


Virtual Private Networks Resource Guide



**A Resource Guide for
Senior Telecommunications Engineers
and Department Managers of
Network Service Providers Worldwide**



Table of Contents

1. Executive Summary	1
2. The VPN Opportunity	4
3. VPN Essentials	15
4. Creating the VPN Infrastructure	24
5. Implementing VPNs	43
6. Appendices	50
VPN Application Example	50
Ascend Product Information	52
Reference Material	59
Comment/Information Request Form	60

Table of Diagrams

Figure 1.	The VPN Concept.....	2
Figure 2.	VPN Driving Forces	5
Figure 3.	The Conventional Private Network	6
Figure 4.	The Population Pyramid	9
Figure 5.	A Mobile Workers Accessing Corporate Resources.....	9
Figure 6.	Accessing Public and Private Resources.....	11
Figure 7.	An IP Multicast Training Class	11
Figure 8a.	Daily Traffic Loads Follow the Sun	13
Figure 8b.	Private Data Network Traffic	13
Figure 8c.	Today's Consumer Oriented Internet Traffic Load.....	14
Figure 8d.	Sum of Business and Consumer Oriented Traffic	14
Figure 9.	VPN Essentials	15
Figure 10.	GRE Packet Format	17
Figure 11.	ATMP Communications Flow	18
Figure 12.	End-to-End Tunneling.....	19
Figure 13.	Encrypted Tunnel Mode Packet.....	21
Figure 14.	Gateway Mode of Tunneling.....	24
Figure 15.	Router Mode of Tunneling	26
Figure 16.	VPN Building Blocks.....	28
Figure 17.	The Network Access Switch Building Block.....	29
Figure 18.	The Piecemeal Approach.....	30
Figure 19.	The Integrated Approach	31
Figure 20.	The CPE Building Block	32
Figure 21.	The Local WAN Services Building Block	33
Figure 22a.	Before xDSL.....	34
Figure 22b.	After xDSL	35
Figure 23.	The NSP Backbone/Internet Building Block.....	36
Figure 24.	The MegaPOP™ Architecture.....	38
Figure 25.	The Management Building Block.....	39
Figure 26.	Proxy RADIUS Configuration	41
Figure 27.	The CPE Building Block	43
Figure 28.	Major Site Connectivity	46
Figure 29.	ISDN Integrated Access Device Application.....	47
Figure 30.	The RADSL Alternative	47
Figure 31.	An End-to-End Virtual Private Network.....	49

Ascend Communications, Inc. is a leading, worldwide provider of remote networking solutions for corporate central sites, Internet Service Providers' points of presence, remote offices, mobile workers, and telecommuters. Ascend develops, manufactures, markets, sells and supports products that utilize bandwidth on demand to extend existing corporate networks for applications such as remote LAN access, Internet access, telecommuting, SOHO connectivity and videoconferencing/multimedia access. Detailed information on Ascend products, news announcements, seminars, service and support is available on Ascend's home page at the World Wide Web site: <http://www.ascend.com>.

Ascend markets the GRF, MAX, Multiband, Pipeline and Security families of products. Ascend products are available in more than 30 countries worldwide.

Ascend and the Ascend logos are registered trademarks and all Ascend product names are trademarks of Ascend Communications, Inc. Other brand and product names are trademarks of their respective holders.

Specifications are subject to change without notice.

1. Executive Summary

The Internet is an unprecedented success. Its user population will skyrocket from 50 million by the end of 1996 to more than 200 million by 2000 by some analysts' estimates. Early efforts by Network Service Providers (NSPs) to capitalize on this explosive growth have focused on consumer users. In terms of sheer numbers, this approach has been a great success: every 30 seconds, a new user joins the Internet. But in terms of revenue and profit, success is elusive with such a strategy. In fact, many Network Service Providers (NSPs) are actually losing money. Why? Because consumers, the typical NSP customers today, do not generate adequate revenue streams. They access the Internet during a relatively small window of time, and are not willing to pay very much when they do.

Most NSPs now realize that the Internet's profit potential lies with business subscribers. Companies, large and small, have started to appreciate just how valuable the Internet is for promoting and selling products and services, supporting customers, exchanging e-mail internally and externally, conducting research, collaborating with business partners and more. Many organizations now view the Internet as a more cost-effective alternative to private data networks. These "virtual private networks," or VPNs, represent a major opportunity – perhaps the most significant opportunity – for NSPs.

The essence of a VPN is its use of the Internet as a "public data network." Sending private data traffic via the public Internet is not much different than sending internal correspondence by mail, or faxing sensitive documents through the public switched telephone network. From the user's perspective, information sent simply arrives at their appropriate destination. Users should not need to, nor do they want to, take responsibility for the intervening infrastructure. It is this need to simplify that makes VPNs such a profound opportunity for NSPs.

A VPN can link all of an organization's offices, telecommuters (also called "telecommuters" outside of North America), traveling employees, and even its customers and suppliers around the globe. Owing to the Internet's worldwide presence, users just about anywhere can connect with a local phone call or leased line service. By eliminating long-distance charges, consolidating equipment needs and minimizing network management responsibilities, Forrester Research estimates companies can achieve a savings of up to 60% over private networks. The VPN also leverages user familiarity with the Internet and enhances overall flexibility. For these reasons, and others, VPNs offer businesses a more attractive solution to corporate data communication needs.

Virtual Private Networking: In Concept




Figure 1 – An Internet-based virtual private network makes corporate data communications more cost effective for companies, and more profitable for Network Service Providers.

A properly structured VPN offering reaps substantial rewards for the NSP as well. Much of the money a company saves with a VPN is internal. Correspondingly, a greater percentage of what companies spend for VPNs is external. With VPNs, NSPs have an opportunity to derive revenue not just from the traffic itself but also from network design consulting, end-user support services, the resale and management of customer premises equipment and more. Best of all, VPNs create long-term strategic relationships between corporate clients and the NSP. These partnerships open the door to many other lucrative revenue opportunities that range from the usual Web site hosting and content design to full-scale implementation of an organization's worldwide intranet.

Corporations will use VPNs to supplement or replace existing private network applications, and to implement new forms of communications. According to the International Data Corporation, 90% of organizations want to use the Internet to give employees remote access to internal information. Many also want to use VPNs for remote office internetworking and even enterprise-wide intranets. And the Internet's multicast capability will enable organizations to implement new collaborative work and distance learning applications. The almost limitless possibilities make VPNs the next logical step in corporate communications.

The enormous potential offered by VPNs has motivated NSPs to take action. Two-thirds of the NSPs surveyed by Infonetix Research are already planning VPN offerings. These NSPs recognize that with VPN-related technologies ready for mission-critical applications today, there is no longer an excuse for waiting. Encryption and authentication security provisions now put the "private" in virtual private networking. Tunneling and encapsulation techniques now allow the Internet Protocol (IP) to carry a wide range of popular non-IP traffic. Performance enhancements in the Internet backbone and access equipment now provide the throughput needed to compete with private networks. And all of these enabling technologies are based on standards that yield end-to-end interoperability. Finally, preparing Points of Presence (POPs) for VPNs is relatively simple and inexpensive. Low costs with high margins – VPNs are good business.

This planning guide can help NSPs understand, implement and sell a successful virtual private network offering. The target audience is any Network Service Provider (NSP), Internet Service Provider (ISP), Public Telephone and Telegraphs (PTTs), Local Exchange Carriers (LECs), InterExchange Carriers (IXCs), Competitive Access Providers (CAPs) and other providers interested in the profit potential of VPNs. Chapter 2 assesses the VPN opportunity with an overview of benefits and applications. Chapter 3 describes the enabling technologies and alternative techniques for implementing VPNs. Chapter 4 outlines the steps and systems needed to create the NSP's VPN infrastructure. And Chapter 5 contains practical information for implementing a customer's very own VPN. The Appendices provide useful supplemental material, including case studies and reference material available from Ascend.

One thing is clear: companies will use Internet-based VPNs extensively. The benefits are quite compelling and even irresistible. NSPs entering the market early stand to become the major beneficiaries of the VPN opportunity.

2. The VPN Opportunity

The success of the Internet is changing the way companies implement private data networking. Private networks and the Internet infrastructure now exist in parallel. The benefits to end users and NSPs alike are causing these “parallel universes” to converge as VPNs. Four forces are driving this convergence:

The increasingly dispersed and mobile workforce is making private networks unmanageable. Personnel who travel need dial-up access around the world, and more employees are now taking work home in the evenings. Such mobility requires at least two network connections for each worker. In addition to universal analog modem access, cellular data services such as the Global System for Mobile communications (GSM) are becoming ubiquitous. Full-time telecommuting arrangements dramatically increase the number of permanent “remote offices” a company must interconnect. Acquisitions, mergers and expansion add even more sites and nodes. As a result, many private networks have become unwieldy – and unmanageable.

The need to interact “on-line” with customers and suppliers adds a new dimension of complexity, where multiple private networks must be interfaced in a delicate balance of integration and isolation. The individual networks normally use different protocols, different applications, different carriers and different network management systems. With so few common denominators, interfacing two private networks can be a major challenge.

The desire to consolidate and simplify the user interface has become almost a business imperative under the relentless onslaught of networked applications. Users are unable to keep up with so many new applications, each with its own style and conventions. Users often find they lose more time becoming proficient in a new application than they can ever hope to gain becoming more productive.

The high cost of implementing and maintaining private networks has no light at the end of the tunnel. Long-distance charges for leased lines and switched services mount daily. The support staff required to manage the often-complex topologies involved continues to grow in both number and expertise. The dependence on networked applications requires separate backup and overflow provisions, further expanding the already burdensome private network infrastructure. And rather than provide relief, most new technology only makes private networks more complicated – and more expensive.

Virtual Private Network Driving Forces




Figure 2 – These driving forces are making Internet-based virtual private networks a compelling alternative to private networks for most businesses.

A Private Network Primer

Most organizations have private networks for at least some of their communications needs. These private networks carry exclusively data traffic, or a mix of voice/video and data traffic. They are constructed using a variety of Wide Area Network (WAN) services based primarily on the Public Switched Telephone Network (PSTN).

The typical private network employs high-speed leased lines that carry voice, fax, video and data traffic between major facilities. The Integrated Services Digital Network (ISDN) Primary Rate Interface (PRI) or conventional T1/E1 services both offer 1.544-2.048 Mbps of throughput. The 24-32 individual 64 Kbps channels integrate seamlessly with the PSTN, and can be used separately or in combination to handle anything from a single voice call to a videoconferencing session. Because leased lines are point-to-point, a mesh topology is needed to interconnect multiple facilities.

Conventional Private Network

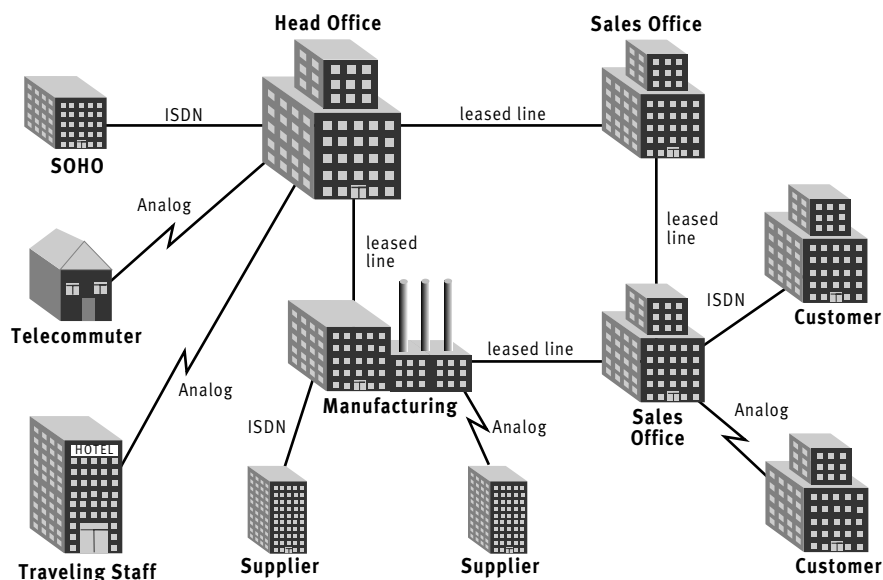


Figure 3 – The typical private network employs a range of WAN services to handle all of an organization's communications needs.

For remote offices, switched services via the PSTN generally offer the most cost-effective solution for voice, fax and video traffic. A sales office with a dozen users or more is often served by a single ISDN Basic Rate Interface (BRI) line at 128 Kbps. An alternative to ISDN is a Frame Relay line at 64 to 384 Kbps or an even higher data rate. Because Frame Relay is a packet-based network, it can handle multipoint communications – a distinct advantage where numerous offices are involved.

For individual users, either analog modems or ISDN BRI offer the optimal in price/performance. Mobile workers or other employees who travel regularly find the universality of an analog modem far outweighs the inconvenience of its lower speed. Tethered users working from their homes generally prefer the productivity boost provided by an ISDN BRI line, which can provide a complete voice/fax/data communications solution for the small office/home office (SOHO).

Leveraging the Internet with Virtual Private Networks

The VPN concept is not new. Indeed, the phrase “virtual private network” is also used to describe integrated voice/data networks offered by some NSPs. Although these offerings are packaged as VPNs, they resemble true private networks in both their costs and capabilities. A Frame Relay network could also be considered a VPN, since multiple users share a “public” network infrastructure. Rather than cover all forms of VPNs, this planning guide addresses only Internet-based VPNs, which hold the greatest potential for users and, therefore, service providers.

Technically, any private network could be considered “virtual” because it uses the public switched telephone network for both leased line and dial-up communications. But such a view is based on semantics and not on network characteristics or requirements, which are quite different for PSTN-based private networks and Internet-based virtual private networks. Requirements unique to Internet-based VPN exist in four key areas – compatibility, security, availability and interoperability – which are covered in detail in Chapter 3: VPN Essentials.

Private networks have been the most cost-effective way to implement enterprise-wide data communications – until now. Today, the Internet-based virtual private network offers many of the same capabilities as a private network, but at a fraction of the cost.

Essentially, a VPN is a private data network that uses a public data network to carry all traffic. The most ubiquitous, least expensive public data network – for now and the foreseeable future – is the Internet. The Internet, with its worldwide presence and affordable access, is indeed the perfect foundation for a VPN.

An Internet-based VPN is virtual because it appears to the organization as a dedicated private network, with exclusive use of the intermediate infrastructure, even though this is far from reality. In reality, of course, traffic from other VPNs and the Internet itself traverses the Internet infrastructure on a packet-by-packet basis. But it does so in such a way that ensures the appropriate traffic, and only the appropriate traffic, arrives at the appropriate, and only the appropriate, destinations. Because all users see only their own traffic, the network appears to be theirs – and theirs alone: a virtual private network.

Identifying VPN Applications

Additional candidates for VPNs:

- *Support full- and part-time telecommuting programs*
- *Handle all branch office interconnectivity on a dedicated VPN*
- *Move an existing application from the private network to a VPN*
- *Add sites not already on the private enterprise network*
- *Provide backup and overflow capacity for private networks using the Internet as a secondary “carrier”*
- *Perform overnight backup or software distribution of applications and/or data*
- *Institute virtual project teams with outside partners using inter-organizational “members only groupware”*

VPNs are suitable for a wide range of commercial data networking needs. An Internet-based VPN can:

- Replace existing private network segments or subnets
- Supplement private networks by offloading certain applications or meeting back-up/overflow needs
- Handle new applications without disturbing the existing private network
- Add new locations, especially international sites

There are three general conditions that favor use of a VPN:

- Numerous locations, particularly individual users and remote office sites
- Widespread users/sites involving long distances, including worldwide locations
- Relatively modest bandwidth and latency requirements

Conversely, there are three conditions where Internet-based VPNs may not be appropriate:

- Any situation where performance is vitally important – at any price
- Applications with unusual protocols that cannot be made compatible with the Internet Protocol (The various techniques for providing IP compatibility are covered in Chapter 3: VPN Essentials.)
- Applications where most traffic is isochronous, such as voice and video (A separate isochronous solution could be offered in parallel to the VPN, however.)

Three applications – remote access, enterprise-wide intranets and collaborative work – are particularly well-suited for Internet-based virtual private networks. Each provides both Internet and VPN access, and can be implemented as a small “pilot” or trial application to gain experience and confidence.

Population Pyramid




Figure 4 – The revenue-generating potential of a VPN increases with the number of sites.

Remote access for mobile workers traveling worldwide is an obvious use of the Internet. Rather than place a long-distance call to a centralized corporate facility, a traveling employee simply calls a local NSP Point of Presence (POP) to access the company's VPN via the Internet. The employee can now exchange e-mail, catch the latest news, update a price list, enter orders, and perform other tasks. The same arrangement can provide remote access for tethered employees as well.

Mobile Workers Accessing Corporate Resources




Figure 5 – Traveling workers using the ubiquitous Internet to access corporate resources.

End-user VPN Benefits

Organizations with VPNs are able to save up to 60% over equivalent private networks, according to a study by Forrester Research. VPNs save money because they:

- Eliminate long-distance leased lines among major facilities, including those needed for alternate “mesh” paths
- Eliminate long-distance switched calls via the PSTN for analog modems and ISDN access equipment
- Allow companies to pay only for actual usage with no idle lines or wasted Frame Relay Committed Information Rate (CIR) commitments
- Require less equipment (a single solution provides both Internet and VPN access, eliminating the need for separate modem banks, terminal adapters, remote access servers, and so on; the consolidation also permits utilization of cost-effective high-speed trunk lines)
- Minimize end-user network design and management responsibilities

VPNs leverage the Internet infrastructure’s built-in robustness to provide a more capable and dependable alternative to private networks:

- NSPs in nearly every city create a worldwide presence
- Local access improves throughput by minimizing line noise
- Mesh redundancy and fault tolerance afford end-to-end reliability
- User familiarity simplifies training needs

The Internet’s global presence also makes VPNs more flexible than private networks. With VPNs, end-user organizations can:

- Add and delete connections instantaneously
- Provide permanent, periodic or temporary connectivity as needed
- Integrate third-party users, such as customers and suppliers, almost effortlessly
- Select optimal data rates ranging from analog modem to T1/E1 speeds, and beyond with Digital Subscriber Line (DSL) technology

Because VPNs offer a more affordable, capable, dependable and flexible alternative to private networks, nearly all organizations surveyed by IDC and other industry analysts expect to leverage the Internet for internal data communications needs. VPNs are the next step in business communications.

Enterprise-wide intranets are a natural and symbiotic fit with the Internet. From a single network connection and a single application interface, users have access to both public and private resources. An employee can search the World Wide Web for publicly-available background information on a current project, then correlate this with private information on one of the company's Web-enabled servers. The result is easily posted for others to review before presentation to management.

Accessing Public and Private Resources




Figure 6 – An Internet-based VPN intranet gives users seamless access to both private and public resources – all from the friendly and familiar Web browser interface.

Collaborative work or distance learning/training applications are perfect for Internet-based VPNs. The Internet's powerful multicast capability can "broadcast" material to any number of sites, allowing users around the world to participate in a meeting or attend a training class from the convenience of their own offices. The multicast material could be as simple as a shared whiteboard, or as sophisticated as full-motion video and audio. Such applications, which are generally cost-prohibitive with private networks, can boost productivity substantially – and inexpensively – with an Internet-based VPN.

Multicast "Broadcast" to Multiple Sites




Figure 7 – By capitalizing on the Internet's worldwide multicast capability, the VPN lets users participate in a virtual meeting or attend a training class.

NSP VPN Benefits

Multicast in the Internet

Multicast capabilities will dramatically expand the applications potential of the Internet. The optimal way to handle IP multicast in the global Internet is a topic of intense scrutiny at this time. The Multicast Backbone (MBone) has proven itself quite successful, but now only covers a portion of the world. Basically, IP multicast employs the equivalent of a “group address” administered by the Internet Group Management Protocol (IGMP). Users wanting to participate in a multicast session register with the nearest multicast-capable router as a member of that session. This router then directs the multicast traffic to the unique IP addresses of its member participants. These “edge” routers must, in turn, register with other routers all the way back to the multicast source using either the Distance Vector Multicast Routing Protocol (DVMRP) or Protocol Independent Multicast (PIM). DVMRP and PIM allow multicast traffic to reach all participants without duplication of traffic on the mesh topology of the Internet itself. The technology is robust enough for an NSP’s IP.

VPNs provide a win/win opportunity for users and service providers alike. The benefits for users are readily apparent, but NSPs also stand to profit from VPNs. The raw revenue potential of VPNs as organizations migrate from private networks is tremendous, and the Return On Investment (ROI) is even better. Most NSPs experience peak traffic during only a limited period of a typical day. As a result, most of the resources needed to meet peak demand sit idle most of the time. A business-oriented VPN offering permits “around the clock” resource utilization – and revenue – with corporate users creating a peak during the day and individual users creating their usual peak in the evening. With the right balance of business and consumer subscribers, an NSP can double ROI with a VPN offering.

Additional benefits VPNs provide NSPs include:

- Ability to attract corporate clients
- Opportunity to establish long-term strategic relationships with large organizations
- A competitive edge against other NSPs – or an opportunity to establish mutually beneficial service agreements with other NSPs
- Leveraging of the existing infrastructure with little special investment
- High margins that offer quick payback on the minimal investment
- Revenue from end-user support and other value-added services
- Supplemental profits from reselling CPE

Daily Traffic Loads Follow the Sun

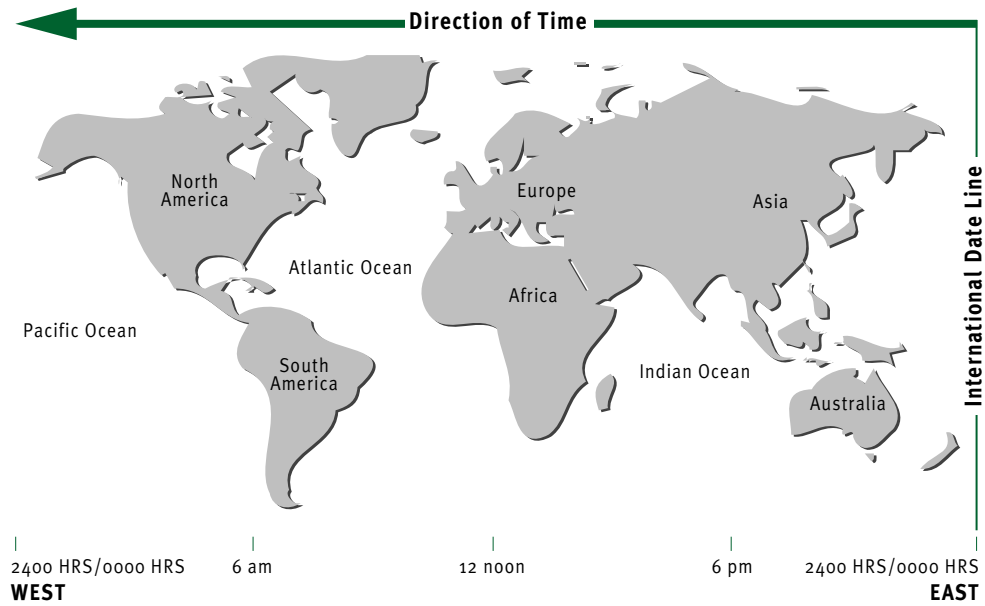


Figure 8a

Private Data Network Traffic (approximate)




Figure 8b

Today's Consumer Oriented Internet Traffic (approximate)




Figure 8c

Sum of Business & Consumer Oriented Traffic (approximate)




Figure 8d – Traffic – and revenue – flow for a full 18 hours a day through all POPs supporting both commercial (VPN) and consumer customers. The remaining six hours can be used for maintenance, or additional revenue-generating applications such as an VPN-based data backup service.

3. VPN Essentials

While there are numerous different requirements for any network, only four present special considerations with Internet-based VPNs: compatibility, security, availability and interoperability. Because private and virtual private networks are so similar, all other network needs are essentially identical. This chapter, therefore, focuses on the recent advents in the four special areas that now make the Internet robust enough for mission-critical virtual private networks.

VPN Essentials

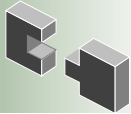




	Compatibility
	Security
	Availability
	Interoperability

Figure 9 – An Internet-based virtual private network has these four special requirements.

Compatibility



To use the Internet for a VPN, the private network must be made compatible with the Internet Protocol, or IP, at the International Standards Organization (ISO) Layer-3. The obvious way to do this is to use officially-assigned Internet addresses. And indeed, private networks with such addresses can use the Internet “as is” for a VPN, provided appropriate security measures are taken. But because most private networks use unofficial or “private” IP addresses, few can be used as is with the Internet. In fact, the vast majority of private networks are non-IP or private IP.

There are three proven options for making these private networks compatible with the Internet:

- Convert to Internet addresses
- Install special IP gateways
- Employ tunneling techniques

Private IP Networks

Many organizations that use IP have “private” IP addresses. The reason is simple: obtaining a block of official Internet addresses large enough to facilitate subnetting is impossible. Subnets simplify address administration and router/switch management, but “waste” precious addresses.

This practice is so common that a standard was published (RFC 1597) to set aside or sanction certain IP addresses for private use. Three such address blocks, or subnets, are available to suit any size organization:

- 10.0.0.0 – 10.255.255.255
(24 bits for nearly 17 million addresses)
- 172.16.0.0 – 172.31.255.255
(20 bits for about 1 million addresses)
- 192.168.0.0 – 192.168.255.255
(16 bits for slightly over 65,000 addresses)

Routers in the Internet block these sanctioned private IP addresses to avoid any ambiguity among private networks. Note that private IP address schemes that do not use the sanctioned private subnets must be carefully administered with similar blocking or filtering techniques. If two nodes have the identical address (one official, one not) serious problems can result for both users.

The next generation IP (IPng, which is also dubbed Version 6 or IPv6) will eliminate the need for private address schemes. IPv6 quadruples the current 32-bit address space. Even with waste from rampant subnetting, 128 bits is generous enough to yield 50,000 registered IP addresses for every square meter of land on Earth!

All three options make Internet-based VPNs more universal; selecting the best option(s) in each case depends on the organization’s current situation and long-term networking goals.

Internet addresses are the “official” IP addresses administered and assigned by the governing body, InterNIC. Of course, any user organization can simply select 32-bit IP addresses at random or as part of a rational scheme, and these addresses will work just fine in a private network. But these private IP addresses will not work in an Internet-based VPN (see sidebar on Private IP Networks).

This option is viable for organizations with existing private IP networks, and may even be suitable for those with non-IP networks. What is needed in either case is a compatible internal internetwork of local routers and switches. The internal internetwork can support multiple protocols, as long as it supports IP.

Converting an organization’s entire network to official Internet addresses is unnecessary for a limited VPN. Even if such enterprise-wide conversion is desirable or inevitable, most organizations will likely choose to wait for the next generation IP Version 6 to avoid double work: converting now with IPv4 and again with IPv6.

A less ambitious approach is to convert only those clients and servers in the VPN. All VPN servers should be given a permanent Internet address, but VPN clients can use temporary ones from a common pool of Internet addresses. The Dynamic Host Configuration Protocol (DHCP) and/or Network Address Translation (NAT) allow organizations to “lease” an Internet address temporarily to any client normally assigned a private IP address. After the VPN session is over, the Internet address is returned to the DHCP or NAT pool for use by others.

The official Internet addresses coexist with private IP addresses on the organization’s internal internetwork of routers and switches. In other words, a “private” IP client can still access an “official” IP server via the local internetwork with no special provisions.

IP Gateways are the second option for making private networks compatible with the Internet Protocol. A gateway works by translating another protocol to IP, and vice versa. Because the typical IP gateway operates at ISO Layer-3, it should be called an IP relay – technically speaking. But the popular terminology is, and will likely remain, the IP gateway.

Normally, an IP gateway supports the clients assigned to a particular server. The server has a Network Operating System (NOS) with a “native” protocol. The gateway converts traffic to/from the native protocol from/to IP. For example, Novell NetWare clients running browser-like applications on Novell’s IPX protocol can access Web servers on the Internet via one of the many gateways available. IP gateways used for Internet access can be used without modification for Internet-based VPNs. The gateway application can run on the NOS-based server itself, or on a separate server, perhaps a dedicated device, with a different operating system.

Tunneling is the third, and generally the best, option for making private networks Internet-compatible. Tunneling protocols and various encapsulation techniques have been used for years to integrate different network protocols on a common backbone. These proven technologies have now been optimized for use with Internet-based VPNs.

GRE Packet Format

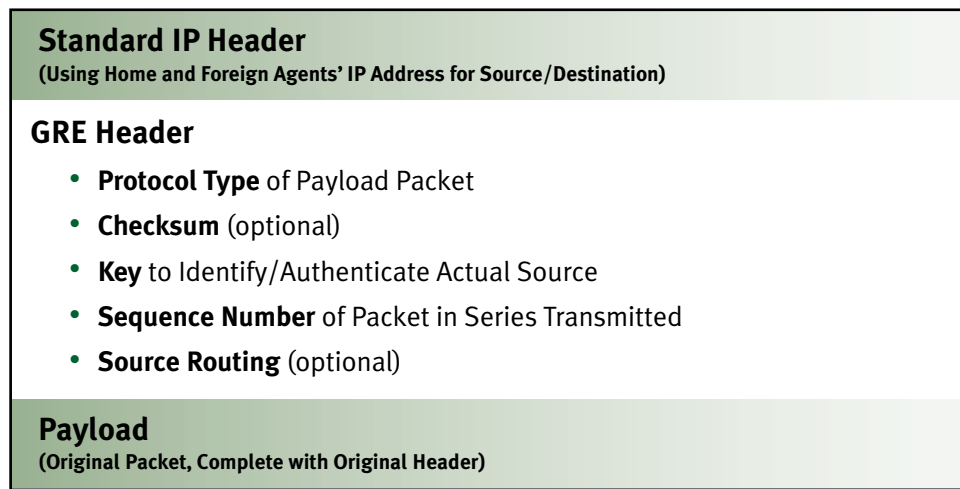


Figure 10 – Packet encapsulation techniques, such as Generic Routing Encapsulation, add special headers to the original packets for tunneling via the Internet.

Tunneling occurs at both ends of a connection. The source end encapsulates the other protocol's packets in IP packets for transit across the Internet. The encapsulation process consists in the addition of a standard IP header; the original packet is the payload. A corresponding process at the destination end decapsulates the IP packet (removes the IP header), leaving the original packet (the payload) intact. (See the sidebar on How Tunneling Works for a more detailed description of the process.)

Because tunneling is relatively simple, it is often the most cost-effective and easiest to manage alternative for making virtually any private network, including a private IP network, operate as an Internet-based virtual private network. Another advantage of tunneling is that it can be implemented in the NSP's POP or in customer premises equipment (CPE) – or in a combination of POP and CPE. Many network access switches, remote access servers and WAN routers already support interoperable tunneling standards (see sidebar on Tunneling Protocols).

How Tunneling Works

ATMP (RFC 2107) Communications Flow

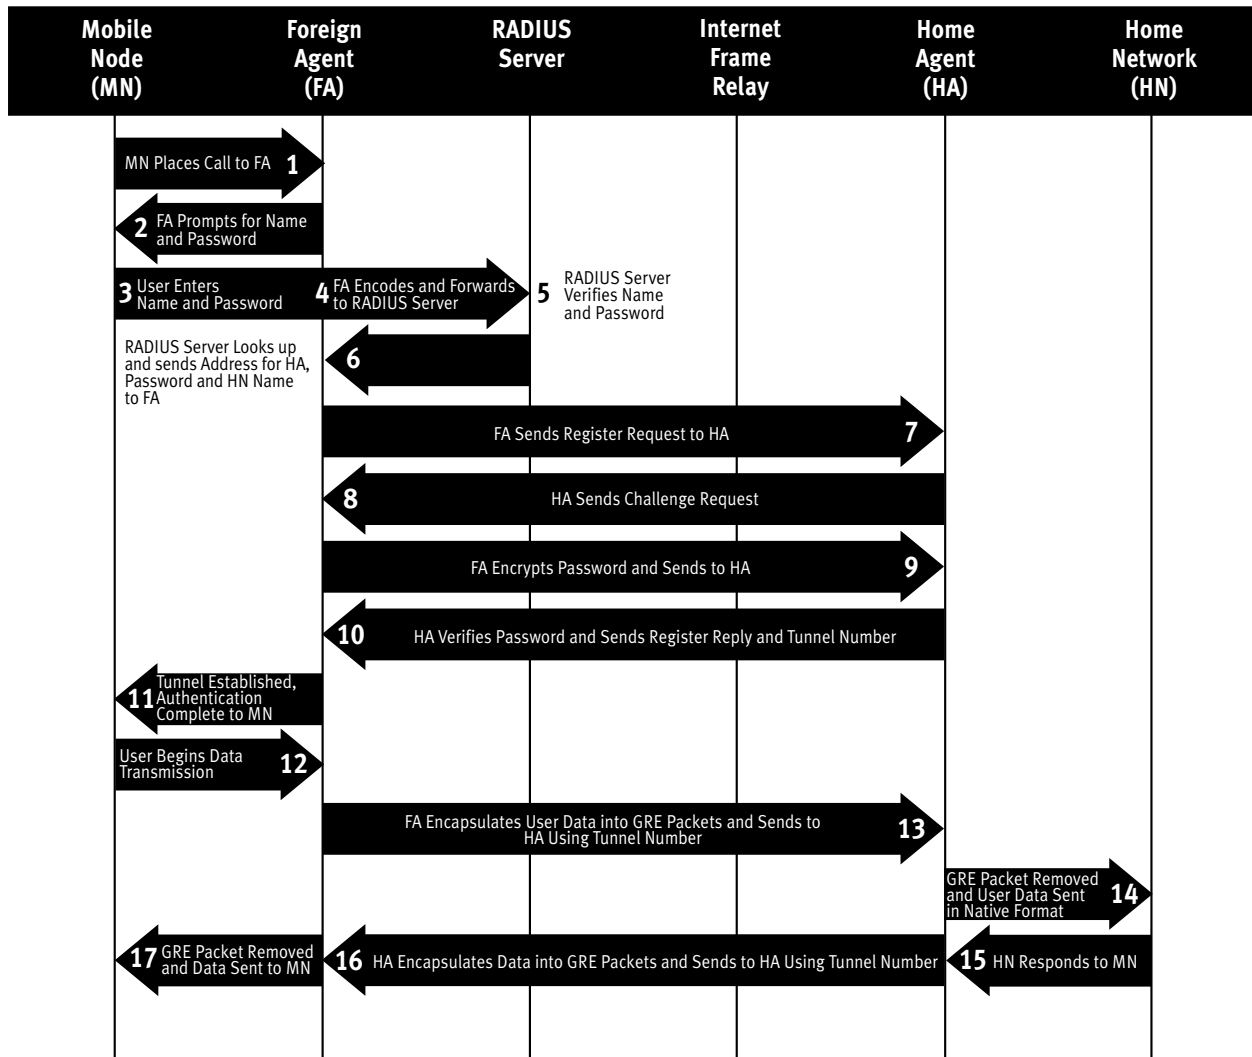


Figure 11 – This flow chart shows the detailed handshaking exchange of the GRE-based Ascend Tunnel Management Protocol (ATMP).

The end-to-end tunneling process has four elements:

- The Mobile Node (MN) is the remote client or server initiating the VPN session. (Mobile Nodes may be stationary, i.e. attached to a LAN, or truly mobile, i.e. traveling employees.)
- The Foreign Agent (FA) resides in the network access equipment at the Mobile Node's site or NSP POP.
- The Home Network (HN) is the private network containing the resources the Mobile Node wishes to access.
- The Home Agent (HA) resides in the network access equipment at the Mobile Node's Home Network site or NSP POP.

How Tunneling Works

Tunneled packets are sent across the Internet from agent to agent using each agent's Internet address in the header to designate source and destination, depending on the direction. The source agent (Home or Foreign) creates the tunnel's header; the destination agent (Foreign or Home) removes the tunnel header, and delivers the original packet to the Mobile Node or the Home Network, respectively. The tunnels can either be static or dynamic. Static tunnels, which remain active for extended periods of time, are acceptable for site-to-site VPNs. Dynamic tunnels are activated only as traffic requires, and are, therefore, more secure.

The location of the agents determines where tunnels originate and terminate. This example shows tunnels originating and terminating in the network access switches, which are located either at the NSP's POPs or at the customer's sites. With the Point-to-Point Tunneling Protocol (PPTP), however, the Home Agent that terminates the tunnel is always in the server (Windows NT or NetWare). The Foreign Agent that originates the tunnel can be in one of two locations: in the client workstation or in the network access switch.

End-to-End Tunneling

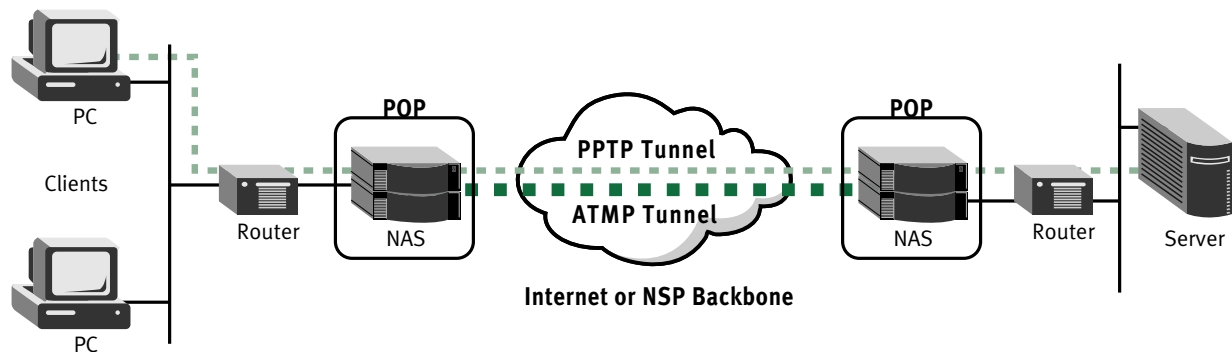


Figure 12 – Tunnels originate and terminate in different places, depending on the location of the Home and Foreign Agents. Shown here is an example of POP-to-POP tunneling with ATMP, and client-to-server tunneling with PPTP.

Mobile IP is similar to PPTP, except that the mobile node is referred to as the Mobile Host, which can also act as its own Foreign Agent. The Home Agent on the Home Network is responsible for forwarding (tunneling) traffic to the Mobile Host at its temporary location.

Tunneling Protocols: Making the Virtual Paths in Virtual Private Networks

- *Generic Routing Encapsulation (GRE), as defined in RFCs 1701/1702, is used by a wide variety of tunneling protocols.*
- *The Point-to-Point Tunneling Protocol (PPTP), created by Microsoft and Ascend Communications, is an extension to the Point-to-Point Protocol (PPP) for Windows NT and NetWare client/server environments.*
- *The Ascend Tunnel Management Protocol (ATMP), defined as RFC 2107, implements both GRE and PPTP for tunneling IP, IPX, NetBIOS and NetBEUI traffic.*
- *Layer-2 Forwarding (L2F) is a tunneling protocol created by Cisco Systems.*
- *The Layer-2 Tunneling Protocol (L2TP) is a proposed industry standard that will combine the best features of both L2F and PPTP.*
- *Data Link Switching (DLSw), originally defined by IBM and now an industry standard, encapsulates SNA traffic (the LU 6.2 protocol) in IP.*
- *Mobile IP is designed to tunnel IP within IP for individuals traveling away from their “home” network, but can also be used to tunnel private IP traffic in a VPN.*
- *IPsec supports tunneling with or without encryption.*

Security



Security puts the “private” in virtual private networks. Providing adequate security is often the primary concern for organizations considering use of an Internet-based VPN. Many Information Systems managers have become accustomed to the inherent privacy afforded by private networks, and may consider the Internet “too” public for private networking. With the proper security provisions, however, the public Internet can be made just as private as the public switched telephone network.

The fundamental concern is that private information could be accessed while in transit or directly from servers/hosts. Several robust security measures are now available to keep transmitted data strictly confidential, as well as to prevent unauthorized users from gaining access to network-attached resources (see sidebar on Techniques for Securing a VPN).

VPN security provisions should meet the following three objectives:

- **Provide adequate security** – *A minimal security system should validate users with passwords to protect VPN-accessible resources from unauthorized access. The addition of encryption protects traffic in transit. Other levels of security are available, and the more levels provided, the more secure the VPN becomes.*
- **Provide ease of administration** – *The VPN security provisions chosen should be easy to both set up initially and maintain over time. The security system’s administrative functions must also be secure from tampering by users.*
- **Be transparent to users** – *Even legitimate users may attempt to circumvent security methods that are difficult to use, so the security system should make logging onto the VPN as easy as logging on at a LAN-attached workstation.*

Techniques for Securing a VPN

Users can choose from a number of options for securing a VPN (relevant standards, where applicable, are listed):

- Password identification affords the minimum level of protection (Password Authorization Protocol, or PAP).
- Challenge Handshake Authentication Protocol (CHAP) is superior because passwords are never transmitted across the line during the authentication.
- Token cards offer virtually “bullet-proof” authentication with single-use passwords.
- Authorization grants authenticated users permitted access only.
- Encryption keeps data strictly confidential during transit through the Internet.
- Dynamic firewalls control access to internal hosts, networks and applications, and are used to isolate private resources from public ones, such as an organization’s World Wide Web server.

Encryption is critical to keeping VPN-transmitted data confidential. The IP Security or IPsec standard embodies several proven technologies to provide robust encryption. For protecting the payload, which is the entire original packet being tunneled, IPsec employs the Encapsulating Security Protocol (ESP) and the Data Encryption Standard (DES). Other encryption algorithms are also supported so long as they are available to both ends of the connection. To provide additional protection IPsec creates an authentication header, which uses Message Digest (MD5) or a Secure Hash Algorithm (SHA), effectively providing a separately encrypted “certificate of authenticity” for the packet. This effectively provides a digital signature on every packet. Packets can be encrypted, authenticated, tunneled, not tunneled, or any combination of these.

For managing security provisions, the industry-standard RADIUS (Remote Authentication Dial-In User Service) database maintains user profiles that contain passwords (authentication) and access privileges (authorization). For managing the keys that “lock” and “unlock” encrypted packets, IPsec uses a combination of the Internet Security Association Key Management Protocol (ISAKMP) and the Oakley Key Determination Protocol.

Tunnel Mode Packet




Figure 13 – Encryption “scrambles” each and every packet transmitted so that only the authorized user, with the appropriate key can read the content. An optional Authentication Header (AH) provides an additional layer of security that prevents an ongoing session from being hijacked, i.e., having someone in the middle of the transmission take over the session by pretending that it was the actual source of traffic at the remote end.

Additional security options may be supported by some NSPs, such as calling line ID and callback. Certain other security standards being proposed for Internet commerce are not designed for VPNs; these include the Secure Sockets Layer, Secure-HTTP, Secure-MIME, Secure Electronic Transaction (SET) and Private Communications Technology (PCT).

Availability



Availability encompasses both up-time and throughput, and private networks assure certain levels of service in both dimensions. Dedicated leased lines operate at their full capacity around the clock; Frame Relay’s Committed Information Rate (CIR) offers trustworthy throughput; and dial-up access via the PSTN is quite dependable in most regions of the world.

The Internet has no such service level assurances today. But despite the occasional outage hyped in headlines, the Internet is nearly as dependable as most private networks. And how could it not be? The Internet is constructed using the very same facilities employed in all private networks: high-speed leased lines, Frame Relay, PSTN dial-up access, routers, IP switches, and so on.

Short-Term VPN Performance Boosters

- *ISDN bandwidth on demand with the Multilink Protocol (MP) and the Multilink Protocol Plus™ (MP+)*
- *Digital modem technology to improve analog modem performance, including use of new asymmetric 56 Kbps modems*
- *Digital Subscriber Lines (DSL) for high-speed continuous access*
- *Standard STAC or other data compression*
- *IP multicast for “multipoint” applications*
- *Frame Relay Direct to channel tunnels through virtual circuits in the NSP’s Frame Relay backbone, rather than route traffic unnecessarily as IP packets*
- *CPE with integral PSTN-based dial backup and overflow provisions*

Of course, the Internet exists on a much larger scale that is being expanded and re-engineered relentlessly without the benefit of a single organization overseeing all of this activity, so it may seem to be less dependable than most private networks. On the other hand, however, the Internet has much more robust end-to-end redundancy and resiliency than a typical private network. Both serve to prevent Internet-wide catastrophes, leaving service problems isolated to the occasional local brownout or outage hyped in headlines. Because the Internet is unregulated, with no formal accounting of its overall up-time, it is impossible to say which is more dependable: the typical private network or a VPN. But it is reasonable to claim that the Internet offers reliability sufficient for the vast majority of business applications.

With an Internet-based VPN, businesses share the existing capacity of the Internet. The throughput assurance of a private network is overkill most of the day, and comes at a premium price. And when throughput is considered in the context of cost, the price/performance advantage of a VPN becomes quite compelling.

In the short-term, organizations can maximize VPN throughput by selecting appropriate WAN services and using data compression techniques. A full ISDN BRI line to the local NSP POP, compressed 4:1, offers a throughput of up to 512 Kbps. In addition, the digital modem technology at most POPs worldwide combined with the superior quality of local vs. long-distance calls, allows analog modems to operate at top-rated speeds, potentially giving dial-in users better performance with the VPN. Finally, NSPs with their own IP backbones may be able to offer throughput assurances for “service area” VPNs. (See sidebar on Short-Term VPN Performance Boosters for other ideas.)

Quality of Service

Quality of Service, or QoS, involves a network’s ability to assure a specified level of performance end-to-end. QoS provides either a guaranteed amount of bandwidth or a never-to-exceed latency, or a combination of the two. The PSTN offers both bandwidth and latency assurances today with a switched virtual circuit that grants exclusive use of a channel or multiple channels (at 64 Kbps each) for the duration of a call. Currently, the Internet delivers data on a “best effort” basis. Within the next few years, the Internet will offer support for QoS through enhancements like RSVP and RTP. In the meantime, NSPs with private backbones can provide QoS for VPNs using Frame Relay’s Committed Information Rate (CIR) or other techniques.

In the long-term, technology advances will permit service level assurances within the Internet infrastructure itself. IP switching combined with higher speed WAN services are expanding the Internet’s overall capacity. Emerging standards like the Resource ReSerVation Protocol (RSVP) will add support for quality of service (see sidebar). Other advances, like the Real Time Protocol (RTP), will broaden a VPN’s capabilities. Although not in place today, these future enhancements provide comfort for organizations considering Internet-based VPNs in their strategic networking plans.

Interoperability



Implementations of the first three VPN requirements raise a fourth one: interoperability. Although standards exist for providing VPN compatibility, security and availability, unspecified details and other nuances preclude a guarantee of multivendor interoperability at this early stage. The tunneling, authentication, encryption and performance-related standards, mentioned earlier, are new or emerging. Many are not fully robust yet, enticing vendors to add some desirable enhancements.

Early adopters of VPN technology must, therefore, pay careful attention to end-to-end interoperability. This responsibility can reside either with the end user organization or with the NSP, depending on how a particular VPN is implemented (see the discussion on VPN architectures in Chapter 4: Creating the VPN Infrastructure).

One way to assure interoperability is to select a single-vendor solution. If no single vendor can meet all requirements, limit the interoperable aspects to an essential few, and use equipment that has been lab-tested or field-proven conformant to the applicable standards. In either case, select a vendor or vendors fully committed to VPN standards. And buy only equipment that can be upgraded inexpensively with software, firmware or plug-in modules to comply with future enhancements to emerging standards.

4. Creating the VPN Infrastructure

Just as all private networks are unique in some respects, so too will all VPNs vary from one organization to another. Creating a robust VPN infrastructure requires an understanding of these potential differences. This chapter outlines the three alternative VPN architectures, presents checklists of considerations for establishing a VPN offering, and describes the five VPN building blocks, including a list of requirements for each.

VPN Architectures

There are three fundamental architectures for VPNs: fully outsourced, in-house, and a hybrid combination of the outsourced and in-house alternatives. The difference involves where the four basic VPN requirements are implemented: at the NSP's POPs and/or at the end-user's facilities.

An outsourced VPN exists when the NSP POP provides the complete VPN solution. NSP POP equipment handles all tunneling (normally used with outsourced VPNs), security, performance and interoperability requirements for the end-user organization. The advantage to end-users is that they can employ the existing network infrastructure “as is,” except for the possible addition of a firewall between the VPN and the existing private WAN/LAN.

ATMP (RFC 2107) Gateway Mode of Tunneling

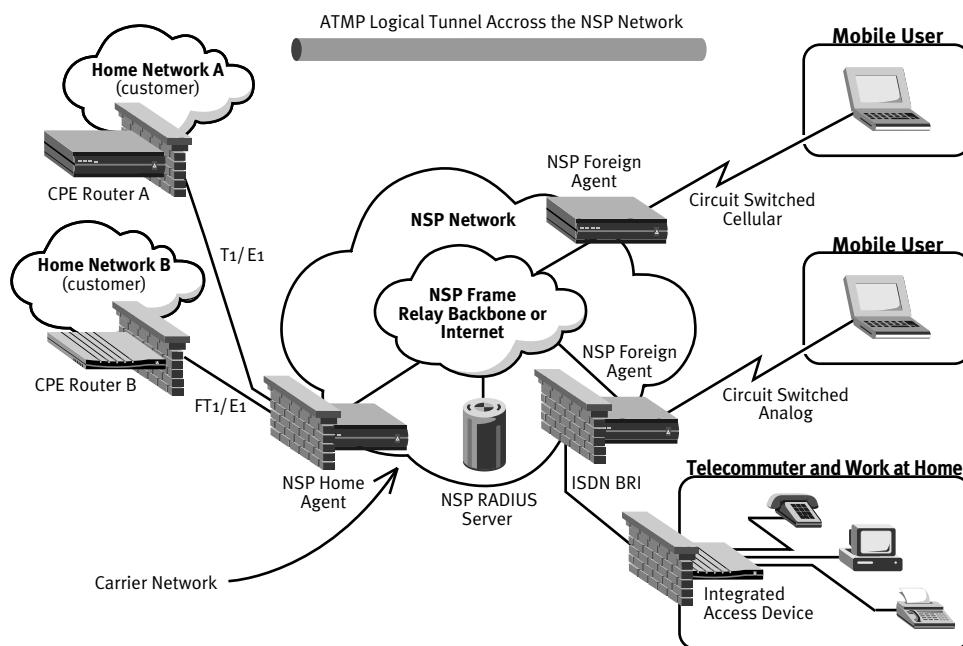


Figure 14 – With an outsourced VPN, the customer can use existing WAN internetworking equipment with the addition of a firewall. When using gateway mode tunneling the tunnel is built across the NSP network, between the NSP edge Routers.

Outsourced VPNs work as follows: All VPN sites interface with their respective local POP using an appropriate WAN service. When a user logs on, the NSP network access equipment queries the customer's RADIUS database to obtain the user's password, access privileges and tunneling parameters. All traffic from and to the user is encapsulated and decapsulated at the local POP, which is sometimes referred to as the Gateway Tunneling Mode. The tunneling process, along with the entire intervening NSP POP and Internet infrastructure, are totally invisible to the user, who sees only native-mode traffic, such as IP, IPX or NetBEUI.

With an outsourced VPN the end-user organization can – and probably should – administer all user security and access capabilities. And, this is just as easily done as said: the NSP's RADIUS server acts as a proxy for the customer's RADIUS server, which contains the database of user profiles. This server is managed and secured by the end user organization, not the ISP. The alternative would be for the customer to turn over all necessary internal information – with regular updates – which may actually reduce the effectiveness of security. Gathering, “publishing” and distributing employee lists with sensitive security-related information is a risk-filled endeavor. Maintaining user profiles on a RADIUS database is a task best kept in-house, even with an outsourced VPN.

Offering a fully outsourced VPN means more responsibility for the NSP, but it will generate substantially more revenue. Outsourcing also creates a long-term strategic relationship with the customer. Much of the VPN-specific needs are available as add-on options to popular POP equipment. The only real requirement is providing POPs in, or POP access for, all current and planned VPN sites. For customer locations where the NSP has no POP, there are two solutions: coordinate the customer's VPN with another NSP in a strategic partnership, or place VPN-capable CPE at the customer's site (see hybrid architecture discussion below).

An in-house VPN is where the end-user organization handles all VPN requirements on its own CPE, relegating the NSP to an Internet “carrier” role. The NSP sees only Internet traffic – and its associated revenue – and does not concern itself whether the traffic is for the Internet alone or for the customer's VPN.

With an in-house VPN, all participating sites exchange IP traffic with the local POP. If tunneling is employed, all traffic is encapsulated and decapsulated at the user site, which is sometimes referred to as the Router Tunneling Mode.

The in-house approach is ideal for organizations that believe VPNs are too new to outsource to an NSP. In-house VPNs allow the customer to maintain direct control over day-to-day operations, and the total costs – both internal and external – may not be that much more than an outsourced VPN. NSPs can, of course, still provide complementary services, and the in-house architecture can migrate to a hybrid one or even to a fully outsourced one over time.

ATMP (RFC 2107) Router Mode of Tunneling

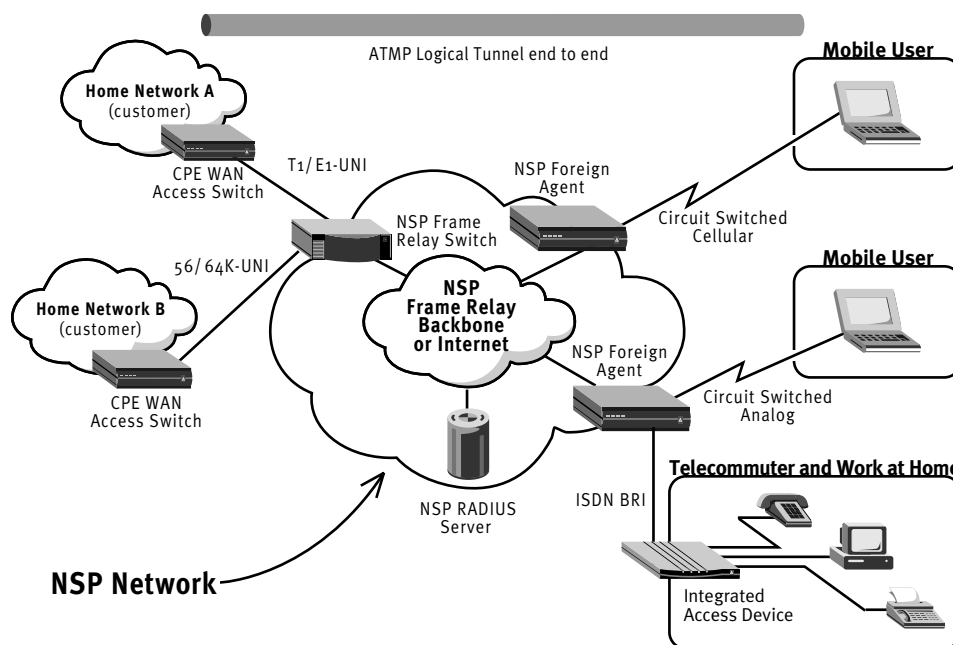


Figure 15 – With an in-house VPN, the NSP's role is that of an Internet "carrier" hauling IP traffic. In Router mode the tunnel is built end to end across the network.

The hybrid VPN involves a combination of outsourced and in-house VPN sites. Each site is one or the other, making the total virtual private network a hybrid. While the hybrid approach is, technically, an architecture, it can also be an element of an outsourced VPN offering. For example, a customer may have some sites beyond the NSP's service area, but wants an outsourced solution. No problem. The NSP installs VPN-capable network access switches at these sites, and the customer's monthly charge is increased to include a lease payment for all of this equipment. The systems can even be managed remotely by the NSP, again as part of a fully outsourced solution. The VPN-equipped sites then access the Internet, and hence the Internet-based VPN, through another service provider's POP. The other service provider handles the traffic just as it would any other IP traffic, and might even send its monthly statement to the "primary" NSP for consolidated billing, so the customer sees only one invoice. From an Engineer's perspective, the result is a hybrid VPN. But from the customer's perspective, it is a fully outsourced solution.

Planning the VPN Offering

Planning a VPN offering has four basic elements; monitoring the results adds a fifth. This checklist itemizes some of the considerations involved in getting a VPN offering off to the right start. Similar experience with other programs will provide additional items for consideration.

1. Determine the Desired Offering

- Outsourced or in-house VPN or both/hybrid
- Up-front network design/implementation consulting services
- Resale or partner with a reseller for CPE
- On-going remote office/user support services
- Other services or value-add
- Limit VPNs to the service area backbone or utilize the full Internet for worldwide reach
- Create a list of equipment/upgrades needed to implement the offering
- Is there an opportunity to implement the offering in phases?

2. Formulate the Rate Structure (NOTE: Because setting rate structures is a strategic and sometimes regulated process, these checklist items are offered merely as ideas for thought.)

- By packet, connect time and bandwidth, fixed monthly fee, or other means
- A fixed monthly fee may allow for highest margins because the customer is replacing or supplementing an expensive private network, and desires an arrangement that can be budgeted easily
- When setting rates keep in mind that customers are replacing expensive and burdensome private networks with the VPN
- But also keep in mind that some customers will need to move voice communications from the private network to the PSTN
- Billing options (distributed or centralized)

3. Conduct an ROI Analysis


- Assess all possible business opportunities in the service area market
- Conversely, evaluate the potential for lost business by not offering certain capabilities and options
- Are the total costs in line with anticipated revenues?
- Does the plan minimize costs by making maximum use of the existing POP/backbone infrastructure?
- If the ROI seems low (it should not be), review the offering itself and/or the rate structure

4. Launch the VPN Offering

- Resolve business integration issues for consistency with mission and objectives
- Implement and test VPN infrastructure or first phase offering
- Publish promotional materials (and place on Web server)
- Create sales tools, such as a slide presentation and a proposal “boilerplate”
- (Optional) implement a beta VPN for an existing customer
- Train sales and support personnel

5. Monitor the Program

- Is the offering attractive to prospective customers?
- Is the rate structure appealing?
- Is the VPN infrastructure working as expected?
- Are revenues in line with the forecast?
- Is the competition taking business away? If so, why?
- Are any changes needed?



VPN Building Blocks

All virtual private networks are constructed using the five fundamental building blocks depicted in the diagram. The requirements for each are discussed, in turn, in this section.

The Five VPN Building Blocks





Figure 16 – Every virtual private network is made up of these five fundamental building blocks.

The Network Access Switch (NAS) is the heart of the NSP Point of Presence (POP). A capable, carrier-class network access switch provides a fully integrated POP solution in a single chassis. Here is a list of capabilities desirable in a network access switch:

- Variety of WAN access options including T1/E1, ISDN PRI/BRI, xDSL, DS-3, analog modems, cellular, Frame Relay and X.25
- Digital modem technology for better performance and compatibility with a broad assortment of analog modems, including new asymmetric 56 Kbps modems
- High-speed PSTN trunk lines for optimal channel utilization with switched traffic provisions for direct connection to the Internet or NSP backbone in POPs that do not require IP switching
- Ability to handle a variety of LAN options, such as Ethernet, Fast Ethernet and the Fiber Distributed Data Interface (FDDI)
- Adequate security provisions, including encryption and authentication
- Support for PPTP, L2TP, Mobile IP and/or other GRE-based tunneling
- Support for IP Direct and Frame Relay Direct to channel tunneled packets through a virtual circuit within the NSP's IP or Frame Relay backbone, rather than route traffic unnecessarily onto the Internet itself
- Firewall technology integrated into CPE or NSP equipment.



Network Access Switch Building Block

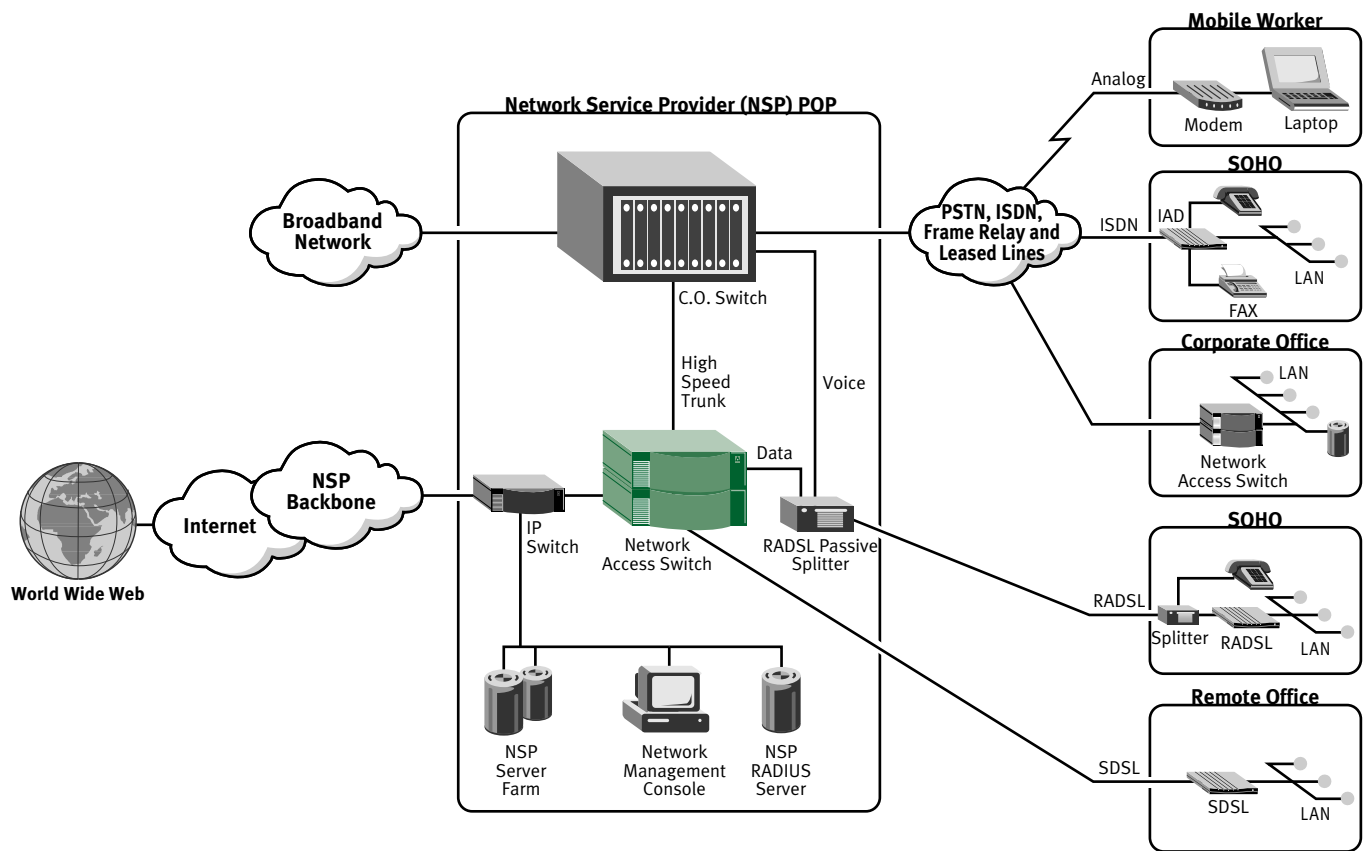



Figure 17 – All five building blocks are represented in this local view, or one end, of a VPN. The network access switch is the core component of the complete end-to-end virtual private network.

- Built-in compression to maximize throughput
- Dynamic bandwidth management for enhanced performance
- Ability to handle IP multicast applications
- Software upgradable to conform with emerging tunneling and security standards
- Remote download of software upgrades, via the Internet or NSP backbone
- Robust local and remote management to maximize uptime at minimal cost
- Call detail reporting (CDR) to track usage for virtually any rate structure, as well as to log attempted security violations
- RADIUS database support for administering security and accounting
- Resiliency with dual power supplies and hot-swappable interface cards
- Adequate capacity to support anticipated traffic volumes
- Compatible family of scalable products to suit a variety of POP sizes
- Certification for operation with local carriers



The Piecemeal Approach vs. the Integrated Approach

Piecemeal Solution of Network Access

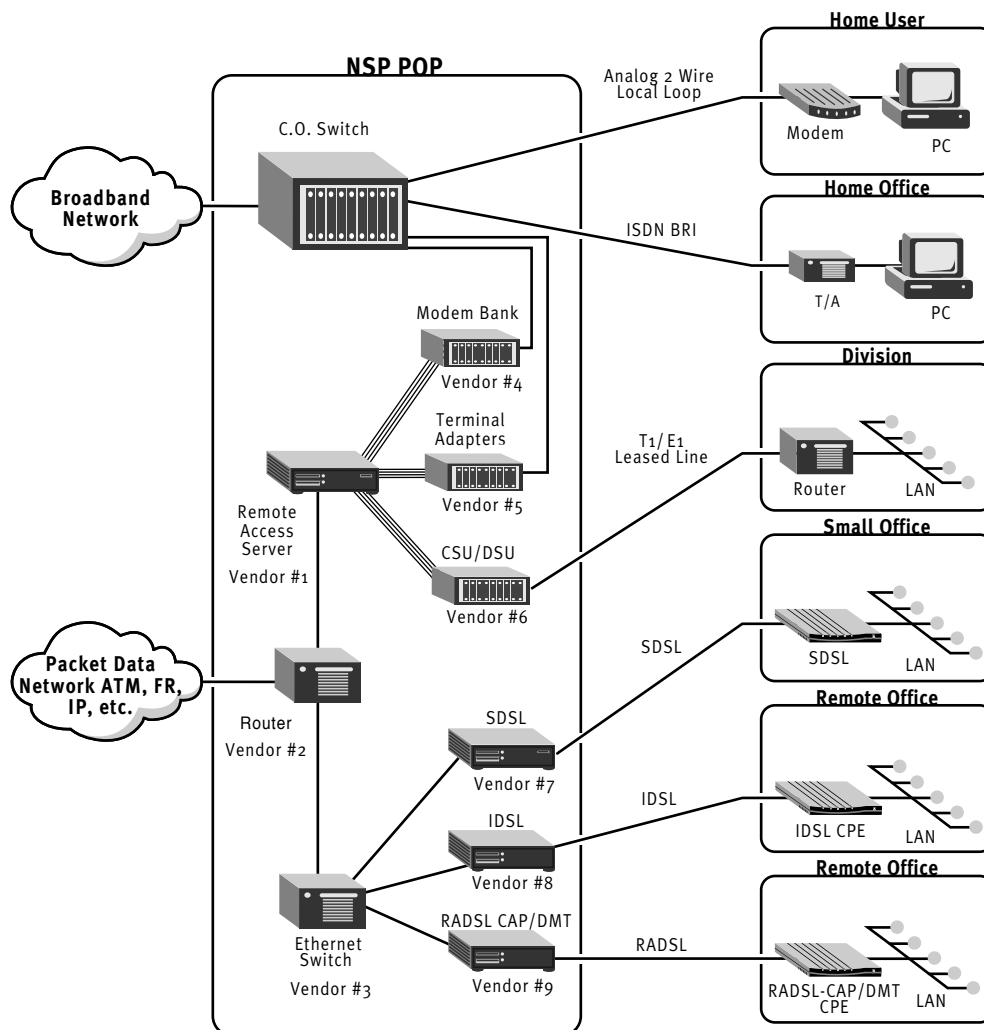


Figure 18 – Multiple access technologies, vendors and solutions make for a considerable management headache.

Integrated Approach




Figure 19 – The network access switch is a more capable, flexible, scalable and manageable replacement for the old-fashioned piecemeal assortment of equipment in the POP.

CPE Building Block

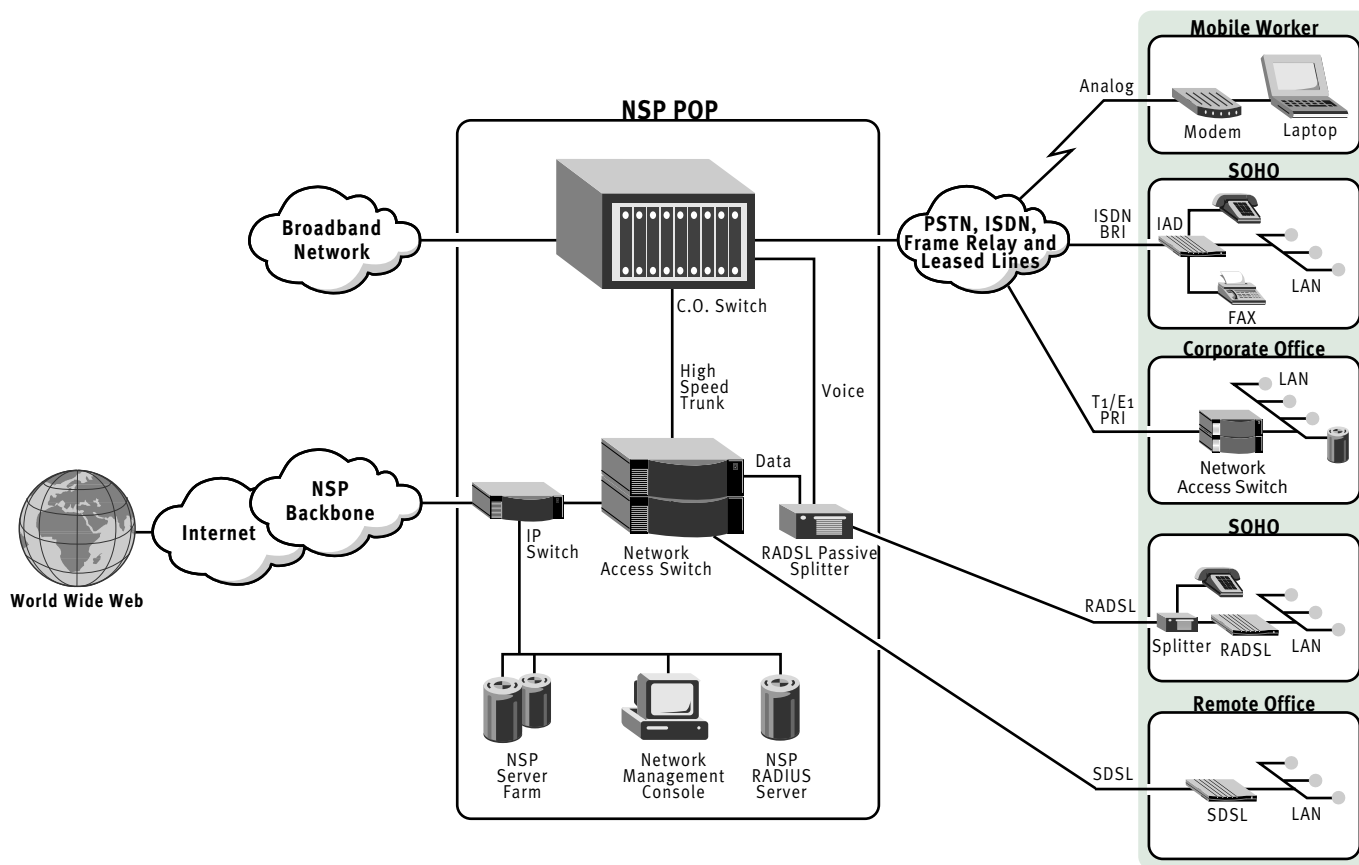


Figure 20 – Customer premises equipment defines the “edge” of the VPN.

Customer Premises Equipment (CPE) is covered in detail in Chapter 5: Implementing VPNs. As a fundamental building block, however, some level of understanding is relevant here. The type of CPE required depends on the VPN architecture. With an out-sourced VPN, the customer can use existing WAN equipment, such as routers or remote access servers. An in-house VPN architecture, on the other hand, requires CPE that supports many of the same features needed on a network access switch at the NSP POP, such as provisions for tunneling and security.

CPE raises two issues for the NSP. The first is how to handle situations where a VPN customer needs CPE. A typical example involves customer sites outside of the NSP's service area. The most expedient way to implement a VPN for these sites is to install CPE compatible with the NSP's network access switch. In these situations, and others, the NSP should either resell (or lease) the CPE directly, or work in partnership with a reseller that does. The first issue creates the second: selecting a compatible solution of both CPE and POP equipment. In other words, when evaluating equipment for the POP, be certain the vendor also offers a family of suitable – and compatible – CPE.

There is a third issue, as well. It is inevitable that NSPs will need to provide at least some level of end-user support in VPNs, even for those with in-house architectures. Why not turn this “problem” into a profitable element of the VPN offering? Selling and supporting CPE can generate substantial revenue. It also limits the permutations and combinations of end-to-end configurations to a manageable few. Take control of the entire VPN, end-to-end, and profit in the process.

Local WAN Services Building Block

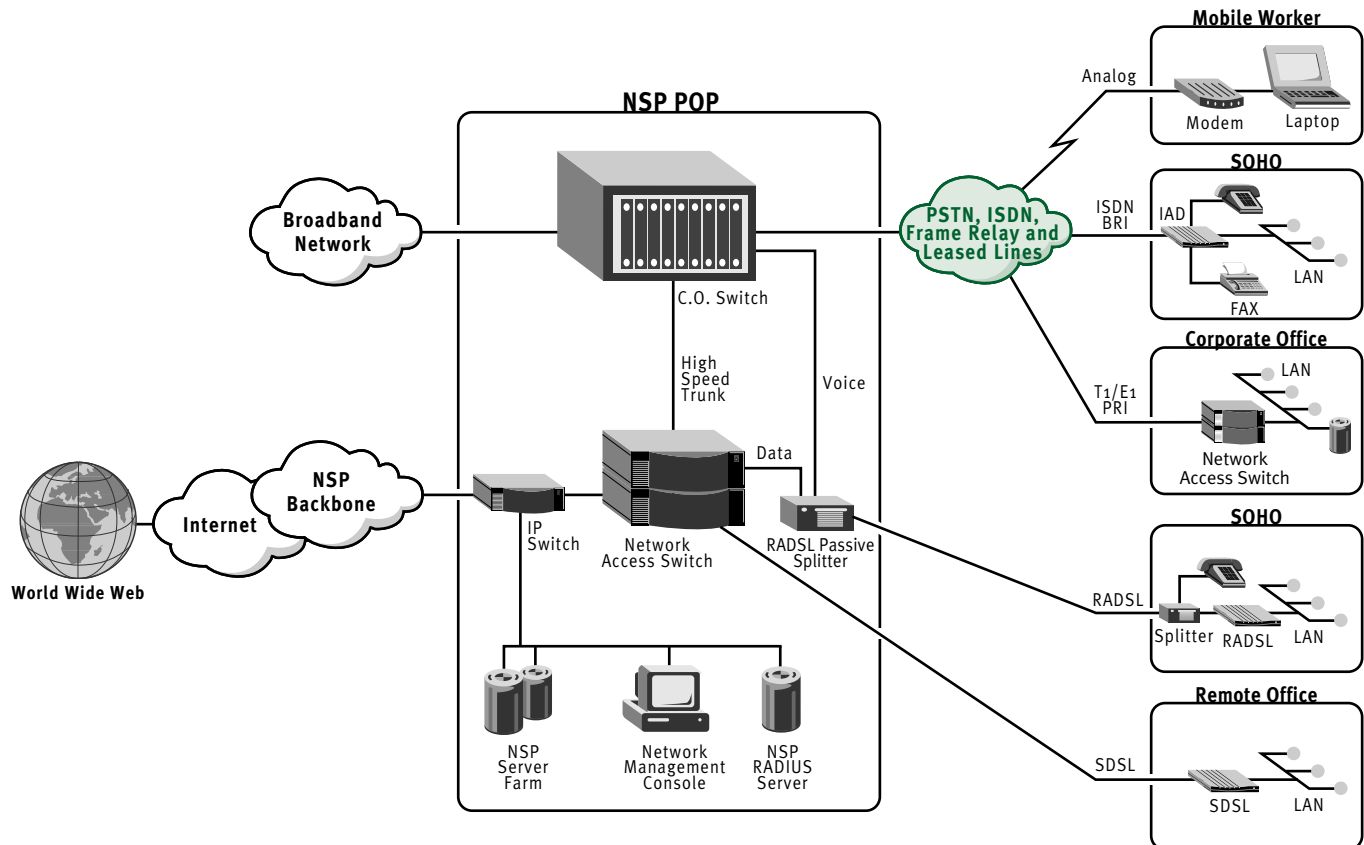


Figure 21 – Local WAN Services connect the customer’s users and sites to the VPN at the POPs.

Local WAN Access Services (also known as WAN Services) connect the CPE with the NAS at the POP. There are essentially two choices:

- Dial-up services, such as analog and ISDN, which are best for traveling employees and telecommuters, respectively
- Continuous forms of access, such as that provided by leased lines or Digital Subscriber Lines, which are best for multi-user offices

Beyond these two broad categories, choosing the best alternative is really only a matter of speed: How much throughput does the site need? When replacing a private network with the VPN, the current WAN access services used should be sufficient. If the private network also handles voice communications, an appropriate lower speed WAN access service option can be used.



Before DSL




Figure 22a – Data traffic in general is contributing to congestion on PSTN switches.

Digital Subscriber Lines

Digital Subscriber Line (xDSL) technology increases the throughput of ordinary twisted pair wiring in the local loop. Voice telephone services use this same wiring, but employ analog signaling methods that severely limit bandwidth. xDSL technologies achieve higher transmission speeds by using line frequencies up to 1.2 MHz, rather than the 3.4 KHz used by ordinary analog voice systems. By utilizing advanced digital signal processing techniques, similar to those used for ISDN and T1/E1 today, xDSL is able to achieve throughput as high as 7 Mbps.

An xDSL link, in effect, creates a high-speed “leased line” between the central office and customer site, which is ideal for a VPN. There are many different variations on the xDSL theme; however, the three versions that can utilize existing twisted pair wiring are:

- ISDN Digital Subscriber Line (IDSL), an Ascend innovation, delivers 128 Kbps performance and offers compatibility with existing ISDN access equipment.
- Symmetric Digital Subscriber Line (SDSL) furnishes 768 Kbps of throughput as a cost-effective alternative to leased lines.
- Rate Adaptive Asymmetric Digital Subscriber Line (RADSL) integrates lifeline analog voice with high-speed digital data for a total communications solution on a single pair of wiring. RADSL is available in Carrier Amplitude/Phase (CAP) and Discrete Multi-Tone (DMT) options that provide 544 Kbps - 1 Mbps in the upstream direction (from the subscriber) and 640 Kbps - 7 Mbps in the downstream direction, where bandwidth is needed the most.

After DSL

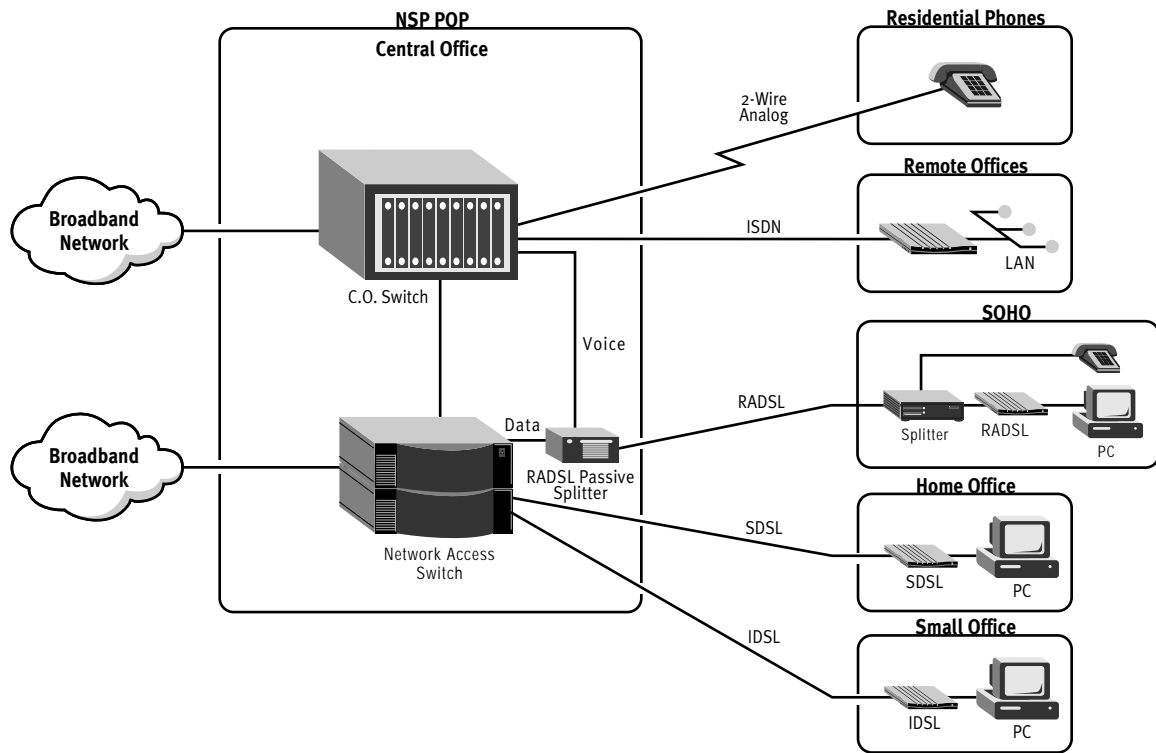


Figure 22b – DSL provides high-speed, continuous access while relieving the PSTN of analog modem and ISDN traffic.

For new applications, the customer will need to supply an estimate of the traffic volume for each VPN site. Be certain to get these estimates in writing. A customer that invests in special CPE only to find out immediately that it is undersized, should have no reason to blame the NSP. Here is another example of where providing a turnkey solution can be to the NSP's advantage. If the unit is too small, merely exchange it for a more powerful one, charging the customer only for the difference. To minimize the occurrence of such opportunities-to-solve-customer-problems, however, it is always good practice to review a customer's estimates up front.

NSP Backbone/Internet Building Block

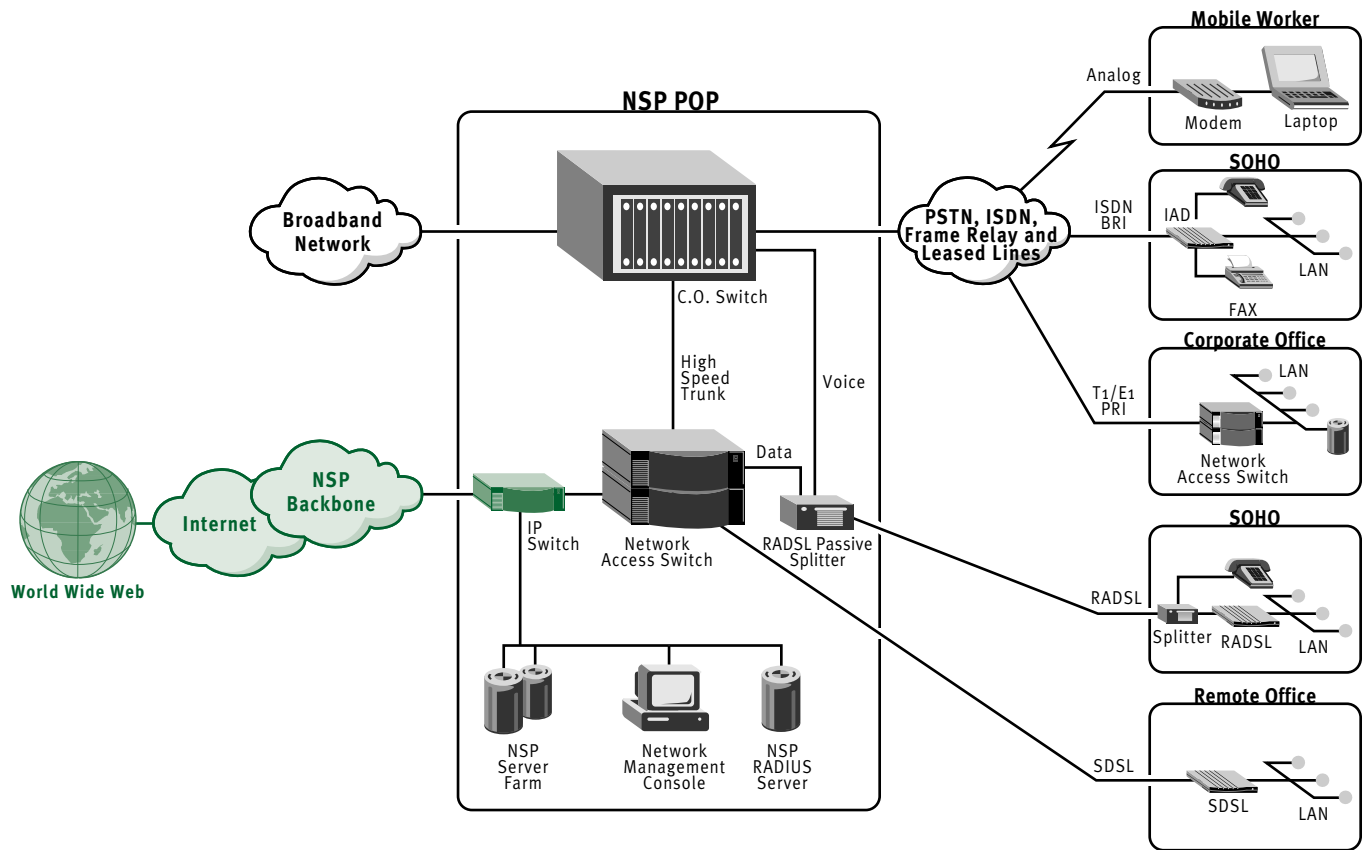



Figure 23 – The NSP’s backbone and the Internet itself provide the “long distance” communications for the virtual private networks.

The Internet and/or NSP Backbone constitute the circulatory system of the VPN, fed by the NAS “hearts” pumping at all POPs. The volume of traffic flow requires either a high-performance router or an IP switch. Because the conventional router architecture is no longer able to accommodate the Internet’s growth, causing both worldwide congestion and local service outages, IP switching technology is now the preferred choice. Small and medium-sized POPs will still be able to interface to the Internet or NSP backbone using the integral routing capability of the network access switch. Larger POPs, however, are better served by an IP switch with the following features:

- Compatibility with existing network infrastructures, including interoperability with conventional routers and LAN switches
- Full compliance with industry routing standards, such as RIP1/2, BGP4, EGP, OSPF, IS-IS and IP multicast to eliminate any need for proprietary gateways or special client software
- Next-hop address lookup fast enough to take advantage of the switching engine’s low latency and high throughput



- Route table capacity of greater than 100,000 next-hop routes to keep pace with anticipated Internet infrastructure growth
- Linear scalability with no performance degradation
- Overall capacity to support a sufficient number of LAN/WAN ports
- Wire-speed performance for all LAN/WAN ports with sustainable throughput that is independent of traffic characteristics, such as flows and cache hits
- Support for a wide range of popular LAN and WAN media, such as 10/100 Mbps Ethernet, FDDI, HIPPI, HSSI, OC-3c ATM (ATM STM-1) and SONET/SDH, and OC-12c ATM (ATM STM-4)
- Packaging in a small chassis able to fit into limited POP space
- Built-in resiliency, including dual power supplies and hot-swappable interface cards
- Bandwidth required to accommodate high-powered applications

High-performance IP switching will solve one of the biggest paradoxes in the Internet. Today, with the limited performance of traditional IP routers, the only way to grow the Internet's overall capacity is to add more and more routers. But the proliferation of routers causes next-hop route tables to grow exponentially. Updating these enormous route tables consumes Internet capacity; processing them brings many routers to a grinding halt. The solution is a more streamlined infrastructure that adds capacity with more ports on IP switches rather than with more next-hop nodes.



The next-generation MegaPOP Architecture

The Internet market today is characterized by a constant need to increase both access port and backbone network capacity to meet the requirements of the increasing numbers of users, while staring in the face of declining subscription revenue. NSPs want to change this unprofitable scenario in both halves of the ROI equation with business-oriented services, such as VPNs, along with a more efficient and cost-effective network architecture. That architecture is what Ascend has dubbed the next-generation MegaPOP™, an integrated solution that combines IP switching with high-capacity remote access for analog modems, ISDN, Frame Relay and DSL technologies.

The next-generation MegaPOP configuration affords two major benefits to NSPs. The first is a new and enduring architecture that allows for the separation of voice and data traffic using Digital Subscriber Lines and other techniques. The split preserves the PSTN for voice, postponing or eliminating costly upgrades, and lets the resulting parallel data-only backbone take advantage of IP switching for dramatically improved throughput.

The second benefit is the increased profit potential generated by a combination of lower costs and higher revenues. NSPs can implement and grow a MegaPOP in profitable steps. And because VPNs substitute for expensive PSTN-based private networks, they can command a premium price.

MegaPOP Architecture

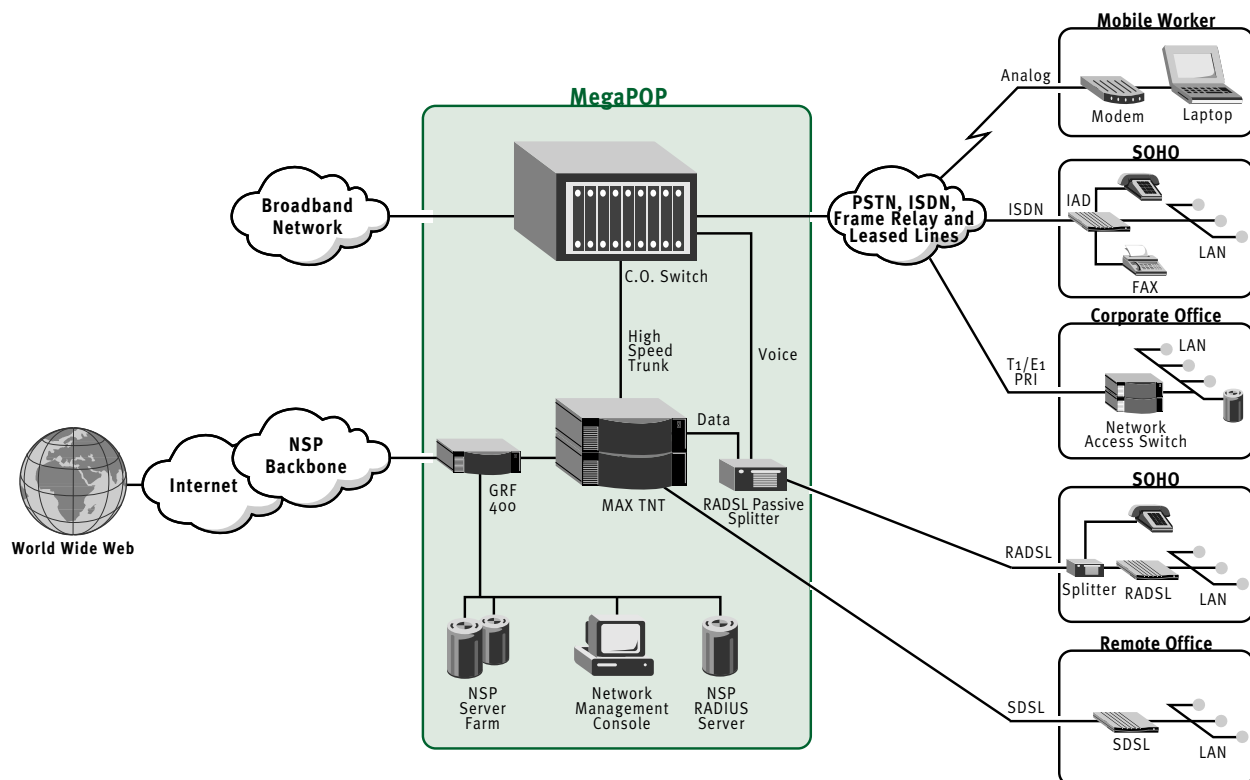



Figure 24 – The next-generation MegaPOP architecture provides the most capable, manageable and profitable solution for large-scale Internet access and Internet-based VPNs. An entire POP supporting thousands of users fits easily into a single rack.



Management Building Block

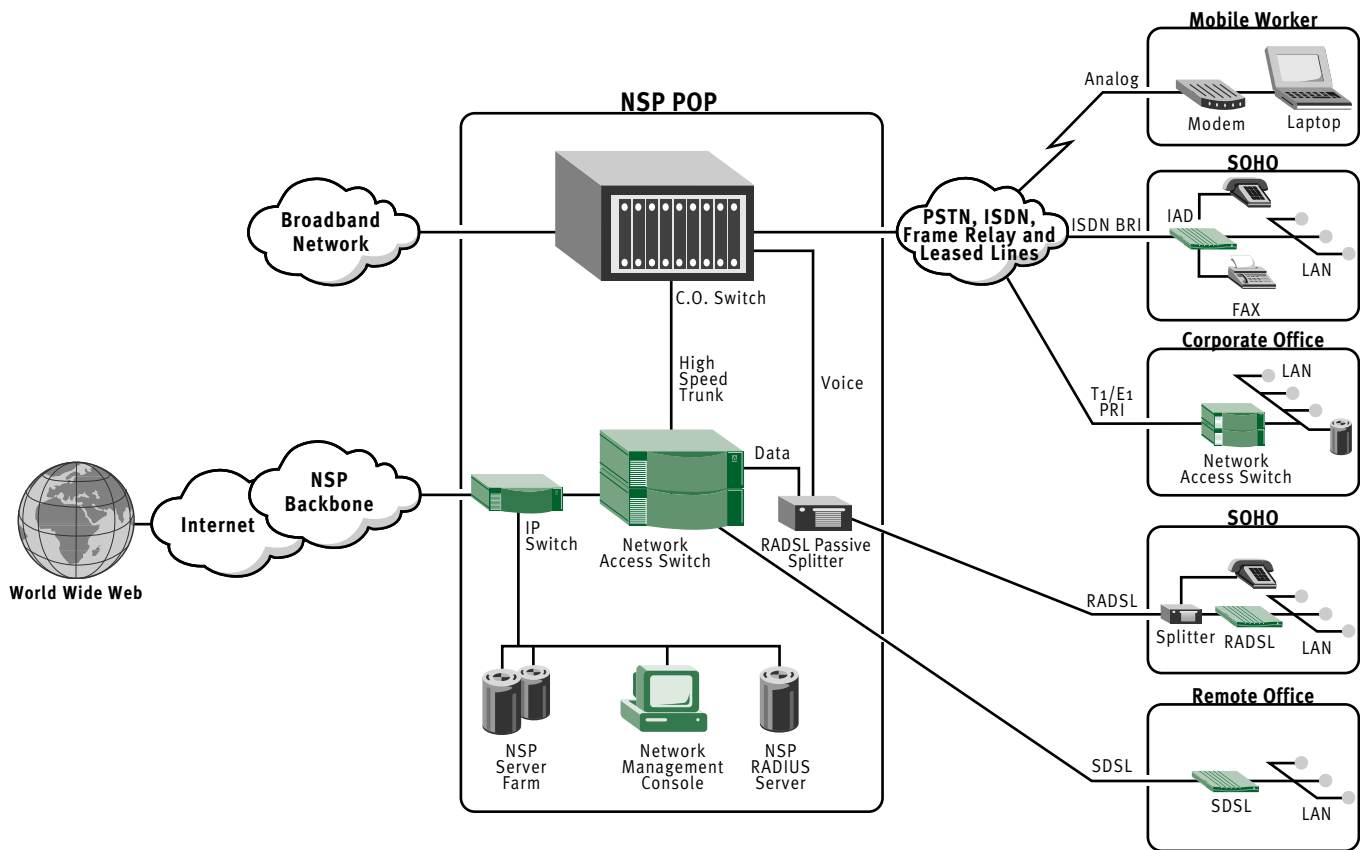



Figure 25 – Various tools are needed to manage the entire VPN infrastructure, as well as each customer's individual virtual private network. Agents are also needed in each piece of networking equipment.

The Management Tools needed by an NSP must encompass many networks: the one at each POP, the private backbone among multiple POPs, the “edge” of those at all customer sites connected in VPNs, the interface to the Internet, and, to some extent, the intervening WAN infrastructure. This seemingly daunting task is vastly simplified by first selecting easily-managed equipment, then adding the right collection of tools. These tools fall into three categories:

- Vendor-supplied utilities for managing the network access switch, IP switch and customer premises equipment
- Diagnostic and troubleshooting aids for both real-time traffic analysis and capacity planning based on usage patterns
- A database for security and accounting administration

The vendor-supplied utilities for managing the end-to-end POP/VPN configuration should be considered when evaluating each vendor's equipment offering. The tasks involved include installing, configuring, monitoring and troubleshooting the network equipment and all interconnections.



Look for the following minimum features:

- Support for industry standards, like the Simple Network Management Protocol (SNMP)
- Applications that operate on the same workstation or management platform
- A means to compare actual vs. intended equipment configurations
- Auto-discovery and mapping of the network topology for the entire enterprise
- Alarm generation based on user-definable thresholds
- An ability to identify a particular customer's VPN
- A way to coordinate network-wide upgrades
- Ability to gather aggregate data for multiple devices in both logical and geographic groups
- Ability to monitor access devices (modems, ISDN devices, remote access servers) throughout the network
- Capability to monitor multiple vendor products

The equipment vendors may also provide diagnostic or capacity planning tools, but be sure to investigate third party solutions as well. An NSP's management tool kit should contain these capabilities, at a minimum:

- Real-time monitoring of traffic conditions to give an early warning of imminent problems
- Traffic monitoring also offers a way to assess actual throughput on WAN Access lines, and helps control delivery of contracted Quality of Service (QoS) levels
- Historical data gathering and reporting to help determine overall network "health" and for capacity planning needs
- An end-to-end trace function that tracks traffic through the network to assist in pinpointing problems and spotting bottlenecks

For administering security and accounting, there is one clear winner: the Remote Authentication User Dial-In Service (RADIUS) standard. RADIUS handles all security, accounting and other administrative needs with individual profiles for all users. The user profiles, combined with call detail reporting from the network access switch, give an NSP all the information necessary to provide adequate VPN security, manage subscriber accounts and generate invoices. RADIUS employs a client/server architecture with network access switches at NSP POPs as the clients, and the RADIUS database as the server. Multiple RADIUS servers, such as one at the customer's site for security and one at a central POP for accounting, are also supported. This ability to handle distributed management with centralized control makes RADIUS ideal for VPNs of any architecture.

The Role of RADIUS in VPNs

- RADIUS (Remote Authentication Dial-In User Service) is a standard for maintaining authentication, authorization, accounting and auditing information for remote access networks, including virtual private networks
- RADIUS employs a client/server architecture: RADIUS is the server; the network access equipment becomes its clients
- “Proxy RADIUS” lets one server become a proxy for accessing another RADIUS server at any location in the network (see diagram)
- Normally both the NSP and customer have RADIUS servers
- The NSP’s RADIUS server is used for accounting and auditing, i.e. it tracks and logs all network usage (both Internet and VPN) by all users
- The customer’s RADIUS server is used for authentication and authorization, i.e. it is the repository for individual profiles of all Mobile Nodes (workstations and hosts) in the VPN
- Each profile identifies a particular node’s Home Agent and Home Network, as well as any applicable user password(s) and access privileges
- For workers who travel (truly mobile Mobile Nodes), separate profiles are used for dial-out and dial-in capabilities
- Support for ODBC-compliant databases such as Oracle and Sybase
- The ability to use authentication schemes such as token cards (Security Dynamics, AssureNet Pathways) as well as other types of authentication schemes such as TACACS+, Kerberos and Windows NT domains
- IP pools so that multiple RADIUS clients can draw from one pool of IP addresses

Proxy RADIUS Configuration





Figure 26 – The primary RADIUS server can be used as proxy for reaching other RADIUS servers, anywhere in the network, that contain the information needed. A common use of Proxy RADIUS is for a server at the NSP POP to access user profiles on a server at the customer’s site.



Creating the VPN Infrastructure

This checklist itemizes the major steps involved in creating the VPN infrastructure. Similar experience with other undertakings can provide additional steps or a deeper level of detail.

- Design the POP/backbone network
 - Create a list of equipment requirements for the POPs, and optionally for eventual CPE, that includes both necessary and desirable features
 - Evaluate vendor offerings and select the equipment to be used
 - Install the network access switches in the POPs
 - Install the IP switches in the backbone or for the backhaul connection
 - Implement the infrastructure:
 - install/configure the tunneling software
 - install/configure the security software for encryption, authentication and authorization
 - activate compression, dynamic bandwidth management and other performance enhancements
 - reconfigure the router/switch internetwork as needed
 - configure the RADIUS server(s) for necessary security/accounting provisions
 - Verify/establish proper operation end-to-end with a test VPN
 - Characterize the VPN's performance both for setting customer expectations now and planning additional capacity in the future
 - Continually monitor the VPN infrastructure, taking corrective action when and where necessary
-

5. Implementing VPNs

When the VPN offering has been created and the infrastructure is in place, the process of implementing VPNs for paying customers begins. This final chapter evaluates options for customer premises equipment and outlines a list of the steps needed to implement each customer's own virtual private network.

CPE Building Block

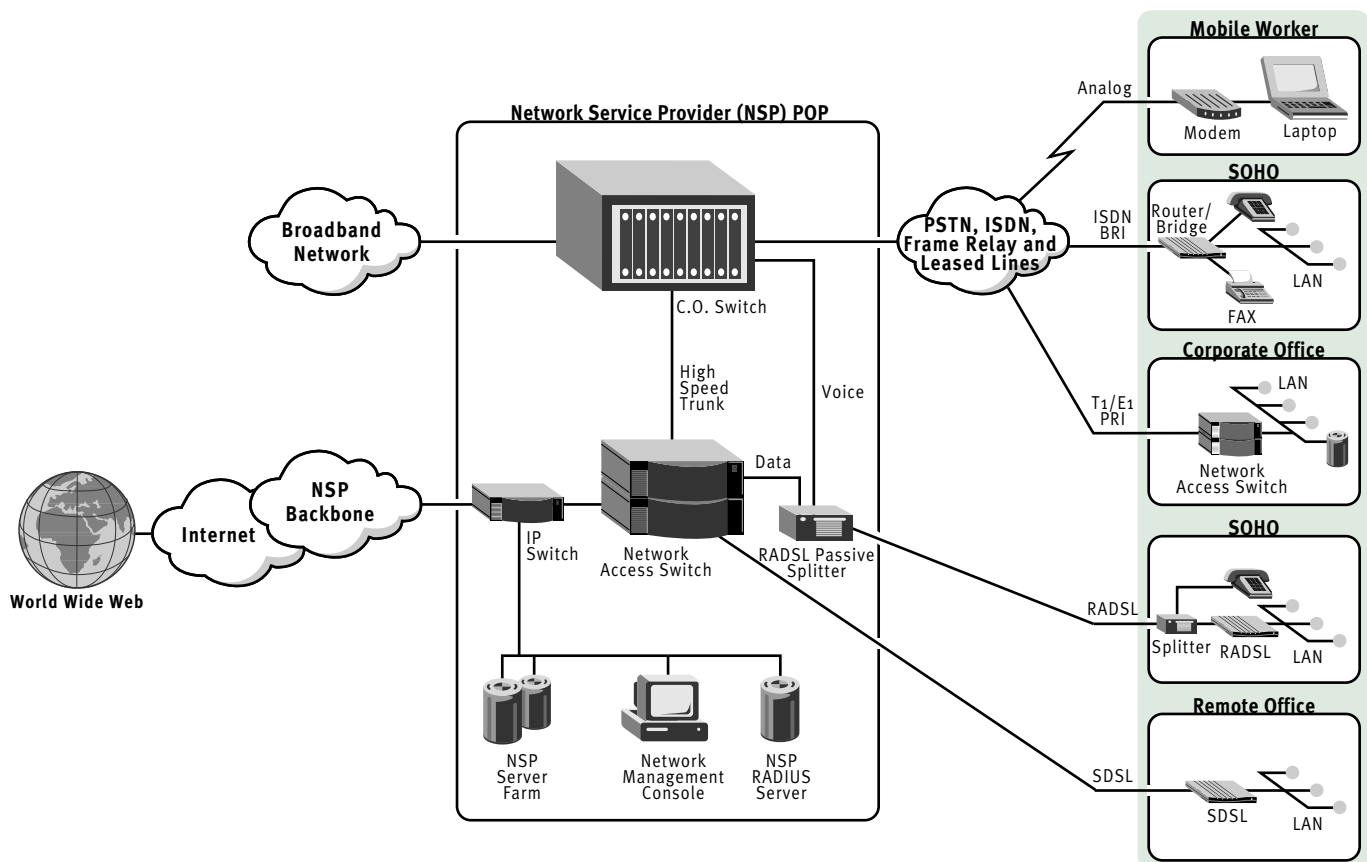


Figure 27 – Customer locations interface to the VPN through on-site WAN access equipment.

Customer Premises Equipment

Customer Premises Equipment (CPE) is the only one of the physical VPN building blocks that exists outside the NSP-Internet infrastructure. From the customer's perspective CPE is where a VPN begins, and ends. The rest is left to the NSP. So CPE is vitally important in any virtual private network.

With an outsourced VPN architecture, CPE requirements are minimal. The customer will probably want to use whatever is already in place at each site in the current private network, whether a WAN router, remote access server or other device.

For any other situation – an in-house or hybrid VPN, new sites not already connected to the private network or new applications – the customer will need new CPE. This section outlines the required and desired features for this CPE in two stages: a list of general requirements, followed by a discussion of site-specific needs.

General requirements for VPN-capable CPE are similar to those of NSP POP equipment:

- A fully-integrated solution
- Support for PPTP, L2TP, Mobile IP and/or other GRE-based tunneling protocols
- Security provisions for encryption, authentication and authorization
- Integral dynamic firewall for protection of local resources
- Software upgradable to conform with emerging tunneling and security standards
- Remote download of software upgrades, via the VPN
- Robust local and remote management to maximize uptime at minimal cost
- Adequate capacity to support anticipated traffic volumes
- An Ethernet LAN interface for attaching to the local network
- Support for the most cost-effective WAN option desired, such as T1/E1, ISDN PRI/BRI, xDSL, Frame Relay or X.25
- Built-in compression to maximize throughput
- Dynamic bandwidth management for enhanced performance
- Ability to handle IP multicast applications
- Compatibility with any advanced capabilities offered by the network access switch at the NSP's POP to magnify the benefits of a VPN
- Compatible family of scalable products to suit a variety of site types and sizes
- Certification for operation with local carriers

The final “general requirement” for CPE is the RADIUS server database. The VPN user profiles are best maintained by the end-user organization (on its RADIUS server) for utilization by its own, as well as the NSP's network access equipment (the RADIUS clients). In this way, the customer assumes full responsibility for the accuracy of the passwords and access privileges for all of their VPN users.

Implementing a Customer's VPN

Each new VPN customer needs a brand new VPN – at least from that customer's perspective. From the NSP's perspective, the same infrastructure applies and the same implementation process is used. Outlined here is a checklist of the process for implementing a customer's very own VPN:

- Plan the VPN working closely with the customer (this step is often best performed during the sales cycle):
 - Create a list of all VPN applications
 - Create a list of all VPN members, both sites and individuals
 - The member list should include the location of and the WAN service option for each site or individual user
 - Determine which sites/users will be “outsourced” and which will be “in-house” make a list of all new CPE needed
- Prepare for the implementation:
 - Check/test all existing CPE, possibly altering its configuration
 - If any existing CPE proves unsuitable for the VPN, add its replacement to the equipment list
 - Order all necessary CPE
 - Obtain a sufficient block of IP addresses (not necessary when tunneling is used)
- Implement the pilot or “Phase I” VPN sites:
 - Install new CPE where needed, and configure the CPE at all sites
 - Configure the customer's RADIUS server with profiles for all VPN sites/users
 - Configure the NSP's RADIUS server for accounting and billing purposes
 - Reconfigure the customer's internal router internetwork, if necessary
 - Install, configure and test each site's firewall, if not already in place
- Test the pilot or “Phase I” VPN:
 - Check/alter equipment configurations
 - Verify each RADIUS database
 - Assess the performance
- Deploy the remainder of the network using the steps listed for both implementing and testing the pilot or “Phase I” VPN
- Perform on-going management responsibilities, which include:
 - Monitoring attempted security violations and taking corrective action
 - Evaluating and tuning end-to-end performance
 - Downloading of any new feature set software to CPE
 - Assisting the customer, as needed, with maintenance of the RADIUS database

Site-by-Site CPE Requirements

A major or central site, such as the headquarters or a division, should be considered a mission-critical installation. These facilities have a large number of users, and some may be Web sites. A single line to such sites, whether in a private network or a VPN, is risky. Dedicated lines, while ideal for such mission-critical applications, regularly become overloaded and can go out of service at any time. Supplemental dial-up bandwidth can handle either situation. A single dial-up ISDN BRI line with 4:1 compression provides up to 512 Kbps of throughput – often enough to handle an overload condition or maintain Internet/VPN access until the primary link comes back on-line.

The best solution is a LAN-attached remote access router or network access switch with both primary and secondary WAN ports. The dual-WAN device should automatically and transparently add the supplemental dial-up ISDN bandwidth when the dedi-

cated primary link is saturated or goes down. It should also automatically terminate the ISDN call when, in either situation, the supplemental bandwidth is no longer needed. No on-duty attendant should be required once the system has been configured for the desired operation.

Major Site Connectivity

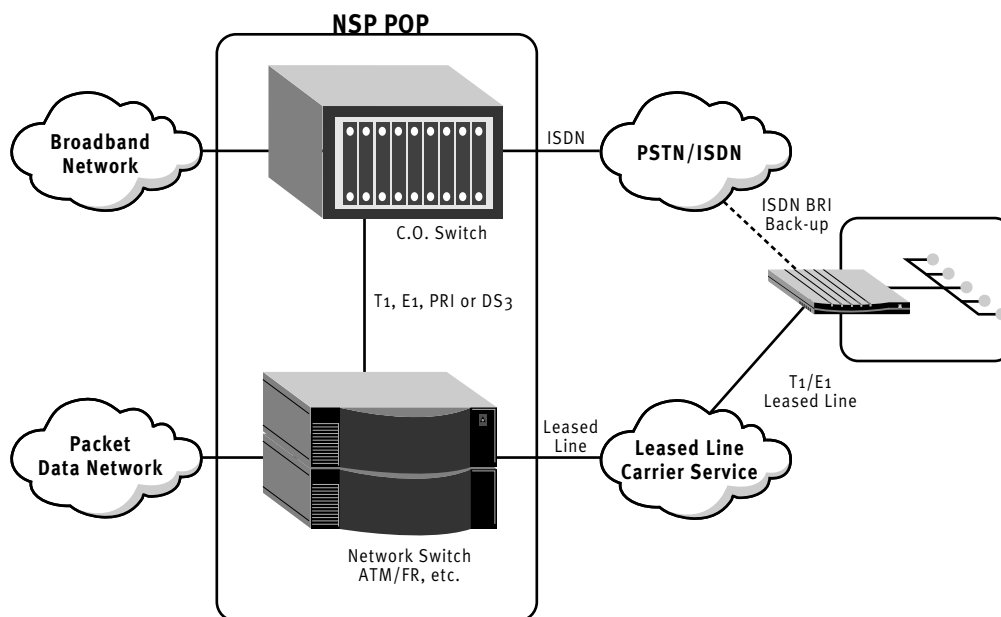


Figure 28 – The dual-WAN configuration offers the mission-critical availability needed by all major facilities in the virtual private network.

For remote offices customers have a choice of ISDN terminal adapters that plug into a card slot on the local Windows NT or NetWare server, or stand-alone remote access routers. The Terminal Adapter (TA) typically lacks the more advanced features such as multiprotocol routing and integrated firewall technology.

For slightly more money, the stand-alone remote access router provides a much more capable and flexible solution. Because it does not depend on the local server, the remote access router eliminates configuration complexities, consumes no server resources and can be managed remotely because it is always on-line. And one of the nicest VPN-related features of a robust remote access router is the integral firewall protection it affords.

Individual users come in two types: traveling and tethered. Traveling workers are stuck with analog modems for the foreseeable future. The analog modem is the only remote access device compatible with the Plain Old Telephone System (POTS), and POTS is the only universally available service. Fortunately, for the daily access usually needed by those on the go, the modem's modest performance is generally quite adequate.

For the tethered telecommuter, however, the better the performance, the better the productivity. ISDN and IDSL at 128 Kbps, in the form of either plug-in interfaces or stand-alone units, offer comparable alternatives for VPN-only applications. But a

common condition in the Small Office/Home Office (SOHO) environment adds an interesting twist. Many SOHO workers need three or more lines: the home line, a business voice line, a data line and, maybe, a separate fax line. The problem is that many homes and apartments are wired for only one or two lines. In these situations there are two solutions. The first is the ISDN Integrated Access Device (IAD), which uses the two BRI channels as needed to handle all data, voice and fax communications on a single line. The second solution is RADSL, which delivers an analog voice channel and a digital data channel on a single line.

ISDN Integrated Access Device Application

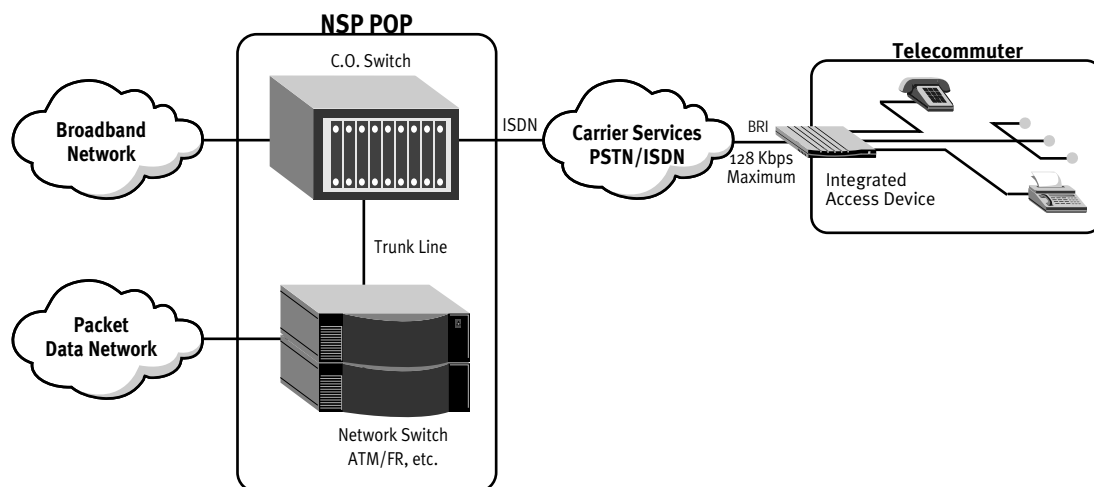


Figure 29 – The ISDN Integrated Access Device (IAD) provides a complete dial-up data/voice/fax communications solution on a single Basic Rate Interface line.

The RADSL Alternative

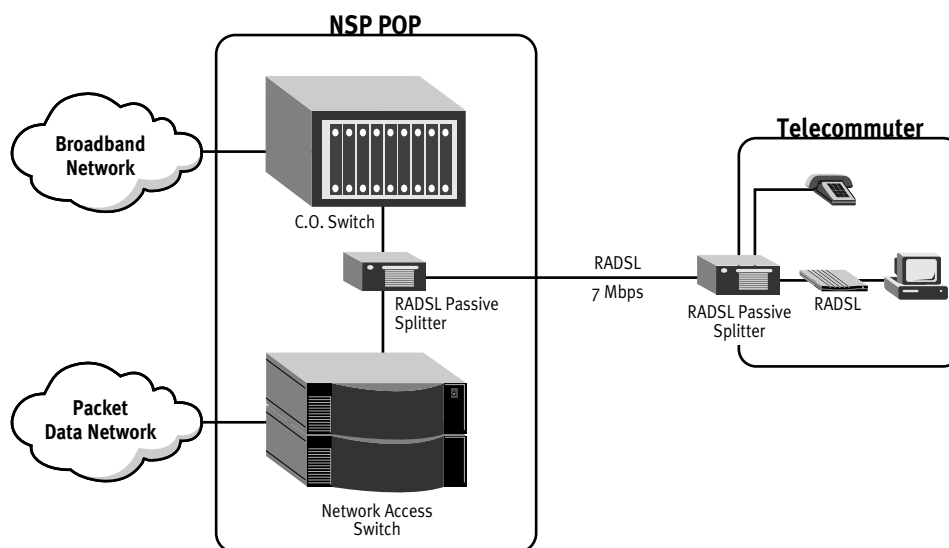


Figure 30 – The Rate Adaptive Asymmetric Digital Subscriber Line (RADSL) is a single-line data/voice/fax alternative for users needing a higher data rate and/or continuous access.

Incorporating “outside members” into a VPN, such as customers, suppliers or business partners, is becoming an increasingly common requirement. Note that VPN membership involves access to private internal resources above and beyond those already available to the public on the World Wide Web. For example, customers may want to check order status, suppliers may need access to the master production schedule, and business partners may be on internal teams that use groupware for project management. If it makes sense to give outsiders access to internal resources – a business decision only the organization can make – then it makes sense to use an Internet-based VPN.

Outside members of the VPN are likely to be equivalent to either a remote office or an individual user. Even though these parties are not part of the organization, they are handled just like all other members of the VPN: with RADIUS profiles to specify “home” networks, passwords and access privileges. The only difference is that outsider access privileges are likely to be quite restrictive. It is useful to install a perimeter firewall to control and monitor access by “outside” members of the VPN. Firewalls integrated directly into the corporate CPE router perform this task well. This service can also be provided from NSPs by installing the firewall directly in their remote access servers.

Defining RADIUS profiles for a limited number of suppliers and business partners is a manageable task, but the effort could become burdensome when a multitude of customers is involved. One way to simplify the job is to define a “generic” customer profile for use by all customers. A customer that requests access to the order processing database, for example, is given the generic user name and password. The access in this case, of course, should be read-only, which must be enforced on the server itself.

One special requirement with outsiders is network interoperability. The obvious, easiest, least expensive and best way to provide this interoperability is to use Internet addresses and IP applications. And the best application by far is the Web’s powerful server/browser combination. Making an internal application or a subset of its information available on a Web server is likely to be a whole lot easier than getting numerous other organizations to convert to the native application, that might be on a mainframe or midrange system, or on a non-IP server. In effect, such an arrangement is similar to an intranet and has, therefore, been dubbed an “extranet”. The IPsec standards for key exchange (ISAKMP/Oakley), when combined with trusted Certificate Authorities, make secure (encrypted) interconnection with other companies via this “extranet” safer than accepting the check of a heretofore unknown third party. With Certificate Authorities and Certificates (such as Level 3 certificates from Verisign), it is as though you receive a check with credit verification already encoded on the check.

Another special requirement for outsiders is firewall protection. A firewall integrated into the router or remote access server can prevent hackers from accessing sensitive or proprietary areas on the network. This added level of protection works in conjunction with the RADIUS server to provide iron-clad security for the entire organization.

The Customer's VPN

All five building blocks come together to establish the customer's complete, end-to-end virtual private network. The VPN can involve numerous NSP POPs, customer sites and third party locations. The result can be a single solution for the customer to access private, public and partner information resources.

Some might ask, "Is an extranet really a virtual private network?" The answer: If it uses the Internet it is! The same application could run on a private network constructed with leased lines or switched WAN services. But as a major computer vendor's sales force always asks, rhetorically, "Why would you want to do that?" Indeed, with the affordability, capability, dependability and flexibility of Internet-based VPNs, why would any organization ever want a purely private data network again? And why would any NSP want to miss out on this golden opportunity?

"End-to-End VPN"




Figure 31 – An Internet-based virtual private network holds the potential to connect an entire corporation, along with its many customers and suppliers.

The diagram shows a VPN that connects a company's headquarters with remote offices and individual workers. Select customers and suppliers are also members of the VPN. Note that traffic can traverse the NSP's backbone alone, or both the Internet and the NSP backbone. A world-wide VPN is also likely to involve the Internet via another NSP's POPs, especially for access by mobile workers and customers. And therein lies the power of an Internet-based virtual private network: wherever your customers locate or travel, their very own VPN is just a local phone call away.

6. Appendices

Actual Case Study:

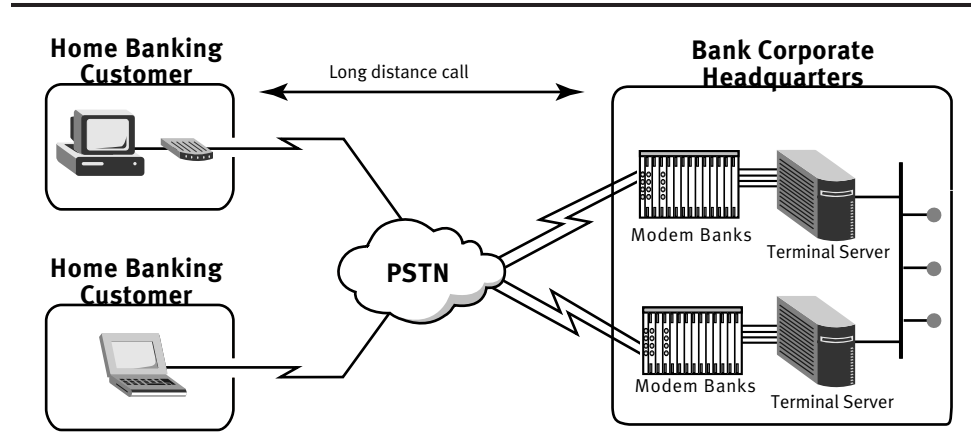
International Carrier Introduces Internet-based Virtual Private Network

Background

An international carrier has been a large systems integrator, offering tailor-made, turnkey telecommunications solutions to small, medium and large businesses. The solutions range from simple networks with several low-speed leased lines and switched connections, to massive WANs that support thousands of users over high-speed packet services such as Frame Relay and ATM.

Over the last few years, the carrier has seen requests for remote networking services skyrocket. Most of these requests come from banks, insurance companies, online service providers and other large organizations that plan to connect remote sites into corporate intranets or offer services such as online banking to customers in geographically dispersed locations.

Carrier's Previous Network



The Solution

But the high cost of leased lines and long distance connections have discouraged many companies from implementing these plans. To meet the demand for remote networking at a price that its customers could afford, the carrier decided to offer an Internet-based Virtual Private Network (VPN). This revolutionary new service uses the infrastructure of the Internet, instead of private lines, as a transport mechanism for establishing secure, point-to-point connections over long distances.

To design and build the VPN; the carrier enlisted the aid of Ascend Communications, the worldwide leader in remote networking solutions and Ascend's local premier VAR. Ascend's MAX™ 400x WAN access switches serve as the foundation of the new VPN service.

Ascend Equipment

- 400 MAX 400x units
- Digital modem cards
- Ascend Access Control and Proxy-RADIUS
- Ascend Tunnel Management Protocol (ATMP)

The Benefits

- Remote users connect to the network with local phone calls.
- Proxy-RADIUS lets customers retain control of password files and access parameters.
- Digital modems can be easily swapped for ISDN cards as callers need high-speed digital access over ISDN BRI lines.
- ATMP enables transmission of packets that would otherwise be unacceptable on the Internet, such as packets that use unregistered IP addresses or IPX packets.
- High-speed Frame Relay port on the MAX allows connection to the Internet backbone at 2 Mbps.

According to the carrier, Ascend is the only company that could provide the technology to build a successful VPN service. Its winning points were the Ascend Tunnel Management Protocol (ATMP) and Access Control, features of the MAX which provide secure tunneling for customer data and foolproof identification and authentication of an unlimited number of users.

How it Works

The Ascend VPN solution consists of one or more MAX 400x units installed in the Carrier 150 points of presence (POPs) and a MAX 400x at each VPN subscriber's corporate headquarters. The MAX in each POP contains five digital modem cards, handles 60 simultaneous calls from analog modem users and is connected to the Internet through a 2 Mbps Frame Relay port. The MAX at a subscriber's corporate headquarters is connected to the Carriers IP backbone via an E1 or fractional E1 circuit.

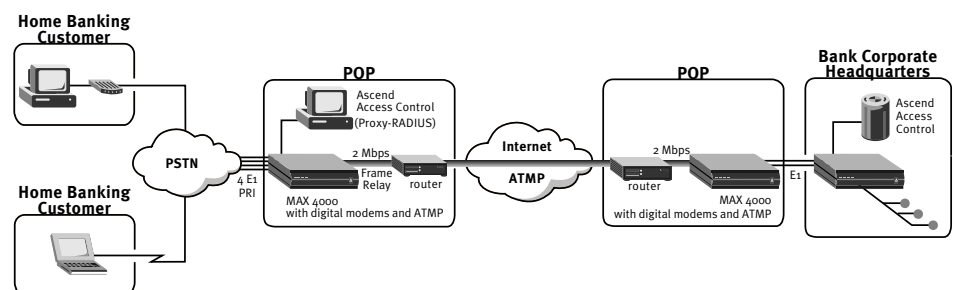
Remote callers such as the home banking customers shown below gain access to their VPN simply by placing a local call to the nearest POP, eliminating any long-distance phone charges. The MAX uses ATMP, an encapsulation protocol developed by Ascend, to establish a temporary "tunnel" between the caller and the MAX at the corporate network. The tunnel ensures the security of the caller's data as it travels across the Internet, and allows the transmission of unregistered and non-unique IP addresses that otherwise could be unacceptable on shared Internet facilities. When the call is terminated, the MAX automatically tears down the tunnel.

To manage VPN users and protect subscribers' network resources, the carrier is using Ascend's network security management system, Ascend Access Control™. In conjunction with Proxy-RADIUS, an authentication server that resides at the POP, Access Control serves to identify legitimate VPN callers, perform user authentication and authorization and determine network destinations. It also lets each VPN subscriber maintain local control of their users and customize their own user access profiles.

The network is attracting strong interest from corporate customers, who recognize that Internet-based VPNs provide enhanced flexibility and tremendous cost savings over other networking alternatives. A large bank with 15,000 home banking customers has already signed up. By the end of the year, the carrier expects the new service, which became operational in February of 1997, to have tens of thousands of subscribers.

To support this influx of subscribers, the Carrier plans to install at least 400 MAX products in its POPs by year's end. To support the additional traffic on its Internet backbone, the company also plans to add hundreds of MAX products.

Bank using the Carrier's VPN Service – Ascend Solution



Ascend Product Information

Ascend's industry-leading remote networking products deliver seamless and secure end-to-end solutions for high-speed connectivity. With the expanding lines of Ascend's integrated products – MAX™, Pipeline®, GRF™, NetWarp™, MultiDSL™ and security families – users are assured of the most comprehensive solutions for high-performance access to central sites or the Internet.

MAX family products are modular WAN access switches that scale to high densities and support digital, analog and cellular carrier services in a single product. The Pipeline family of products are high-speed, easy-to-use remote access devices for telecommuters, remote offices and home users. The high-performance GRF family products are scalable, high-bandwidth Layer-3 IP switches for network access and backbone services. The affordable NetWarp family of terminal adapters are Plug-and-Play solutions for individual PC users that need high-speed digital links. The MultiDSL family products are the first fully integrated Digital Subscriber Line (DSL) solutions that can be immediately deployed for high-speed access ranging from 128 Kbps to 6 Mbps. And integrated security products like Secure Access Firewall™ and Ascend Access Control furnish comprehensive, dynamic security management for your remote network – including central sites, remote offices and telecommuters' home offices.

Ascend's products are designed for telecommuting, Internet access, remote office connectivity and multimedia access. Ascend is the market share leader, with 54 percent of all access concentrator ports around the world using Ascend's award-winning products to access the Internet, link branch offices, transmit video, telecommute and perform hundreds of other remote networking tasks. Ascend's products deliver the most cost-effective, high performance and scalable solutions for a full range of remote access applications from Internet access, to remote office connectivity, to high-density MegaPOP™ solutions.

Ascend's end-to-end solutions allow customers to leverage both the economies of an easy-to-manage, single-vendor platform and the adaptability of Ascend's modular, scalable architecture to meet the remote access needs of a changing workplace.

Applications

MegaPOP

Ascend's MegaPOP application combines the GRF family of IP switches with the MAX TNT™ to give network service providers unprecedented aggregation and packet forwarding performance. This high-density combination supports hundreds and thousands of users in less than eight feet rack space. With packet forwarding rates of up to 2.8 million packets per second, Ascend's MegaPOP solution eliminates network brownouts and access denials caused by conventional solutions. Together, a GRF IP switch and the MAX TNT let network service providers increase network capacity, reduce equipment costs and eliminate the need for multivendor installations.

Internet Access

To get connected to the Internet, corporations need flexible, reliable equipment that provides high-speed digital performance, conforms strictly to industry standards and guarantees compatibility with existing network protocols. Ascend's products meet, and exceed, these requirements.

Ascend's support for analog and digital services means all users enjoy high-speed, error-free connections, no matter where they're located or what carrier service they use. Ascend's products deliver advanced Internet-specific feature sets, integrated firewall security, superior technical support and complete compatibility within Ascend's own product families and those of other vendors. These are just some of the reasons why 28 of the world's 30 largest Internet Service Providers have standardized on Ascend solutions.

Remote Office Connectivity

Ascend's high-performance products give you everything you need to construct a rock-solid remote networking solution. They support ISDN, IDSL and other DSL technologies, Switched 56, Frame Relay as well as leased lines and analog circuits. Ascend's user authentication schemes and integrated, dynamic firewall capabilities implement security even before the corporate LAN is accessed. Sophisticated bandwidth optimization features transparently manage digital connections, giving the highest throughput at the lowest possible cost.

At your corporate site, Ascend's products replace multiple terminal servers, terminal adapters, and sluggish modem banks that no longer meet user performance requirements. At branches and remote offices, Ascend's products link users at lightening-fast speeds to key information resources on corporate networks and the Internet. Robust support for all standard protocols and network management schemes ensures efficient connectivity and multivendor interoperability.

Telecommuting/Small Office/Home Office Access

Ascend's award-winning products provide security, scalability and manageability for enterprise wide telecommuting programs and fit seamlessly into current network security architecture. They can be configured and managed either remotely or locally, so technical support staff can take care of telecommuter needs without ever leaving the office.

Ascend's end user products integrate analog and digital capabilities. They connect computers, phones and fax machines to a single high-speed line, so users reap the benefits of digital connectivity without leaving the analog world behind. With Ascend's end-to-end telecommuting solution, employees work as productively from home as when they're in the office.

Products

MAX Family of WAN Access Switches

Ascend's MAX™ products are powerful remote access concentrators that help corporations build remote networks of any size, with virtually any combination of analog and digital carrier services. All MAX products combine the functionality of a router, a terminal server, an ISDN switch and a Frame Relay concentrator in a single box. MAX products represent the industry's most adaptable, secure and interoperable solutions for remote networking. No wonder there are more than 15,000 MAX units installed, managing over 1 million calls each day.

With industry-leading manageability and plenty of room for expansion, there is a MAX product to precisely fit you or your customers' needs – whether you are servicing a small company with several offices across town or a corporation with thousands of users around the world.

MAX family features include:

- Integrated support for digital, analog and cellular services
- MAX family supports from 8 to 672 callers in a single platform
- Dynamic Bandwidth Allocation™ (DBA), inverse multiplexing, Multilink Protocol Plus™ (MP+)
- Remote management
- Scalable and stackable units
- Series56™ Digital Modem modules that are V.34 - and K56flex - compatible
- Extensive standards-based security
- Optional integrated Secure Access Firewall

Across the board, MAX products supply their rich functionality and high performance for the industry's lowest cost per port – once again explaining why over 54 percent of all access concentrator ports installed by corporations, carriers and service providers are Ascend's MAX products.

MAX Hardware and Software Options

Several hardware and software options available for Ascend's MAX products help you optimize MAX performance and customize it to fit your individual needs. Used separately or together, these options provide you with a complete, end-to-end remote networking solution.

MAXLink Pro

With MAXLink Pro™ client software, remote users can access enterprise network resources. These include Internet and Intranet access, host connectivity, electronic mail and messaging, and access to file servers, printers and other resources. MAXLink Pro includes a complete Novell client and the most popular TCP/IP applications.

- Easy to install and use
- TCP/IP and IPX protocol stacks
- Client software for Windows 3.1x, WIN95, Windows NT, Macintosh and MS-DOS
- Bundled with all Pipeline products and the MAX 200Plus
- Optional for all other MAX products

MAXDial

MAXDial™ is client software that lets LAN users place outgoing modem calls and send faxes utilizing the digital modem cards in their MAX WAN access switch. Using MAXDial saves you money because you won't need to install separate dial-out servers or direct lines and desktop modems in every users' office.

- Outgoing modem and fax calls from desktop via a MAX
- AT command set for V.34 modems
- MS-DOS, Windows 3.x and WIN95 support
- Novell IPX LAN support
- TCP/IP LAN support (WIN95 only)

Digital Modem Cards

Ascend's Series56 high-density digital modem cards let users with analog modems dial into a MAX product over digital access lines. They provide full access to 8, 12, or 48 analog callers, enhancing modem call performance and reducing operating costs. Not all modem cards are supported on all platforms.

- Available in models that support 8, 12, 16 or 48 modem-based users per card
- Support all modem speeds up to 56 Kbps (K56flex) and 33.6 Kbps (V.34)
- Data compression (V.42) and error control (V.42bis)
- Cards plug directly into expansion slots

Pipeline Family of Remote Access Products

Ascend's award-winning family of Pipeline® products represent the largest range of remote access devices for Internet access, remote offices, home offices and telecommuters. Key applications such as corporate LAN access, Internet/intranet access, management of large numbers of telecommuters, and mission-critical remote networking all benefit from the Pipeline family's high-speed, industry-leading performance.

The Pipeline 25-*Px*, 25-*Fx*, 50, 75 and 130 are Ethernet-to-ISDN routers and bridges. The Pipeline 15 is an easy-to-use terminal adapter for inexpensive, high-speed access to a corporate LAN or the Internet. The Pipeline supports high-speed ISDN BRI communications for stand-alone PCs. Analog devices (modems, phones, faxes) are supported on the Pipeline 15, 25-*Px*, 25-*Fx* and 75 models.

High-performance Pipeline 130 models support growth from a 56 Kbps Frame Relay dial-up or leased line connections at up to T1/-E1 speeds. They provide a dedicated WAN connection and an integrated switched connection for mission-critical back-up and overflow applications.

Pipeline family features at a glance:

- Multiprotocol bridging and routing*
- Dynamic Bandwidth Allocation (DBA), inverse multiplexing, Multilink Protocol Plus (MP+)
- Data compression on Pipeline 50, 75 and 130
- Remote management
- Extensive standards-based security
- Integrated Secure Access Firewall optional on Pipeline 50, 75 and 130

In addition to their rich features, Pipeline products are convenient solutions with modem-sized footprints and no external wiring – exactly what you'd expect from Ascend, the ISDN market share leader.

**Note that Pipeline 15 is a terminal adapter that does not include bridging and routing.*

The Pipeline Companion CD-ROM

The Pipeline Companion CD-ROM includes all software for key applications programs in addition to the Java-based Pipeline Configurator (JBPC), a graphical setup and configuration utility for the Pipeline family. Software included on the Pipeline Companion CD-ROM includes MAXLink Pro Client software for remote access to the corporate LAN, and MAXDial software that allows digital dial-out for LAN-attached users.

The JBPC allows Pipeline 25-*Px*, 25-*Fx*, 50 and 75 units to be configured from a PC over an Ethernet LAN connection. Running as a stand-alone program, the Pipeline Configurator offers new levels of set-up convenience and efficiency – including a comprehensive Quick Start utility to guide users or network administrators through configuration by application, as well as a complete HTML-based help utility for the Java program.

- Java-based Pipeline Configurator (JBPC) for graphical Java-based set-up over Ethernet LAN
- Quick Start set-up utility for JBPC
- Key applications including MAXLink Pro, MAXDial and Internet Explorer 3.0
- Free demo software programs from Microsoft, NetManage, Adobe, and Funk Software
- Bundled FREE with Pipeline products and the MAX 200Plus

GRF Family of IP Switches

Ascend's breakthrough GRF™ family of high-performance IP switches are the ideal solution for the congestion problem on the Internet. In combination with the MAX TNT™ WAN access switch, carriers and ISPs can support an entire metropolitan area in less than eight feet of rack space.

GRF IP Switches can be easily integrated into service provider networks. Using IP Forwarding Media Cards, the scalable GRF family of products provide the bandwidth needed for next-generation high-throughput networks. The hot-swappable cards are intelligent, self-contained IP forwarding engines that provide exceptional performance, configuration flexibility, and port density far beyond the range of conventional routing solutions.

- Layer-3 IP switching and distributed routing provide up to 16 Gb/s bandwidth
- Hardware-assisted, full route table lookup enables wire-speed performance
- High-density, redundant modular chassis ideal for the POP
- Open architecture supports concurrent HSSI, 10/100Base-T, ATM OC-3c/STM-3, FDDI, IP SONET OC-3c with Frame Relay and PPP framing, ATM OC-12c/STM-12

MultiDSL Family of Digital Subscriber Line Solutions

Rapid Deployment of DSL Services over existing wiring

The Ascend MultiDSL family of products deliver and implement xDSL solutions immediately. Ascend's MultiDSL products work on a single, scalable platform and supports all xDSL services in addition to all other carrier services, such as analog, Frame Relay and ISDN. Ascend's new MultiDSL family includes integrated Central Office Equipment (COE) and Customer Premise Equipment (CPE) for implementing a wide range of Digital Subscriber Line (DSL) solutions with speeds from 128 Kbps to 7 Mbps over existing twisted pair wiring. Supported technologies include Ascend's own IDSL (ISDN Digital Subscriber Line), in addition to SDSL (Symmetric Digital Subscriber Line), RADSL-CAP (Rate Adaptive Asymmetric Digital Subscriber Line – Carrierless Amplitude Phase) and soon RADSL-DMT.

Central Office Solutions (MAX 4000, 4002, 4004 and the MAX TNT)

MultiDSL COE products allow carriers to leverage their copper wiring investments to quickly launch high-speed access services from 128 Kbps (IDSL) to 7 Mbps (RADSL). By installing MultiDSL line cards into one of the MAX products, users can cost-effectively support a wide range of services and advanced software functionality with a single manageable platform.

Customer Premise Solutions (DSLPipe, Pipeline)

For deployment at subscriber premises, MultiDSL CPE products complement the COE offerings with IDSL capabilities for Ascend Pipeline products, and specialized DSLPipe™ remote access solutions for SDSL, RADSL-CAP and soon RADSL-DMT.

Ascend's Security Family

Working hand-in-hand with Ascend's MAX and Pipeline products, Ascend's security products bring affordable, bullet-proof security to your network — from central sites to the smallest remote user locations. Most corporate networks are only secured at the corporate LAN or not at all, leaving your company's assets vulnerable to unauthorized users. With Ascend's integrated firewall security, any size company can implement a secure, end-to-end, remote network without resorting to expensive stand-alone firewalls. And for easy, cost-effective manageability, all security measures can be managed and configured remotely from a central site.

Secure Access Firewall

Secure Access Firewall is NCSA certified and integrates "stateful inspection" based dynamic firewall technology with Ascend's Pipeline (Pipeline 50, 75 and 130 models) and all MAX products, providing a single-vendor security solution for remote access networks of all sizes. The dynamic firewall technology prevents intruders from attacking the router and accessing the corporate LAN. It also protects the network at break-in points, including remote offices and telecommuters' homes.

The Secure Access Firewall option offers:

- State-of-the-art dynamic firewall technology
- A record of attempted break-ins as well as an audit trail
- Comprehensive control over TCP/IP applications and non-IP protocols
- Transparency for authorized users

Secure Access Manager

Free of charge with the Secure Access Firewall is Secure Access Manager. This point-and-click graphical user interface is Windows-based (supports Windows 3.x, WIN95 and NT) and allows network managers to install and fine tune the firewall functions of an entire network from the corporate site.

Ascend Access Control

Ascend's Access Control™ is an authentication server that provides a single-point solution for secure remote networking. It encompasses extensive identification, authentication, authorization and accounting functions. It also features extensive per-user filters for authorization and access. Based on the industry-standard RADIUS server, Ascend Access Control analog and digital (ISDN) dial-up lets you manage users from central or remote sites.

Ascend Access Control offers the following:

- An easy to install and manage interface with Java-based GUI features
- Support for Internet and virtual private networking
- Detailed accounting for billing record generation
- Support of ODBC-compliant databases
- Secure Dynamic Bandwidth Allocation
- Comprehensive security when combined with Secure Access Firewall
- Ascend TelNet Manager

Reference Material

The following documents have provided valuable input to the development of the VPN cookbook.

These reference documents are available from Ascend Communications, Inc.

Document	Content	Reference	Author
Ascend Tunnel Management Protocol	ISP Application Note	05-5	Ascend
Ascend Tunnel Management Protocol	Carrier Application Note	05-5	Ascend
Remote Access Dial In User Services (RADIUS)	Technical Overview	05-6	Ascend
Networking for Remote LAN Access and Internet Connectivity	Remote LAN Access		Telechoice
Access Denied – Integrated Firewall Security	Technical Overview	05-10	Ascend
Multiprotocol VPN Protocols	Technical Backgrounder	05-13	Ascend
Remote Access Network Security	Resource Guide	06-2	Ascend
20 Profit making tips for ISP's	Resource Guide	06-4	Ascend
Money saving tips for Remote Access Networks	Resource Guide	06-03	Ascend
Corporate Remote Access Guide	Business Tools	06-06	Ascend
MultiDSL Profit Opportunities for Network Service Providers	Resource Guide	06-8	Ascend

Virtual Private Networks Resource Guide

Fax this Form For Additional Information

To: **ASCEND COMMUNICATIONS INC.**
Worldwide Headquarters
Attn.: PreSales Technical Consulting
Phone: 800-621-9578, option 4
Fax: 510-337-2668

From: Name: _____
Title: _____
Company: _____
Address: _____
City: _____
State/Zip: _____ / _____
Country: _____
Telephone/Fax: _____ / _____
E-mail: _____



Please send information on the following products:

☐ GRF

☐ Pipeline

☐ MAX

☐ Multiband

☐ xDSL

☐ Security

☐ Network Management



Please have an Ascend sales representative call.

