




The Basic Guide to **Frame Relay** Networking

Your Complete Guide

to Frame Relay

from the

Frame Relay Forum

**Frame
Relay
Forum** 
www.frforum.com

The Basic Guide to **Frame Relay** **Networking**

**Your Complete Guide
to Frame Relay
from the
Frame Relay Forum**



Notice

Contributors

Todd Bahner, ADC Kentrox

Skip Carlson, Cabletron Systems

Anne Exter, Bell Atlantic

Mark Kaplan, Newbridge Networks

Chris Nicoll, Current Analysis, Inc.

Cheryl Vandegriff Hyon, Sync Research

Additional contributing companies include

MCI, Visual Networks, Netrix

Editor

Jan Thibodeau

JT Communications LLC

Design

Alan Greco Design

Copyright ©1998 Frame Relay Forum

Frame Relay Forum

39355 California Street Suite 307

Fremont, California 94538

No part of this publication may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopy, recording, or information storage and retrieval systems now known or to be invented, without permission in writing from the Frame Relay Forum.

TABLE OF CONTENTS

Introduction	iii
Chapter 1	
Basic Gear	7
Chapter 2	
How Frame Relay Works	17
Chapter 3	
Frame Relay Signaling Mechanisms	26
Chapter 4	
Frame Relay Standards and Interoperability	40
Chapter 5	
Where Frame Relay is Used	46
Chapter 6	
Planning Your Frame Relay Network	70
Frame Relay Glossary	75
Appendix	85

INTRODUCTION

How to Use This Guide

Understanding a new technology is a lot like taking a hike on an unfamiliar trail. It helps to have a guide. Well, since we can't hike along with you, we thought we'd do the next best thing: provide a trail guide that will help you chart your course.

We're assuming, by the way, that you're not a completely uninitiated hiker. That is, you've got some basic familiarity with data communications and basic internetworking concepts, as well as data communications devices and their functions.

Like any good trail guide, we've tried to present the information in an easy-to-read format. Stop at the **Base Camp** in each chapter. That's where you'll find out what information is in that chapter and how it's organized. Typically, the **Basic Trail** will offer the fundamental information about the chapter topic. On the **Advanced Trail**, you'll get more challenging information. Here, we'll go into a little more depth on the topic and use a few more technical terms.

The Base camp will also list the **View Points** in each chapter. You guessed it...the things to look at. Tables, diagrams, figures.

Finally, if you've only got 20 minutes or so, take the **Shortcut**.

We believe that most network planners will benefit by reading this guide from cover to cover. But we also know that if you're involved in networking today, your hectic schedule may force you to skim through this booklet for highlights until you can find a quiet moment to read more thoroughly. The Shortcuts contain summaries of the major points you should know about frame relay networking. A quick scan can tell you which sections are most relevant to your planning needs. Here's how the content breaks out:

- In the first chapter, you'll get your basic gear: a background on why frame relay was developed and what the benefits are. The



Base Camp



Basic Trail



Advanced Trail

Advanced Trail will discuss the ways in which circuit switching and X.25 switching are less than ideal and how they measure up to the demands of certain new high speed network applications.

- Chapter 2 defines the simple way in which frame relay data is switched, based on the address at the beginning of the frame.
- Chapter 3 describes the various mechanisms that are used by the frame relay network to communicate with the user device in order to avoid congestion on the network, to recover from an overload situation, and to convey the status of various connections.
- Chapter 4 discusses frame relay standards and interoperability and the work of the Frame Relay Forum.
- Chapter 5 examines several common applications used over frame relay networks.
- Chapter 6 discusses the steps you need to take and the questions you'll want to consider if you're planning a frame relay network.
- In the back of the book, you'll find a comprehensive frame relay glossary.



View Points



Shortcut

So fill up your canteen, lace up your hiking boots, and let's set out on the trail.

CHAPTER 1

BASIC GEAR



Base Camp

In this chapter, we will introduce you to frame relay and how it works. In other words, we'll give you the basic gear you'll need to continue your exploration.

Basic Trail: The Basic Trail will give you a definition of frame relay as a technology. Then, we'll explore the networking trends that combined to create a market need for frame relay. Finally, we'll talk about the benefits of using frame relay in your network.

Advanced Trail: If you take the advanced trail, you'll find a comparison of the characteristics of frame relay and other network switching technologies, namely Time Division Multiplexing (TDM), circuit switching and X.25 packet switching. (For a more detailed comparison of frame relay processing and X.25 processing, see Chapter 2.)



View Points:

- Figure 1: Frame relay network
- Figure 2: Opens Systems Interconnection (OSI) Model
- Table 1: Comparison chart of TDM circuit switching, X.25, and frame relay

Shortcut: If you have only a few minutes, take the shortcut. It will give you a quick overview of how frame relay was developed and the benefits of frame relay.



Basic Trail

What is Frame Relay?

Frame relay is a high-speed communications technology that is used in hundreds of networks throughout the world to connect LAN, SNA, Internet and even voice applications.

Simply put, frame relay is a way of sending information over a wide area network (WAN) that divides the information into frames or packets. Each frame has an address that the network uses to determine the destination of the frame. The frames travel through a series of switches within the frame relay network and arrive at their destination.

Frame relay employs a simple form of packet switching that is well-suited to powerful PCs, workstations and servers that operate with intelligent protocols, such as SNA and TCP/IP. As a result, frame relay offers high throughput and reliability that is perfect for a variety of today's business applications.

A Quick Look at a Frame Relay Network

A frame relay network consists of endpoints (e.g., PCs, servers, host computers), frame relay access equipment (e.g., bridges, routers, hosts, frame relay access devices) and network devices (e.g., switches, network routers, T1/E1 multiplexers).

Accessing the network using a standard frame relay interface, the frame relay access equipment is responsible for delivering frames to the network in the prescribed format. The job of the network device is to switch or route the frame through the network to the proper destination user device.

(See Figure 1.)

A frame relay network will often be depicted as a network cloud, because the frame relay network is not a single physical connection between one endpoint and the other. Instead, a logical path is defined within the network. This logical path is called a virtual circuit. Bandwidth is allocated to the path until actual data needs to be transmitted. Then, the bandwidth within the network is allocated on a packet-by-packet basis. This logical path is called a virtual circuit.

We'll be talking a lot more about virtual circuits – both Permanent Virtual Circuits (PVCs) and Switched Virtual Circuits (SVCs) in the next chapter. And we'll also discuss how frames or packets are "relayed" across the network.

But, before we get too technical, let's turn our attention to how and why frame relay got its start.




Figure 1: Frame Relay Network

Why was Frame Relay Developed?

From the beginning, frame relay was embraced enthusiastically by users because it was developed in response to a clear market need, namely the need for high speed, high performance transmission. Frame relay technology also made cost-effective use of widespread digital facilities and inexpensive processing power found in end user devices. Developed by and for data communications users, frame relay was simply the right technology at the right time. Let's explore the network trends that contributed to the development of frame relay.

As the 1980's came to a close, several trends combined to create a demand for and enable higher speed transmission across the wide area network:

- The change from primarily text to graphics interaction
- The increase in "bursty" traffic applications
- Intelligent end-user devices (PCs, workstations, X-Windows terminals) with increased computing power
- The proliferation of LANs and client/server computing
- Widespread digital networks

Need for Increased Speed

Today, rapid storage and retrieval of images for interactive applications is as common as transmitting full screens of text was in the 1970s and 1980s. Early graphics applications users who were accustomed to rapid information transfer over their LANs expected similar response times when transmitting data over the wide area. Since peak bandwidth requirements for graphics were substantially higher than for text transactions, increased bandwidth and throughput were clearly needed if response time expectations were to be met.

Dynamic Bandwidth Requirements

This type of LAN user required high bandwidth in bursts, followed by periods of idle time. "Bursty" traffic, as we call it, is well-suited for statistical sharing of bandwidth, which is a characteristic of frame relay technology.

Smarter Attached Devices

As networking requirements were changing, computing power was changing as well. Decreasing cost of processing power resulted in the proliferation of intelligent PCs and powerful workstations and servers all connected by LANs. These new end-user devices also offered the possibility of performing protocol processing, such as error detection and correction. This meant that the wide area network could be relieved of the burden of application layer protocol processing – another perfect fit for frame relay.

End-user equipment was becoming more sophisticated in its ability to recognize errors and retransmit packets at the same time as digital facilities were reducing error rates within the network. In addition, industry-standard high layer protocols, such as TCP/IP, added intelligence to end-user devices.

Without the overhead associated with error detection and correction, frame relay could offer higher throughput than other connectivity solutions, such as X.25.

Higher Performance

More LANs in general and Internet Protocol (IP) LANs in specific, fueled the need to internetwork LANs across the wide area network, another factor that drove the growth of public frame relay services.

Some users tried to solve the internetworking challenge by

simply hooking LAN bridges or routers together with dedicated lines. This approach worked for simple networks, but as complexity increased, the drawbacks became apparent: higher transmission costs, lower reliability, limited network management and diagnostics, and hidden inefficiencies.

It soon became apparent that a better approach to LAN internetworking was to connect bridges and routers into a reliable, manageable WAN backbone designed to make the best use of facilities and offer the high performance users demanded.

Frame relay technology offered distinct advantages for the wide area network. First, it was a more efficient WAN protocol than IP, using only five bytes of overhead versus 20 for IP. In addition, frame relay was easily switched. IP switching was not widely available in the WAN, and IP routing added unnecessary delays and consumed more bandwidth in the network.

Widespread Digital Facilities

As the public telecommunications infrastructure migrated from analog facilities to high quality digital facilities, bandwidth availability increased and error rates decreased. The error-correcting capabilities of X.25 and SNA, which were developed to cope with the inherent errors of analog lines, were no longer necessary in digital wide area networks.

In the Beginning

While telecommunications managers contemplated the task of how to manage growing user requirements and increased network complexity, frame relay was being conceived in Bell Labs as part of the ISDN specification. Soon, frame relay evolved into a network service in its own right.

In 1990, four companies collaborated to refine the frame relay specification. "The Gang of Four," as they were known, later formed the Frame Relay Forum, which was incorporated in 1991. Since its inception, the Frame Relay Forum has grown to more than 300 members, evidence of widespread acceptance of frame relay as the method of choice for high-speed networks.

We'll discuss the work of the Frame Relay Forum in more detail in Chapter 4. Now, let's focus on the benefits of frame relay as a technology.

Banking on Frame Relay

The success of a new technology is often dependent upon compelling economic reasons for implementation. In the years since its inception, users of frame relay have found that it provides a number of benefits over alternative technologies:

1. lower cost of ownership
2. well-established and widely-adopted standards that allow open architecture and plug-and-play service implementation
3. low overhead, combined with high reliability
4. network scalability, flexibility and disaster recovery
5. interworking with other new services and applications, such as ATM

Cost of Ownership

Frame relay provides users a lower cost of ownership than competing technologies for a number of reasons:

- It supports multiple user applications, such as TCP/IP, NetBIOS, SNA and voice, eliminating multiple private lines to support different applications at a single site.
- It allows multiple users at a location to access a single circuit and frame relay port, and it efficiently uses bandwidth, thanks to its statistical multiplexing capability.
- Because only a single access circuit and port are required for each site, tremendous savings are often realized in recurring costs of transmission facilities.
- Customers realize a significant reduction in hardware, such as the number of router cards and DSU/CSUs required, reducing up-front costs and on-going maintenance costs when compared with point-to-point technologies.

Standards

Well established, widely-adopted standards are key to equipment interoperability and efficient use of capital acquisition funds.

With frame relay, users can relax knowing that frame relay standards are in place both in the United States and around the world. This ensures that equipment and services provided today will be functional for the long run, with constantly evolving standards to support new applications and to meet dynamic market place needs.

Low Overhead and High Reliability

By using only two to five bytes of overhead, frame relay makes efficient use of each frame. This means that more of the frame relay bandwidth is used for sending user data and less for overhead. Bandwidth utilization of frame relay is nearly equivalent to leased lines and better than numerous other technologies, such as X.25 or IP switching.

When the effects are spread over a large network with numerous sites, results improve exponentially:

- Simplified switching means less delay.
- Statistical multiplexing leads to more efficient bandwidth use.
- Low overhead means bandwidth is used only for user data, not for data transport.

Network Scalability, Flexibility and Disaster Recovery

To the end user, a frame relay network appears straightforward: one user simply connects directly to the frame relay cloud. A frame relay network is based on virtual circuits which may be meshed or point-to-point, and these links may be permanent or switched. (See Chapter 2 for more details.)

Because of this structure, frame relay is more easily scalable than a fixed point-to-point network. This means that additions and changes in a network are transparent to end users, giving telecommunications managers the flexibility to modify network topologies easily and scale networks as applications grow and sites are added.

This inherent flexibility lends itself equally well to the provision of alternate routes to disaster recovery sites, which are, in many cases, transparent to the end user.

Interoperability with New Applications and Services

Compared with using point-to-point leased lines, frame relay suits meshed networks and hub and spoke networks equally well. This means that frame relay easily accommodates new applications and new directions of existing networks, for example, SNA migration to APPN.

In addition, frame relay standards have been developed to interwork with newly evolving services such as ATM. As new applications emerge and/or bandwidth requirements increase, networks can gracefully migrate to the appropriate technology without stranding existing network equipment.



Advanced Trail

If you're an experienced hiker, we've provided a little more of a challenging trail in this section. Here, you'll read why frame relay offers advantages over other technologies.

Frame Relay: The Right Mix of Technology

Frame relay combines the statistical multiplexing and port sharing features of X.25 with the high speed and low delay characteristics of TDM circuit switching. Defined as a "packet mode" service, frame relay organizes data into individually addressed units known as frames rather than placing it into fixed time slots. This gives frame relay statistical multiplexing and port sharing characteristics.

Unlike X.25, frame relay completely eliminates all Layer 3 processing. (See Figure 2.)




Figure 2: Open Systems Interconnection (OSI) Model

Only a few Layer 2 functions, the so-called "core aspects," are used, such as checking for a valid, error-free frame but not requesting retransmission if an error is found. Thus, many protocol functions already performed at higher levels, such as sequence numbers, window rotation, acknowledgments and supervisory frames, are not duplicated within the frame relay network.

Stripping these functions out of frame relay dramatically increases throughput (i.e., the number of frames processed per second for a given hardware cost), since each frame requires much less processing. For the same reason, frame relay delay is lower than X.25 delay, although it is higher than TDM delay, which does no processing.

In order to remove this functionality from the frame relay network, end devices must ensure error-free end-to-end transmission of data. Fortunately, most devices, especially those attached to LANs, have the intelligence and processing power to perform this function.

Table 1 summarizes the characteristics of TDM circuit switching, packet switching and frame relay

	TDM Circuit Switching	X.25 Packet Switching	Frame Relay
Time-slot multiplexing	yes	no	no
Statistical (virtual circuit) multiplexing	no	yes	yes
Port sharing	no	yes	yes
High speed (per \$)	yes	no	yes
Delay	very low	high	low

Table 1: Comparison of TDM circuit switching, packet switching, and frame relay

Frame relay uses a variable length framing structure, which, depending on user data, ranges from a few to more than a thousand characters. This feature, similar to X.25, is essential for interoperability with LANs and other synchronous data traffic, which requires variable frame size. It also means that traffic delays (although always lower than X.25) will vary, depending on frame size. Some traffic types do not tolerate delay well, especially variable delay. However, frame relay technology has been adapted to carry even delay-sensitive traffic, such as voice.



Shortcut

As the 1980's came to a close, several network trends combined to create a need for a new form of wide area network switching:

- Growth in high speed, high throughput applications
- Proliferation of end-user devices
- Increased availability of error-free, high-speed transmission lines.

This new wide area switching technology required high speed, low delay, port sharing, and bandwidth sharing on a virtual circuit basis. While existing TDM circuit switching and X.25 packet switching had some of these characteristics, only frame relay offered a full complement. These characteristics make frame relay an ideal solution for the bursty traffic sources found in LAN-WAN internetworking.

Frame relay provides a number of benefits over alternative technologies:

- lower cost of ownership
- well-established and widely-adopted standards that allow open architecture and plug-and-play service implementation
- low overhead , combined with high reliability
- network scalability, flexibility and disaster recovery
- interworking with other new services and applications, such as ATM

Frame relay offers users the ability to improve performance (response time) and reduce transmission costs dramatically for a number of important network applications. In order to be effective, frame relay requires that two conditions be met:

1. the end devices must be running an intelligent higher-layer protocol
2. the transmission lines must be virtually error-free

Other wide area network switching technologies, such as X.25 packet switching and TDM circuit switching, will remain important where line quality is not as good, when the network itself must guarantee error-free deliver, or when the traffic is intolerant of delay.

CHAPTER 2

HOW FRAME RELAY WORKS



Base Camp

In this chapter, we will discuss in more detail how frame relay works. We'll concentrate on the basic flow of data within a frame relay network.

Basic Trail: The basic trail will give beginners an overview of virtual circuits. Next, we'll show you the frame relay frame, how it's constructed and how it moves across the frame relay network. Finally, on the Basic Trail, we'll introduce the concept of discarding frames.

Advanced Trail: If you take the advanced trail, you'll find a comparison of X.25 and frame relay processing and a more detailed discussion of error recovery by higher layer protocols.



View Points:

- Figure 3: Basic frame structure of popular synchronous protocols
- Figure 4: Frame structure and header format of the frame relay frame
- Figure 5: DLCI path through the network
- Figure 6: X.25 versus frame relay processing flow chart

Shortcut: The shortcut summarizes the basic flow of data in a frame relay network.



Basic Trail

Virtual Circuits in Frame Relay

Frame relay technology is based on the concept of using virtual circuits (VCs). VCs are two-way, software-defined data paths between two ports that act as private line replacements in the network. While today there are two types of frame relay connections, switched virtual circuits (SVCs) and permanent virtual circuits (PVCs), PVCs were the original service offering. As a result, PVCs were more commonly used, but SVC products and services are growing in popularity. A more detailed discussion of SVCs and their benefits occurs in Chapter 3. For now, let's discuss the basic differences between PVCs and SVCs.

Using PVCs

PVCs are set up by a network operator – whether a private network or a service provider – via a network management system. PVCs are initially defined as a connection between two sites or endpoints. New PVCs may be added when there is a demand for new sites, additional bandwidth, alternate routing, or when new applications require existing ports to talk to one another.

PVCs are fixed paths, not available on demand or on a call-by-call basis. Although the actual path taken through the network may vary from time to time, such as when automatic rerouting takes place, the beginning and end of the circuit will not change. In this way, the PVC is like a dedicated point-to-point circuit.

PVCs are popular because they provide a cost-effective alternative to leased lines. Provisioning PVCs requires thorough planning, a knowledge of traffic patterns, and bandwidth utilization. There are fixed lead times for installation which limit the flexibility of adding service when required for short usage periods.

Using SVCs

Switched virtual circuits are available on a call-by-call basis. Establishing a call by using the SVC signaling protocol (Q.933) is comparable to normal telephone use. Users specify a destination address similar to a phone number.

Implementing SVCs in the network is more complex than


using PVCs, but is transparent to end users. First, the network must dynamically establish connections based on requests by many users (as opposed to PVCs where a central network operator configures the network). The network must quickly establish the connection and allocate bandwidth based on the user's requests. Finally, the network must track the calls and bill according to the amount of service provided.

Although SVCs were defined in the initial frame relay specifications, they were not implemented by the first carriers or vendors of frame relay. Today, applications well-suited to SVCs are driving its deployment. While PVCs offer the statistical bandwidth gain of frame relay, SVCs deliver the any-to-any connectivity that can result in network savings and flexibility. (See Chapter 3 for a more complete discussion of SVC applications and benefits.)

The Frame Relay Header and DLCI

Now that we know about virtual circuits, and the fundamental differences between PVCs and SVCs, let's take a look at the basic structure of a frame relay frame and how it accommodates other technologies.

In the most popular synchronous protocols, data is carried across a communications line in frames which are similar in structure, as shown in Figure 3.



In a frame relay frame, user data packets are not changed in any way. Frame relay simply adds a two-byte header to the frame. Figure 4 shows the frame relay frame structure and its header in more detail.




Figure 4: Frame structure and header format for frame relay.

For now, let's look at the largest portion of the header, the DLCI. The remaining six bits of the frame relay header are discussed in the next chapter.

The frame relay header contains a 10-bit number, called the Data Link Connection Identifier (DLCI). The DLCI is the frame relay virtual circuit number (with local significance) which corresponds to a particular destination. (In the case of LAN-WAN internetworking, the DLCI denotes the port to which the destination LAN is attached.) As shown in Figure 5, the routing tables at each intervening frame relay switch in the private or carrier frame relay network route the frames to the proper destination.

Note: In the figures illustrating frame relay networks, the user devices are often shown as LAN routers, since this is a common frame relay application. They could also be LAN bridges, hosts, front end processors, FRADs or any other device with a frame relay interface.




Figure 5: The DLCI denotes the port on which the destination is located

The DLCI allows data coming into a frame relay switch (often called a node) to be sent across the network using a simple, three-step process, which is shown as a flow chart in Figure 6 in this chapter.

1. Check the integrity of the frame using the Frame Check Sequence (FCS) — if it indicates an error, discard the frame.
2. Look up the DLCI in a table — if the DLCI is not defined for this link, discard the frame.
3. Relay the frame toward its destination by sending it out the port or trunk specified in the table.

Simple Rule: If there is a problem, discard the data

In order to simplify frame relay as much as possible, one simple rule exists: if there is any problem with a frame, simply discard it. There are two principal reasons why frame relay data might be discarded:

- detection of errors in the data
- congestion (the network is overloaded)

But how can the network discard frames without destroying the integrity of the communications? The answer lies in the existence of intelligence in the endpoint devices, such as PCs, workstations, and hosts. These endpoint devices operate with multilevel protocols which detect and recover from loss of data

in the network. Incidentally, this concept of using intelligent upper layer protocols to make up for a backbone network is not a new idea. The Internet relies on this method to ensure reliable communication across the network.

If you're interested in a more detailed discussion of how the upper layer protocols recover from the loss of a frame and the causes of discarded frames, continue on to the advanced trail.

If you prefer to go right to Chapter 3, you'll find a discussion of how the frame relay network handles congestion and frame discards.



Advanced Trail

Processing: Frame Relay Versus X.25

The frame relay node processes data in a relatively simple manner compared to more fully-featured protocols like X.25. Figure 6 contrasts the simplicity of frame relay with the more complex processing of X.25. (For the sake of simplicity, the diagram reflects the path of a valid data packet. Showing the steps for error recovery and non-information frame processing for X.25 would make it much more complicated.)

Recovery by Higher Layer Protocol

As you can see in Figure 6, frame relay technology simplifies the processing task, and it relies on the endpoint devices to compensate for frame loss.

How does an upper layer protocol recover from the loss of a frame? It keeps track of the sequence numbers of the various frames it sends and receives. Acknowledgments are sent to let the sending end know which frame numbers have been successfully received. If a sequence number is missing, the receiving end will request a retransmittal after waiting for a "time-out" period.

In this manner, the two end devices ensure that all of the frames eventually are received without errors. This function occurs at Layer 4, the Transport Layer, in protocols like TCP/IP and OSI Transport Class 4. By contrast, X.25 networks perform this function at Layers 2 and 3, and the endpoints need not duplicate the function in Layer 4.

While the higher layers will reliably recover from frame discards, end-to-end recovery is costly. A single lost frame will result in retransmitting all unacknowledged frames. Such recovery takes extra cycles and memory in the endpoint computers, and it uses extra network bandwidth to retransmit multiple frames.




Figure 6: Simplified model of X.25 and frame relay processing

Worst of all – it causes large delays due to the higher layer time-outs (the time spent waiting for the frame to arrive before declaring it lost) and the time spent retransmitting.

Even though the higher layers can recover when discards occur, a major factor in the overall performance of a network is the ability of the network to minimize frame discards.

Two causes of frame discards are bit errors and congestion.

Frame Discards Caused by Bit Errors

When an error occurs in a frame, typically caused by noise on the line, it is detected upon receipt of the frame using the Frame Check Sequence (FCS). (See Figure 4.)

Unlike X.25, the frame relay node detecting the error does not request the sender to correct the error by retransmitting the frame. The node simply throws the frame away and moves on to receive the next frame. It relies on the intelligence of the PC or workstation that originated the data to recognize that an error has occurred and to resend the frame. Because the cost of having the higher layers recover is great, this approach would have a disastrous effect on network efficiency if the lines are noisy, generating many errors.

Fortunately, most backbone lines are based on fiber optics and experience extremely low error rates. This lowers the frequency of error-induced endpoint data recovery on lines and effectively eliminates the problem. Thus, frame relay is useful with clean, digital lines that have low error rates, while X.25 may be required for good performance on lines with higher error rates.

Frame Discards Caused by Congestion

Network congestion occurs for two reasons. First, a network node receives more frames than it can process. This is called receiver congestion. Second, a network node needs to send more frames across a given line than the speed of the line permits, which is called line congestion.

In either case, the node's buffers (temporary memory for incoming frames awaiting processing or outgoing frames lining up to be sent) are filled and the node must discard frames until the buffers have room.

Since LAN traffic is extremely bursty, the probability of congestion occurring occasionally is high unless, of course, the user excessively overconfigures both the lines and the switches – and thereby overpays on network costs. As a result, it is very important that the frame relay network have excellent congestion management features both to minimize the occurrence and severity of congestion and to minimize the effect of the discards when they are required. Congestion management features are discussed in more detail in the following chapter.



Shortcut

The basic flow of data in a frame relay network can best be described in a series of key points:

- Data is sent through a frame relay network using a data link connection identifier (DLCI), which specifies the frame's destination.
- If the network has a problem handling a frame due to line errors or congestion, it simply discards the frame.
- The frame relay network does no error correction; instead, it relies on the higher layer protocols in the intelligent user devices to recover by retransmitting the lost frames.
- Error recovery by the higher layer protocols, although automatic and reliable, is costly in terms of delay, processing and bandwidth; thus, it is imperative that the network minimize the occurrence of discards.
- Frame relay requires lines with low error rates to achieve good performance.
- On clean lines, congestion is by far the most frequent cause of discards; thus, the network's ability to avoid and react to congestion is extremely important in determining network performance.

CHAPTER 3

FRAME RELAY SIGNALING MECHANISMS



Base Camp

In this chapter, we'll discuss how frame relay technology handles interface signaling for control. If that sounds too complicated, think of it this way: interface signaling provides information about what is happening on the network so that users can get the response time they expect and the network will have the greatest efficiency possible. Signaling mechanisms can also provide options for building different types of frame relay networks to match your applications and the performance they demand.

Basic Trail: The basic trail will acquaint you with three types of signaling mechanisms used in frame relay:

- congestion notification mechanisms
- status of the connections
- SVC signaling

Advanced Trail: If you take the advanced trail, you'll find more information about the Local Management Interface (LMI) specification, which is a connection status mechanism.



View Points:

- Figure 7: The importance of congestion management
- Figure 8: Frame relay frame showing the FECN, BECN and DE bits
- Figure 9: The use of FECN and BECN in explicit congestion notification Table 2: LMI Specifications
- Figure 10: PVC Signaling using the LMI specification
- Table 2: LMI specification

Shortcut: The shortcut will quickly cover the three types of congestion management. It will also give highlights of the two other types of interface signaling discussed in the chapter, PVC status and SVC signaling.



Basic Trail

The Need for Signaling Mechanisms

When frame relay was first proposed, it was based on a simple rule: keep the network protocol simple and let the higher layer protocols of the end devices worry about the other problems. But upon further study, it became apparent to the standards organizations that practical implementation of frame relay in real-world networks would need to specify signaling mechanisms to address three important issues:

- Allowing the network to signal that congestion exists
- Telling the status of connections (PVCs)
- Setting up new calls (SVCs)

Although these mechanisms add complexity to frame relay, the standards have an important provision which allows basic frame relay to remain simple: the use of signaling mechanisms is optional. That is, a vendor is not required to implement these features.

Without the signaling mechanisms, the resulting frame relay interface would still be compliant with the standard and data will still flow. With the signaling mechanisms, however, the throughput of the network, the response time to users, and the efficiency of line and host usage are improved.

Let's look at how these frame relay signaling mechanisms work.

Congestion Notification Mechanisms

Congestion management mechanisms, like the other signaling mechanisms, are optional for compliance, but they will affect performance. The importance of congestion management is illustrated in Figure 7.

The traffic entering the network is called the "offered load." As the offered load increases, the actual network throughput increases linearly. The beginning of congestion is represented by Point A, when the network cannot keep up with the entering traffic and begins flow control.

If the entering traffic continues to increase, it reaches a state of severe congestion at Point B, where the actual effective throughput of the network starts to decrease due to the number of re-




Figure 7: The importance of congestion management

transmissions. This causes a given frame to be sent into the network multiple times before successfully making it through.

In severe congestion, the overall network throughput can diminish, and the only way to recover is for the user devices to reduce their traffic. For that reason, several mechanisms have been developed to notify the user devices that congestion is occurring and that they should reduce their offered load.

The network should be able to detect when it is approaching congestion (Point A) rather than waiting until Point B is reached before notifying the end devices to reduce traffic. Early notification can avoid severe congestion altogether.

The ANSI specifications are very clear about the mechanisms used to indicate the existence of congestion in the network. There are two types of mechanisms to minimize, detect and recover from congestion situations, in effect providing flow control:

- Explicit Congestion Notification
- Discard Eligibility

Another mechanism that may be employed by end user devices is implicit congestion notification.

These mechanisms use specific bits contained within the header of each frame. The location of these specific bits (FECN,

BECN and DE) are shown in Figure 8.

Let's look at how each of these mechanisms function.




Figure 8: Frame relay frame showing the FECN, BECN and DE bits

Explicit Congestion Notification (ECN) Bits

The first mechanism uses two Explicit Congestion Notification (ECN) bits in the frame relay header. They are called the Forward Explicit Congestion Notification (FECN) and the Backward Explicit Congestion Notification (BECN) bits. Figure 9 depicts the use of these bits.




Figure 9: The use of FECN and BECN in explicit congestion notification

Let's suppose Node B is approaching a congestion condition. This could be caused by a temporary peak in traffic coming into the node from various sources or by a peak in the amount of traffic on the link between B and C. Here is how forward congestion notification would occur:

- Node B would detect the onset of congestion based on internal measures such as memory buffer usage or queue length.

- Node B would signal Node C (the downstream node, toward the destination) of the congestion by changing the forward ECN (FECN) contained within the frames destined for Node C from 0 to 1.
- All interim downstream nodes, as well as the attached user device, would thus learn that congestion is occurring on the DLCI(s) affected.

Depending upon the protocols used and the capabilities of the CPE device and the network switches, it is sometimes more useful to notify the source of the traffic that there is congestion, so the source can slow down until congestion subsides. (This assumes that the source is capable of responding to receipt of the congestion notification signals.) This is called Backward Congestion notification.

This is how backward congestion notification occurs:

- Node B watches for frames coming in the other direction on the connection.
- Node B sets the backward ECN bit within those frames to signal the upstream node(s) and the attached user device.

The FECN and BECN process can take place simultaneously on multiple DLCIs in response to congestion on a given line or node, thus notifying multiple sources and destinations. The ECN bits represent an important tool for minimizing serious congestion conditions.

Implicit Congestion Notification

Some upper layer protocols, such as Transport Control Protocol (TCP), operating in the end devices have an implicit form of congestion detection. These protocols can infer that congestion is occurring by an increase in round trip delay or by detection of the loss of a frame, for example. Reliance on network traffic characteristics to indicate congestion is known as implicit congestion notification.

These upper layer protocols were developed to run effectively over networks whose capacity was undetermined. Such protocols limit the rate at which they send traffic onto the network by means of a "window," which allows only a limited number of frames to be sent before an acknowledgment is received.

When it appears that congestion is occurring, the protocol can reduce its window size, which reduces the load on the network. As congestion abates, the window size is gradually increased.

The same window-size adjustment is also the normal way for the end-user devices to respond to explicit congestion notification – FECN and BECN. The ANSI standards state that implicit and explicit congestion notification are complementary and can be used together for best results.

Discard Eligibility

Frame relay standards state that the user device should reduce its traffic in response to congestion notification. Implementation of the recommended actions by the user device will result in a decrease in the traffic into the network, thereby reducing congestion. If the user device is incapable of responding to the signaling mechanisms, it might simply ignore the congestion signal and continue to transmit data at the same rate as before. This would lead to continued or increased congestion.

In this case, how does the network protect itself? The answer is found in the basic rule of frame relay: if there is a problem, discard the data. Therefore, if congestion causes an overload, more frames will be discarded. This will lengthen response times and reduce overall network throughput, but the network will not fail.

When congestion does occur, the nodes must decide which frames to discard. The simplest approach is to select frames at random. The drawback of this approach is that it maximizes the number of endpoints which must initiate error recovery due to missing frames.

A better method is to predetermine which frames can be discarded. This approach is accomplished through the use of the Committed Information Rate (CIR). The CIR is the average information capacity of the virtual circuit. When you subscribe to or buy a frame relay service from a carrier, you specify a CIR depending on how much information capacity you think your network will need.

In each frame header, there is a bit called the Discard Eligibility (DE) bit (see Figure 8). A DE bit is set to one (1) by the CPE device or the network switch when the frame is above the CIR. When the DE bit is set to 1, it makes the frame eligible for discard in response to situations of congestion. A frame with a DE bit of 1 is discarded in advance of non-discard-eligible data (those frames with a DE bit set to zero (0)). When the discard of

DE-eligible data, by itself, is not sufficient to relieve severe congestion, additional incoming frames are discarded without regard to the setting of the DE bit.

Status of Connections (PVCs and SVCs)

The next type of optional signaling mechanism defines how the two sides of a frame relay interface (e.g., the network and the router) can communicate with each other about the status of the interface and the various PVCs on that interface.

Again, these are optional parameters. It is possible to implement a frame relay interface and pass data without implementing these parameters. This signaling mechanism simply enables you to retrieve more information about the status of your network connection.

This status information is accomplished through the use of special management frames with a unique DLCI address which may be passed between the network and the access device. These frames monitor the status of the connection and provide the following information:

- Whether the interface is still active — this is called a "keep alive" or "heartbeat" signal
- The valid DLCIs defined for that interface
- The status of each virtual circuit; for example, if it is congested or not

The connection status mechanism is termed the Local Management Interface (LMI) specification. There are currently three versions of the LMI specification:

Protocol	Specification
LMI	Frame Relay Forum Implementation Agreement (IA) FRF.1 superceded by FRF.1.1
Annex D	ANSI T1.617
Annex A	ITU Q.933 referenced in FRF.1.1

Table 2: LMI Specifications

While LMI was used colloquially for the FRF.1 IA, it may also be used as a generic term to refer to any and all of the protocols.

The revised Frame Relay Forum IA FRF.1.1 calls for the mandatory implementation of Annex A of ITU Q.933.

Each version includes a slightly different use of the management protocol. Virtually all equipment vendors support LMI and most support Annex D, while Annex A is supported by fewer vendors. To ensure interoperability when your network consists of equipment from different vendors, the same version of management protocol must be at each end of the frame relay link.

For a little history of LMI and more detail on the functions of the different versions, see the advanced trail later in this chapter.

Switched Virtual Circuits

The final signaling mechanism we will discuss is SVC signaling. Unlike the previous two signaling mechanisms discussed – congestion status and connection status – SVC signaling does not offer the network operator information about the network. Rather, SVC signaling specifications allow an alternative to permanent virtual circuits. In turn, SVC signaling must provide call setup and call disconnect. Call setup includes information about the call, such as measuring data sent, acceptance, addresses and bandwidth parameters.

SVCs can also provide opportunities for new applications and new network usage. This section will discuss those alternatives and opportunities.

SVC Implementation Agreement

Implementation Agreement FRF.4 defines the needed messages and procedures to establish an SVC. Basically, the network alerts the requested destination of the incoming call and the destination chooses whether or not to accept it. If the destination accepts, the network builds the SVC across the network switches. Once the network establishes the SVC, the two endpoints can transfer information. When the endpoints no longer need the connection, either one notifies the network to terminate the call.

SVC Benefits

While current provisions for PVCs are adequate for the vast majority of near-term applications, SVC capability is beginning to gain momentum for use in public frame relay networks and for very large private networks. With SVCs, users can request set up of virtual connections only when needed and negotiate through-

put rate and burst size depending on the application.

SVC Network Applications

Some of the benefits become clearer as we look at the various applications where SVC technology is well-suited.

Remote Connectivity

At the fringes of the network or in sites where there is little need to contact other locations, SVCs are an excellent way to provide basic connectivity cost effectively. The customer pays only for the use of the network when needed, without requiring PVCs at the user-to-network interface. This application holds a great deal of promise for remote locations accessing high-speed frame relay implementations.

Overflow Traffic

There may be times of the day or night when using the burst capability of the main PVC alone cannot satisfy the need for excess capacity. Since SVCs can be set up on an adhoc basis, they can fulfill the demands of seasonal, sporadic or finite-use traffic and offer true bandwidth on demand.

Intranets and Extranets

These two applications are compelling because they allow frame relay (with SVCs) to access the Internet territory. For customers uncomfortable with the variations in quality of the Internet, building an intranet or extranet using frame relay may be a good alternative. This represents a whole new set of services carriers could offer.

Dial Access

To access a frame relay service from a carrier, a "local loop" connects the user premises to the carrier's nearest point-of-presence (POP). This local loop can be either a leased line or a dial line. Users dialing into the frame relay network can connect to either a PVC network or an SVC network.

Disaster Recovery or Alternate Network Paths

For networks using back-up or recovery sites and alternate network paths, SVCs can provide an economical alternative to leased lines, switched services or PVCs. They provide the net-

work flexibility required when leased lines are not available or there is no time to provision PVCs.



Advanced Trail

The advanced trail will discuss two topics in more depth: the LMI specification and the SVC Implementation Agreement.

In this section, we will refer to the two major standards-setting organizations, the American National Standards Institute (ANSI) and the International Telecommunications Union - Telecommunications Services Sector (ITU-T). For more information on standards, please read Chapter 4.

LMI Specification

As you may recall if you went through the basic trail, there are three versions of the LMI specification:

FRF. 1 superseded by FRF.1.1, ANSI T1.617 and ITU Q.933 referenced in FRF.1.1.

The first definition for PVC status signaling was in the LMI specification. The protocol defined for the LMI provides for a "status inquiry" message which the user device (e.g., router) can send, either simply as a "keep alive" message to inform the network that the connection to the router is still up, or as a request for a report on the status of the PVCs on that port.

The network then responds with a "status" message, either in the form of a "keep alive" response or in the form of a full report on the PVCs. (See Figure 10.) An additional optional message, "status update," is also defined which enables the network to provide an unsolicited report of a change in PVC status.

Notice that the LMI status query only provides for one-way querying and one-way response, meaning that only the user device (e.g., router) can send a "status inquiry" message, and only the network can respond with a "status" message. While this approach was simple to implement, it resulted in some limitations in functions. Using status inquiries in this manner, both sides of the interface are unable to provide the same commands and responses. Most notably, it addressed only the User Network Inter-

face (UNI) and would not work in a Network-to-Network Interface (NNI) due to the one-way communications of the interface. UNI provides the end device interface to the network.

NNI provides the ability for networks to query and respond to one another. When only UNIs are available, this could lead to problems within hybrid private/public networks, where a private network node would have a frame relay NNI interface to a public frame relay service.




Figure 10: PVC Status Signaling per LMI specification

Therefore, just before final approval of the standard for frame relay signaling, ANSI extended the standard to provide a bi-directional mechanism for PVC status signaling that is symmetric. The bi-directional mechanism provides the ability for both sides of the interface to issue the same queries and responses. This mechanism is contained in Annex D of T1.617, known simply as Annex D. Annex D works in both the UNI and NNI interfaces.

In contrast to the LMI (which uses DLCI 1023), Annex D reserves DLCI 0 for PVC status signaling. The current requirement in FRF.1.1, Annex A signaling, is similar to Annex D and also uses DLCI 0.

To insure interoperability in a multi-vendor network environment, the same version of management protocol must be at each end of the frame relay link.

SVC Implementation Agreement

The SVC Implementation Agreement is based on existing SVC

standards in ANSI and ITU-T. The current SVC standards are T1.617 in ANSI and Q.933 in ITU-T. These two documents lay the basis for Q.2931, the standard for access signaling for ATM (Asynchronous Transfer Mode), as well as for the PVC management procedures for frame relay.

The SVC Implementation Agreement can enable expanded service in frame relay networks. Use in internal networks involves implementing SVCs that are internal to a public or private network. The SVCs would remain transparent to the users who maintain their user-to-network interface PVCs, for example, in the case of disaster recovery. In wide area networks SVCs may be used over large geographic areas such as transatlantic applications, which have been traditionally cost-prohibitive.

ISDN and Switched Access for PVCs and SVCs

Access on demand for PVCs and SVCs, whether via Integrated Services Digital Network (ISDN) or switched access is another method to reach the frame relay network. Access on demand holds a great deal of promise for remote locations accessing high-speed frame relay implementations.

In switched access, a circuit-switched connection to the frame relay switch can be established using the existing voice network. An indication is then sent to the switch that a frame relay call is being established; the switch makes the connection and bills the call appropriately. The customer pays only for the use of the local loop when needed, without requiring PVCs at the user-to-network interface. The same benefits are true for ISDN access and E.164 addressing and lead to true, any-to-any connectivity through ISDN or switched access.



Shortcut

Interface signaling mechanisms provide information about the frame relay network so that network operators can improve efficiency.

Signaling mechanisms also provide optional ways of configuring your frame relay network to match applications usage.

There are three types of signaling mechanisms used in frame relay:

- congestion notification mechanisms
- status of the connections
- SVC signaling

The ANSI standard defines a method for the network to signal the existence of congestion called the Explicit Congestion Notification (ECN) bits.

Frame relay uses FECN (Forward ECN) and BECN (Backward ECN) bits to notify end user devices about network congestion.

Although the frame relay protocol does not respond to congestion, some higher layer protocols for end-user devices may respond to Implicit Congestion Notification by recognizing that end-to-end delays have increased or that frames have been dropped.

The use of the "discard eligibility" (DE) bit can be a powerful tool for managing throughput, including the ability to meter traffic and to guarantee a level of service.

The ANSI and ITU standards define a mechanism for communicating the status of PVCs on a frame relay interface based on a modification of the method in the LMI specification.

SVC signaling allows an alternative to permanent virtual circuits which can improve the efficiency of the network. SVCs can also provide opportunities for new applications and new network usage.

CHAPTER 4

FRAME RELAY STANDARDS AND INTEROPERABILITY



Base Camp

No discussion of frame relay would be complete without mentioning standards.

This chapter will discuss frame relay standards – those that currently exist and how they developed. We'll also acquaint you with the Frame Relay Forum and the Implementation Agreements the Forum develops to ensure frame relay interoperability.

Basic Trail: The basic trail will discuss the development of ANSI and ITU standards for frame relay. This section also gives an overview of the Frame Relay Forum and lists the current Implementation Agreements.

There is no advanced trail in this chapter.



View Points:

- Table 3: Frame relay standards
- Table 4: List of frame relay Implementation Agreements (IAs)

Shortcut: The shortcut summarizes current frame relay standards and the work of the Frame Relay Forum.



Basic Trail

How Did Frame Relay Standards Develop?

The remarkable degree of industry consensus about the need for frame relay to supplement existing switching technologies resulted in rapid development of industry standards. There are two major standards organizations which are active in this area:

- American National Standards Institute (ANSI)
- International Telecommunications Union - Telecommunications Services Sector (ITU-T), which was formerly called the Consultative Committee for International Telephone and Telegraph known as CCITT

To understand how frame relay standards developed, we need to go back to 1988. That year, ITU-T (then called CCITT) approved Recommendation I.122, "Framework for additional packet mode bearer services."

I.122 was part of a series of ISDN-related specifications. ISDN developers had been using a protocol known as Link Access Protocol - D channel (LAPD) to carry the signaling information on the "D channel" of ISDN. (LAPD is defined in ITU Recommendation Q.921.)

Developers recognized that LAPD had characteristics that could be very useful in other applications. One of these characteristics is that it has provisions for multiplexing virtual circuits at level 2, the frame level (instead of level 3, the packet level as in X.25). Therefore, I.122 was written to provide a general framework outlining how such a protocol might be used in applications other than ISDN signaling.

At that point, rapid progress began, led by an ANSI committee known as T1S1, under the auspices of the Exchange Carrier Standards Association (ECSA). This work resulted in a set of standards defining frame relay very clearly and completely. The principal frame relay standards are shown in Table 3.

Description	ANSI Standard	Status	ITU Standard	Status
Service Description	T1.606	Standard	I.233	Approved
Core Aspects	T1.618 (previously known as T1.6ca)	Standard	Q.922 Annex A	Approved
Access Signaling	T1.617 (previously known as T1.6fr)	Standard	Q.933	Approved

Table 3: Frame relay standards

T1.606 was approved early in 1990. Thanks to the hard work of the ANSI committee, coupled with a clear mandate from the market, the remaining ANSI standards sped through the stages of the standards process to receive complete approval in 1991.

Frame Relay Standards

The fast pace of frame relay standards work at ANSI was matched by an outstanding degree of cooperation and consensus in the international arena. As a result, the ITU-T recommendations for frame relay are in alignment with the ANSI standards and have also moved rapidly through the approval process. (Authors' note: although we refer to ANSI standards throughout this book, most of the discussion applies equally to the ITU-T standards.)

Interoperability and Standards Compliance

With the number of options in the standards and the range of design choices faced by vendors, what does it mean to a customer interested in interoperability?

Minimum requirements: Basic Data Handling

In order to achieve interoperability, frame relay network equipment must comply with the basic data transport method specified in the ANSI standard, which states that frame relay takes place using the DLCI in the two-byte frame relay header. This subject is covered in Chapter 2. With that relatively simple re-

quirement met, there is interoperability. The remaining requirements determine how well the network performs and whether it can be managed.

Required for real-world networks: Interface signaling

The interface control mechanisms described in Chapter 3 are optional. Data flows without them, and ignoring them is not a violation of the standard.

In real networks, however, you may find interface signaling essential to ensure that the network operates with adequate performance. Otherwise, there is no way for a network to control congestion. This means that as the traffic increases, network throughput may decrease. And throughput may continue to decrease as congestion is further exacerbated by more discards and retransmissions.

The Frame Relay Forum

The Frame Relay Forum is a non-profit organization dedicated to promoting the acceptance and implementation of frame relay based on national and international standards. Established in 1991, the Forum now has more than 300 member companies worldwide.

The Forum develops and approves Implementation Agreements (IAs) to ensure frame relay interoperability and facilitates the development of standard protocol conformance tests for various protocols. Since the earliest frame relay IAs, additional features, such as multicast, multiprotocol encapsulation and switched virtual circuit signaling, have been defined in subsequent IAs to increase the capabilities of frame relay.

Work by the Frame Relay Forum has resulted in the completion of several implementation agreements, which are listed in Table 4. Work on implementation agreements and standards is ongoing to add enhancements and broaden the applications for frame relay.

An updated listing of IAs can be found on the Frame Relay Forum web site at <<www.frforum.com>> and in the Forum's quarterly newsletters.

FRF.1.1	User-to-Network (UNI) Implementation Agreement
FRF.2.1	Frame Relay Network-to-Network (NNI) Implementation Agreement
FRF.3.1	Multiprotocol Encapsulation Implementation Agreement (MEI)
FRF.4	Switched Virtual Circuit Implementation Agreement
FRF.5	Frame Relay/ATM PVC Network Interworking Implementation Agreement
FRF.6	Frame Relay Service Customer Network Management Implementation Agreement (MIB)
FRF.7	Frame Relay PVC Multicast Service and Protocol Description Implementation Agreement
FRF.8	Frame Relay/ATM PVC Service Interworking Implementation Agreement
FRF.9	Data Compression over Frame Relay Implementation Agreement
FRF.10	Frame Relay Network-to-Network Interface SVC Implementation Agreement
FRF.11	Voice over Frame Relay Implementation Agreement
FRF.12	Frame Relay Fragmentation Implementation Agreement

Table 4: Frame Relay Forum Implementation Agreements (IAs)



Shortcut

There are two major standards organizations:

- American National Standards Institute (ANSI)
- International Telecommunications Union- Telecommunications Services Sector (ITU-T)

The initial frame relay standard was approved in 1990 by ANSI, and the remaining standards were approved by 1991. ITU recommendations for frame relay are in alignment with the ANSI standards.

In order to achieve interoperability, frame relay network equipment must comply with the basic data transport method specified in the ANSI standard, which states that frame relay takes place using the DLCI in the two-byte frame relay header.

Although interface control mechanisms are optional, they are essential to ensure that the network operates with adequate performance.

The Frame Relay Forum is a non-profit organization dedicated to promoting the acceptance and implementation of frame relay based on national and international standards.

The Forum develops and approves Implementation Agreements (IAs) to ensure frame relay interoperability. Since the earli-

est frame relay IAs, additional features, such as multicast, multiprotocol encapsulation and switched virtual circuit signaling, have been defined to increase the capabilities of frame relay.

CHAPTER 5

WHERE FRAME RELAY IS USED



Base Camp

Previous chapters covered the what and how of frame relay. By focusing on applications, this chapter provides insight into the practical benefits of frame relay. We'll discuss frame relay applications that are widely deployed and others that are emerging.

Basic Trail: The basic trail will give you an overview of four popular and growing applications: LAN peer-to-peer networking over frame relay, SNA over frame relay, voice over frame relay (VoFR) and frame relay-to-ATM interworking.

Advanced Trail: On the advanced trail, we'll go into more detail about three of the applications covered on the basic trail. Specifically, you'll also read how FRF.3.1 provides interoperability in SNA networks and how traffic is managed in SNA over frame relay applications. We'll also talk about how voice over frame relay works and the associated trade-offs. Finally, you'll find a discussion of the Interworking Function (IWF) and FUNI or ATM frame-based UNI.



View Points:

- Figure 11: Traditional Solution for LAN or Client/Server Networking
- Figure 12: Frame Relay Solution for LAN or Client/Server Networking
- Figure 13: Parallel bank branch networks
- Figure 14: Consolidated bank network

- Figure 15: Integrated voice and data network
- Figure 16: Frame/ATM Network Interworking
- Figure 17: Frame/ATM Service Interworking
- Figure 18: Typical Multidrop SNA Network
- Figure 19: SNA Network Migrated to Frame Relay
- Figure 20: NCP direct FRE.3.1 Network
- Figure 21: FRAD FRE.3.1 Network
- Figure 22: Normal Speech Components
- Figure 23: Frame/ATM Network Interworking (Encapsulation)
- Figure 24: Frame/ATM Service Interworking (Transparent)
- Figure 25: Frame/ATM Service Interworking (Translation)
- Figure 26: ATM DXI and ATM FUNI
- Table 5: Comparison of Frame Relay/ATM Interworking with FUNI and ATM DXI

Shortcut: The shortcut will give the highlights of the four applications discussed and the associated benefits.



Basic Trail

Initially, frame relay gained acceptance as a means to provide end users with a solution for LAN-to-LAN connections and to meet other data connectivity requirements. Frame relay's

compelling benefit is that it lowers the cost of ownership compared to competing technologies:

- Frame relay supports multiple user applications, such as TCP/IP, NetBIOS, SNA and voice and thus eliminates the need for multiple private line facilities supporting different applications at a single site.
- Because it statistically multiplexes, frame relay allows multiple users at a location to access a single circuit and frame relay port, making efficient use of the bandwidth.
- Since only a single access circuit and port are required for each user site, users often realize tremendous savings in the cost of transmission facilities.
- Customers realize a significant reduction in the number of router cards and DSU/CSUs required, reducing up-front costs as well as ongoing maintenance compared with point-to-point technologies.

Application #1: Meshed LAN Peer-to-Peer Networking

In a traditional solution for LAN or client/server networking across a WAN, meshed network implementations can be costly. Since private line pricing is distance sensitive, the price of the network increases as geographic dispersion increases. Changes in network design normally require physical reconfigurations in addition to software changes, which increases the time to administer the changes. (See Figure 11.)

Frame Relay for LAN or Client/Server

By moving to frame relay for LAN or client/server applications, additional VCs between locations can be provisioned for minimal incremental cost. Most public frame relay pricing is distance insensitive. Virtual connections are software configurable. Changes to VCs can be done relatively quickly. This makes frame relay ideal for meshed configurations. (See figure 12.)




Figure 11: Traditional Solution for LAN or Client/Server Networking




Figure 12: Frame Relay Solution for LAN or Client/Server Networking

Application #2: SNA Over Frame Relay

Over the past few years there has been a migration of legacy traffic, such as BSC (binary synchronous communications) and Systems Network Architecture (SNA), from low speed leased lines onto frame relay services. Ratified standards from the Internet Engineering Task Force (IETF) and the Frame Relay Forum enable the encapsulation of multiple protocols, including SNA, over frame relay networks. Together, they provide a standard method of combining SNA and LAN traffic on a single frame relay link. This enables FRADs and routers, which provide network connectivity, to handle time-sensitive SNA and bursty LAN traffic simultaneously.

The integration of legacy and LAN-to-LAN traffic provides network administrators with a more efficient, flexible and cost-effective network as well as a number of other benefits:

- Simplify the network
- Leverage investment in capital equipment
- Move in SNA's stated direction, with migration strategies to distributed and peer-to-peer enterprise networks
- Dramatically lower line costs – a potential of 30 to 40 percent compared to dedicated links
- Provide up to a 40 percent increase in network utilization through frame relay's multiprotocol support
- Experience no disruption of operations – integrity and control of the network are sustained with NetView and SNMP management
- Offer high performance networking for Advanced Peer-to-Peer Networking (APPN)

Let's explore a few of these benefits. SNA installations have expanded their use of frame relay because it is a mature, proven and stable technology. Moreover, users can leverage existing capital equipment and maintain existing network management practices.

Frame relay can be integrated into SNA networks with little or no disruption. Users may migrate to frame relay without any changes to Front End Processor (FEP) hardware or software, SNA naming or network topologies.

Frame relay allows the familiar NetView tools and practices to be maintained, so there is no need to retool network operations or retrain operations staff. This allows users to migrate at their own pace to multi-vendor enterprise network management such as SNMP.

Because frame relay is consistent with SNA's stated direction, end users have the comfort of adopting a migration strategy to distributed and peer-to-peer enterprise networks to advance and optimize their SNA network. Users benefit from improved response times and session availability with higher performance than the original multidrop network.

A typical frame relay SNA application will illustrate these benefits.

Frame Relay in a Banking Application

In a large bank with many branch locations, SNA and BSC devices are co-located with LANs, resulting in parallel branch networks and high monthly WAN costs. (See Figure 13.)




Figure 13: Parallel SNA, BSC, Alarm, and LAN Branch networks

Consolidating traffic on frame relay customer premise equipment (CPE) combines serial protocol networks and LAN networks in each branch. The CPE provides an integration of legacy devices typically found at a branch with emerging devices that support client/server applications.

How does it work? The frame relay CPE consolidates SNA/SDLC (Synchronous Data Link Control) and BSC data and LAN data onto the frame relay-based WAN. This eliminates multiple single protocol leased lines connecting the branch to its host resources. It also exploits the advantages of the higher performance LAN internetwork to consolidate serial and LAN traffic, compared with low speed analog leased-line networks.

The result is better performance, greater reliability and lower costs. Because one frame relay access line can be used to reach the same number of sites as multiple leased lines, the amount of networking equipment may be reduced. Monthly telecommuni-

cations charges are also reduced and the network is simplified. Other benefits include efficient bandwidth utilization, predictable and consistent response times, enhanced session reliability and simplified network topologies and troubleshooting.

For example, a BSC-to-frame relay conversion can be performed by the frame relay CPE for BSC Automated Teller Machine locations that are supported by an SNA host at the data center. This leverages the bank's investment in capital equipment. (See Figure 14.)




Figure 14: Consolidated Bank Network

SNA Over Frame Relay Pays

Frame relay enables mission-critical SNA networks to improve performance and reduce operating costs. These savings are available because frame relay meets the response time, availability and management requirements of mission-critical applications.

If you're interested in the details of how frame relay can be used as a replacement for SDLC point-to-point networks and how FRF.3.1 provides interoperability, take a look at the SNA section in the advanced trail.

Application #3: Voice Over Frame Relay (VoFR)

Today, non-traditional uses for frame relay are beginning to emerge. One new application, voice over frame relay (VoFR), offers telecommunication and network managers the opportunity to consolidate voice and voice-band data (e.g., fax and analog modems) with data services over the frame relay network. (See Figure 15.)




Figure 15: Integrated Voice and Data Network

In migrating leased line networks to frame relay, many network administrators found a cost-effective solution for their data needs. However, since many leased-line networks also carried voice, a solution was needed to address corporate voice requirements.

With the ratification of FRF.11, a standard was established for frame relay voice transport. Among other things, FRF.11 defines standards for how vendor equipment interoperates for the transport of voice across a carrier's public frame relay network.

Maximizing Frame Relay Networks

FRF.11 enables vendors to develop standards-based equipment and services that interoperate. It also enables network managers seeking to reduce communications costs and maximize their frame relay network to consider VoFR as an option to standard voice services.

In some cases, users may find they have excess bandwidth in

their frame relay network that could efficiently support voice traffic. Other telecommunications managers may find that the incremental cost of additional frame relay bandwidth for voice traffic may be more cost-effective than standard voice services offered by local or long distance carriers.

VoFR can provide end users with a cost effective option for voice traffic transport needs between company locations. For instance, the network manager may integrate some voice channels and serial data over a frame relay connection between a branch office and corporate headquarters. By combining the voice and data traffic on a frame relay connection already in place, the user has the potential to obtain cost-effective intracompany calling and efficient use of the network bandwidth.

Because it does not significantly increase network architecture, link speeds or CIR, the integration of voice, fax and data traffic over a single access link provides a viable option for network managers and adds to the growing list of new and non-traditional applications for frame relay.

For a discussion of how VoFR works, refer to the advanced trail in this chapter.

Users are finding that frame relay offers another significant advantage: the ability to interwork with other advanced services, such as ATM.

Application #4: Frame Relay-to-ATM Interworking

Frame Relay/ATM Interworking is a viable solution that provides users with low cost access to high speed networks. Ratified by both the ATM and Frame Relay Forums, the Frame Relay/ATM PVC Interworking Implementation Agreements (IAs) provide a standards-based solution for interworking between existing or new frame relay networks and ATM networks without any changes to end user or network devices.

Why do users want to interwork frame relay and ATM? While frame relay is well suited for many applications including LAN internetworking, SNA migration and remote access, other applications, such as broadcast video and server farm support, may be better suited for ATM networks.

Users are also interested in interworking frame relay and ATM networks to protect their capital investment in existing frame relay networks and to support planned migrations from frame relay to ATM.

Frame Relay/ATM SVC Interworking IAs are currently being developed.

Frame Relay-to-ATM Interworking Standards

There are two Frame Relay/ATM Interworking IAs for PVCs, each encompassing two different types of interworking. The first one, Frame Relay/ATM Network Interworking for PVCs (FRF.5) allows network administrators to scale the backbone beyond the 45 Mbps trunks supported by frame relay. In other words, it provides the standards for ATM to become a high speed backbone for frame relay PVC users. The second one, Frame Relay/ATM Service Interworking for PVCs (FRF.8) defines the standard for frame relay PVC and ATM PVC end users or systems to communicate seamlessly.

Frame Relay/ATM Network Interworking for PVCs can be thought of as encapsulation while Frame Relay/ATM Service Interworking for PVCs is translational between the two protocols. Let's take a closer look at each standard.

Frame Relay/ATM Network Interworking for PVCs

Frame Relay/ATM Network Interworking allows Frame Relay end-user or networking devices such as FRADs or routers to communicate with each other via an ATM network employed as the backbone.

For example, SNA terminal users connected to FRADs in branch offices communicate with frame relay-attached IBM 3745 Communications Controllers located in corporate headquarters locations using a high speed ATM network as the backbone.




Figure 16: Frame Relay/ATM Network Interworking

The frame relay devices interact as if they are using frame relay for the entire connection without knowing that an ATM network is in the middle. An ATM backbone connecting multiple frame relay networks can provide scalability and high speed support

for a large number of locations and end-user devices, without requiring changes to the devices themselves.

Frame Relay/ATM Service Interworking for PVCs

Frame Relay/ATM Service Interworking enables communication between an ATM and frame relay network or end user devices. Frame Relay/ATM Service Interworking allows existing frame relay devices in the remote branch offices to communicate with end users at headquarters who are using ATM-based applications.

By enabling existing devices to access new ATM-based applications, Frame Relay/ATM Service Interworking protects the investment in existing equipment. This promotes the decoupling of client and server sides of the network, allowing each to use the resources that best meet bandwidth requirements and budget constraints.




Figure 17: Frame Relay/ATM Service Interworking

A Quick Look at IWF and FUNI

Before we leave frame relay to ATM interworking, there are two more topics we need to touch upon briefly. Then, if you're interested in more details on these topics, you can proceed to the advanced trail. The topics are Interworking Function (IWF) and Frame-based User-to-Network Interface (FUNI).

An important advantage of Frame Relay/ATM Interworking is that it provides solutions to support communications between frame relay and ATM environments without modifications to end-user devices. However, successful support of end-to-end communications in a Frame Relay/ATM Interworking environment requires performing technical functions to compensate for the differences between frame relay and ATM. These functions are defined within the Frame Relay/ATM Service and Network Interworking IAs and are provided by the IWF generally located on the switch at the boundaries of the frame relay and ATM

services. The advanced trail will discuss the responsibilities of the IWF and how it works.

FUNI was defined by the ATM Forum to provide frame-based access to ATM networks. It is an alternative to Frame Relay/ATM Service Interworking and it is most viable where the wide area infrastructure uses ATM. FUNI enables ATM quality of service levels for network throughput and delay to be maintained end-to-end, despite the fact that the access method is frame-based, rather than native or cell-based ATM.

Approved by the ATM Forum in 1995, the FUNI specification provides improved efficiency of access line bandwidth. FUNI enables users to transmit variable length (low overhead) frames rather than fixed length cells to the ATM network. The advanced trail will discuss how FUNI differs from ATM DXI (Data Exchange Interface) and the benefits of FUNI.



Advanced Trail

On the advanced trail, we'll go into more detail about three of the applications covered on the basic trail: SNA over frame relay, Voice over Frame Relay and Frame Relay/ATM Interworking.

Frame Relay as an SNA SDLC Point-to-Point Line Replacement

Traditional SNA networks are based on leased lines which connect multiple controllers to the Front End Processor (FEP). These are typically low-speed analog lines, which represent a single point of failure between user and host, as in Figure 18.




Figure 18: Typical multidrop SNA network

Multidrop leased-line networks are a familiar evil, subjecting network managers to the complexities of a multitude of leased lines. Despite advances in networking technology, many organizations continue to base their SNA mission-critical applications on multidrop private lines.

SNA networks using point-to-point, non-switched lines can be migrated from SDLC to frame relay without any changes to the existing applications or hardware. Frequently, all that is needed is an upgrade to the communications software in the controllers.

Controllers that cannot be upgraded to support frame relay may be connected to a FRAD or router for frame relay connectivity. Frame relay uses the same hardware framing as SDLC, so all SDLC line interface couplers, modems and DSU/CSUs can be used with frame relay networks.

Another item to be considered when connecting controllers to a frame relay network is how the FRADs or routers connect to the remote and host sites on the SNA network. SNA has a form of maintenance communications called polling. An SNA device responsible for a sub-area network regularly polls each downstream controller for status, inquiring if it has data to send. At the remote site, each local controller responds.

Frame relay access devices can provide local polling or a decoupled polling capability in order to provide optimal networking conditions. A frame relay access device can poll its downstream devices or respond on their behalf. This process, called spoofing, eliminates polls on the network because only data is passed end-to-end. Extracting polls increases usable network bandwidth, directly impacting network performance.

Frame relay, as a virtual private line replacement, offers straightforward migration from the complexities of multidrop leased lines to a higher performance and more cost effective network.

As shown in Figure 19, migrating SNA networks to frame relay can occur without any change to FEP hardware or software. Users can realize significantly lower monthly WAN costs, which can pay for a frame relay migration within months.




Figure 19: SNA Network Migrated to Frame Relay

Upgrading to frame relay allows fully meshed topologies for redundancy and backup without managing a large number of dedicated lines. Adding and deleting virtual connections is done via network management and service subscription versus adding and deleting hardware. For high traffic volumes at a data center, frame relay supports access speeds up to 45 Mbps (e.g., T3/E3).

Frame relay supports "one-to-many" and "many-to-many" connections over a single line, where SDLC requires a multidrop line for a "one-to-many" configuration. SNA multipoint hardware configurations must be changed to point-to-point hardware configurations to use frame relay. The changes can be chosen to provide the best economic solution by combining the positioning of frame relay switches and frame relay terminal equipment. For example, frame relay switches may be used to provide the best use of point-to-point line tariffs, and FRADs may be used to provide frame relay to SDLC interworking, where SDLC multipoint lines are less expensive.

If users want additional cost savings, other migration paths are possible. For example, the FEP can be upgraded to allow direct connections to NCP (Network Control Program) from an FRR3.1-compatible FRAD, as shown in Figure 20. This eliminates a FRAD or router at the host, which reduces hardware costs and complexity.




Figure 20: NCP-direct FRR3.1 Network

Alternatively, the FEP may be upgraded to a token ring connection from an SDLC line, and a FRAD provides connectivity to the host, as shown in Figure 21, eliminating hardware on the host.




Figure 21: FRAD FRR3.1 Network

With the additional bandwidth available from frame relay, overall performance including session availability and user response time is improved as users migrate from multidrop lines which are typically 4.8/9.6 Kbps to frame relay connections of 56/64 Kbps to T1/E1.

How FRF.3.1 Provides Interoperability

Recognizing the ability of frame relay networks to carry multiple protocols, members of the Internet Engineering Task Force (IETF) developed a standardized method to encapsulate various protocols in frame relay. This multiprotocol encapsulation technique is called RFC 1490 after its IETF designation. ANSI and the Frame Relay Forum enhanced the multiprotocol encapsulation method to include support of the SNA protocols (FRF.3.1). FRF.3.1 was adopted and implemented by numerous vendors and is invaluable in multi-vendor environments. FRF.3.1 is used to carry SNA traffic across a frame relay network and may also be used to transport IP.

Protocol Encapsulation and Practical Implementation

Typically, SNA controllers, routers and FRADs encapsulate SNA as multiprotocol data as described in the Frame Relay Forum FRF.3.1 IA. SNA topologies supported across a frame relay network include:

- Intermediate Network Node (INN)
- Boundary Network Node (BNN)
- SNA Network Interconnect (SNI)
- Advanced Peer-to-Peer Networking (APPN), including High Performance Routing (HPR)
- Boundary Access Node (BAN)

FRF.3.1 specifies how to encapsulate SNA Subarea, SNA/APPN with and without HPR within the RFC 1490 multiprotocol framework. Because data is transparent to the frame relay network, it allows multiple distinct protocols to be multiplexed across a single frame relay interface. Frame relay network access nodes are responsible for converting the user data into an appropriate FRF.3.1 format for SNA and LAN traffic.

There are other alternatives to FRF.3.1 for transporting SNA over frame relay. One method uses routers to encapsulate SNA data within TCP/IP using a standard such as Data Link Switching

(DLSw) for link layer transport. The transport method a user selects depends on the application involved and the type of network equipment used.

Traffic Management Considerations

The mission-critical nature of SNA applications requires prioritization and bandwidth allocation mechanisms to avoid poor response times and SNA session failures caused by large bursts of other data traffic. One solution is to assign a higher priority to SNA data than LAN IP/IPX data if both are multiplexed over the same virtual connection. Another alternative is to send the data streams over two separate virtual connections and use the frame relay CIR mechanism to allocate bandwidth dynamically to each virtual connection.

Bandwidth allocation by percentage of CIR is a feature supported by some FRADs and routers, and it may not require separate PVCs. A smaller amount of bandwidth may be allocated (e.g., 20 percent) to the LAN traffic connection, giving the SNA traffic more frequent transmission opportunity. Further, it is recommended that both the frame relay service provider and the frame relay equipment support explicit congestion management indicators such as FECN/BECN and Discard Eligibility (DE). If these mechanisms are supported, SNA traffic flow is adjusted properly and packet discards are minimized. As with other applications carried over frame relay, congestion management plays an important role in supporting SNA applications.

Please visit the Frame Relay Forum's Web site (www.frforum.com) for a white paper on SNA over Frame Relay. This paper goes into more detail on the many options available for SNA over Frame Relay.

Voice over Frame Relay (VoFR)

Unlike most data which can tolerate delay, voice must be handled in near real time. This means that transmission and network delays must be kept small enough to remain imperceptible to the user. Until recently, packetized voice transmission was unattainable due to the voice bandwidth requirements and transmission delays associated with packet based networks.

Human speech is burdened with a tremendous amount of redundant information that is necessary for communications to occur in the natural environment, but which is not needed for a

conversation to occur. Analysis of a representative voice sample shows that only 22 percent of a typical dialog contains essential speech components that must be transmitted for complete voice clarity (see Figure 22). The balance is made up of pauses, background noise, and repetitive patterns.

Packetized voice is possible and low-bit rates are attained by analyzing and processing only the essential components of the voice sample, rather than attempting to digitize the entire voice




Figure 22: Normal speech components

sample with all its associated pauses and repetitive patterns. Current speech processing technology takes the voice digitizing process several steps further than conventional encoding methods.

VoFR Trade-offs

There are potential trade-offs when implementing VoFR. These include:

- loss of the quality commonly associated with toll traffic due to VoFR's use of voice compression
- loss of management and administrative benefits associated with carrier voice services (i.e., the loss of consolidated voice billing and invoice itemization, end user charge back capabilities, and other advanced features such as ID and accounting codes)
- lack of equipment interoperability between customer premise equipment vendors
- lack of standards defining the acceptable levels of quality for voice transport over a carrier's frame relay network

These trade-offs do not necessarily negate the value and promise of VoFR. Significant advances in digital signal processors and compression algorithms often provide voice at a level approaching

toll quality, for a fraction of the cost of public service. VoFR vendors continue to add advanced capabilities in management and administration capabilities. In addition, future industry work will also seek to develop standards which define acceptable levels of quality and performance metrics for voice transport through carriers' frame relay networks.

Please visit the Frame Relay Forum's Web site (www.frforum.com) for a white paper on Voice over Frame Relay. This paper goes into more detail on how VoFR works and the mechanics of voice compression.

More on the IWF

As we discussed on the basic trail, support of end-to-end communications in a Frame Relay/ATM network requires performing technical functions to compensate for the differences between frame relay and ATM. These functions are provided by the IWF generally located on the switch at the boundaries of the frame relay and ATM services.

Primary responsibilities for services provided by the IWF include mapping various parameters or functions between frame relay and ATM networks. These include:

- Frames or cells are formatted and delimited as appropriate.
- Discard eligibility and cell loss priority are mapped.
- Congestion indications are sent or received appropriately (frame relay's FECN is mapped to ATM's EFCI (Explicit Forward Congestion Indicator)).
- DLCI to VPI/VCI (Virtual Path Identifier/Virtual Circuit Identifier) mapping is performed.

The IWF also supports traffic management by converting ATM and frame relay traffic conformance parameters, supporting PVC management interworking via status indicators and providing upper layer user protocol encapsulation.

Frame Relay to ATM Network Interworking for PVCs may be thought of as encapsulating frame relay in ATM, since the ATM transport is transparent to the two frame relay users. The end user protocol suite remains intact. The IWF provides all mapping and encapsulation functions necessary to ensure that the service provided to the frame relay CPE is unchanged by the presence of an ATM transport. This is also sometimes referred to as frame relay transport over ATM. (See figure 23.)

Figure 24 and 25 illustrate the Interworking Function in a




Figure 23: Frame/ATM Network Interworking (Encapsulation)

Frame Relay/ATM Service Interworking environment. To enable communications between a frame relay desktop device and the ATM based application, the IWF performs all tasks associated with mapping the frame relay User-to-Network Interface (UNI) Q.922 core-based message in the frame relay network to the ATM UNI adaptation layer in the ATM network.

Figure 24 shows Transparent Protocol Support. For encapsulation methods other than FRE.3.1 (RFC 1490) and 1483 or when a single protocol is used, the IWF forwards the data unaltered.




Figure 24: Frame/ATM Service Interworking (Transparent)

Transparent mode can be used when the terminal equipment on one side of the IWF uses the encapsulation method of the terminal equipment on the other side. For example, an ATM CPE may use the RFC 1490 encapsulation method which is directly compatible with the frame relay equipment and no translation is required by the IWF.

Figure 25 shows Protocol Translation Mode. Encapsulation methods for carrying multiple upper layer user protocols (e.g. LAN-to-LAN) over a frame relay PVC and an ATM PVC conform to the standard FRF.3.1 and RFC 1483, respectively. The IWF performs mapping between the two encapsulation methods. Translation Mode supports the interworking of routed and/or bridged protocols (e.g., ARP translation). For more details, please refer to the Frame Relay and Frame-Based ATM white paper on the Frame Relay Forum web site (www.frforum.com).




Figure 25: Frame/ATM Service Interworking (Translation)

How Does FUNI Work?

FUNI requires FUNI-compatible software in the user equipment and a complementary frame-based interface, as well as FUNI software in the switch to which user equipment connects. Within the switch interface, the frames are segmented into cells and sent into the network. Cells coming from the network are reassembled into frames and sent to the user. Thus, hardware costs

of segmentation and reassembly are moved from the user equipment to the switch where it can be shared across a large number of users.

How Does FUNI Differ from ATM DXI?

Both the ATM DXI (Data Exchange Interface) and the ATM FUNI specifications translate frames of up to 2000 bytes into 53-byte ATM cells, but they differ as to where the translation takes place. The DXI standard requires a DXI-enhanced DSU to convert frames sent over an access line into cells and DXI software in the user equipment as well.

In contrast, the FUNI specification allows frames to be sent directly to the ATM switch where they are divided into cells, an approach which reduces processing and memory overhead in the remote server or workstation and makes more efficient use of the access line bandwidth. Figure 26 illustrates the two approaches.




Figure 26: ATM DXI and ATM FUNI

The major benefit of ATM FUNI is that it is "ATM ready." Although limited to VBR (variable bit rate) services, it uses the same schemes as cell-based ATM UNI in the following areas:

- upper layer multiprotocol encapsulation and address resolution
- traffic parameters
- ILMI (Interim Length Management Interface)

- OAM (Operations, Administration and Maintenance) cells (future requirement)
- ATM SVC signaling (future requirement)

By contrast, frame relay requires Frame Relay/ATM

Interworking functionality to achieve the same interoperability.

The current FUNI specifications address T1/E1 and Fractional T1/E1 (256 VCs per interface), whereas DXI supports full T1/E1, but not fractional T1/E1. Like frame relay and DXI, FUNI support for CPE routers consists of a software option.

Comparing Frame Relay/ATM Interworking with FUNI and ATM DXI

The following table helps to compare frame relay/ATM interworking with frame-based ATM.

	Frame Relay ATM Service Interworking	FUNI	ATM DXI
Access Transport	Frame based	Frame based	Cells
Software Requirements	IWF in frame relay or cell relay switch	FUNI software in user device and ATM network switch	DXI software in user device and DSU
Hardware Requirements	None	None	May require enhanced DSU

Table 5: Comparison of Frame Relay/ATM Interworking with FUNI and ATM DXI

When should you consider FUNI over Frame Relay-ATM service interworking? Clearly, one technology is not superior to the other. Rather, it amounts to selecting the solution which best addresses the network requirements, current topology and future network needs.

Since frame relay dominates the wide area architecture and remote site connectivity, Frame Relay/ATM Service Interworking is usually the most logical solution for most applications today. As the deployment of ATM approaches 50 percent of the wide area infrastructure, the case for deploying FUNI-based access becomes more compelling.



Shortcut

This chapter discussed four popular applications for frame relay: meshed LANs over frame relay, SNA over frame relay, voice over frame relay (VoFR) and Frame Relay/ATM Interworking.

- Frame relay enables networks to improve performance and provide cost reductions. These savings are available because frame relay meets the response time, availability and management requirements of business applications.
- Frame relay enables peer-to-peer, meshed LAN internetworking without the expense of a fully meshed leased line networking.
- Frame relay enables mission-critical SNA networks to improve performance and reduce costs. These savings are available because frame relay meets the response time, availability and management requirements of mission-critical applications.
- With compatible frame relay network access devices, most branch office equipment can connect to frame relay without hardware or software changes. With configuration changes or upgrades, additional savings are possible.
- The hidden costs of operation and maintenance are reduced by allowing network management staff to use the tools they are familiar with while providing a migration path to enterprise network management using SNMP.
- Voice over frame relay (VoFR) technology consolidates voice and voice-band data (e.g., fax and analog modems) with data services over the frame relay network. It has the potential to provide end users with greater efficiencies in the use of access bandwidth and cost-effective voice traffic transport for intra-company communications.
- Wide area network savings are possible because of the interworking of LAN, ATM, SNA and other legacy protocol, and voice traffic over frame-relay using industry standards. This may reduce branch office CIRs and port access speeds while always lowering CIRs and port access speeds at the data center. This improves user response times and reduces WAN costs.
- Frame Relay/ATM Network Interworking allows frame relay end-user or networking devices such as FRADs or routers to communicate with each other via an ATM network. Thus,

frame relay devices interact as if they are using frame relay for the entire connection without knowing that an ATM network is in the middle.

- Frame Relay/ATM Service Interworking enables communication between ATM and frame relay network or end user devices. Enabling existing devices to access new ATM based applications allows low cost access to high speed networks while protecting the investment in existing equipment.

CHAPTER 6

PLANNING YOUR FRAME RELAY NETWORK



Base Camp

In this chapter, we will discuss the steps you need to take and the questions you should consider if you're planning a new frame relay network or planning to deploy new applications over an existing frame relay network.

Basic Trail: The basic trail will discuss four steps and several considerations to help you plan your frame relay network.

Advanced Trail: The advanced trail presents more detailed information to consider in planning a frame relay network.



Basic Trail

It doesn't matter where in the world you're located or how large your telecommunications network is – you're probably trying to do more with less and leverage your telecommunications devices and services to extend your reach.

Frame relay offers compelling advantages over today's leased line networks, including flexibility, reduced WAN costs and scalability. Migrating your leased-line network, however, requires careful planning to assess network requirements because WAN savings at the cost of network performance and reliability is not a savings at all.

Let's look at four steps and several considerations to help you plan your frame relay network.

Assess Your Network Requirements

Before you migrate your leased-line applications to a frame relay network or add new applications to your existing frame relay network, consider these questions:

1. What is the average and peak bandwidth required by your target applications?
2. What is the maximum network latency your applications can tolerate before having an impact on users?
3. What are your objectives for application and network availability?

Addressing these questions can help you assess the proper access trunk speed, Committed Information Rate (CIR), Excess Information Rate (EIR) or burst capabilities, and the optimal Service Level Agreement (SLA) from your carrier or service provider.

There are many service level management products available to enable your network to gather this type of information. These systems help you to establish your current network baselines, and they also report on ongoing network performance which can be compared against your SLAs.

Assess the Impact on Your Management Procedures

Because you are "outsourcing" a large part of your network, your network management procedures will change. Understanding these changes and how they affect application availability is critical to the success of your frame relay network. Consider these questions:

1. How will problem identification, tracking and resolution procedures change?
2. What are the responsibilities of your organization and your carrier or service provider?
3. Do you or your service provider have tools to isolate and diagnose frame relay related problems?

Examine Your Service Level Agreement

The Service Level Agreement (SLA) between you and your service provider states network performance and availability commitments. These are some guidelines:

1. Maximum network transit latency (delay): influences application response times.
2. Network availability: defining measurement of service reliability.
3. Mean time to restoral: how fast service is restored after an outage.
4. Measurement intervals: how often the service provider measures these metrics.
5. Reporting: in what form the SLA metrics are reported and how often.
6. Data delivery rate (throughput): what percentage of your data actually is delivered at the destination side.

Conduct Ongoing Capacity and Performance Planning (Service Level Management)

Launching your frame relay network is only the beginning. Changes in the organization, applications and user population necessitate a continual assessment of your networking needs. Plan ahead for network changes. Consider how these factors can impact your network:

1. Enhancements to existing applications
2. Deployment of new applications
3. Increase in user populations
4. Acquisitions or reorganization
5. Service provider and switch loading factors



Advanced Trail

Designing a Frame Relay Network

When designing a frame relay network, you should ask yourself:

Is Frame Relay Service Right for this Network?

Here are some general rules that help identify applications that are suited for frame relay service.

- **Connecting Multiple Sites:** frame relay service will most likely be beneficial when multiple sites must be connected, not just a pair of sites.
- **High Speed:** if a network is using X.25 or a large number of analog private lines and is approaching the limits of existing bandwidth, frame relay may prove to be a cost-effective way to gain speed and efficiency.
- **Multi-Vendor, Multi-Protocol Environment:** If the network has a multi-vendor, multi-protocol environment, frame relay service may be a good choice because of its network transparency.
- **A Goal to Reduce Networking Costs:** If a network has been over-engineered to meet connectivity requirements, frame relay may offer a cost-effective solution.
- **Interactive or Bursty Traffic:** frame relay is a better choice when the traffic pattern between sites is interactive or bursty.
- **Widely Separated Locations:** frame relay is a better choice to link fairly widely separated locations, because its pricing structure is usually insensitive to distance; that is, it does not have the “mileage rates” usually attached to the private line tariffs.

If the network passes the preliminary qualification audit, the next step is to diagram the proposed frame relay network. This should include your diagramming current configuration, labeling the locations, listing the CPE, and noting the WAN connections. This will give you a better sense of the benefits of frame relay.

Specific Network Design

Frame relay network design consists of two steps:

- Diagram your existing network. For example, is it hubbed or meshed? What are the speeds of existing connections?
- Identify traffic patterns and flow characteristics. This will help determine the bandwidth and logical port connectivity requirements.

When you have completed a first draft frame relay network design, consult with your technical support team to ensure that the diagram and applications are well matched. Keep in mind that frame relay is an interface, not an architecture, and the applications to be run on interconnected LANs or via legacy protocols must conform to a distributable architecture.



Shortcut

The major challenge in migrating your leased-line network to a frame relay service is achieving the reliability, performance and network availability your users and applications require while maximizing your WAN networking budget. With careful planning, you can achieve these goals. Four steps are helpful in planning a frame relay network:

- Assess your network requirements
- Assess the impact on your management procedures
- Examine your Service Level Agreement
- Conduct ongoing capacity and performance planning

To help identify applications that are and are not well suited for frame relay service are, examine your network for these characteristics:

- Connecting multiple sites
- High speed
- Multi-vendor, multi-protocol environment
- A goal to reduce networking costs
- Interactive or bursty traffic
- Widely separated locations

FRAME RELAY GLOSSARY

Access Line A communications line (e.g. circuit) interconnecting a frame-relay-compatible device (DTE) to a frame-relay switch (DCE). See also Trunk Line.

Access Rate (AR) The data rate of the user access channel. The speed of the access channel determines how rapidly (maximum rate) the end user can inject data into a frame relay network.

American National Standards Institute (ANSI) Devises and proposes recommendations for international communications standards. See also Comite Consultatif International Telegraphique et Telephonique (CCITT) and International Telecommunications Union-the Telecommunications Services Sector (ITU-T).

Asynchronous Transfer Mode (ATM) A high-bandwidth, low-delay, connection-oriented packet-like switching and multiplexing technique. Usable capacity is segmented into 53-byte fixed-size cells, consisting of header and information fields, allocated to services on demand. Also referred to as cell relay.

ATM Forum An industry organization which focuses on speeding the development, standardization and deployment of Asynchronous Transfer Mode (ATM).

Backward Explicit Congestion Notification (BECN) A bit set by a frame relay network to notify an interface device (DTE) that congestion avoidance procedures should be initiated by the sending device.

Bandwidth The range of frequencies, expressed in Kilobits per second, that can pass over a given data transmission channel within a frame relay network. The bandwidth determines the rate at which information can be sent through a channel - the greater the bandwidth, the more information that can be sent in a given amount of time.

Bridge A device that supports LAN-to-LAN communications. Bridges may be equipped to provide frame relay support to the LAN devices they serve. A frame-relay-capable bridge encapsulates LAN frames in frame relay frames and feeds those frame relay frames to a frame relay switch for transmission across the network. A frame-relay-capable bridge also receives frame relay frames from the network, strips the frame relay frame off each LAN frame, and passes the LAN frame on to the end device. Bridges are generally used to connect local area network (LAN) segments to other LAN segments or to a wide area network (WAN). They route traffic on the Level 2 LAN protocol (e.g., the Media Access Control address), which occupies the lower sub layer of the LAN OSI data link layer. See also Router.

Burstiness In the context of a frame relay network, data that uses bandwidth only sporadically; that is, information that does not use the total bandwidth of a circuit 100 percent of the time. During pauses, channels are idle; and no traffic flows across them in either direction. Interactive and LAN-to-LAN data is bursty in nature, because it is sent intermittently, and in between data transmissions the channel experiences idle time waiting for the DTEs to respond to the transmitted data user's input of waiting for the user to send more data.

Channel Generically refers to the user access channel across which frame relay data travels. Within a given T1 or E1 physical line, a channel can be one of the following, depending on how the line is configured.

Unchannelized:

The entire T1/E1 line is considered a channel, where:

- The T1 line operates at speeds of 1.536 Mbps and is a single channel consisting of 24 T1 time slots.
- The E1 line operates at speeds of 1.984 Mbps and is a single channel consisting of 20 E1 time slots.

Channelized:

The channel is any one of N time slots within a given line, where:

- The T1 line consists of any one or more channels. Each channel is any one of 24 time slots. The T1 line operates at speeds in multiples of 56/64 Kbps to 1.536 Mbps, with aggregate speed not exceeding 1.536 Mbps.

- The E1 line consists of one or more channels. Each channel is any one of 31 time slots. The E1 line operates at speeds in multiples of 64 Kbps to 1.984 Mbps, with aggregate speed not exceeding 1.984 Mbps.

Fractional:

The T1/E1 channel is one of the following groupings of consecutively or nonconsecutively assigned time slots:

- N T/1 time slots (NX56/64Kbps where N = 1 to 23 T1 time slots per FT1 channel).
- + N E1 time slots (NX64Kbps, where N = 1 to 30 time slots per E1 channel).

Channel Service Unit (CSU)

An ancillary device needed to adapt the V35 interface on a frame relay DTE to the T1 (or E1) interface on a frame relay switch.

The T1 (or E1) signal format on the frame relay switch is not compatible with the V35 interface on the DTE: therefore, a CSU or similar device, placed between the DTE and the frame relay switch, is needed to perform the required conversion.

Committed Burst Size (Bc) The maximum amount of data (in bits) that the network agrees to transfer, under normal conditions, during a time interval T_c . See also Excess Burst Size (Be).

Comite Consultatif International Telegraphique et Telephonique (CCITT) International Consultative Committee for Telegraphy and Telephony, a standards organization that devises and proposes recommendations for international communications. The CCITT is now known as the ITU-T, the International Telecommunications Union-the Telecommunications Services Sector. See also American National Standards Institute (ANSI) and the International Telecommunications Union (ITU-T).

Committed Information Rate (CIR) The committed rate (in bits per second) at which the ingress access interface trunk interfaces, and egress access interface of a frame relay network transfer information to the destination frame relay end system under normal conditions. The rate is averaged over a minimum time interval T_c .

Committed Rate Measurement Interval (Tc) The time interval during which the user can 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 the subscription parameters of CIR and Bc) as $Tc = Bc/CIR$. Tc is not a periodic time interval. Instead, it is used only to measure incoming data, during which it acts like a sliding window. Incoming data triggers the Tc interval, which continues until it completes its committed duration. See also Committed Information Rate (CIR) and committed Burst Size (Bc).

Cyclic Redundancy Check (CRC) A computational means to ensure the accuracy of frames transmitted between devices in a frame relay network. The mathematical function is computed, before the frame is transmitted, at the originating device. Its numerical value is computed based on the content of the frame. This value is compared with a recomputed value of the function at the destination device. See also Frame Check Sequence (FCS).

Data Communications Equipment (DCE) Term defined by both frame relay and X.25 committees, that applies to switching equipment and is distinguished from the devices that attach to the network (DTE). Also see DTE.

Data Link Connection Identifier (DLCI) A unique number assigned to a PVC end point in a frame relay network. Identifies a particular PVC endpoint within a user's access channel in a frame relay network and has local significance only to that port.

Discard Eligibility (DE) A bit indicating that a frame may be discarded in preference to other frames if congestion occurs to maintain the committed information rate. See also Excess burst Size (Be) and CIR.

Egress Frame relay frames leaving a frame relay network in the direction toward the destination device. Contrast with Ingress.

End Device The ultimate source or destination of data flowing through a frame relay network sometime referred to as a Data Terminal Equipment (DTE). As a source device, it sends data to an interface device for encapsulation in a frame relay frame. As a destination device, it receives de-encapsulated data (i.e., the frame relay frame is stripped off, leaving only the user's data) from the interface device. Also see DCE.

NOTE: An end device can be an application program or some operator-controlled device (e.g., workstation). In a LAN environment, the end device could be a file server or host.

Encapsulation A process by which an interface device places an end device's protocol-specific frames inside a frame relay frame. The network accepts only frames formatted specifically for frame relay; hence, interface devices acting as interfaces to a frame relay network must perform encapsulation. See also Interface device or Frame-Relay-Capable Interface Device.

Excess Burst Size (Be) The maximum amount of uncommitted data (in bits) in excess of Bc that a frame relay network can attempt to deliver during a time interval Tc. This data (Be) generally is delivered with a lower probability than Bc. The network treats Be data as discard eligible. See also Committed burst Size (Bc).

E1 Transmission rate of 2.048 Mbps on E1 communications lines. An E1 facility carries a 2.048 Mbps digital signal. See also T1 and channel.

File Server In the context of frame relay network supporting LAN-to-LAN communications, a device servicing a series of workstations within a given LAN.

Forward Explicit Congestion Notification (FECN) A bit set by a frame relay network to notify an interface device (DTE) that congestion avoidance procedures should be initiated by the receiving device. See also BECN.

Frame Check Sequence (FCS) The standard 16-bit cyclic redundancy check used for HDLC and frame relay frames. The FCS detects bit errors occurring in the bits of the frame between the opening flag and the FCS, and is only effective in detecting errors in frames no larger than 4096 octets. See also Cyclic Redundancy Check (CRC).

Frame Relay Access Device (FRAD) A device that is responsible for framing data with header and trailer information (control information) prior to presentation of the frame to the frame relay switch. On the receiving end, the FRAD strips away the frame relay control information so that the target device is presented with the data in its original form. A FRAD may be a standalone device or it may be embedded in a router, switch, multiplexer or similar device.

Frame-Relay-Capable Interface Device A communications device that performs encapsulation or a device with an integral FRAD. Frame-relay-capable routers and bridges are examples of interface devices used to interface the customer's equipment to a frame relay network. See also Interface Device and Encapsulation.

Frame Relay Forum Worldwide organization of frame relay equipment vendors, service providers, end users and consultants working to speed the development and deployment of frame relay. Web site address: www.frforum.com.

Frame Relay Frame A variable-length unit of data, in frame-relay format that is transmitted through a frame relay network as pure data. Contrast with Packet. See also Q.922A.

Frame Relay Network A telecommunications network based on frame relay technology. Data is multiplexed. Contrast with Packet-Switching Network.

High Level Data Link control (HDLC) A generic link-level communications protocol developed by the International Organization for Standardization (ISO). HDLC manages synchronous, code-transparent, serial information transfer over a link connection. See also Synchronous Data Link Control (SDLC).

Hop A single trunk line between two switches in a frame relay network. An established PVC consists of a certain number of hops, spanning the distance from the ingress access interface to the egress access interface within the network.

Host Computer A communications device that enables users to run applications programs to perform such functions as text editing, program execution, access to data bases, etc.

Ingress Frame relay frames from an access device toward the frame relay network. Contrast with Egress.

Interface Device Provides the interface between the end device(s) and a frame relay network by encapsulating the user's native protocol in frame relay frames and sending the frames across the frame relay backbone. See also Encapsulation and Frame-Relay-Capable Interface Device.

International Telecommunications Union-the Telecommunications Services Sector (ITU-T) Formerly known as the Comité Consultatif International Télégraphique et Téléphonique (CCITT), the ITU-T is a standards organization that devises and proposes recommendations for international communications. See also Comité Consultatif International Télégraphique et Téléphonique (CCITT).

Latency The time it takes for information to get through a network, sometimes referred to as delay.

Link Access Procedure Balanced (LAPB) The balanced-mode, enhanced, version of HDLC. Used in X.25 packet-switching networks. Contrast with LAPD.

Link Access Procedure on the D-channel (LAPD) A protocol that operates at the data link layer (layer 2) of the OSI architecture. LAPD is used to convey information between layer 3 entities across the frame relay network. The D-channel carries signaling information for circuit switching. Contrast with LAPB.

Local Area Network (LAN) A privately owned network that offers high-speed communications channels to connect information processing equipment in a limited geographic area.

LAN Protocols A range of LAN protocols supported by a frame relay network, including Transmission Control Protocol/Internet Protocol (TCP/IP), Apple Talk, Xerox Network System (XNS), Internetwork Packet Exchange (IPX), and Common Operating System used by DOS-based PCs.

LAN Segment In the context of a frame relay network supporting LAN-to-LAN communications, a LAN linked to another LAN by a bridge. Bridges enable two LANs to function like a single, large LAN by passing data from one LAN segment to another. To communicate with each other, the bridged LAN segments must use the same native protocol. See also Bridge.

Local Loop The physical wires that run from the subscriber's telephone set or PBX to the telephone company central office.

Open Systems Interconnection (OSI) Model The only internationally accepted framework of standards for communication between different systems made by different vendors. Developed by the International Standards Organization (ISO).

Packet A group of fixed-length binary digits, including the data and call control signals, that are transmitted through an X.25 packet-switching network as a composite whole. The data, call control signals, and possible error control information are arranged in a predetermined format. Packets do not always travel the same pathway but are arranged in proper sequence at the destination side before forwarding the complete message to an addressee. Contrast with Frame Relay Frame.

Packet-Switching Network A telecommunications network based on packet-switching technology, wherein a transmission channel is occupied only for the duration of the transmission of the packet. Contrast with Frame Relay Network. Typically refers to an X.25 packet network.

Parameter A numerical code that controls an aspect of terminal and/or network operation. Parameters control such aspects as packet size, data transmission speed, and timing options.

Permanent Virtual Circuit (PVC) A frame relay logical link, whose endpoints and class of service are defined by network management. Analogous to an X.25 permanent virtual circuit, a PVC consists of the originating frame relay network element address, originating data link control identifier, terminating frame relay network element address, and termination data link control identifier. Originating refers to the access interface from which the PVC is initiated. Terminating refers to the access interface at which the PVC stops. Many data network customers require a PVC between two points. Data terminating equipment with a need for continuous communication use PVCs. See also Data Link Connection Identifier (DLCI).

Point-of-Presence (POP) Physical place where a long distance carrier interfaces with the network of the local exchange carrier (LEC).

Q.922 Annex A (Q.922A) The international draft standard that defines the structure of frame relay frames. Based on the Q.922A frame format developed by the CCITT. All frame relay frames entering a frame relay network automatically conform to this structure. Contrast with Link Access Procedure Balanced (LAPB).

Q.922A Frame A variable-length unit of data, formatted in frame-relay (Q.922A) format, that is transmitted through a frame relay network as pure data (i.e., it contains no flow control information). Contrast with Packet. See also Frame Relay Frame.

Router A device that supports LAN-to-LAN communications by connecting multiple LAN segments to each other or to a WAN. Routers route traffic on the Level 3 LAN protocol (e.g., the Internet Protocol (IP) address). Routers may be equipped to provide frame relay support to the LAN devices they serve. A frame-relay-capable router encapsulates LAN frames in frame relay frames and feeds the frame relay frames to a frame relay switch for transmission across the network. See also Bridge.

Statistical Multiplexing Interleaving the data input of two or more devices on a single channel or access line for transmission through a frame relay network. Interleaving of data is accomplished using the DLCI.

Synchronous Data Link Control (SDLC) A link-level communications protocol used in an International Business Machines (IBM) Systems Network Architecture (SNA) network that manages synchronous, code-transparent, serial information transfer over a link connection. SDLC is a subset of the more generic High-Level Data Link Control (HDLC) protocol developed by the International Organization for Standardization (ISO).

Switched Virtual Circuit (SVC) A virtual circuit connection established across a network on an as-needed basis and lasting only for the duration of the transfer.

T1 Transmission rate of 1.544 Mbps on T1 communications lines. A T1 facility carries a 1.544 Mbps digital signal. Also referred to as digital signal level 1 (DS-1). See also E1 and channel.

Time Division Multiplexing (TDM) A method of transmission which relies on providing bandwidth based on fixed time slots or channels, also called circuit switching. See also channel.

Trunk Line A communications line connecting two frame relay switches to each other.

Virtual Circuit A communications link – voice or data – that appears to the user to be a dedicated point-to-point circuit. A virtual circuit is referred to as a logical, rather than a physical path, for a call.

APPENDIX

RELEVANT DOCUMENTS

- IBM Frame Relay Guide (IBM GG24-4463-00)x
- Systems Network Architecture - Format and Protocol Reference Manual: Architectural Logic (IBM SC30-3112-2)
- System Network Architecture - Advanced Peer to Peer Networking: Architecture Reference (IBM SC30-3422-03)



Frame Relay Forum
39355 California Street, Suite 307
Fremont, CA 94538

