White Paper

A Discussion of Voice over Frame Relay

August 2000

The Frame Relay Forum
Market Development & Education Committee
and
Technical Committee

A Discussion of: Voice over Frame Relay

Special thanks to the following individuals who contributed to this document:
Todd Bahner, ADC Kentrox
Ed Bursk, FastComm Communications
John Casadonte, Nortel
Richard Cornetti, Hewlett Packard
Bernard da Costa, Bell Canada
Micaela Giuhat, Newbridge Networks
Larry Greenstein, NUERA
Ross Kocen, ACT Networks
Terence La Du, Bellcore
Christopher Nicoll, Visual Networks
Douglas O'Leary, Bell Atlantic
Michelle Olesiejuk, FRF Management
Chuck Reiss, IBM Corporation
John Scarmalis, Paradyne Networks Inc.
Martin Sendyk, Newbridge Networks
Dave Sinicrope, IBM Corporation

INTRODUCTION

Additional Possibilities and Benefits Associated With Frame Relay

The Frame Relay Forum, a non-profit organization of 300-plus member companies dedicated to promoting the acceptance and implementation of frame relay, along with over 10,000 worldwide users, is witnessing the evolution of frame relay from a single-application technology to one with a broad spectrum of uses. Network managers are constantly seeking new ways to make their company networks more efficient through the use of new and innovative services that are continually being introduced into the market. Often, they are faced with the need to connect remote offices to the corporate backbone to enable access to corporate
e-mail, local area networks, mainframe computers, and other corporate services. Frame Relay is frequently the technology of choice used to meet these needs.

Initially, frame relay gained acceptance as a means to provide end-users with a solution for LAN-to-LAN connections and other data connectivity requirements. Besides providing a flexible and efficient data transport mechanism, frame relay lowered the cost of bandwidth for tying together multi-protocol networks and devices. Over the past few years there has been a migration of legacy traffic such as bisync and SNA from low-speed leased lines onto frame relay. The integration of this so-called legacy traffic with today's LAN-to-LAN connectivity needs provided network managers with a more efficient, flexible, and very cost-effective network. More recently, non-traditional uses are beginning to emerge. Due to advances in areas such as digital signal processing, end-users are beginning to see viable methods being developed that incorporate non-data traffic such as voice and video over frame relay.

Voice over frame relay (VoFR) technology offers telecommunication and network managers the possibility of consolidating voice and voice-band data (i.e. fax and analog modems) with data services over frame relay. The Frame Relay Forum Technical Committee has developed an Implementation Agreement [FRF.11] in order to allow vendors to interconnect their VoFR-capable equipment. It is anticipated that this work will lay the groundwork for future deployment of VoFR capabilities in multi-vendor and public network environments. Prior to the development of this Implementation Agreement, many equipment vendors developed proprietary methods for implementing voice over frame relay, thus enabling end-users to successfully deploy voice over their frame relay networks.

Frame relay will continue to see explosive growth. The acceptance and use of ATM (Asynchronous Transfer Mode, a technology designed with the intent to transport voice, data, and video) will also increase. Both end-users and network service providers will increasingly find that frame relay and ATM do not only coexist, but are complementary; both frame relay and ATM access is offered to, and used by end-users. In addition, service providers have begun to migrate their frame-based networks to ATM-based backbones.. The continued use and increased acceptance of frame relay and ATM technologies will bring greater bandwidth and high performance networking to a wider variety of user applications.

In addition to giving the reader some insight into how VoFR works, this paper will provide an overview of a few of the potential applications of VoFR, some of the considerations faced by end-users, and an overview of the Voice over Frame Relay Implementation Agreement [FRF.11]. The intent of this paper is not to promote or dissuade frame relay users from incorporating voice into their frame relay networks. Instead, it is meant to provide a balanced perspective and information so that users may have more information to decide for themselves as to whether they may benefit from voice over frame relay technology.

VOICE OVER FRAME RELAY
Theory Of Operation

Over the years communication networks have become more reliable. Older, low speed, analog connections, which are often susceptible to network-induced errors, are being replaced with higher speed digital links offering relatively error-free performance. In addition, devices communicating between sites have become more intelligent, allowing them to more readily accommodate network delay, and recover from and re-transmit lost data. Unlike most data communications, which can tolerate delay, voice communications must be performed 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 not possible due to the requirements of voice bandwidth, and the transmission delays associated with packet-based networks.

Packetized voice is now possible; low bit rates are attained by analyzing and processing only the essential components of a voice sample, rather than attempting to digitize the entire voice sample (with all the associates pauses and repetitive patterns). Current speech processing technology takes the voice digitizing process several steps further than conventional encoding methods.

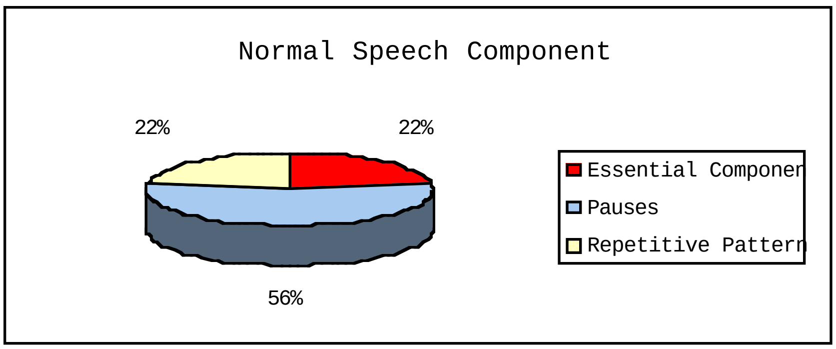

Figure 1. Only 22% of normal speech needs to be sent for high-quality voice communications.

Human speech is burdened with a tremendous amount of redundant information that is necessary for communications to occur in a natural environment, but which is not needed for a conversation to occur over a communications network. Analysis of a representative voice sample shows that only 22 percent of a typical conversation consists of essential speech components that need to be transmitted for complete voice clarity (Figure 1). The balance is made up of pauses, background noise, and repetitive patterns.

Removal of Repetitive Speech Sounds
Repetitive sounds are inherent in human speech, and are caused by vibration of the vocal cords. These repetitive sounds (like the 's' in the word *snake* or long 'o' in the word *loan*, are easily compressible. While traveling through the natural environment, perhaps only half of what is spoken will reach the listener's ear. However, in a typical communications network, all of the speech content is transmitted. Transmission of these identical sounds is not necessary; their removal can increase bandwidth efficiency.

Removal of Pauses (Silence Suppression)
A person speaking does not provide a continuous stream of information (regardless of how fast they speak). Pauses between words and sentences, and the gaps that occur at the end of one person talking but before the other begins, can also be removed. The pauses may be represented in compressed form and can be re-created at the destination side of the call in order to maintain the natural quality of the spoken communication. The suppression and removal of silent periods can also significantly improve bandwidth utilization.

Voice Frame Formation
The removal of silent periods and redundant information through advanced techniques enables voice to be efficiently “compressed”. After the removal of repetitive patterns and silent periods, the remaining speech information may then be digitized and placed into voice packets suitable for transmission over a frame relay network. These packets or frames (both terms are often used interchangeably) also tend to be smaller than the average data frame. The use of smaller packets helps to reduce transmission delay across a frame relay network. The concepts introduced above provide the basis for efficiently using the smallest amount of bandwidth possible for voice transmission over a frame relay network.

USING VOICE OVER FRAME RELAY

Potential End-User Applications

Telecommunication managers continue to explore alternatives for obtaining the most efficient use of their corporate network resources. Many network managers have migrated their point-to-point leased communication networks (built in the 1980’s with TDM (Time Division Multiplexers) equipment) to public and private frame relay networks. Since many of these point-to-point leased line networks carried both voice and data, these network managers are interested in meeting not only their data communications needs, but also their voice communications needs. VoFR (voice over frame relay) offers a potential alternative for carrying voice communications over a frame relay network in order to meet intra-company communications needs. Current users of frame relay may find that they have “excess” bandwidth available even with the tremendous expansion of applications and increase in data traffic. And, even when existing bandwidth is efficiently utilized, some network managers might find that the incremental cost for the additional frame relay network bandwidth needed for voice transport is more cost-effective than some of the standard voice services offered by local and long distance carriers. In other cases, some end-users might find that VoFR is a viable option to be used in place of Off-Premises Extension (OPX) and Private Line Auto Ringdown (PLAR) lines.

Of course, the motivation behind the interest in VoFR will vary. VoFR has the potential to provide end-users with greater efficiencies in the use of access bandwidth by functionally integrating voice, data, and fax over a single access link. In addition, VoFR has the potential to provide end-users with a cost-effective option for the transport needs of voice traffic between their company locations.

As an example, a network manager may choose to integrate a few voice channels and serial data over the frame relay connection between a branch office and corporate headquarters. By transmitting the voice traffic over the frame relay connection, which is already carrying data traffic (Figure 2), the user has the potential to obtain cost-effective intra-company calling and efficient use of network bandwidth.

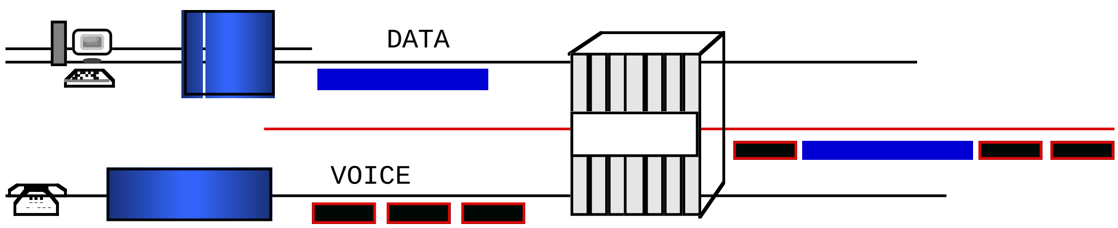

Figure 2: Voice and data being integrated by customer premises equipment.

The examples provided do not necessarily reflect all the potential possibilities of VoFR. An exploration of the full possibilities and potential for the implementation of VoFR is outside the scope of this paper. There are many reasons used, and possibilities explored by network managers in their attempt to justify more efficient, flexible, and cost-effective networking capabilities. VoFR represents one of many possible methods that enable users to increase the flexibility and efficiency of their company’s network resources.

There might, however, be potential trade-offs that the network manager may face when implementing VoFR. Some of the potential trade-offs may include some loss of the quality commonly associated with toll traffic due to VoFR's use of voice compression; the 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 Caller ID and accounting codes); and the lack of standards defining the acceptable levels of quality for voice transport over a carrier's frame relay network.

In addition, carriers offering public frame relay service cannot always guarantee the quality or performance of voice transported over their frame relay networks. In the absence of standards there are not any specifications which define the quality of a voice conversation (i.e. delay, tonal and pitch qualities) occurring over the carrier's frame relay service. Since the quality of VoFR is subjective, it is troublesome for the carrier to guarantee complete user satisfaction. The lack of specific voice over frame relay service guarantees and full carrier troubleshooting capabilities is a result of the fact that in today's environment the implementation of voice over frame relay occurs in equipment on the end-user's premises and outside of the carrier's frame relay network.

The potential trade-offs do not necessarily negate the value and promise of VoFR. Significant advances in digital signal processors and compression algorithms often provide voice approaching toll quality. VoFR vendors continue to add advanced capabilities in management and administration capabilities. Future industry work will seek to define standards which define acceptable levels of quality and performance metrics for voice transport through data networks.

Some end-users might not be concerned with the potential trade-offs noted; some may find the trade-offs unacceptable in particular situations; others may find that if trade-offs exist, they are outweighed by the potential for costs savings and efficiencies gained by integrating voice and data over frame relay. In the end, it will be up to the customer to decide.

VOICE OVER FRAME RELAY EQUIPMENT

Common Considerations Faced By Equipment Vendors

Vendors offering equipment capable of integrating voice and data traffic over frame relay must consider how they will address issues such as compression, echo cancellation, delay and delay variation, frame loss, and traffic prioritization. Each of these, and other considerations, can affect voice quality. While vendors offering voice over frame relay-capable equipment may have similar objectives regarding quality and performance, each vendor may choose to pursue these objectives through different hardware and software implementations. Common considerations and a few of the many potential methods used to provide voice over frame relay are presented below.

Voice Compression:

Compression of voice is a result of the removal of the silent periods and redundant information found in human speech. Voice compression is used to reduce the amount of information needed to recreate the voice at the destination end. Uncompressed digitized voice and fax require a large amount of bandwidth. This often makes it impractical to transmit these signals over low-speed access links The use of low bit rate voice compression algorithms can make it possible to provide high quality speech while using bandwidth efficiently.

Various algorithms are used to sample speech patterns and reduce the information sent - all while retaining the highest possible level of voice quality. A relatively simple ADPCM (Adaptive Delta Pulse Code Modulation) algorithm can reduce the speech data rate to half that of PCM (Pulse Code Modulation), an ITU standard for digital voice coding, which consumes 64 Kbps and is optimized for speech quality. PCM is the voice algorithm that is commonly used in telephone networks today. ADPCM may be used in place of PCM, while maintaining about the same voice quality. In addition to ADPCM, there are a number of standard low bit rate voice compression algorithms (e.g., ITU G.729) as well as proprietary algorithms implemented by various vendors, which provide more significant reductions (ie; to 10% or less than that of PCM) in the amount of information required to recreate speech at the receiver.

Other voice compression algorithms model speech more efficiently (i.e., with fewer bits) by using advanced predictive techniques. These algorithms further reduce the bandwidth required to maintain good voice quality. Implementation of these advanced compression techniques, and meeting their processing demands, is made possible by the use of Digital Signal Processors (DSPs). A DSP is a microprocessor that is designed specifically to process digitized signals such as those found in voice and video applications. In the last ten years significant advances in the design of DSPs have occurred. This development has allowed manufacturers to bring to market even higher quality digitization algorithms that consume very little bandwidth.

The general function of these strategies is to scrutinize the speech signal more carefully in order to eliminate the redundancies in the signal more completely, and to use the available bits to code the non-redundant parts of the signal in an efficient manner. As the available bit rate is reduced from 64 Kbps to 32, 16, 8, and 4 Kbps or below, the strategies for redundancy removal and bit allocation need to be ever more sophisticated. Low cost general purpose DSP processors and other advanced compression algorithms allow the possibility of accomplishing voice compression within VoFR-capable devices at lower and lower bit rates.

Echo Cancellation:
Echo is a phenomenon found in voice networks. Echo occurs when the transmitted voice is reflected back to the point from which it was transmitted. In voice networks, echo cancellation devices are used within a carrier's network when the propagation delay increases to the point where echo results. The longer the distance, the more the delay, and the more likely that echo will result. Voice transmitted over a frame relay network will also face propagation delays. As the end-to-end delay increases, the echo will become noticeable to the end-user if it is not canceled. Since carriers do not use echo cancellation equipment in their frame relay networks, it is up to the CPE vendor to address echo cancellation in the CPE.

Delay and Delay Variation:
The bursty nature and variable frame sizes of frame relay may result in variable delays between consecutive packets. The variation in the time difference between each arriving packet is called "jitter".

Jitter can impede the ability of the receiving end CPE to smoothly regenerate voice. Since voice is inherently a continuous waveform, a large gap between the regenerated voice packets will result in distorted sound. Equipment vendors can contribute to the mitigation of jitter across the network by employing fragmentation of data packets in order to transmit uniform packet sizes into the network. FRF.12 provides an interoperable technique for fragmenting data prior to its transmission into a FR network."

To avoid dropping speech samples, data can be buffered sufficiently at the speech decoder to account for the worst case delay jitter through the network. Equipment vendors look to incorporate this capability within their equipment.

Frame Loss:
Compressed voice can usually withstand infrequent packet loss better than data can. If a voice packet is lost, the user will most likely not notice. If excessive frame loss occurs, it is equally unacceptable for VoFR and for data traffic.

Traffic Integration - Fax and Modem Support
Vendors implementing VoFR technology appear to be mimicking switched public voice services. Since VoFR supports fax and data modem services as well, end-users who have high fax traffic volumes between branches and headquarters will find this ability beneficial.

Voice band fax signals are demodulated at the locally connected CPE equipment and transmitted over the network as digital data in a standard packet format. In effect, the local Voice FRAD tricks the fax machine into thinking it is connected to a remote fax machine across an analog network. However, it is difficult to reliably compress fax and data modem signals to achieve the low bandwidth utilization often necessary for the most efficient integration over frame relay. Some vendors have implemented schemes where voice is compressed to a low bit rate, but upon detection of a fax tone, the bandwidth is reallocated to a higher bit rate to allow for faster fax transmission.

Prioritization:
Voice, fax and some data types are delay-sensitive. This means that if the end-to-end delay or the delay variation exceeds a specified limit, the service level will get degraded. To minimize the potential for service degradation, vendors can employ a variety of mechanisms and techniques.

To minimize voice traffic delay, a prioritization mechanism that provides service to the delay-sensitive traffic can first be employed. Vendors offering equipment capable of integrating voice and data over frame relay may choose to use a variety of proprietary mechanisms to ensure a balance between voice and data transmission needs. Although they may differ, the concept remains essentially the same. For example, each input traffic type may be configured into one of several priority queues. Voice and fax traffic can be placed in the highest-priority queue, for expeditious delivery to the network. Lower-priority data traffic can be buffered until the higher-priority voice and fax packets are sent (Figure 3 below).

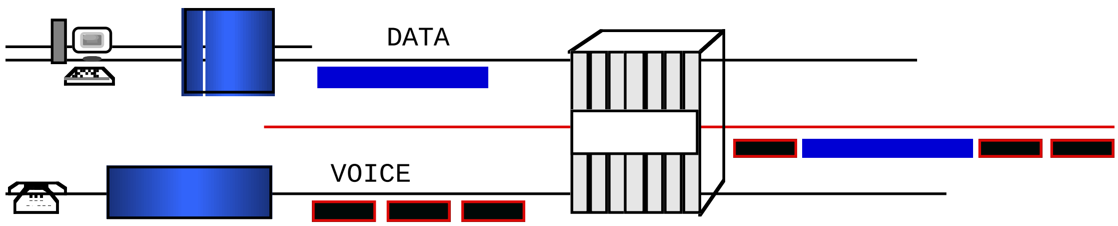

Figure 3. Prioritization places delay-sensitive traffic, such as voice, ahead of lower priority data transmissions.

Fragmentation
Fragmentation is used to break up large blocks of data into smaller, less delay-creating frames. This is another means used to ensure the highest level of voice quality possible. Fragmentation attempts to ensure an even flow of voice frames into the network, thus minimizing delay jitter across circuits that carry both packet-voice and data.

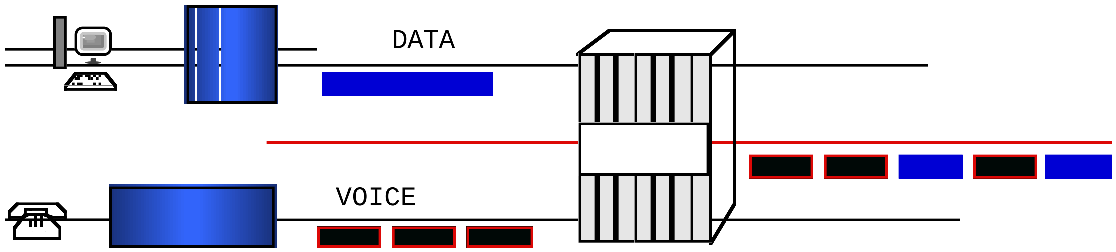

Figure 4. Fragmentation ensures that high priority traffic such as voice does not have to wait to be sent. Long data packets can be interrupted to send a voice packet.

Fragmentation often involves all of the data in the network, to retain consistent voice quality. This is because even if the voice information is fragmented, delay will still occur if a voice frame is held up in the "middle" of the network, behind a large data frame. This fragmentation of data packets (as shown in Figure 4) ensures that voice and fax packets are not unacceptably delayed behind large data packets. Additionally, fragmentation reduces jitter because voice packets can be sent and received more regularly. Fragmentation, especially when used with prioritization techniques, is used to ensure a consistent flow of voice information. The objective of this and other techniques is to enable VoFR technology to provide service approaching toll quality voice. The Frame Relay Forum recommends the use of the Fragmentation Implementation Agreement [FRF.12] when employing fragmentation for VoFR.

Digital Speech Interpolation:

Digital speech interpolation addresses silence suppression. The nature of speech communication entails pauses between words and sentences. Advanced voice compression algorithms, which identify and remove these redundant patterns, effectively reduce the amount of speech information to be transmitted. DSI uses advanced voice processing techniques to detect silence periods and suppress transmission of this information. By taking advantage of this technique, bandwidth consumption may be reduced.

Multiplexing Techniques

Some equipment vendors offering voice FRADs (Frame Relay Access Devices) use different bandwidth optimization multiplexing techniques such as Logical Link Multiplexing and Subchannel Multiplexing. Logical Link Multiplexing allows voice and data frames to share the same PVC (Permanent Virtual Circuit). This can provide savings on carrier PVC charges and increase the utilization of the PVC.

Subchannel Multiplexing is a technique used to combine multiple voice conversations within the same frame. By allowing multiple voice payloads to be sent in a single frame, packet overhead is reduced. This may offer increased performance on low speed links. This technique can allow slow speed connections to transport small voice packets efficiently across the frame relay network.

Other Considerations:

In addition to providing basic services such as encapsulation of data traffic for transport over the frame relay network, voice capable FRADs may sometimes provide connectivity between PBXs and other voice equipment. As a result the voice FRAD would have to manage different traffic types and accommodate their different needs.

When voice is carried over a Frame Relay network that employs ATM in the backbone, there is no impact due to the use of the ATM backbone since ATM functions purely as a transport medium.

VOICE OVER FRAME RELAY IMPLEMENTATION AGREEMENT (FRF.11)

Overview

Frame Relay Forum Implementation Agreements provide an agreed upon basis for vendors and service providers to develop equipment and services that interoperate. In the case of VoFR, as with many emerging technologies, vendors are often able to develop and deploy capabilities before the various industry and user organizations achieve consensus on uniform standards and implementations.

FRF.11 provides an outline for an agreed-upon basis for VoFR so that companies may build equipment and offer services that will be capable of functionally interoperating with each other.

The IA addresses the following:

- Transport of compressed voice within the payload of a frame relay frame, via the support of a diverse set of voice compression algorithms such as CS-ACELP, LD CELP, MP-MLQ, PCM, etc.
- Effective utilization of low bit rate frame relay connections
- Multiplexing of up to 255 sub-channels on a single frame relay DLCI, such that a single DLCI may contain both voice and data payloads
- Support of multiple voice payloads on the same or different sub-channel(s) within a single frame

Reference Model

The reference model for VoFR is shown in Figure 5. Using the VoFR feature, it is possible for any type of VFRAD on the left-hand side of Figure 5 to exchange voice and signaling information with any type of VFRAD on the right-hand side of Figure 5.

Three types of device are shown in Figure 5. The top layer shows end-system devices similar to telephones or fax machines; the middle layer shows transparent multiplexing devices similar to channel banks; the bottom layer shows switching system devices similar to PBX's.

A VFRAD connects to a frame relay UNI via physical interfaces as defined in [FRF.1.1].

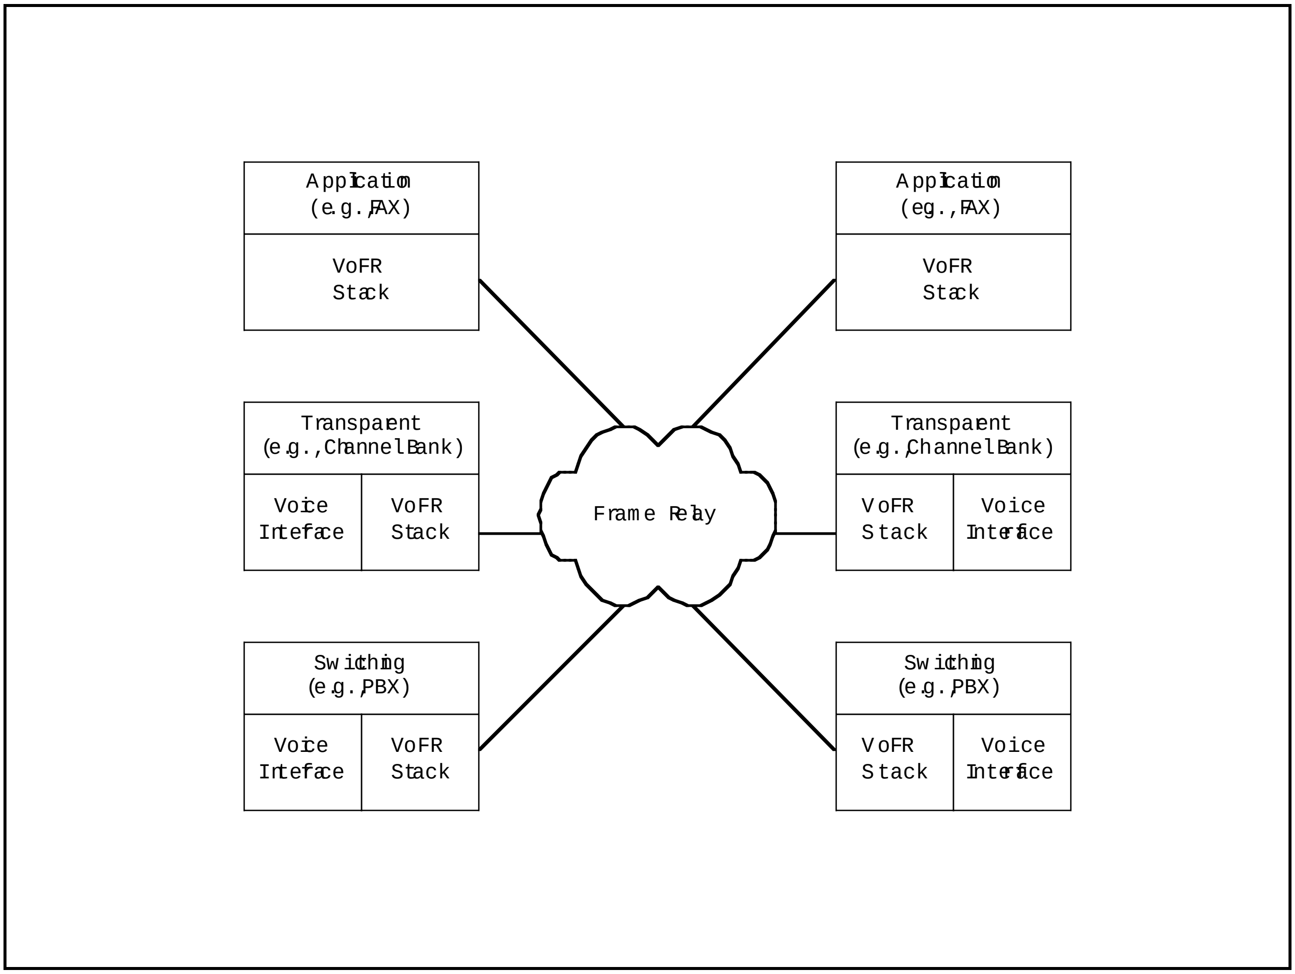

Figure 5. Voice over Frame Relay Network Reference Model

SUMMARY

Frame Relay is beginning to evolve from a single, data-only application technology to one with a broad spectrum of uses. Integration of voice and data over frame relay represents one of many promising areas of development that could not only benefit end users, but one that could continue to fuel the continued growth of frame relay services and applications.

Voice over frame relay (VoFR) technology consolidates voice and voice-band data (i.e. 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 to provide end-users with cost-effective transport of voice traffic for intra-company communications.