WHITE PAPER

Frame Relay or Internet VPNs

The Landscape for Next Generation Business Applications

Escalating network operating costs and the scarcity of trained management personnel are two main drivers behind the corporate outsourcing of private wide area networks (WANs) to public service providers. Frame Relay gained its success as the first outsourcing opportunity for enterprise customers, but now there's new opportunity in a competing WAN alternative: Internet/IP services promise to deliver a rich set of business and e-commerce solutions that significantly enhance or improve on Frame Relay's answer.

The Internet is, in many respects, becoming the world's business network. With flexibility, diversity, and reach beyond traditional network solutions, Internet-based Virtual Private Networks (VPNs) are right in step with the arriving generation of business-to-business, E-commerce, and remote access applications.

Technically the Internet has come a long way. Its performance now matches and even surpasses Frame Relay along many dimensions, in part because the rules of the game are changing. New, more explicit service level requirements are being defined by business customers who expect Quality of Service (QoS) to be the measure of WAN solutions. Along with QoS demand there is also the greater scope of network connectivity, flexibility, and performance required by new business applications.

Both Frame Relay and Internet services are equally at a new threshold - neither technology yet has the lead. But the Internet holds a very large edge with its potential for VPN services.

Background: Frame Relay and IP in Perspective

A technology assimilation is underway between Frame Relay and IP. While Frame Relay attempts to inherit the flexibility and ubiquity of IP's world, the Internet is moving to gain the predictability associated with Frame Relay networks. Both Frame Relay and IP technologies are evolving to meet more aggressive QoS requirements. 1999 is in many respects "Year 0" for both technologies as they press forward on next-generation QoS capabilities. As the two agendas progress they will have great potential as strong partners in building the New Public Network.

Frame Relay service developed in the early 1990's as a simpler and more cost-effective alternative to leased lines for interconnecting large wide area networks. Its success proved the case for cost-savings and reliability that corporate networks can gain through outsourced WAN services.

However, Frame Relay, as currently deployed, presents a relatively static network environment that has not achieved global proportion. Most of the existing services also assume low speed (T1 and below) access. As a result, Frame Relay service remains best suited for branch office-to-HQ connectivity - it has not been extended to the small office and remote access community.

By comparison, the Internet is rooted as a tool for broadly distributed communications among diverse communities. From the start, the Internet was defined to enable low cost, any-to-any access to a global base of users and information. More recently the Internet has come of age for corporate applications. ISPs have made significant investments to add bandwidth, expand their Points of Presence (POPs), and establish international positions. They have also defined aggressive Service Level Agreements (SLAs), with several ISPs making 100% network uptime and explicit round-trip latency commitments. These advances have been instrumental in preparing the Internet for the new wave of business applications.

While both Frame Relay and Internet VPN services will grow over the next few years, the rate of growth is expected to be dramatically higher on the Internet VPN services side. Initial VPN growth will come as corporations build new networks connecting their remote employees, partners, and customers. Over time, the existing base of corporate applications will require more bandwidth, greater connectivity, and increased cost reduction. At that point they are likely to migrate to a common Internet-based WAN service.

Existing and emerging carriers recognize these dynamics and are aggressively moving toward a high performance, QoS capable IP foundation in their New Public Networks. The enterprise should prepare for the ability to deploy key IP-enabled business applications over these public Internet VPN services.

The Business Applications Driving VPN Service

The most commonly identified applications for Internet VPNs include corporate Intranets and Extranets, remote access, and e-commerce.

- **Corporate Intranets** enable secure communications among enterprise business centers, small offices and mobile or remote users. Their primary objective is to increase productivity by enabling more rapid access to critical corporate information. The result can be shorter sales cycles, lower cost of sales, better efficiency in delivering customer service, and more accurate and timely financial reporting. While many corporate Intranets were initially built on top of existing network infrastructures, they are now being outsourced to service providers, adding the advantages of low capital cost, lower operating cost, and dramatically improved network reach.
- **Extranets** connect businesses with their trading partners, suppliers, and agents. An Extranet is typically comprised of business affiliates that do not share any common network infrastructure. There can be significant productivity and profitability gains as partner Extranets are used to automate or refine business processes for faster time to market, lower cost or better efficiency. Since they include multiple business entities, Extranets have very demanding requirements for flexible access, broad reach and low cost.
- **Remote Access** to many enterprise networks today is delivered via direct-dial to a corporate 800 number, which connects users to a bank of modems managed by the enterprise. This is a high cost solution that does not easily scale to a large number of remote users. Therefore, VPNs for small offices, mobile workers, and tele-

commuters have been targeted as an immediate priority. The Internet with its widely deployed dial access infrastructure is a natural for this application. Frame Relay networks, because they are designed for dedicated access, do not support this application.

- **E-commerce** applications and services must provide a secure environment for exchanging customer orders and payment information. E-commerce is a specialized type of Extranet that can reduce cost of sales through lower order processing and procurement overhead, as well as increase revenue via on-line access to an expanded customer base. E-commerce is expected to dominate Internet business service revenues in the very near future. The Internet is the only viable vehicle for building these customer centered and cross-enterprise services.

The common thread among all of these applications is their need for flexible and diverse connectivity. They stress the ability to reach any site using via either dial-up or dedicated access from anywhere in the world. Performance and security are critical, and these applications also typically require rapid deployment with potentially global distribution of service. Internet VPNs are often the only service that can satisfy all of these criteria.

Underlying Technologies Compared: Frame Relay and IP Capabilities

Cost alone is not the most important factor when comparing Internet and Frame Relay VPN services. The overall impact of VPN features lands at a more strategic level where network services can help to realize new business methods.

	Internet VPN	Frame Relay
Flexibility	Any-to-any networking across VPN or public access	Statically provisioned PVCs between sites
Ubiquity	Global access with well established peering architecture	Regional or national, with limited interconnectivity between service providers
Diversified access	Dial-up or dedicated using many access technologies	Primarily dedicated access with limited technology options
Remote access	Yes, everywhere	no
Site-to-site	High-speed via OC3 and faster access links	Bandwidth up to T3 available, but predominantly T1 or less
Business-to-business services	Yes, globally possible	No model exists
Ease of deployment	Global access to/from any site	Must be provisioned per-PVC
Low cost	Ideal for high speed connectivity and large meshed topologies	More economical for low speed traditional hub and spoke topologies
End-to-end SLAs	Currently available on backbone only	Yes, based on throughput

FRAME RELAY ADVANTAGES

Frame Relay packet switching is generally characterized as a Layer 2 connection-oriented technology. Connections are pre-configured in the form of Permanent Virtual Circuits (PVCs), which are fixed logical channels between sites.

The Frame Relay value proposition is simple. First, it offers lower cost than leased line services. Second, Frame Relay delivers performance guarantees in the form of Committed Information Rates (CIR), which specify the bandwidth characteristics of a PVC. In a Frame Relay service environment, each PVC has a cost that varies with its CIR. Increasing the number of PVCs or the CIR of a PVC will result in higher cost. As a result, the Frame Relay model is best suited for network topologies with limited connectivity and low speed requirements. Indeed, the greatest success for Frame Relay service has been with traditional hub and spoke branch office networks supporting predominantly IP-based LAN traffic.

Since Frame Relay operates at Layer 2, it is independent of the routing protocols that run on top of it. As a result, it has been a good solution for multiprotocol environments. At this time, however, most traffic on Frame Relay networks, including SNA traffic, is running the IP protocol.

Frame Relay providers today offer SLAs that include throughput guarantees and are a step towards network QoS. But these SLAs are complex to measure, and collecting information to validate an SLA is still a bit of an art. A new breed of DSU/CSU-like monitoring devices have been developed to augment Frame Relay managed services with SLA monitoring information. Detailed statistics are collected, typically at 5 minute intervals, to measure SLA compliance. However, even with these tools in place, SLAs in general have a long way to go to become more enforceable. The next step for Frame Relay services is to offer tighter latency and jitter guarantees to support the more demanding requirements of SNA and even voice applications. This will require the implementation of more advanced QoS techniques within the backbone.

FRAME RELAY LIMITATIONS

One of the most pressing Frame Relay limitations is its lack of remote access capability. First, there is a very limited Frame Relay dial-up infrastructure; it uses predominantly ISDN and is targeted for either dial-backup or lower cost local loop access from fixed locations. In addition, Frame Relay PVCs create a static definition of the network, which does not easily adapt to on-demand remote access applications. While Switched Virtual Circuits (SVCs) will eventually allow connections to be dynamically established, implementations remain few and far between.

Secondly, Frame Relay services are not as widely deployed as Internet services. There are fewer Frame Relay providers, and even the largest have only 400 to 500 POPs. This means that local access charges can be very high for networks of broad reach. International coverage is also limited, which complicates deployment of a multinational Frame Relay service.

Ubiquity of service is further constrained by the lack of Frame Relay “peering points.” In general, the Frame Relay Network-to-Network Interface (NNI) is not widely implemented - mainly because it doesn’t include the mechanisms necessary to protect SLAs across provider boundaries. The result is very little provider-to-provider connectivity, which again constrains the geographic coverage of a Frame Relay-based VPN. There is hope that this will improve as providers leverage more advanced ATM and PNNI (Public Network to Network Interface) technology within the core of their Frame Relay networks, but this will take time.

At a more basic level, because Frame Relay is a Layer 2 technology its service-level control is strictly at Layer 2, not at the level of the applications using the Frame Relay network. A single PVC may have specific QoS associated with it, but the applications sharing that PVC do not have priority or differentiated service capability. And because each PVC has a cost, it is very expensive to create multiple PVCs between two points to accommodate different QoS requirements.

In summary, Frame Relay services do not offer the ubiquity, flexibility, and breadth of services offered by the Internet. Multi-national Frame Relay access is still difficult to achieve; and because network-to-network connections are not widely supported, the basic infrastructure needed to support SLAs between Frame Relay carriers is not even in place. While Frame Relay has provided a valued service, there is a very real need for the Internet to support the next wave of network applications.

INTERNET ADVANTAGES

The Internet is comprised of a collection of networks based on standard IP networking technologies and well established peering arrangements that are unique to the Internet. The result is ubiquity of service - a VPN user connecting to one ISP can reach a user on any other connection to the Internet. This global, any-to-any model of communication is fundamental to delivering VPN services.

The Internet also offers great diversity of access technologies. The dial access infrastructure is widely deployed, and high-speed dedicated connections already reach OC3 rates. Aggressive cable and DSL installations will further expand both business and residential access options at very competitive prices.

Internet cost of service is often lower than Frame Relay, particularly as connectivity requirements increase. This lower cost is achieved in part by the larger scale of the Internet. Many national ISPs have between 600 and 800 POPs along with regional access partnerships and a large international presence. The important benefits are lower U.S. access costs and easier international connection when compared to Frame Relay services. Cost is also lower given the connectionless nature of IP services - the customer buys a single IP interface that offers the flexibility to reach any number of locations on the Internet. There is no requirement to pay for separate connections between users, unlike the case with PVCs. Flexibility is further enhanced with the ability to use a single Internet interface both for secure VPN traffic and for open Internet communications as well.

Internet providers are being very aggressive in their service level commitments. The most advanced providers offer 100% availability guarantees, round trip latency guarantees and, most recently, packet loss commitments. Of great importance to Internet providers is the emergence of new, high performance Internet access routers, e.g., Xedia's Access Point, that provide integrated service level measurement, monitoring, and control capabilities required to extend SLAs from the backbone to the customer premises.

INTERNET LIMITATIONS

The Internet is often criticized for the best-effort nature of its service. All traffic is treated equally, which means high priority applications have no better QoS guarantee than lower priority applications. However, new standards like Differentiated Services (DiffServ) and Multiprotocol Label Switching (MPLS) will be used with the new generation of QoS-enabled products to render very effective Internet QoS.

For example, implementing DiffServ with advanced IP classification and control technologies such as Class-Based Queuing (CBQ) will enable absolute rates of IP service comparable to Frame Relay CIRs. Other technologies such as MPLS will allow backbone traffic engineering that can allocate network resources based on QoS attributes.

Another often cited limitation is that Internet SLAs do not yet cross network provider boundaries, making it difficult to assign end-to-end QoS characteristics to Extranet applications among business partners. The good news is that several of the required mechanisms, such as a well established peering infrastructure, are in place. Standards like DiffServ will enable providers to indicate service level requirements at these well defined boundaries, meaning that end-to-end Internet SLAs should be within sight.

Perhaps more importantly, there are established business models for Internet SLAs such as the one established for the Automotive Network Exchange (ANX). The ANX is an on-line procurement network (an Extranet) operated for the benefit of the "big three" U.S. auto makers and their suppliers. It is managed by a separate agency that establishes and monitors service interworking rules. Participating ISPs must agree to a specific set of performance criteria, and equipment vendors must be certified to meet standards for ANX service applications.

TECHNOLOGY FOCUS: NETWORK QOS FOR FRAME RELAY AND IP

FRAME RELAY QOS

Today, Frame Relay network performance is more consistent and predictable than Internet performance. This stems mainly from the ability to define and monitor end-to-end throughput or bandwidth characteristics at the level of individual PVCs. It is implemented as part of the network configuration process, where each PVC is provisioned on an end-to-end basis with a CIR of anywhere from zero (for best effort service) up to the full rate of the link.

There are two levels at which provisioning needs to occur: 1) the access link into the network, and 2) across the backbone. The access link can be oversubscribed in order to reduce the cost of the local circuit. This means that the sum of all CIRs on the access link can exceed the maximum rate of the link. If all PVCs burst to their full CIR, traffic is typically lost. When this happens, traffic is discarded without regard to the underlying applications - because Frame Relay has no packet by packet classification capability, it cannot differentiate business critical traffic from best-effort traffic.

Across the backbone, bandwidth is provisioned to meet the aggregate demand of PVCs for all network users. This part of the process is supported by Frame Relay provisioning tools that provide an instantaneous view of the bandwidth available to accommodate each new PVC. Again, the aggregate demand typically exceeds the bandwidth available in the backbone. Network managers generally apply a specific over-subscription ratio (e.g., 1 Mbps of bandwidth is available for every 4Mbps of bandwidth provisioned) in order to gain the cost efficiencies of statistically multiplexing many customer flows.

This model is exactly like other “packet switching” networks, such as the Internet, which take advantage of the bursty nature of data traffic to efficiently share expensive wide area resources. However, Frame Relay has a real advantage in that it is very easy to configure and monitor the over-subscription level. Since Frame Relay is a connection-oriented service, trunk over-subscription is a static factor that is easily tracked and measured. The same cannot be said for connectionless IP services.

While Frame Relay has gone a long way toward achieving good performance, it still falls short in delivering the explicit QoS often associated with technologies like ATM. In fact, the preferred track for achieving Frame Relay QoS is to map Frame Relay PVCs to ATM PVCs. Frame Relay traffic can then inherit the more explicit real time and non-real time QoS attributes of ATM services. Given that most Frame Relay networks within the United States have migrated to ATM backbones, the challenge will be to take advantage of ATM QoS without adding unnecessary complexity.

INTERNET QOS

Internet providers currently rely on traffic engineering - with continuous monitoring of traffic patterns and regular re-engineering of the network - to best meet service level commitments. Because the IP protocol is connectionless, IP traffic engineering tools do not pre-allocate traffic flows over specific paths through the Internet. However, ISPs have found that traffic patterns do not change dramatically from one cycle to the next. They can therefore engineer network resources to deliver stronger latency and availability commitments in their SLAs. These commitments will strengthen as new Internet QoS technologies and products enable more explicit service levels.

Again, the QoS discussion can be divided into two parts: access and backbone. Access QoS applies from the customer premises to the backbone. Historically, Internet users have had no ability to control the access link. Traffic has competed for a single pipe, over-subscription has been the rule, and the effect has been unpredictable performance for all applications. In addressing this problem, Internet users will benefit tremendously from a new class of QoS-capable Internet access routers that can assign committed access rates and prioritize traffic. Operating at the level of individual applications, users, or any policy-based group, Class-Based Queuing (CBQ) is a standards-compliant technology defined within the IETF community. CBQ enables explicit policing, shaping, and bursting policies at the network access point. It can offer even better control and recovery than current Frame Relay access devices. This control is based on the ability to shape traffic on a per-application basis, which is much more granular than the per-PVC definition of Frame Relay.

DiffServ, also defined by the IETF, takes this control one step further by providing a standards-based way to define service requirements at key network boundary points, e.g., the customer premises or network-to-network boundary. CBQ controlled classes can be mapped to DiffServ service descriptions for backbone delivery of end-to-end service levels.

DiffServ and MPLS will be a strong team in this environment. MPLS promises to provide a means of taking DiffServ flows and mapping them to aggregated flows with common service level requirements. In this scenario, the “label” or “tag” associated with the aggregated flow is analogous to a PVC identifier. Because this process brings a measure of connection-oriented transport to the IP network core, ISPs will be able to control over-subscription levels with more advanced traffic engineering. Network provisioning tools will also need to evolve to make it easier to add new users, etc.

Comparing The Cost

In general, the cost difference between a Frame Relay and Internet VPN solution widens in the Internet's favor as connectivity and bandwidth requirements increase. Local access and international access charges are typically lower for the Internet given its broad, global distribution of service. And because each new Frame Relay connection or PVC has a cost, the more diverse the connectivity required, the more cost effective the Internet will be.

In addition, some customers push Internet traffic back to headquarters, where a single corporate Internet access connection is maintained. This creates additional traffic on the backbone, raising the cost of more expensive corporate backbone links. As Internet usage increases, this will become a more noticeable portion of the private network cost.

It remains to be seen how cost equations will change as more explicit QoS is added to both Frame Relay and Internet services. However, customers should expect to pay differentiated pricing for differentiated services. It should also be noted that cost can vary with the application, and cost-saving is not the primary reason for a customer to choose Internet or Frame Relay VPNs.

Summary

The business rationale for VPNs is proven. The resulting new deal in wide area networking will require flexible business connectivity, universal customer access and diverse dial-up and dedicated service options. Only the Internet holds the cards to be the unified network that can meet all of these business needs.

While Frame Relay has become a relatively established solution for replacing private WANs, it will reach its limits with the emerging generation of business applications. It is expected that existing Frame Relay applications will remain with Frame Relay networks in the near term. This will change as network administrators must expand or enhance these applications. At some point it will be important to eliminate the cost and complexity of managing multiple discreet services, and these applications will converge with other applications already leveraging the power of the Internet.

In the final analysis, both Frame Relay networks and the Internet are taking large steps to accommodate new VPN service requirements. However, the more significant service provider investment in products and new technologies is happening in the Internet. In many respects the Frame Relay technology investment is now incremental, attempting to better solve existing problems. The Internet investment is fundamental to enabling a new networking paradigm that serves a global community of corporate users, e-commerce and business partner applications… all on the same network with the same technology!

Motivation is high, and so are the stakes. Network providers and customers alike can see large returns on the long curve ahead.

XEDIA CORPORATION
119 RUSSELL STREET
LITTLETON, MA USA
TEL: 978-952-6000
FAX: 978-952-6066
WWW.XEDIA.COM

93-1001463-01