facility, operating at 1.544 Mbps.

DSR

See *Data Set Ready.*

DSU

See *Data Service Unit.*

DSX-1

A T1 specification that indicates the physical and electrical characteristics of the standard T1 cross-connection.

Data Terminal Equipment (DTE)

The devices in a category that includes terminals and computers. It also refers to the interface to users' equipment, as opposed to the DCE interface to the network.

Data Terminal Ready (DTR)

An RS-232 modem interface control signal that indicates the DTE is ready for data transmission.

DXI

See *Data Exchange Interface.*

dynamic routing

A routing technique that allows a message's route to change "en route" through the network.

E1

The European counterpart to the North American T1 transmission speed. Adopted by the Conference of European Posts and Telecommunications Administrations, the E1 standard carries data at the rate of 2.048 Mbps.

encapsulation

The wrapping of data in a particular protocol header. For example, Ethernet data is wrapped in a specific Ethernet header before being transmitted. Also, when bridging dissimilar networks, the entire frame from one network is simply placed in the header used by the data link layer protocol of the other network.

error rate

In communications, the ratio between the number of bits received incorrectly and the total number of bits in the transmission.

ESF

See *Extended Superframe Format.*

Ethernet address

A 48-bit number physical address. Each Ethernet address is unique to a specific network card or PC on a LAN, which forms the basis of a network-addressing scheme. Compare with *Internet address.*

Excess Burst

The maximum allowed amount of uncommitted data (in bits) in excess of Bc that the network attempts to deliver during time interval Tc. In general, this data (Be) is delivered with a lower probability than Bc. See *Be.*

Extended Superframe Format

In Frame Relay, a frame structure that extends the DS1 superframe structure from 12 to 24 frames, for a total of 4632 bits. This format redefines the 8-Kbps channel consisting of framing bits previously used only for terminal and robbed-bit signaling synchronization.

FCS

See *Frame Check Sequence.*

FDDI

See *Fiber Distributed Data Interface.*

FDM

See *Frequency-Division Multiplexing.*

FECN

See *Forward Explicit Congestion Notification.*

Fiber Distributed Data Interface

An ANSI standard for fiber-optic links with a data transmission rate up to 100 Mbps.

File Transfer Protocol

A method of transferring information from one computer to another, either over a modem and telephone line, or over a network. FTP is a TCP/IP application utility.

Forward Explicit Congestion Notification (FECN)

A bit in the Frame Relay header that indicates the frame has passed through a node that is experiencing congestion in the same direction in which the frame is traveling.

fractional T1

One channel of a T1 circuit. T1 circuits consist of 24, 64-Kbps channels. Customers can lease as many of these channels as needed; they are not required to lease all 24 channels in one circuit.

FRAD

See *Frame Relay Assembler/Disassembler.*

frame

In Frame Relay, a block that can be transmitted as a single unit.

Frame Check Sequence

In a frame, a field that contains the standard 16-bit cyclic redundancy check used to detect errors in HDLC and LAPD frames. See also *Cyclic Redundancy Check.*

Frame Relay

A type of data transmission based on a packet-switching protocol, with transmission rates up to 2 Mbps. Frame Relay provides for bandwidth-on-demand.

Frame Relay Assembler/Disassembler

A function that enables a logical port to perform Frame Relay encapsulation/de-encapsulation for HDLC/SDLC-based protocols. The FRAD function encapsulates HDLC/SDLC traffic entering a Cascade Frame Relay network and de-encapsulates it upon exiting the network. This function is restricted to one point-to-point PVC.

Frame Relay RFC1294 Multi-protocol Encapsulation

A specification allowing for a single circuit to be established between two devices.

Frequency-Division Multiplexing

A method of sharing a transmission channel by dividing the total bandwidth of the circuit into several smaller channels. This is accomplished by allocating specific frequency ranges to each channel. All signals are carried simultaneously. Compare with *Time Division Multiplexing.*

Generic Flow Control

The field in the ATM cell that controls the flow of traffic across the User-Network Interface (UNI), and into the network.

GFC

See *Generic Flow Control.*

globally significant DLCI

A feature of the Local Management Interface (LMI) enhancement to Frame Relay, which enables DLCIs to use the same connection identification scheme across the network (global values), to specify individual end devices.

graceful discard

When enabled, this function turns red frames into best-effort frames. When disabled, this function discards frames.

green frames

Cascade's own class of packet frames used to identify packets as they travel through the network. Green frames are never discarded by the network except under extreme circumstances, such as node or link failure.

group addressing

The ability to send a single datagram/packet to multiple locations.

guaranteed packets

Data delivered according to some time constraint and with high reliability.

HDLC

See *High-level Data Link Control.*

header

The initial part of a data block, packet, or frame, that provides basic information about the handling of the rest of the block, packet, or frame.

Header Error Control

In ATM, a feature that provides protection against misdelivery of cells due to addressing errors.

HEC

See *Header Error Control.*

High-level Data Link Control

An international protocol defined by ISO. In HDLC, messages are transmitted in variable-length units known as frames.

High-Speed Serial Interface

A high-speed interface (up to 52 Mbps full duplex) between a DTE and a DCE. The DCE provides the timing for the interface. HSSI can connect over a 50 ft. (15m) shielded twisted-pair cable.

hop (count)

The number of links that must be "jumped" to get from a source node to a destination node.

hot swappable

A feature that allows the user to add, replace, or remove interface processors in a Cascade switch without interrupting switch operations.

HP OpenView

The UNIX-based network management application used with CascadeView/UX on an NMS to manage a Cascade Network

HSSI

See *High-Speed Serial Interface.*

ICMP

See *Internet Control Message Protocol.*

Input/Output Adapter

A module that connects the various IOP and IOP Plus modules in a switch. IOA configurations vary according to the specific IOP module they support.

Input/Output Processor

A module in a switch that manages the lowest level of a node's trunk or user interfaces. An IOP performs physical data link and multiplexing operations on external trunks and user links.

Integrated Services Digital Network (ISDN)

A CCITT standard for a worldwide digital communications network, intended to replace all current systems with a completely digital transmission system.

internal clocking

A hardware function of the Cascade switch that provides the transmit and receive clocks to the user equipment.

International Standards Organization

An international standards group based in Geneva, Switzerland that establishes global standards for communications and information exchange.

International Telecommunication Union Telecommunication Standard Sector

An advisory committee established under the United Nations to recommend worldwide standards for voice and data. One of the four main organizations of the International Telecommunications Union.

Internet Control Message Protocol

The IP portion of TCP that provides the functions used for network layer management and control.

Internet Protocol

The TCP/IP session-layer protocol that regulates packet forwarding. See also *ICMP.*

Internet Protocol address

A 32-bit address assigned to hosts using TCP/IP. The address is written as four octets separated with periods (dotted decimal format) that are made up of a network section, an optional subnet section, and a host section.

ISDN call setup

A procedure that establishes an ISDN backup trunk.

ISO

See *International Standards Organization.*

ITU-T

See *International Telecommunication Union Telecommunication Standardization Sector.*

jitter

A type of distortion found on analog communications lines, resulting in data transmission errors.

Kbps

Kilobits per second.

keep-alives

A series of polling messages used in the Local Management Interface (LMI) of a Frame Relay port to verify link integrity between devices.

LAN

See *Local Area Network.*

Link Access Protocol

The link-level protocol used for communications between DCE and DTE devices.

Link Management Interface

A set of enhancements to the basic Frame Relay specification. LMI dynamically notifies the user when a PVC is added or deleted. The LMI also monitors each connection to the network through a periodic heartbeat “keep alive” polling process.

Link Management Interface Rev 1

A synchronous polling scheme used for the link management of a Frame Relay channel where the user polls the network to obtain status information of the PVCs configured on the channel. LMI exchanges this information using DLCI 1023.

link-state routing protocol

A sophisticated method of determining the shortest paths through the network. See also *OSPF.*

LMI

See *Link Management Interface.*

LMI Rev 1

See *Link Management Interface Rev 1.*

load balancing

A technique that distributes network traffic along parallel paths to maximize the available bandwidth while providing redundancy at the same time.

Local Area Network

Any physical network technology that connects a number of devices and operates at high speeds (10 Mbps through several gigabits per second) over short distances. Compare with *Wide Area Networks.*

locally significant DLCI

In Frame Relay, an identifier or address that specifies a local router, PVC, SVC, or endpoint device. It is re-usable at non-overlapping end points and allows for scalability. Compare with *globally significant DLCI*.

logical port

A configured circuit that defines protocol interaction.

loss of frame

A T1 error condition when an out-of-frame condition exists for a normal period of 2 1/2 seconds.

management DLCI

A value that specifies a PVC or SVC from a LAN connected via a router to a Cascade switch over a Frame Relay network.

Management Information Base (MIB)

The set of variables forming a database contained in a CMIP or SNMP-managed node on a network. Network management stations can fetch/store information from/to this database.

Mbps

Megabits per second.

mild congestion

In Frame Relay, the state of a link when the threshold is exceeded.

multicast

A type of broadcast transmission that sends copies of the message to multiple stations, but not to all possible stations.

multiplexer (mux)

A device that merges several lower-speed transmission channels into one high-speed channel at one end of the link. Another mux reverses this process at the opposite end.

multiplexing

A technique that transmits several signals over a single communications channel.

name server

A server connected to a network that converts network names into network addresses.

network address

A network layer address refers to a logical, rather than a physical network device; also called protocol address.

Network Interface Card (NIC)

A card, usually installed in a PC, that enables you to communicate with other users on a LAN; also called adapter.

Network-to-Network Interface (NNI)

The standard that defines the interface between ATM switches and between Frame Relay switches. In an SMDS network, an NNI is referred to as Inter-Switching System Interface (ISSI).

node

Any device such as a pc, terminal, workstation, etc., connected to a network and capable of communicating with other devices.

node number

A unique number that identifies a device on the network.

Open Shortest Path First (OSPF)

A routing protocol that takes into account network loading and bandwidth when routing information over the network. Incorporates least-cost routing, equal-cost routing, and load balancing.

Open Systems Interconnection (OSI)

An international standard program created by ISO and ITU-T to develop standards for data networking, such as the OSI model, to facilitate multivendor operating environments.

OPTimum PVC trunk

A logical port configuration that optimizes interoperability in performance and throughput in networks where both ends are connected by Cascade switches.

OPTimum trunking

A software function that allows public data networks based on Frame Relay, SMDS, or ATM to be used as trunk connections between Cascade switches.

out of frame

A T1 error condition where two or three framing bits of any five consecutive frames are in error.

packet

Any block of data sent over a network. Each packet contains sender, receiver, and error-control information, in addition to the actual message; sometimes called payload or data bits.

Packet Assembler/Disassembler

A device connected to a packet-switched network that converts a serial data stream from a character-oriented device (e.g., a bridge or router) into packets suitable for transmission. It also disassembles packets into character format for transmission to a character device.

packet processor

The Cascade switch module that performs the frame format validation, routing, queuing and protocol conversion for the switch. This module is not hot swappable.

packet-switched network

A network that consists of a series of interconnected circuits that route individual packets of data over one of several routes and services.

packet switching

Type of networking in which nodes share bandwidth with each other by intermittently sending logical information units (packets). In contrast, a circuit-switching network dedicates one circuit at a time to data transmission.

PAD

See *Packet Assembler/Disassembler.*

path

The complete location of a directory or file in the file system. See *define path* and *alternate path.*

payload

The portion of an ATM cell that contains the actual user data.

PCR

See *Peak Cell Rate.*

Peak Cell Rate

In ATM transmission, the maximum transmission rate that cells are transmitted. Equivalent to Be for Frame Relay, PCR is measured in cells per second and converted internally to bits per second. PCR defines the shortest time period between two cells.

PDN

See *Public Data Network.*

Permanent Virtual Circuit (PVC)

A logical connection across a packet-switched network that is always in place and always available along a predetermined network path. See also *Virtual Circuit.*

Point-to-Point Protocol

A protocol that provides router-to-router and host-to-network connections.

PPP

See *Point-to-Point Protocol.*

Primary Rate Interface (PRI)

An ISDN interface to primary rate access. Primary rate access consists of a single 64-Kbps D channel plus 23 (T1) or 30 (E1) B channels for voice or data.

protocol

A set of rules governing communication between two entities or systems to provide interoperability between services and vendors. Protocols operate at different layers of the network, e.g., data link, network, and session.

Public Data Network

Any government-owned or controlled commercial packet-switched network, offering WAN services to data processing users.

Quality Of Service (QoS)

A statistical report that specifies certain characteristics of network services, sessions, connections, or links. For example, the CascadeView Statistics report describes the lost packets and round-trip delay measurements.

rate enforcement

A process used to measure the actual traffic flow across a given connection and compare it to the total admissible traffic flow for that connection. Traffic outside of the acceptable level can be tagged and discarded en route if congestion develops. ATM, Frame Relay, and other types of networks use rate enforcement.

red alarm

A T1 alarm condition indicating a loss of signal or loss of frame at the device's local termination point.

red frames

In Frame Relay, a type of frame to be discarded. Color designators green, amber, and red identify packets as they travel through the network.

redundancy

The duplication of hardware or software within a network to ensure fault-tolerant or back-up operation.

remote connection

A workstation-to-network connection made using a modem and telephone line or other WAN services equipment. Remote connections enable you to send and receive data over greater distances than you can with conventional cabling methods.

Request For Comment (RFC)

A series of notes and documents available on-line that describe surveys, measurements, ideas, techniques, and observations, as well as proposed and accepted Internet protocol standards, such as Telnet and FTP.

Request To Send

A hardware signal, defined by the RS-232-C standard, that a device sends to request permission to transmit.

RFC1294

A specification documenting multiprotocol access over Frame Relay.

RIP

See *Routing Information Protocol.*

route recovery

In Frame Relay, an OSPF routing function in the Cascade switch. When a tandem node or trunk is down, new shortest-path routes for those affected PVCs are recalculated immediately at the ingress nodes due to fast convergence of the link-state updates. The PVCs are then rerouted to the new route. Recovery time is typically under four seconds. The network reports PVC rerouting as an event/alarm.

router

An intelligent LAN-connection device that routes packets to the correct LAN segments destination address(es). The extended LAN segments may or may not use the same protocols. Routers link LAN segments at the ISO/OSI network layer.

Routing Information Protocol

A routing protocol that maintains a list of accessible networks and calculates the lowest hop count from a particular location to a specific network.

routing protocol

A protocol that implements routing using a specific routing algorithm. Routing protocols include IGRP, OSPF, and RIP.

SCR

See *Sustainable Cell Rate.*

SEAL

See *Simple and Efficient Adaption Layer.*

Serial Line over Internet Protocol

A protocol that enables point-to-point serial communication over IP using serial lines or telephone connections and modems.

serial management port

A management port on the Packet Processor card in a Cascade switch.

severe congestion

In Frame Relay, a state or condition that occurs when the queue size is greater than a second predetermined threshold (32 buffers full). In this state, the continued forwarding of amber and red packets jeopardize the successful delivery of green packets.

shortest path routing

A routing algorithm that calculates the path distances to all network destinations. The shortest path is then determined by a cost assigned to each link. See also *OSPF.*

SIG

See *SMDS Interest Group.*

Simple and Efficient Adaption Layer

In ATM, an extension of the Type 3 AAL. It simplifies the SAR portion of the Adaption layer to pack all 48 bytes of the cell information field with data. This AAL makes ATM look like high-speed Frame Relay. It also assumes that only one message is crossing the UNI at a time. That is, multiple end users at one location cannot interleave messages on the same virtual circuit, but must queue them for sequential transmission.

Simple Network Management Protocol

A standard network management protocol used to manage and monitor nodes and devices on a network.

SIP

See *SMDS Interface Protocol.*

SLIP

See *Serial Line over Internet Protocol.*

SMDS

See *Switched Multimegabit Data Services.*

SMDS In-Band Network Management

The NMS manages the SMDS network traffic using SMDS In-Band Network Management. To be managed from this NMS, all SMDS Access Servers/Switches must be in the same IP subnet.

SMDS Interest Group

A consortium of vendors and consultants committed to advancing worldwide SMDS as an open, interoperable solution for high-performance data connectivity.

SMDS Interface Protocol

The protocol defined at the network and end-user interface connection.

SNMP

See *Simple Network Management Protocol.*

static route

A route or path that is manually entered into the routing table. Static routes take precedence over routes or paths specified by dynamic routing protocols.

subnet address

An extension of the Internet addressing scheme that allows a site to use a single Internet address for multiple physical networks.

subnet mask

A 32-bit address mask used in IP to specify a particular subnet. See also *address mask.*

Sustainable Cell Rate

The average cell transmission rate in ATM transmission. Equivalent to CIR for Frame Relay, SCR is measured in cells per second and converted internally to bits per second Usually, SCR is a fraction of the peak cell rate. Cells are sent at this rate if there is no credit.

Switched Multimegabit Data Services

A high-speed WAN service based on the 802.6 standard for use over T1 or T3 circuits.

Switched Virtual Circuit (SVC)

A logical connection across a packet-switched network providing as-needed connections to any other node in the network. See also *Virtual Circuit.*

synchronization

The timing of separate elements or events to occur simultaneously. In communications, hardware and software must be synchronized so that file transfers can occur.

synchronous transmission

A data transmission method that uses a clock signal to regulate data flow.

T1

A long-distance, point-to-point circuit that provides 24 channels at 64 Kbps each (for a total of 1.544 Mbps). See also *E1*.

T3

A long-distance, point-to-point circuit that provides up to 28 T1 channels. T3 can carry 672 channels of 64 Kbps (for a total of 44.736 Mbps).

Tc

See *Committed Rate Measurement Interval.*

TCP

See *Transmission Control Protocol.*

TDM

See *Time Division Multiplexing.*

telnet

The Internet standard protocol for remote terminal-connection services.

throughput

The actual speed of the network.

Time Division Multiplexing

Technique that allocates bandwidth for multiple channels onto one channel based on preassigned time slots.

time interval "T"

The time interval over which the number of bits used to average the number of bits transmitted, is averaged. To calculate **T,** use the following formula: Bc/CIR=T.

Transmission Control Protocol

The Internet standard, transport-level protocol that provides the reliable, full duplex, stream service on which many application protocols depend.

Transmit Data

A hardware signal (defined by the RS-232-C standard) used by the DTE to transmit data to the DCE. See also *Receive Data.*

trunk

The communications circuit between two switches.

trunk backup

A configuration setting specified by a network operator via the NMS. The network operator can initiate or terminate primary trunk backups at any time via the NMS. Trunk backups take over a connection should the primary trunk fail.

trunk failure

A condition (alarm) that occurs when the Cascade switch status indicates that a trunk is no longer available.

trunk restoration

A process that reroutes the PVCs carried on the backup trunk, and frees up the circuit on the backup trunk.

TXD

See *Transmit Data.*

UIO module

See *Universal Input Output module.*

UDP

See *User Datagram Protocol.*

unshielded cable

Any cable not protected from electromagnetic or radio frequency interference.

UNI

See *User-to-Network Interface.*

UNI DCE

See *User Network Interface Data Communications Equipment.*

UNI DTE

See *User Network Interface Data Terminal Equipment.*

Universal Input Output Module

A module for the Cascade switch that has three 80-pin connectors and is used for redundancy. It is also used as an I/O module for the following interfaces: X.21, RS449, V.35, EIA530, and EIA530A.

User Network Interface Data Communications Equipment

A device that performs the Frame Relay DCE functions for link management and expects a Frame Relay DTE device (e.g., Cascade switch) to be attached to it.

User Network Interface Data Terminal Equipment

A device that performs the Frame Relay DTE functions for link management. The user specifies this option on the NMS to connect to a Frame Relay DCE where the Cascade switch acts as the DTE.

User-to-Network Interface (UNI)

A standard defined by the ATM Forum for public and private ATM network access. UNI connects an ATM end system (such as a router) and an ATM switch, and is also used in Frame Relay. UNI is called SNI (Subscriber Network Interface) in SMDS.

V.35

A standard module used for communication between a network access device and a packet network. It provides clocking from 19.2 Kbps to 4.0966 Mbps.

VC

See *Virtual Channel.*

VCI

See *Virtual Circuit Identifier.*

virtual bandwidth

Channel capacity calculated to allow for oversubscription of channel usage.

Virtual Channel

A connection between two communicating ATM networks.

Virtual Circuit

A logical circuit set up to ensure reliable communication between two network devices. See also *PVCs* and *SVCs.*

Virtual Circuit Identifier

A 16-bit field in the header of an ATM cell. VCI is an addressing identifier used to route cell traffic.

Virtual Path

A group of VCs carried between two points. A VP provides a way to bundle traffic headed in the same direction.

Virtual Path Identifier

An 8-bit field in the header of an ATM cell. A VPI is an addressing identifier used to route cell traffic.

VPI

See *Virtual Path Identifier.*

VP

See *Virtual Path.*

Wide Area Network (WAN)

A network that covers a large geographic area.

yellow alarm

A T1 alarm that is generated when the interface receives a red alarm signal from the remote end.

Index