What is a Virtual Private Network (VPN)? First of all, a VPN is a *network*; that is, it provides inter-connectivity to exchange information among various entities that belong to the VPN. Secondly, it is *private*, that is, it has all the characteristics of a private network.

One might ask, "What characterizes a private network?" A private network supports a closed community of authorized users, allowing them to access various network-related services and resources. The traffic originating and terminating within a private network traverses only those nodes that belong to the private network. Further, there is traffic isolation. That is, the traffic corresponding to this private network does not affect nor is it affected by other traffic extraneous to the private network.

The final characteristic of a VPN is that it is *virtual*. A virtual topology is built on an existing, shared physical network infrastructure. However, the virtual topology and the physical network are usually administered by different administrative bodies.

VPNs can be formally defined as *a communication environment constructed by controlled segmentation of a shared communications infrastructure to emulate the characteristics of a private network*. The access to the communication environment is controlled to permit interconnections for a defined community; even though, the underlying shared communications infrastructure provides services on a non-exclusive basis.

The following terminology will be used. The term *VPN* will refer to the private network of a company or enterprise. The term *shared network infrastructure* will be used to describe the underlying infrastructure on which the VPN is constructed. This can either be the public Internet or a network consisting of one or more service providers.

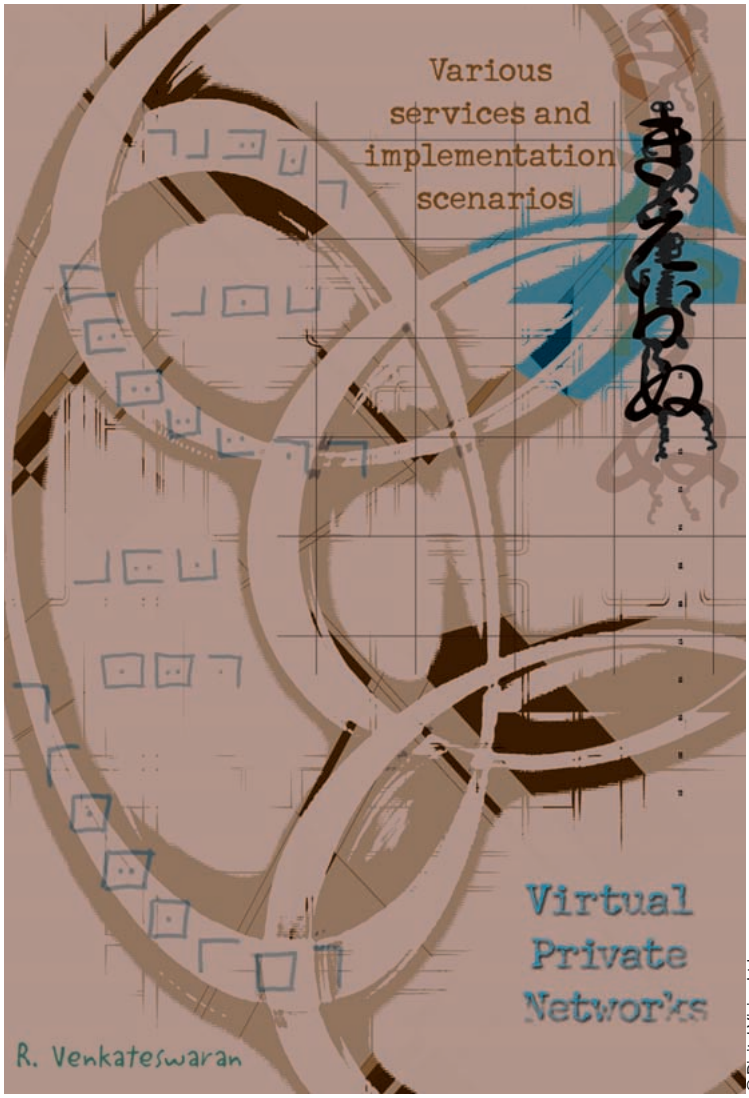

Motivations for VPNs

Why are VPNs so popular today? Traditional private networks facilitate connectivity among various network entities through a set of links, comprising of dedicated circuits (T1, T3 etc.). These are leased from public telecommunication carriers like MCI-Worldcom or Regional Bell Operating Companies (RBOCs) as well as privately installed wiring. [Note: T1 and T3 (used in the United States) have transmission rates of 1.544 and 44.736 Mbps, respectively.] The capacity of these links is available at all times, albeit fixed and inflexible. The traffic on these private networks belongs only to the enterprise or company deploying the network. Therefore, there is an assured level of performance associated with the network.

Such assurances come with a price. Traditional private networks are not cheap to plan and deploy. The costs of dedicated links are especially high when they involve international locations. The planning phase of such networks involves detailed estimates of the applications, their traffic patterns and their growth rates. The planning periods are long because of the work involved in calculating these estimates. Further, dedicated links take time to install. It is not unusual for telecommunication carriers to take about 60 to 90 days to install and activate a dedicated link.

Given the rapidly evolving network technologies and network applications, such a long waiting period can adversely affect a company's ability to react to quick changes in these areas. Network planners have to anticipate all possible scenarios that may arise due to these technological changes. The lack of network agility implies that a minor misreading of the market forces can have a magnified negative impact on the outcome of the company's businesses. Since accurately predicting the market is extremely difficult, a flexible Information Technology (IT) infrastructure that can quickly adapt is crucial.

Another recent trend is the mobility of today's workforce. Many companies are increasing employee's productivity by equipping them with portable computing facilities. Affordable laptops and various palm-based devices have made it easy for people to work without being physically present in their offices. Besides increased productivity, companies are encouraging telecommuting to reduce their investments in real estate. Also, it reduces traffic and pollution from automobiles.

To support the increase in home-offices, companies need to provide a reliable IT infrastructure so employees can access company information from remote locations. This has resulted in the deployment of large modem pools for employees to dial-in remotely. The

0278-6648/01/$10.00 © 2001 IEEE

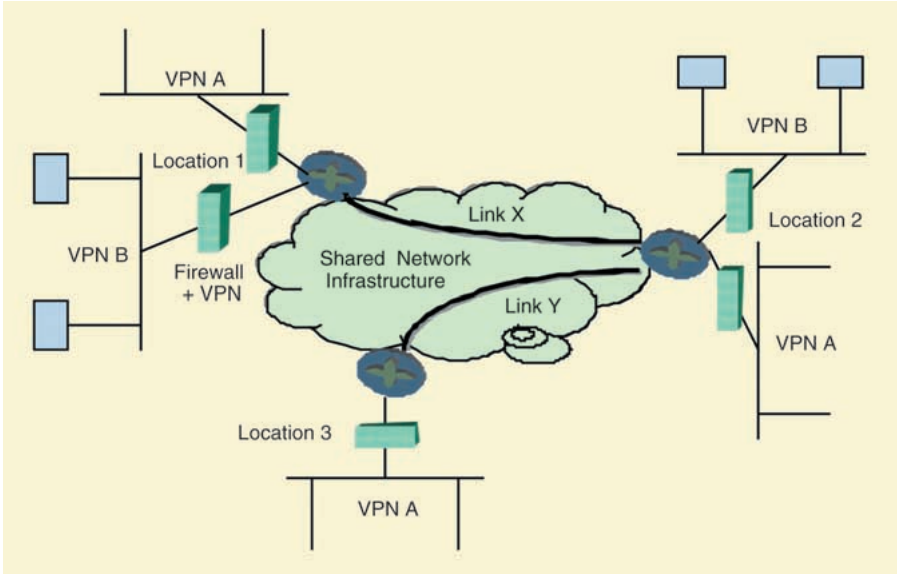

Fig. 1 LAN Interconnect VPN services

cost keeps increasing due to the complexity of managing and maintaining the large modem pools.

An additional cost with mobile users is the long-distance calls or toll-free numbers paid for by the company, which typically costs 5-7 cents per minute. For a US employee who puts in eight hours of work each week from remote locations, the dial-in cost itself adds up to $30 a week or about $130 (USD) a month. The costs are much higher if we consider international calling. Note that this cost does not include the cost associated with the modem pool itself. For companies with a large, mobile workforce, these expenses quickly add up to significant numbers.

What's more, the dial-in connection limits the remote user to a maximum access speed of 56 Kbps for analog modems and 128 Kbps for Integrated Services Digital Network (ISDN) modems. These limitations hamper the day-to-day activities that require high-speed access to the Intranet as available from a regular office. The home-office, therefore, becomes a poor substitute for the regular work environment.

Advent of high-speed access media like cable modems and Digital Subscriber Lines (DSLs), that are now becoming more available and affordable, can overcome the access speed limitations. But, the service providers offering high-speed accesses cannot have easy access to a company's Intranet due to firewall and security restrictions. Companies restrict access to prevent unauthorized intruders or "hackers" from stealing proprietary information. There is an urgent need for a reliable mechanism to authenticate valid users and restrict their accesses based on their access privileges.

E-commerce applications provide considerable advantages over traditional brick-and-mortar operations. Such applications are deployed around inventory management, supply-chain management, electronic data interchange (EDI) etc. For example, suppliers having electronic access to the company's inventory database helps the suppliers to schedule additional supplies based on the demand and current inventory levels. This helps to efficiently manage the inventory, eliminating the need to store large quantities of unused inventory.

But, in a traditional private network, this kind of special access is very difficult to incorporate because it is not easy to install dedicated links to all the suppliers. Further, this infrastructure is not flexible because any change in the supplier involves de-installing dedicated links and installing new links to the new vendor. Quickly replacing a supplier result in enormous cost-savings, but an inflexible infrastructure makes it difficult to take advantage of these savings.

A Virtual Private Network (VPN) can help resolve many of the issues associated with today's private networks. As we will see, a VPN facilitates an agile IT infrastructure. Global VPNs enable connectivity to all locations anywhere in the world at a fraction of the cost of dedicated links. VPN services enable remote access to the Intranet at significantly lower cost, thus enabling support for a mobile workforce. Additionally, the VPN architecture supports a reliable authentication mechanism to provide easy access to the Intranet from anywhere using any available access media including analog modems, ISDN, cable modems, DSL and wireless.

Types of VPN services

There are primarily three types of VPN services:

1) Local Area Network (LAN) Interconnect VPN services,

2) Dial-up VPN services,

3) Extranet VPN services.

LAN Interconnect VPN. LAN Interconnect VPN services help to interconnect local area networks located at multiple geographic areas over the shared network infrastructure. Typically, this service is used to connect multiple geographic locations of a single company. Several small offices can be connected with their regional and main offices. This service provides a replacement for the expensive dedicated links.

A simple LAN interconnect example is shown in Fig. 1 . VPN A has sites in geographic locations 1, 2 and 3, while VPN B has sites in geographic locations 1 and 2. Both VPNs A and B are implemented on top of a shared network infrastructure. The advantage is the flexibility it offers.

For example, it is easy to increase the capacity of any of the links depending on the applications supported on the VPN. As applications change with time, the architecture can be adapted to meet the needs. Further, additional geographical sites can be connected to the VPN with very little effort.

These advantages come with a reduction in cost as well. Dedicated private lines are expensive. For example, in October 2000, a well-known telecommunications carrier quoted a price of $1300 a month for a dedicated 1.544 Mbps T1 link spanning a reasonable distance, plus a one-time installation charge of $3000. These costs are much higher if international sites are involved. On the other hand, using a shared infrastructure is cheap because of the economies of scale. The costs of the links are borne by the different VPNs that are supported on the infrastructure.

For example, in Fig. 1, the cost of

virtual link X is borne by VPNs A and B. Suppose the shared link is used more by VPN A in the mornings, while VPN B uses more of the link in the evenings. If a dedicated link had been used for private network A, its capacity would have to be at least as much as to meet the demands of the morning traffic. This capacity is not needed in the evenings and therefore, remains unused and wasted. Company A has to pay for this unused capacity. Therefore, sharing the infrastructure helps companies A and B to reduce their individual costs. Keep in mind that even though the link is shared, the underlying shared network infrastructure retains the characteristics of a VPN by providing mechanisms to isolate and secure the traffic of each VPN.

Dial-up VPN services

The Dial-up VPN service supports mobile and telecommuting employees in accessing the company's Intranet from remote locations. A typical VPN is shown in Fig. 2. The remote employee (user) dials into the nearest Remote Access Server (RAS, the technical term for modem pool). This is typically a local Point-of-Presence (PoP) of an Internet Service Provider (ISP) or the shared network infrastructure.

In one dial-up VPN model, called the Layer 2 Tunneling Protocol (L2TP), the RAS automatically establishes a secure connection to a prespecified location inside the company's Intranet, usually through a firewall enhanced with VPN capabilities. Contingent upon successful authentication of the user, the secure connection enables the user to transparently connect to the Intranet. The L2TP model is also known as a "static" VPN connection and is usually aimed at home-offices and telecommuters who dial-in to a specific local RAS.

On the other hand, an alternate Point-to-Point Tunneling Protocol (PPTP) model focuses on the mobile user, who may dial-in to any local ISP. After connectivity has been established to the ISP, the user initiates a connection to any of the VPN servers located inside the company's Intranet. A remote authentication mechanism validates the user and establishes the access privileges. The successful establishment of the user-initiated connection enables the user to access the Intranet. In contrast to the L2TP model, the RAS does not participate in the establishment of the VPN connection. Therefore, no specific configuration of RAS is needed for the PPTP model. The PPTP model is also used for VPN connections based on high-speed access media like cable modems and DSL.

The dial-up VPN service results in considerable cost-savings to a company. It eliminates the need for managing large modem pools and uses the RASes that belongs to the local ISPs. There is also a significant reduction in long-distance charges because dialing in to the RAS is, in most cases, a local call. The monthly cost per user for a local ISP is about $20 for unlimited access. This cost is only 15% of the monthly long distance charges of the $130 per user that we had computed earlier. In addition, the dial-in VPN service takes advantage of high-speed access media, thus, eliminating some of the access limitations of the home-office.

Extranet VPN services

An extranet VPN service, shown in Fig. 3, combines the architecture of LAN Interconnect VPN services and dial-in VPN services. This infrastructure enables external vendors, suppliers and customers to access specific areas of the company's Intranet. The allowed specific area is denoted as the Demilitarized Zone (DMZ). When a

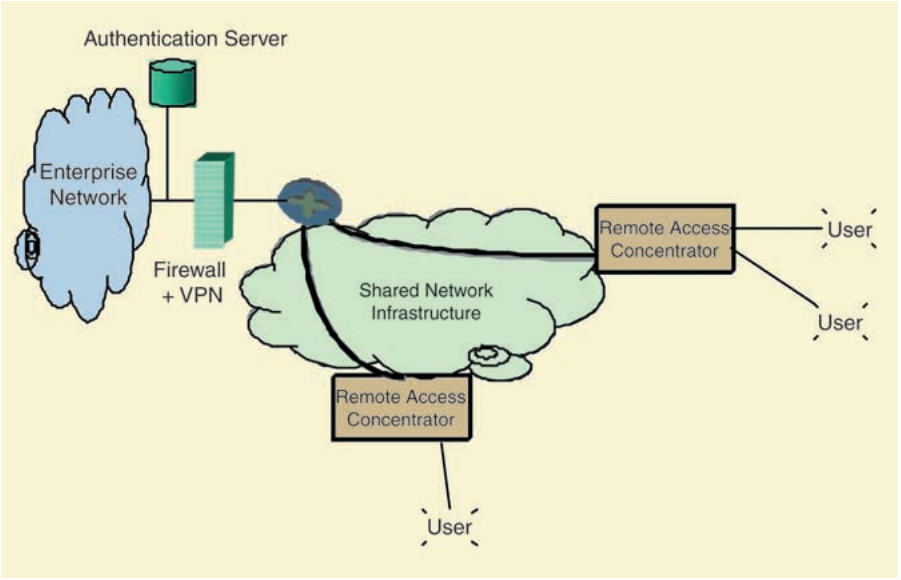

Fig. 2 Dial-up VPN services

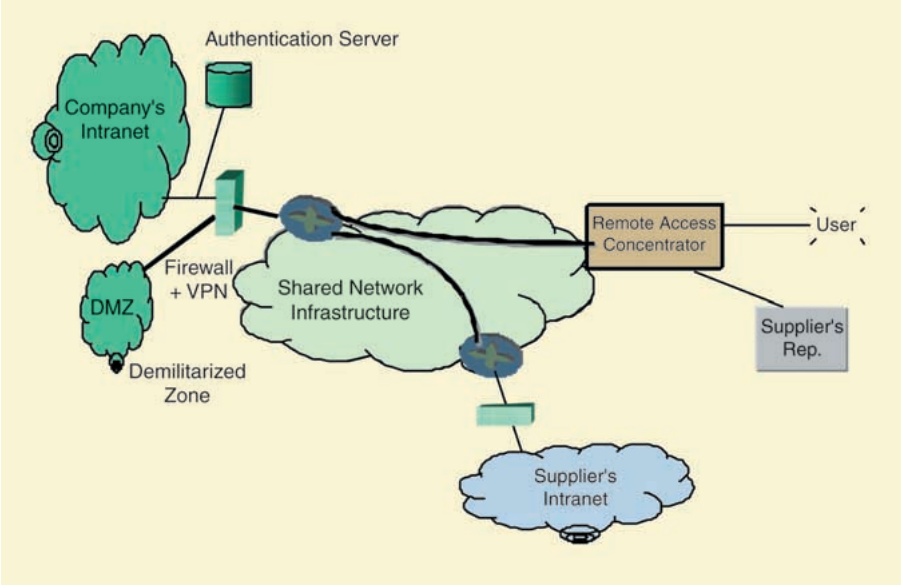

Fig. 3 Extranet VPN services

supplier's representative connects to the company's Intranet, either from the supplier's Intranet or dialing in remotely, the firewall and authentication mechanisms ensure that the connection is directed to the DMZ. A company employee (user), on the other hand, has full access to the company's Intranet.

The flexibility of the extranet services helps to provide connectivity to new external suppliers and customers within a short period of time. The fast communications facilitated by the extranet helps in several e-commerce areas including efficient inventory management and electronic data interchange (EDI). This provides significant savings in cost and the ability to effectively compete in the rapidly growing market.

VPN implementations

We have seen the advantages of VPN services. Now, we will focus on the capabilities of the underlying infrastructure that will help in implementing VPNs. Specifically, we will look at the various architectures for implementing VPNs. Typically, VPNs are implemented either at the network layer or at the link layer. Though application layer VPNs are also possible, we do not focus on such VPNs here.

Network layer VPNs

The network layer VPNs are usually based on Internet Protocol (IP) at the network layer. These VPNs can be implemented either by *tunneling* or by *network layer encryption*. A tunnel connects two points of a VPN across the shared network infrastructure. In the tunnel mode, the end-points of the tunnel are common nodes of the VPN and the shared network infrastructure.

Architecturally, the VPN is a collection of tunnels established over the shared network infrastructure. The network layer packets leaving a VPN node at one end of the tunnel are appended with an additional IP header whose destination address reflects the other end of the tunnel (remote VPN node). The packets are then routed based on this modified destination address through the shared network infrastructure to the other end of the tunnel. At this point, the additional IP header is stripped away to re-create the original packets reaching the remote VPN node at the other end of the tunnel.

Note that the original packets could be based on any Layer 3 protocol (like IP, AppleTalk or Novell's IPX) and still be carried across the shared infrastructure to the remote VPN node. Tunneling, therefore, helps to route multiple protocols across the shared network infrastructure. Further, the VPN and the shared network infrastructure may use different routing protocols without hindering the routing process. Typically, the network-layer protocol within the shared infrastructure is IP.

Another advantage of the tunneling mechanism is that the address mechanisms of the VPN and that of the shared infrastructure can be completely separated. Therefore, address overlaps across multiple VPNs can be easily handled without the need for Network Address Translation (NAT). Note, however, that when a VPN needs to connect to the external Internet, the NAT functionality is needed to translate the private VPN addresses to a set of global IP addresses.

A disadvantage of the tunnel architecture is that it is difficult to manage a large number of tunnels. Therefore, it does not scale well to a large number of VPN nodes. Further, the packets on the unencrypted tunnels can be eavesdropped by others attached to the shared network infrastructure. This tunnel is especially vulnerable at tunnel end-points where the extra headers are stripped away and the packets are visible in their original form. Since tunnels represent only the end-points and not the path taken to reach the other end of the tunnel, the paths taken across the shared network infrastructure may not be optimal. This can create serious performance problems for the VPN.

Network layer encryption provides a secure mechanism for implementing VPNs. The Internet Engineering Task Force (IETF) has standardized on a secure IP architecture, IPSec, which is a collection of protocols, authentication and encryption mechanisms. IPSec is an extension to the standard IP protocol. Parts deal with managing the encryption keys, key exchange protocol and protocol negotiations. These mechanisms can be used for encrypting other tunnels (like L2TP and PPTP) as well.

An IPSec packet has an IP header and, therefore, can be routed by current IP routers from one VPN node to another. Using an encryption algorithm and a set of encryption keys, the IPSec packet is created by encapsulating and encrypting the original IP packet. The encryption keys and the algorithm parameters are negotiated and exchanged between the two VPN nodes using the Internet Key Exchange (IKE) protocols (a part of IPSec protocol specifications) In addition, the IPSec packet may also have an authentication header, which authenticates the validity of the entire IPSec packet. This enables the receiver to verify that the packet has not been modified en route.

Link-layer VPNs

If the shared network infrastructure is based on a switched link-layer technology [like Frame Relay (FR) or Asynchronous Transfer Mode (ATM)], the VPNs implemented directly on these technologies are called Link-Layer VPNs. With advents in switch-routers and protocols like MultiProtocol Label Switching (MPLS) and MultiProtocol over ATM (MPOA), the line distinguishing the Network-layer VPNs and link-layer VPNs is getting hazy.

The links belonging to the VPNs are implemented as virtual circuits at the link-layer. The FR frames or ATM cells are switched across the shared network infrastructure from one node of the VPN to the other. The advantage of virtual circuits is that they are cheaper than dedicated links and they are very flexible. The virtual circuits also come with some Service Level Agreements (SLAs). They provide guarantees on the performance levels of the virtual circuits.

Link-layer VPNs are appropriate for LAN Interconnect VPN services. Link-layer VPNs are not ideally suited for dial-up VPN services because most ISPs provide connectivity through IP. Since dial-up VPN services offer the most cost reductions, IP-based network-layer VPNs look more attractive to IT managers than link-layer VPNs.

As with tunnels, there are scaling concerns when link-layer virtual circuits adopt a full-mesh architecture to connect each pair of VPN nodes. To help scale better, other kinds of architectures like partial meshes, hub-and-spoke may be considered. The trade-off with these architectures is that they may be sub-optimal.

VPN concerns

So, why aren't all companies rush-

ing to deploy VPNs? A survey in "Telechoice VPN Market Report" (www.telechoice.com) showed that more than 25% of the 501 companies surveyed already have some form of VPNs in place. What about the remaining 75%? What are their concerns?

One major concern is security. Traditional private networks are very secure because all the links and the nodes the company deploying the networks the company owns. In a VPN, the data traverses links that are owned by service providers and are shared with other VPNs. Many companies are not comfortable with this idea that their packets could be eavesdropped by someone with malicious intent. The tunnel end-points are the most vulnerable locations on a tunnel-based VPN.

Advances in security features and the standardization of IPSec have reduced some of the security concerns. But, at the time of this writing, these standards are not fully mature . Though there are several IPSec-compliant vendor products , the standards are not mature enough today to guarantee interoperability across vendors. It is important for the protocols to mature and vendor products to become interoperable before we see a wider deployment of IPSec architectures.

Another cause for concern is the cost of deploying VPNs. Security features are not cheap. The encryption algorithms of IPSec are very compute-intensive. Therefore, use of encryption can have a negative impact on network performance. The various key management protocols associated with IPSec are also expensive to implement and manage. Given these facts, the hidden costs of deploying VPNs are not as attractive as the cost reductions advertised for VPNs.

Another issue of concern is the performance of the network and the quality of service (QoS). The public Internet is based on best-effort services. Therefore, no service level guarantees can be made if the public Internet is used as the shared network infrastructure for VPNs. A service provider implementing a VPN may be able to offer some service level agreements (SLAs), but these SLAs add up to the cost of deploying the VPN. Dedicated links must be established from the VPN service provider's network to the company's Intranet, which further increases the cost. The 85% cost reductions that are associated with VPNs will not hold true any longer.

There are different kinds of SLAs that service providers advertise. For example, some service providers advertise guarantees on network availability and end-to-end latency. Living up to these SLAs is a big challenge for the service providers. VPN network administrators must be able to demonstrate that the SLAs are indeed met at all times. On a shared network infrastructure, traffic separation, isolation and prioritization are necessary to assure that each VPN gets its requested share of the service.

Recent advances in IP router technology that enable class-based queuing, support for differentiated services and policy-based services have contributed to the success of meeting these demands. But, these advances are still in the preliminary stages. There is no widespread deployment of these technologies to test vendor interoperability and performance. Once again, designing networks with these capabilities come with added cost, which is reflected in the cost of deploying VPNs.

Using a VPN service provider has another drawback for VPNs because it limits the number of Points-of-Presence (PoPs). All service providers may not have PoPs at all the company locations and providing connectivity to PoPs can add to the cost of the VPN. Given these additional costs, it is no surprise that companies are waiting for the technology to mature before they jump on the bandwagon.

The future

VPN technology is still in its infancy. But the general belief is that in a couple of years VPNs will evolve and demonstrate all the advantages that they have promised. VPN will be a global technology linking geographic regions around the world. Adoption of standards in security and QoS technology will help vendors to minimize interoperability problems among their products. The service providers will then be able to offer and deliver precise SLAs. These SLAs will transcend multiple service provider networks.

As implementing and managing VPNs become increasingly complex, companies will favor outsourcing them, thereby, reducing their IT costs. At the same time, VPN service providers will be able to carve out a profit margin by offering different kinds of services.

Read more about it

• P. Ferguson and G. Huston, "What is a VPN," a whitepaper available from <http://www.clark.net/timw/vpn/Tech/vpn.pdf>.

• D. Fowler, *Virtual Private Networks: Making the Right Connection,* Morgan Kaufmann Publishers, San Francisco, California, 1999.

• Numerous VPN-related web sites. A good starting point with several links to relevant pages can be found at <http://kubarb.phsx.ukans.edu/~tbird/vpn.html>.

• Internet Engineering Task Force (IETF) publishes several documents: RFCs and Internet drafts: describing the various Internet standards. See <http://www.ietf.org>.

About the author

R. Venkateswaran has been a Member of Technical Staff at Lucent Technologies since 1997. He earned his PhD from Washington State University, Pullman, WA in 1997. His work and areas of interest include VPNs, QoS in packet-switched networks and multicasting. He is a member of the IEEE and can be contacted at venki@ieee.org.

Acronyms

ATM	Asynchronous Transfer Mode
DMZ	Demilitarized Zone
DSL	Digital Subscriber Lines
EDI	Electronic Data Interchange
FR	Frame Relay
IETF	Internet Engineering Task Force
IKE	Internet Key Exchange
IP	Internet Protocol
IPSec	IP Security
IPX	Internetwork Packet Exchange
ISDN	Integrated Services Digital Network
ISP	Internet Service Provider
IT	Information Technology
L2TP	Layer 2 Tunneling Protocol
LAN	Local Area Network
MPLS	MultiProtocol Label Switching
MPOA	MultiProtocol over ATM
NAT	Network Address Translation
PoP	Point of Presence
PPTP	Point-to-Point Tunneling Protocol
QoS	Quality of Service
RAS	Remote Access Server
RBOC	Regional Bell Operating Company
SLA	Service Level Agreement
VPN	Virtual Private Network