Gartner

Research

Publication Date: 10 June 2005 **ID Number: G00127825**

An Introduction to IP VPNs

Kathleen M. Adams

Internet Protocol virtual private networks are attractive replacement solutions to legacy data networks, such as frame relay, asynchronous transfer mode and private line. In particular, managed network-based IP VPNs are poised for significant growth during the next four years.

TABLE OF CONTENTS

Gartner

1.0 Technology Basics

Virtual private networks (VPNs) enable enterprises to build secure private communications over public network infrastructure typically at a lower cost than traditional private-line networks and many frame relay and asynchronous transfer mode (ATM) networks. IP VPNs are managed or unmanaged Layer 3 (network or customer premises/customer premises equipment [CPE]-based) offerings, providing a full range of metropolitan-area network/local-area network (MAN/LAN) functionality using IP backbone transport technology — either a private carrier-operated network or the Internet IP network for a closed (or at least well-defined) community of interest.

1.1 Types of IP VPNs

The equipment and software used to run an IP VPN can either be located in the carrier's network or on the customer premises (or both in a hybrid situation). The location of the VPN functionality (that is, encryption, tunneling, authentication, quality-of-service assignment) determines whether an IP VPN is considered a network-based, CPE/premises-based or hybrid solution.

1.1.1 Network-Based IP VPN

Also known as non-Internet IP VPNs, network-based IP VPNs transport traffic over a private IP or Multiprotocol Label Switching (MPLS)-based network. Intelligence and VPN functionality are performed by software/equipment deployed at the edge (a central office or point of presence [POP]) of the service provider's network and extended out to many end-user locations over an ordinary access link. Security is enabled from service provider edge (PE) to PE — thus the access line is not secured by encryption or tunneling nor is there quality of service (QOS) applied to the services unless these features are incorporated into the access router. Although a network-based IP VPN requires only a standard router at the customer premise, customers requiring encryption over the access link may want to consider a hybrid network-based/CPE-based solution or installing IPsec-compliant devices at the customer premise. Geographic availability can sometimes be restricted, since network-based IP VPNs are limited to the reach of the carrier's network, especially for MPLS-based networks. Presently, there are a few techniques that provide cost-effective solutions for interoperability of different carrier’s MPLS networks (for example, solutions by Vanco and Virtela), but in practice, this area is very immature.

Preference for the management of network-based IP VPNs varies by geography — to date, larger enterprises have been inclined toward unmanaged IP VPN in the U.S., although both managed and unmanaged IP VPNs are growing. In Canada and Western Europe, the provider typically manages network-based IP VPNs, although a small number of enterprises have chosen to manage their network-based IP VPN offerings.

MPLS is the dominant architecture for network-based IP VPNs. MPLS-based IP VPNs are fully meshed solutions using Layer 3 routing, with a mix of various vendor-specific and standard-based protocols. MPLS creates virtual circuits between MPLS-enabled endpoints on the network, providing nearly instantaneous automated provisioning capabilities. MPLS provides the ability to establish traffic classes (classes of service), which allow tiered levels of QOS. Although still a less common practice today, compared to the use of DiffServ-enabled routers, MPLS handles multiple protocols (such as Layer 2 [ATM and frame relay] and Layer 3 [IP]) over the network and enables the migration of many ATM control plane functions to Layer 3. MPLS does not encrypt traffic, and therefore end users wanting the predictability of MPLS with the security of an encrypted IPsec flow should consider adding an IPsec overlay. Almost all operators are offering MPLS-based IP VPN services.

1.1.2 CPE/Premises-Based IP VPN

This type of VPN is enabled at the customer edge and transports traffic over a private IP network or the Internet via secure tunnels. Enterprises can choose from a variety of CPE, depending on their requirements: IP VPN routers with tunneling capabilities, standard routers with VPN intelligence in software, IP VPN gateways that work with external routers, IP VPN appliances for smaller offices or firewalls that act as VPN gateways. The CPE creates and maintains a tunnel through a private or public IP network and to a desired endpoint. Encryption is applied to the traffic before it leaves the customer premise, ensuring security from end to end, including the local loop. Standards-based approaches for providing CPE-based IP VPNs currently include PPTP, L2TP and IPsec.

CPE-based IP VPNs typically have just one class of service (COS) — "best effort" — and therefore are designed for networks where QOS is not important. The geographic coverage of CPE-based IP VPNs can be very wide because they can be implemented anywhere the Internet is available as long as the provider has an Internet POP. CPE-based IP VPNs can be managed by the customer, provider or third party, such as a managed services provider or systems integrator, and the enterprise or the provider can own the CPE. Enterprises deploying CPE-based solutions are able to switch service providers relatively easily, since the functionality of the IP VPN is processed by the CPE, not the network.

1.1.2.1 Secure Sockets Layer

Some providers are offering a Web browser-based alternative that supports remote access to a company's server-based information and applications. Secure Sockets Layer (SSL) is an encryption technology originally developed by Netscape Communications for encrypting data transported between browsers and Web servers. SSL-based IP VPNs provide access to browsers or devices that are not IPsec-compliant. SSL enables users to access Web-based applications from any location with a Web browser and an Internet connection without adding software to user systems. Certificates, authentication and encryption are handled through the Web browser. SSL-based VPNs are used primarily for extranets and casual remote access and are not designed for site-to-site connections. They are not a replacement to IPsec remote-based VPNs, but rather a complementary service.

1.1.3 Hybrid IP VPN

A hybrid IP VPN is a combination network-based/CPE-based (edge) IP VPN solution. A sample configuration would consist of a network-based IP VPN for the majority of sites, combined with a CPE-based IP VPN software client for off-net or highly security-conscious sites, remote locations and traveling users (with only dial or broadband access to the Internet).

1.2 Layer 2 vs. Layer 3 VPN

1.2.1 Layer 2 VPN

The distinction between Layer 2 and Layer 3 VPNs is blurring. Frame relay and ATM are legacy Layer 2 services, and Ethernet is an emerging Layer 2 service. Layer 2 VPNs allow the transport of IP and non-IP traffic across a common router infrastructure — they are multiprotocol in nature. Several Layer 2 techniques (including the Internet Engineering Task Force's [IETF's] Martini Draft) have been developed to enable packet-switched traffic (frame relay, ATM and Ethernet) and time division multiplexing traffic (such as voice and leased line) to be transported across an MPLS-enabled network. Another type, based on Cisco'sL2TPv3, enables Layer 2 traffic to be transported across MPLS and pure IP backbones. Layer 2 VPNs leave the carrier out of the

enterprise's IP layer, which is attractive for enterprises with special requirements or if trust in the provider is an issue. The latter is more of a U.S. view.

1.2.2 Layer 3 VPN

A Layer 3 VPN is not multiprotocol – it is a VPN transported across IP only. It is based on techniques defined by the IETF and can be either CPE-based or network-based.

Benefits of a Layer 3 VPN over a Layer 2 VPN include:

- Access Independent: Layer 3 VPNs are access-independent. Layer 2 VPNs are typically limited to certain access media.
- Layer 3 VPNs cope with large or meshed networks better than Layer 2 VPNs, making QOS easier and making it easier to deliver managed/hosted services.
- Convergence: combining Internet and internal (voice and data) traffic onto one network connection reduces complexity and potentially the cost of managing multiple connections.

IP-Enabled Frame Relay/ATM

IP-enabled frame relay/ATM services use a frame relay or ATM interface to access a fully-meshed IP (typically MPLS) transport network. This type of VPN is attractive, as the enterprise network needs a meshed, flexible, more-scalable architecture. With a typical frame relay or ATM network, a separate permanent virtual circuit (PVC) is required between each pair of sites that require direct communication. As the number of sites requiring connectivity increases, the number of PVCs grows as well, and this topology can become expensive. With IP-enabled frame relay/ATM, only one PVC per site is required for access to the service provider's IP-based network, where fully meshed any-to-any connectivity among sites is provided. Additionally, this option may also save some capital upfront since businesses can reuse their edge equipment and less retraining of personnel may be required. On the downside, IP-enabled frame relay/ATM involves two networks in the enterprise communications, and this adds configuration complexity, latency and increased reliability issues. It also limits the access options to those of the frame/ATM network.

1.3 Security

Three key functions are required to ensure security within an IP VPN: encryption, tunneling and authentication.

1.3.1 Encryption

Encryption is the process of encoding data on transmission so that only the intended recipient can read it using a secret decryption "key." Encrypting the data protects it from being intercepted and interpreted by hackers as it passes through the network. The most common encryption algorithms are Data Encryption Standard (DES) (56-bit) and Triple Data Encryption Standard (3DES) (168-bit), but there is growth in vendor support for Advanced Encryption Standard at various key lengths.

There are two types of encryption: symmetric and asymmetric. Symmetric encryption uses the same key to encrypt and decrypt a message. Asymmetric encryption, also known as public key encryption, provides each user with two keys: a public key, shared by many users and used to encrypt the message, and a private key, known only to the intended recipient and used to decrypt the message. Using public-key infrastructure (PKI), public keys can be managed securely for use

by widely distributed users or systems. Public keys are distributed using digital certificates issued by a Certificate Authority, a trusted third-party organization.

1.3.1.1 IPsec

The IETF's IPsec protocol is the leading standard for encryption. IPsec defines a framework of protocols that establishes security at the network layer. IPsec can be used in tunnel mode (entire packet is encrypted) or transport mode (only data is encrypted). The most common encryption algorithm associated with IPsec is 3DES, although others may be used. IPsec is relatively easy to implement for site-to-site connections and for individual remote access. The majority of individual IPsec VPN access solutions are implemented with proprietary value-added clients supplied by the VPN vendor or a third party.

1.3.1.2 Tunneling/Encapsulation

Tunneling is a technology that enables a network transport protocol to carry information for other protocols within its own packets. Specifically, an encrypted packet is encapsulated inside an IP packet with a new header and then delivered unmodified to a remote device. The most common tunneling specifications are:

- PPTP: developed jointly by Microsoft and U.S. Robotics, PPTP is an enhancement to the PPP for use over the Internet. Microsoft added 40- and 128-bit encryption and has embedded PPTP in all Microsoft VPN clients since Windows 98, including full OS and mobile OS (that is, Pocket PC). Third-party versions are available for other OSs. PPTP VPNs are typically used for individual remote access connections.

- L2TP: was jointly developed by Cisco and Microsoft and combines the best features of Cisco's proprietary Layer 2 Forwarding protocol with Microsoft's version of PPTP. L2TP is not commonly seen in use for individual remote access connections, but it has proven popular for carrier-based IP VPN WAN services because Layer 2 and PPP services can be managed on different devices in a packet-switched network, resulting in better performance. L2TP support is included in Windows 2000 and Windows XP. Note that L2TP relies on IPsec for message integrity and encryption.

1.3.1.3 Authentication

Authentication is the process of verifying the identity of a user (or host) trying to access corporate resources or verifying the origin of a transmitted message. Verification can be accomplished through several mechanisms, including Remote Authentication Dial-In User Service (RADIUS), digital certificates or one of the latest methods — digital fingerprinting. RADIUS is a client/software protocol and software that enables remote access servers to communicate with a central server to authenticate dial-in users (via user names and passwords) and to authorize their access to the network. Stronger authentication methods involve digital certificates issued via a token-based solution or PKI. PKI issues and manages certificates for authentication, signatures and encryption.

1.3.1.4 Firewall

No discussion about security would be complete without mentioning firewalls, since they are complementary to encryption — in fact, many firewalls today have encryption embedded. A firewall is a system or combination of systems, comprised of hardware and software, that enforces a boundary between two or more networks. The firewall is placed at the gateway between the enterprise's secure internal network and any insecure external resources. A network-based firewall secures the perimeter of a network from unwanted ingress or egress and

can perform a variety of functions, including content filtering management, site authentication, virus scanning, anti-spoofing, intrusion detection and prevention, and denial-of-service protection.

1.4 Traffic Engineering/ QOS

QOS is often used interchangeably with COS, but they are not one in the same. QOS is the foundation for policy-based networking — determining how to use network resources under specific conditions and how much bandwidth to allot to certain types of applications. QOS for IP VPN is measured using three parameters: jitter (change in time between packets), latency (delay) and packet delivery. COS identifies specific application traffic (such as high-priority/low-priority voice, video, data) as requiring particular QOS treatment and marks each of these packets with a COS mark to ensure they are routed/treated in the network with special emphasis to meeting QOS goals. Following are three prominent standards that allow for traffic engineering.

1.4.1 Resource Reservation Protocol

Resource Reservation Protocol (RSVP) aims to reserve resources, such as bandwidth and buffer space, so that applications can secure their needs. Resources for traffic are reserved specifically to an application's request before the data is transmitted. The protocol relies on end devices to request their priority so that resources can be reserved for flows to facilitate the requested QOS. Each router along the path to the end device attempts to honor the RSVP request by maintaining connection state information. If this cannot be honored, service is denied and packets are dropped. RSVP can be mixed with other protocols, such as DiffServ and MPLS. While RSVP is implemented today on some switches and routers, it is considered more of an enterprise protocol, since its complexity can impair performance on backbone networks. RSVP requires all routers to support it, and this can prove challenging on the Internet.

1.4.2 DiffServ

DiffServ is a prioritization-based Layer 3 protocol for specifying and controlling network traffic by class, enabling certain class types to receive precedence over others. Using prioritization, network traffic is classified into a set of classes, and network nodes provide preferential treatment to those classes deemed as a priority. DiffServ aims to divide traffic into classes and treat the classes differently, especially when there is a shortage of resources, such as bandwidth, as it prioritizes important traffic at the expense of lower-priority traffic. The protocol by itself does not guarantee QOS, but the priority scheme translates into higher throughput for high-priority classes by classifying traffic into aggregate traffic flows to receive better or worse treatment relative to each other.

1.4.3 MPLS

MPLS is a labeling scheme that is used by a network edge router to create paths across the network with specific constraints, such as acceptable packet delay. It is an IETF Committee-specified framework that provides for the efficient designation, routing, forwarding and switching of traffic flows through the network.

1.5 IP VPN Management

There are several ways an IP VPN can be managed, depending on the preference of the customer. Within each of these choices, a variety of CPE ownership options are typically available, depending on the provider. Management options include:

- Customer-managed: Under this scenario, the customer handles all aspects of managing the IP VPN equipment and service. This option is typically attractive to customers with

extensive internal capabilities and staff that prefer internal control to the convenience of having a vendor manage or partially manage their solution.

- Vendor-managed: With this type of service, the vendor handles all aspects of IP VPN management, including fault, configuration, security and CPE management, among others. This option is attractive to organizations wanting to rid themselves of the burdens of managing an IP VPN service and that are not concerned about handing the control of their network over to a provider. It would also be a good fit for businesses that do not have adequate or trained internal capabilities.
- Hybrid: This combined customer-managed and vendor-managed solution is attractive to enterprises that want to retain partial control of their IP VPN.
- Phased: With this type of approach, the customer initially wants to control all aspects of the IP VPN, but over time may become comfortable outsourcing more and more of the IP VPN package.

1.5.1 Portal

A portal is an Internet-based entrance point to various applications and services enabling IP VPN management and monitoring. Portals are an area where providers are differentiating themselves, especially with IP VPN services. Basic portal functions include the ability to view service-level agreements (SLAs) (and possibly check for SLA compliance); view, open and update trouble tickets; and view billing information. Certain providers are enabling users to create VPNs; increase performance parameters of the VPN (such as bandwidth); add, delete or modify user sites; or sign up for additional security features (such as PKI) or authentication.

1.6 SLAs

An SLA is a contract between the enterprise and provider that covers the services and equipment being provisioned and under management. Following are some of the typical components of an IP VPN SLA:

Provisioning Time: consists of the time it takes for local access provisioning; service provisioning (configuration and availability of IP VPN functionality); and in some cases, CPE delivery, configuration and installation.

Coverage: providers will either cover the service from customer edge (CE) to CE or from provider edge PE to PE.

Availability: measurement of the time that the network, connection or site is active, denoted in percentage. Availability figures can be quoted with or without local loop and CPE included. Savvis offers a Service Availability of 100 percent for the network core, 99.9 percent for a single edge; and 99.99 percent for a redundant edge.

Quality of Service (network-based IP VPNs only) levels are determined based on the following:

- **Jitter**: time between packets/variation of latency. Jitter is especially crucial in time-sensitive applications, such as voice and video. Sprint offers a 2-ms jitter guarantee for COS 1, COS 2 and COS 3 offered with its MPLS VPNs. AT&T offers a 1-ms jitter guarantee (in the U.S.) with the highest class of service for its MPLS Private Network Transport Service. Masergy Communications offers a 5-ms jitter guarantee for its Voice and Video COS offered with its Private IP and VPLS services.
- **Latency (delay)**: the total time for a packet to be sent from one location to another. Latency is typically measured in millisecond round-trip time. Round-trip delay will vary

depending on location and COS. AT&T offers a 39-ms node-to-node round-trip delay with its MPLS Private Network Transport Service (in the U.S.) BT Infonet's IP VPN Secure latency guarantee starts at 10 ms.

- **Packet Delivery**: throughput measurement indicating the percentage of packets delivered within a particular interval. Masergy offers a 100 percent packet delivery guarantee for voice and video traffic with its Private IP and VPLS services.

Fault notification: this is typically denoted as a time frame (within five minutes, 15 minutes, one hour and so on) and by method (pager, e-mail, telephone). Qwest offers notification within 10 minutes of an outage via e-mail, page or fax as part of its Private Routed Network service.

Mean Time to Restore (MTTR): the length of time it takes to fix a particular problem and restore service. MCI offers a MTTR of two hours (higher with access that is either partially provided [or not provided] by MCI) with its Private IP service.

Credits for noncompliance of an SLA metric can be either reactive or proactive. With reactive credits, the customer must request a refund by submitting a trouble ticket or similar notification within a certain amount of time following the violation. With proactive credits, the credits are applied automatically and do not require provider notification by the customer. In the U.S., credits are typically reactive; however, most international operators offer proactive credits.

2.0 Technology Analysis

2.1 Business Use

Enterprises are deploying IP VPNs in the following configurations:

- Remote access VPNs enable remote and mobile workers to securely access the corporate network using a variety of access methods, including DSL, cable modem, dial-up or wireless.
- Site-to-site (intranet) VPNs connect internal enterprise sites.
- Extranet VPNs connect a "community of interest" — a company, its partners, suppliers and customers — to an enterprise network.

2.2 Benefits and Risks

2.2.1 Benefits

- Access/remote access: Many enterprises are becoming decentralized, with workforces consisting of remote workers, mobile employees and branch-office locations. All of these employees may require access to central sites or applications. IP VPNs are an affordable and secure option for distributed enterprises.
- Cost: IP VPNs may reduce operating costs significantly on remote dial (local call vs. long distance call) site-to-site bandwidth connection charges, by eliminating redundant WAN connections, and in scaling of the network. Savings depend on a variety of factors, including the number of locations on the network (and the degree of meshing needed between sites) and port speed.
- Security (remote access): Remote access IP VPNs are more secure than Internet dial-up connections. Additionally, Layer 3 network-based IP VPNs are as and can be more secure as private-line or other Layer 2 packet networks as a result of authentication, encryption and tunneling features.

- Complexity: Scaling/adding WAN connections to a Layer 3 IP VPN is much less complex than with private-line or frame relay. Following a merger or acquisition, networks can be integrated quicker and sometimes easier using an IP VPN vs. legacy data technologies.
- Ubiquity: The Internet's vast reach enables IP VPN access from almost anywhere in the world.
- Consolidation: IP VPN architecture enables the enterprise network to be collapsed onto one platform for all applications (remote access, site-to-site and extranet) vs. multiple platforms and possibly multiple providers that may be required with legacy technologies.
- Convergence. Combining various types of application (voice, video, data) and Internet traffic onto one network connection reduces complexity and potentially the cost of managing multiple connections.
- Applications: Server-based applications that require enterprisewide distribution or access (such as customer relationship management, enterprise resource planning and supply-chain management) are transported more easily over an IP VPN than over legacy technologies that require protocol conversion.
- Extranets: IP VPNs enable the ability to securely connect to outside organizations (using IPsec or SSL connections.)
- Investment protection: Most carriers are offering enterprises the option of connecting their current access modality (frame relay, ATM, X.25) to the IP network, thus protecting their established CPE investment and minimizing the cost of IP migration/convergence.

2.2.2 Risks

- Interoperability issues: Extranet members using different firewalls or CPE may not be supported by some IP VPN packages. Interoperability with other MPLS networks is difficult due to the many different ways of implementing QOS/COS, management and traffic engineering.
- Capital expenditures: Installing an IP VPN may require a significant upfront investment in equipment and software and staff retraining. This may deter some enterprises that already have made significant capital investments in data networks.
- Resistance to change: Business may be hesitant to change due to the disruption that is likely to occur from changing/upgrading technologies.

3.0 Technology Alternatives

The primary alternatives to IP VPNs include:

- ATM
- Remote dial access to corporate resources via the long-distance network

4.0 Insight

There are an increasing number of enterprises that have found an IP VPN solution to be a better fit for their organization than their current legacy data services. According to Gartner’s forecasts, managed network-based IP VPNs (IP and MPLS) will enjoy a compound annual growth rate (CAGR) of 30.1 percent between 2004 and 2009, and managed CPE/premise-based IP VPNs will

have a CAGR of 8 percent during the same period. Unmanaged IP VPN services have a forecasted 7.35 percent CAGR for the 2004 to 2009 time frame. Understanding and selecting the correct IP VPN solution is a time-consuming task, but may be well worth the effort in terms of application performance, bandwidth use, cost reduction and network complexity. Enterprise customers should periodically conduct a thorough analysis of their business needs and technological requirements to ensure that their choice of network service(s) is meeting their strategic goals in terms of return on investment, cost containment and reduction, network resource accessibility and employee satisfaction.

RECOMMENDED READING

- Selection Process Critical for North American IP VPN Services
- IP VPN Services: Comparison Columns
- Quality of Service Over IP Networks
- An Introduction to MPLS

REGIONAL HEADQUARTERS

Corporate Headquarters
56 Top Gallant Road
Stamford, CT 06902-7700
U.S.A.
+1 203 964 0096

European Headquarters
Tamesis
The Glanty
Egham
Surrey, TW20 9AW
UNITED KINGDOM
+44 1784 431611

Asia/Pacific Headquarters
Level 7, 40 Miller Street
North Sydney
New South Wales 2060
AUSTRALIA
+61 2 9459 4600

Latin America Headquarters
Av. das Nações Unidas 12.551
9 andar—WTC
04578-903 São Paulo SP
BRAZIL
+55 11 3443 1509

Gartner