Shiva Corporation
Corporate Headquarters
28 Crosby Drive
Bedford MA 01730 USA
800.977.4482
508.788.3061
508.788.1539 FAX
sales@shiva.com
www.shiva.com

The Value of VPN to Your Company

A Guide to Evaluating and Implementing the New Remote Access Technology

Authors: David Allan, Gordon Burnes
Technical Contributors: Ray Beaulieu, Steve Elgar, Mark Tuomenoksa, Jesse Walker

Contents

Overview

This paper examines the technical, administrative and organizational factors that dictate which combination of direct dial and VPN components will produce the optimum return on your remote access investment.

Because VPN implementation must be tailored to your company's particular needs, almost every remote access solution will be different. The following discussion addresses the important questions you should consider before implementing VPN technology as part of a remote access solution. It also outlines the steps necessary for optimizing a remote access solution using VPN technology.

Choices For Information Access describes how the remote access environment of today will evolve to a data dial tone future. The economic and technological implications of this evolution define the available remote access choices.

Steps To Implementing VPN discusses the criteria that you need to consider when planning a VPN implementation: users, data and VPN architectural models.

The Components Of VPN provides a brief discussion of the protocols and technology that make up VPN.

Making VPN Work presents practical examples of working with service providers as partners, and provides examples of cost analysis for direct dial versus VPN and hybrid solutions.

Introduction

Today, remote access to information is a strategic corporate necessity. It:

- Arms employees with up-to-date information, enabling them to make the most informed decisions possible
- Streamlines access to information through Internet and intranet connections
- Reduces networking costs by using the Internet to replace in-house WAN and dial-up networks
- Extends the workplace beyond the office walls to allow people to be fully productive at home and on the road
- Provides an edge in recruiting employees looking for flexible work styles such as telecommuting and job sharing
- Fosters competitive advantage by creating closer links with customers, suppliers and employees

Two complementary technologies can help IS managers maximize their remote access capabilities while minimizing the underlying costs:

- Remote node direct dial
- Virtual private networking

Usually, the most flexible and cost effective remote access solution merges the two technologies, combining direct dial using the Public Switched Telephone Network (PSTN) with a Virtual Private Network (VPN) using the Internet and the PSTN for remote LAN access.

While many corporations already use direct dial for remote LAN access, adding VPN technology could result in lower overall costs.

Choices For Information Access

Today: Data Over Voice

Current remote LAN access commonly uses the existing voice infrastructure, as illustrated in Figure 1. Here, a remote laptop user dials into a corporate central site. A voice circuit then carries the data transmission from the laptop modem to the central site's remote access server modem.

This connection uses a 64 Kbps voice circuit for the duration of the call.

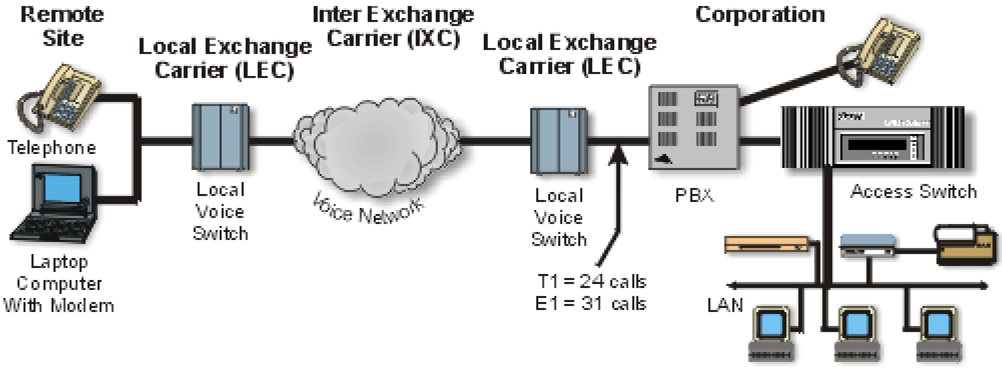

Figure 1: Data over voice

Cost Structure

Because voice tariffs drive the economics of this approach, charges are typically based on connection time, not data transmitted. These fees vary widely throughout the world and represent a major portion of remote access costs.

It is worth noting that connection time charges vary depending on the distance between the remote user and the corporate network. In the US, for instance, local calls may be included in a toll-free zone, but the cost of a long distance call depends upon the service provider's rate. The geographic spread of remote access users is a key factor in determining the cost of the service and the economics of VPN. This is discussed in more detail in *Steps To Implementing VPN, Evaluate User Locations*.

"Last Mile"

The remote site relies on the voice network infrastructure. Thus, the "last mile" of copper wire – usually the weakest link – limits the available data bandwidth.

At the corporate site, the call can terminate as a digital call over a high bandwidth connection such as a T1, E1 or Primary rate ISDN connection. However, each connection monopolizes a complete DS0.

Tomorrow: Data Dial Tone Networks

Data dial tone provides a reliable, universal data service, connecting every user to a high-speed data network.

Currently, data over voice networks are digital only after they pass the local voice switch. In contrast, data dial tone brings the digital signal to all end points on the network.

Digital services such as ISDN are only the start of a trend. The full digital dial tone will provide the same standards for data as consumers now expect from the voice network—100% up-time and universal availability.

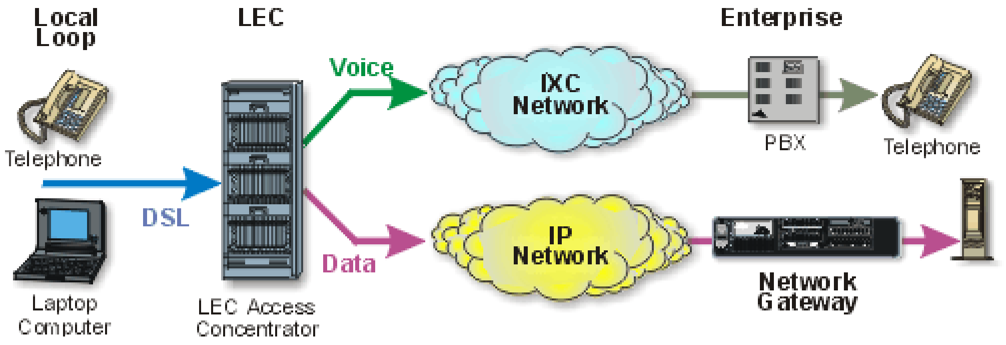

Figure 2. The Future Data Dial Tone Network

Cost Structure

The cost structure for data dial tone service will be different than that for data over voice lines.

Instead of voice tariffs based on distance and time, charges will depend on bytes transferred. Thus, while in today's environment, applications are optimized to limit connection times, future remote access applications will be optimized to limit the number of packets transferred between the remote user and the central site.

"Last Mile"

Bandwidth-hungry users will demand faster access to the data dial tone network.

Replacing the copper network appears to be prohibitively expensive. Instead, service providers will turn to high speed transport services that operate over copper wire to fulfill the need for speed. Digital subscriber line (DSL) technologies can deliver high bandwidth (500 Kbps to 8 Mbps) over the "last mile" of copper to the end user. In addition, residential users may use cable modems to satisfy their demand for higher bandwidth.

Driving Forces for Change

There are three main forces behind the trend toward data dial tone networks:

- The economics of data transport,
- The use of shared resources, and
- The need for higher bandwidth.

Data Transport

Data dial tone will deliver more efficient use of available bandwidth. For example, in the US, a channelized T1 used to transport digital information has a maximum capacity of 24 channels. Since conventional direct dial allocates a full channel for each call, T1 can accommodate 24 calls. Likewise, under this scenario, European E1 lines can accommodate 31 calls.

However, the same digital pipe can carry packetized data on a multiplexed connection, allowing the clearchannel T1 to carry a higher number of user sessions. The magnitude of this increment depends on the type and volume of traffic. If there are many idle or low volume users, then a T1 could support up to and perhaps more than 200 sessions.

Consider, for example, a classic pattern of remote email usage:

1) log in
2) replicate email
3) check email
4) send urgent replies
5) log off

This sequence of steps generates a low volume of traffic. A fully digital connection could, therefore, support a very high number of email users.

In contrast, running graphical packages such as Windows or X Windows over the connection generates a high volume of data traffic. In this case, a clearchannel T1 will handle a much smaller number of sessions, though more than the 24 calls supported by a channelized T1.

Shared Resources

A data dial tone network allows many different users to share network access devices. Thus, because ports are used more efficiently, per port costs drop. In a competitive environment, as the cost of the data dial tone service decreases relative to data over voice, service providers will pass these savings on to their customers.

Higher Bandwidth

The data dial tone network will deliver higher bandwidth all the way to the network end points.

While data over voice networks provide only limited bandwidth connections, data dial tone will, as previously noted, rely on high-speed data transport over copper. DSL technology bypasses the local voice switch to provide high bandwidth over the existing copper infrastructure. This architecture has the added benefit of relieving the local voice switch of data calls.

Remote Access Choices

The basic architecture for all remote access is a connection from a remote or branch site through a network to a central or other branch site. This basic architecture may include two different remote access implementations: direct dial and VPN.

The difference between direct dial and VPN is where the call terminates. The end point of a direct dial connection is the central site. For VPN-based remote access, however, the data call terminates at a local POP. From there, a data session is routed to the central site.

In most cases, the optimal choice for remote access is a combination of both direct dial and VPN to best serve the unique needs of each user, while minimizing the cost to the corporation.

Direct Dial

In a direct dial scenario, a remote user dials from a modem in their laptop through the PSTN, and the call terminates in a corporate LAN-based modem in an access server. Direct dial allows the IS Manager to optimize the connections to the PSTN according to users' needs.

With this solution, the connection depends entirely on the voice network.

Virtual Private Networking

A VPN-based remote access connection begins with a data-over-voice call that terminates at the local POP. From there the data flows through a VPN session over the Internet (or other IP network) and ends at the corporate network gateway.

The differences between direct dial and VPN are illustrated below:

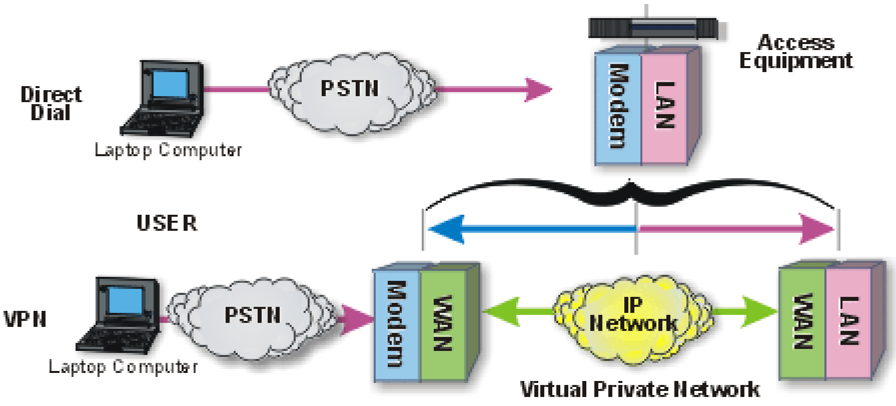

Figure 3. VPN versus Direct Dial Differences

In essence, with a VPN connection, the Internet becomes an extended PCI bus.

Compared to direct dial solutions, VPNs can reduce the long distance phone charges associated with remote LAN access, and potentially eliminate them if remote users can access the public IP network with a toll-free call.

Since public IP network access charges are typically flat rates, VPNs are often more expensive than direct dial when usage is high and most users are in the same calling area as the central site.

While providing low cost access, VPN over the Internet cannot offer the same guaranteed performance as a direct dial solution. The unpredictable nature of Internet traffic places loads on the Internet infrastructure which can greatly alter connection latency. In addition, from time to time, parts of the Internet infrastructure experience temporary bottlenecks and unpredictable outages that can cause some sites to become unreachable. While this may be acceptable for most data communications, it is not an option for mission-critical data transfer.

Optimized Remote Access Solutions

VPNs and direct dial are not mutually exclusive. In fact, a combination of the two technologies usually provides the most cost effective solution.

Merging the advantages of direct dial with the lower costs of VPNs can create an optimized solution. VPNs can help cut long distance charges, while direct dial solutions make more sense for local usage. Further, while a VPN might be used for applications with lower availability and response priorities, direct dial may be the first choice for mission-critical access. Direct dial also provides corporations a backup if the quality of service over the Internet degenerates.

Steps To Implementing VPN

To achieve an optimal return on your remote access investment, the IS manager must review data and usage patterns, and match this information to the appropriate VPN solution. This section examines the evaluation process.

Determine the Number of Ports

In any remote access solution, users and their work patterns are a major factor in defining the system requirements. The key measure for direct dial access is the number of ports deployed; the VPN equivalent is the number of concurrent sessions supported. In either case, the value of this metric determines the necessary capital equipment and associated maintenance and support costs.

Deciding on a ratio of users to ports involves a tradeoff between cost and service level. A higher ratio increases the frequency of failed connection attempts, while a lower ratio increases equipment costs.

A reasonable starting point is to use a ratio of one port or session to ten users. However, the quality of service can be affected by work patterns and remote site bandwidth (the latter issue is discussed below under *Determine Data Rates*).

Users' work patterns can create data "rush hours." For example, if all remote users typically log on at 9 am to pick up a daily message, with a ten to one ratio, most of them will get a busy signal. In this instance, work patterns may need to change to accommodate the limitations of remote access.

Classify Users

Shiva segments remote users into the following classifications:

Mobile workers/business travelers are the classic road warriors who need access to the corporate network and the Internet while traveling. They are either full-time field staff (such as sales people) or office-based workers on business trips. Email is typically their greatest remote access requirement.

Mobile users require self-contained technology on their laptop: a built-in modem (e.g., PCMCIA Card) and integrated dial up client software. A standard modem allows them to plug into any telephone line.

Access should be easy, such as limiting the authentication steps to user name and password. If higher levels of security, such as third party token-based authentication, are required, then this process should be as seamless as possible.

The slow speed of the analog phone line means that any tools that increase data rates – or decrease the amount of data that needs to be sent – are extremely useful to mobile workers.

Telecommuters work one or more full days a week from home. These workers are also more likely to have a higher bandwidth connection such as ISDN.

Home Workers work from home for short periods of time, such as evenings or weekends, in addition to their main activity at the work site.

While telecommuters and home workers can use analog connections, a range of digital options, such as ISDN, DSL, cable modem, etc., offer increased reliability and bandwidth.

The telecommuter requires corporate and Internet access, and may use a separate phone/fax line for work. As with a mobile worker, telecommuters need their remote access system to be easy to use. However, because of their static locations, the types of applications are likely to be more varied than for mobile workers.

In addition to email, telecommuters and home workers are likely to require database and server access. The data-intensive nature of these applications makes bandwidth more important and implies that the total connection times may be longer.

A remote user may fall into different classifications at different times:

- A home worker who picks up email in the evenings or weekends
- A mobile worker who travels
- A worker who telecommutes for a number of days, perhaps while on sick leave

Branch Offices. A final category comprises small branch offices that require access to a central site or to each other. The number of individuals in the office and the data rates required by the office's network connection vary widely depending upon the nature of the work. However, Basic Rate ISDN is likely the minimum connection technology.

Figure 4 illustrates a typical distribution of user types based on a 1996 survey (Merrill Lynch) of 100 US-based LAN Managers and MIS directors.

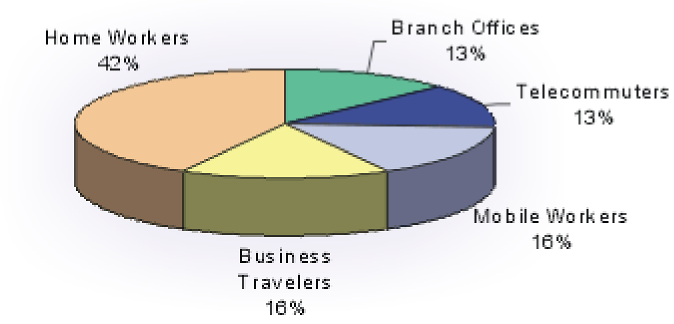

Figure 4. Typical users of remote access

Evaluate User Locations

When comparing VPN to direct dial, it is important to understand how the geographical spread of users affects costs.

Users must dial to a point of presence (POP): a corporate site for direct dial or an ISP POP for a VPN. The distance to the point of presence determines the per minute tariff which, when multiplied by the length of access time, determines the total cost. Even within a U.S. or Canadian free local calling plan, some proportion of users may be outside the free zone.

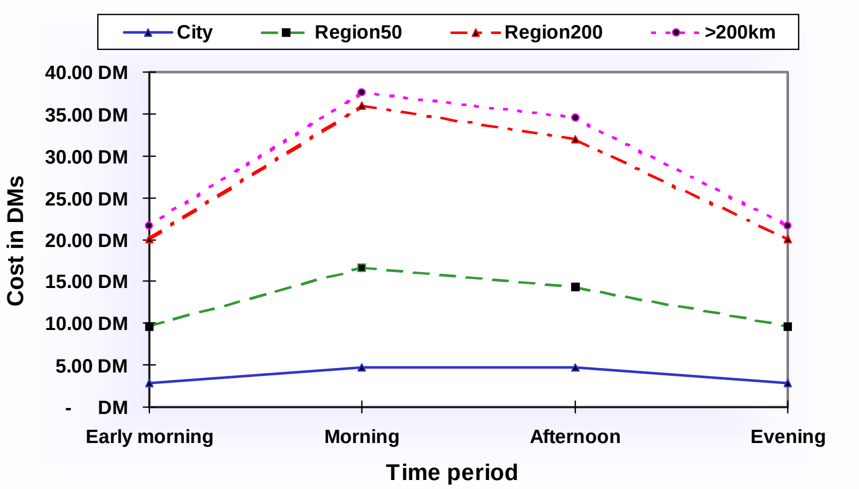

Figure 5. Deutche Telekom rates for one hour of access time

Figure 5 shows an example of remote access costs compared to calling distance in Germany. The cost of one hour of access ranges from $1.62 US to $21.21 US. The main financial argument for VPN comes from the distance-based cost. Using VPN instead of direct dial can greatly reduce the proportion of users who must place expensive long distance calls.

Determine Data Rates

As with the weakest link in a chain, the smallest bandwidth in a data path sets the limit on the data rate for a connection. With a direct dial solution, bandwidth is normally predictable. For example, a modem creates a 33.6 Kbps pipe for the duration of the call.

When all remote sites have the same access technology and the calls arrive at the central site via a channelized connection, the user-to-port ratio gives an accurate measure of the quality of the remote access service. However, when the central site has a non-channelized connection, remote sites that have a high bandwidth can swallow all of the available central site bandwidth. If, for example, the remote site connects by cable modem to a POP, it could use up all the bandwidth of a central site T1 connection.

With VPN, the access session is a data stream passing through the shared network and arriving at the corporate site multiplexed with many other sessions. The bandwidth taken up by each data stream varies and is limited by the indeterminate path taken through the Internet. This makes it difficult to predict the number of concurrent sessions that can be supported and, therefore, the quality of service provided by a given capital outlay.

Assess Data Security Requirements

A key deployment decision is the degree of security required. There are a number of security parameters that should be considered:

1) **Authentication**. Is the data really from the stated source? Is the source a valid user for this corporate network? This is an important issue with all remote access. With VPN, traffic arrives from the Internet; it is, therefore, expecially important to authenticate the connection when establishing the tunnel.

2) **Integrity**. Has anyone tampered with the data? The data session could have been altered as it passed though the network. Integrity checking ensures that the data continues to come from the authenticated source and has not been altered.

3) **Privacy**. Does the data need to be kept confidential? Some data must be concealed from the view of others. This is a more important issue with VPN than direct dial as the data passes through a public and (from the corporation’s point of view) uncontrolled network. Privacy is provided by encryption.

Application traffic is typically segmented into four groups, although some companies further subdivide it into as many as 12 groups. The following categories start with the lowest level of security:

1) **Level 1: Don’t Care**. This data does not need security (general Web surfing is an example) and, therefore, does not require authentication or encryption.

2) **Level 2: Integrity, Non-refutability**. Information authentication is required. The data is not sensitive, but it must be protected against changes during transmission. It requires authentication, but not encryption.

3) **Level 3: Data Integrity and Privacy**. This data requires some degree of privacy. Both authentication and weak encryption are used.

4) **Level 4: Strong Encryption**. This data needs maximum privacy and, thus, requires authentication and strong encryption.

Often different users or groups of users demand different degrees of authentication and security. To make the implementation practical, the required security level should be negotiated when establishing the VPN connection. Network managers also require the tools necessary to pre-assign security levels by user or user group.

Figure 6 shows a typical traffic segmentation. (The percentages noted here are typical customer statistics provided by a major firewall vendor.)

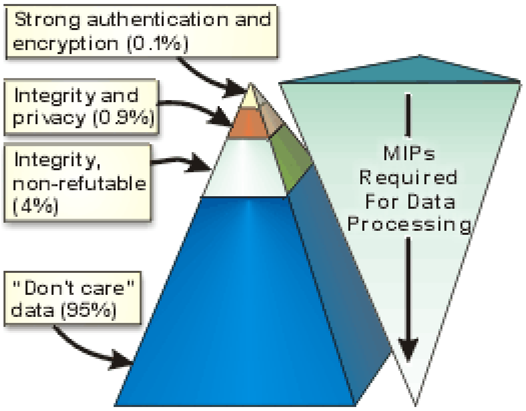
Strong authentication and encryption (0.1%)
Integrity and privacy (0.9%)
Integrity, non-refutable (4%)
"Don't care" data (95%)
MIPs Required For Data Processing

Figure 6. Types of traffic and their security and processing requirements

Select a VPN Solution

There are two basic VPN architectural choices:

- ♦ Service provider independent
- ♦ Service provider dependent

The difference between the two is where the VPN tunnel starts.

Service Provider Independent Model

In a service provider independent model (Figure 7), a VPN-enabled client initiates the tunnel through the public network to the central site.

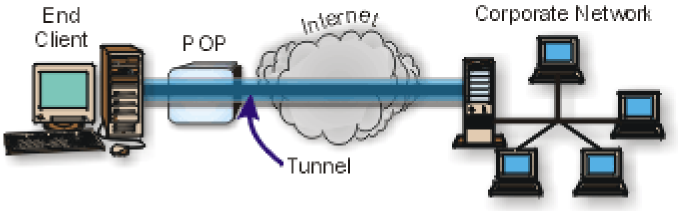

Figure 7. Service provider independent tunneling

To access the corporate network, the client first establishes a PPP (Point-to-Point Protocol) session to a local ISP for Internet access. The client then connects across the Internet to the central site and establishes a tunnel to carry the data traffic.

To the ISP, the tunnel is simply data, and there is no requirement for special processing. The advantage to the corporation is that it can use any POP anywhere in the world, as long as it provides Internet access. The disadvantage is that the client must be VPN-enabled, which could be prohibitively expensive to deploy for a large number of remote users.

Service Provider Dependent Model

With a service provider dependent model (Figure 8), the corporation enters into an agreement with a service provider such as an ISP.

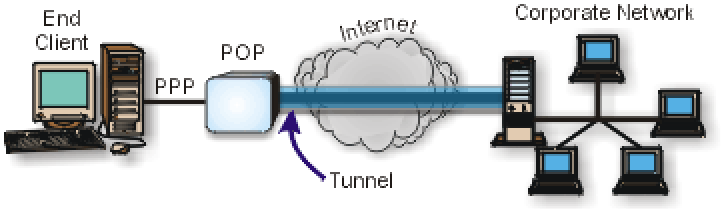

Figure 8. Service provider dependent tunneling

The corporate user dials into a local POP with a PPP client, and the tunnel session is initiated at the POP. The crucial difference is that the client can be any PPP client, such as ShivaRemote™ or Microsoft's Dial-up Networking.

This arrangement can be combined with quality of service agreements to guarantee a level of VPN performance, although few service providers offer true guarantees today.

Deployment is limited by the existence of VPN-enabled POPs. Until standards such as L2TP become widely adopted, a corporation will find it difficult to set up large scale, and especially international, VPN deployments through a service provider.

Where data security is critical, there is also the disadvantage that VPN encryption does not occur until the POP, thus leaving the enterprise's communication unprotected between the remote PC and the POP.

Estimate VPN Performance

In a stable environment with a LAN that is not congested, a LAN-based CPE typically generates data throughput rates equal to the connect speed as boosted by compression. By comparison, VPN users have to contend with issues relating to Internet latency, port availability and connection reliability. In addition, security functionality such as encryption can limit performance because of the extra processing required.

Because the packet overhead associated with carrying a VPN tunnel is minimal, performance problems within the shared network will be caused by congestion on the shared network itself. Some service providers guarantee latency limits on their networks, but these guarantees do not exist when data transits multiple networks.

There is no performance guarantee for the Internet. The latency and nature of the path followed by a tunnel is indeterminate and will vary from session to session. This means that the performance of VPN sessions can vary greatly.

The processing power available at the points of origin and termination of the VPN tunnel also limits performance.

Encryption and decryption are MIPS-intensive operations that can overburden under-powered equipment. Whereas a router usually must process only the packet header, encryption processes each bit of each byte in the packet. For example, triple DES Encryption requires approximately 50-100 times the processing power of straight IP routing.

When data encryption is required, hardware-based solutions are recommended to ensure that origin and destination points will not become bottlenecks. Dedicated hardware solutions also feature tighter physical and logical security. Software-based encryption products (i.e., the main CPU does the work) are generally efficient at data rates of 128Kbps or less.

Infonetics summarize the performance issue by stating:

> *"Because of the computational power required to implement VPNs, hardware-based VPN products deliver the best performance. They also offer tighter physical and logical security. Software-based solutions are best-suited for lower-volume connections for small and medium businesses that have lower security requirements."*

- Copyright © 1997 Infonetics Research, Inc.,
Virtual Private Networks

The Components Of VPN

Tunnels

Tunneling consists of encapsulating packets for secure travel over the shared medium, allowing different protocols to travel though a public IP network.

As the shared network is IP-based, the tunnel must carry traffic over IP. The data inside the tunnel can be either IP-based or some other protocol if the tunnel technology allows it. Multi-protocol tunneling is important where there is a mix of data. For example, previous investments in IPX may mean there are still a number of IPX file servers that remote users need to access through a tunnel.

There are a number of tunneling standards:

- ♦ **L2F (Layer 2 Forwarding)** was developed by Cisco, Nortel and Shiva. It is detailed in an Internet draft, but is unlikely to become a Request for Comment (RFC) or progress any further towards becoming an international standard.

- ♦ **PPTP (Point to Point Tunneling Protocol)** Like L2F, PPTP is detailed in an Internet draft, but is unlikely to become a RFC or standard. (Microsoft has committed to L2TP going forward.) PPTP works at layer 2 of the OSI model. It was designed for client/server operation and, thus, allows only a single point-to-point connection.

- ♦ **L2TP (Layer 2 Tunneling Protocol)** is a hybrid of L2F and PPTP. Like PPTP, it works at layer 2 of the OSI model and was designed for a single point-to-point client/server connection. Multiple protocols can be encapsulated within the tunnel.

  L2TP is currently an Internet draft and is being proposed as an RFC before the Internet Engineering Task Force (IETF). It is moving to become the accepted tunneling standard. The target for L2TP as an RFC is early 1998.

  L2TP eliminates the PPTP’s dependence on IP by abstracting the underlying transport protocol to any packetized protocol, such as IP, X.25 or Frame Relay. The specific network protocol is left up to the vendor implementation.

  L2TP also helps reduce network traffic and enables servers to handle congestion by implementing flow control between the network access system and the home gateway. L2TP is further expected to offer optimal performance over non-managed data networks, such as the Internet, because of its relatively low overhead.

- ♦ **IPSec** defines security on top of standard IP networking. It contains enough functionality to encrypt, authenticate and carry IP-only data through a shared network. It can, therefore, be considered on an equal footing with the other tunneling protocols.

  Compared to PPTP and L2TP, IPSec operates at layer 3 of the OSI model in host-to-host rather than client/server mode. This allows multiple simultaneous tunnels. For example, IPSec would allow a connection to both a central corporate site and the Internet.

None of the tunneling standards include integral encryption. However, the L2TP draft standard recommends the use of IPSec for encryption and key management in IP environments.

Security

Given a tunnel that defines a method for carrying private information across the shared network, encryption can be used to protect the contents of the tunnel. VPNs support standard forms of cryptography, including public key cryptography and DES (Data Encryption Standard) cryptography.

While DES is CPU intensive, public key cryptography can be 1000 times more expensive in terms of CPU cycles. Hence, public key technology is best used as a secure alternative to password-based authentication, and for key distribution, i.e., for distributing DES keys. DES should be used for bulk data encryption.

Compared to public key cryptography, DES is a symmetric cryptosystem. When used for communication, both sender and receiver must know the same secret key, which is used both to encrypt and decrypt the message.

With public key cryptography, keys are generated in pairs known as the public key and the private key. Each party knows its own private key and must never reveal it. Each party publishes its public key. A party can prove its identity by encrypting a challenge with its private key; only its corresponding public key can properly decrypt the encrypted challenge.

Similarly, data can be made private “for your eyes only” by encrypting under a public key. The data can then be decrypted only with the corresponding private key.

If another party learns a valid user’s private key, then they can mimic the valid user. Because of this and the requirement to control and distribute keys, a significant management infrastructure is required to control the technology:

- Certification Authorities (CA) are used to vouch for the validity of keys.
- X.509 Certificates provide a standardized way of representing names and their associated public keys.
- Certificate Revocations Lists (CRLs) list compromised certificates and their keys.
- An X.500 server provides a publicly accessible database for storing certificates and CRLs.

The scale and cost of the infrastructure required for public key cryptography could be so high that small and mid-sized organizations may find it more cost effective to use shared secrets (passwords).

It can also be argued that public key methods do not meet all the needs of remote access because a dial-up termination point, such as a remote access or tunnel server, cannot be proved secure. The dial-in client needs unobstructed access to the latest CRL but, by definition, the remote access server is the dial-up user’s single point of network contact.

The remote access server and the user must mutually authenticate the X.500 database, and the CRL can only be accessed via the unauthenticated and, therefore, as yet untrusted remote access server.

The user needs the CRL to validate the public key presented by the access server. However, if the access server has been compromised, the X.500 server will be unreachable.

Note that the U.S. Government restricts the export of any symmetric cryptography products using keys longer than 40 bits. It also restricts the export of public key cryptography products.

IPSec

IPSec is a collection of security measures that addresses data privacy, integrity, authentication, key management and tunneling. Within IPSec, an authentication header (AH) is used to authenticate the packet, thus ensuring data integrity and authenticity. A technique known as Encapsulation Security Payload (ESP) encrypts the payload to ensure privacy for sensitive data. IPSec version 1 is defined by RFCs 1825 through 1829. IPSec version 2 is likely to attain RFC status during 1998.

Client Technology

In the service provider independent model, the VPN client consists of three basic features: PPP, tunneling and encryption. Deployment of the client generally involves some add-on to a standard dialer. Typically this software comes from the vendor that supplied the central site VPN equipment. However, Microsoft has plans to add an L2TP client with IPSec support to a future release of Windows.

Tunneling and encryption support can be added to a standard PPP client to enable VPN on the client. In the service provider dependent model, the client need only support PPP. A standard Win95 operating system with Dial-up Networking is sufficient.

Making VPN Work

Can VPNs save a company money? The answer lies in the details of a particular company's remote access needs. To deploy the optimized remote access solution, it is necessary to perform a number of tasks:

- Evaluate and analyze all existing telecommunications costs.
- Analyze users. Determine who is making long distance calls and the duration of the calls.
- Establish performance and security requirements.
- Assess costs relating to administration, equipment acquisition, maintenance, management, support and help desk.
- Decide on the degree of direct control required, and find the best partners.
- Finally, decide on the architecture and technology.

The following sections discuss key items to be considered as part of these tasks and present some cost examples.

VPN Deployment—With and Without Service Providers

Relying on a service provider's infrastructure in the Service Provider Dependent model (see *Select a VPN Solution*) eliminates the need for expensive deployment of new software to remote clients. This makes the Service Provider Dependent model particularly attractive for large companies with large numbers of remote workers and sites.

The drawback to the service provider scenario is that some of their networks offer unacceptably slow access. In addition, travelers may find themselves unable to access the provider's network due to limited geographical reach.

Authentication is particularly critical when using a service provider. In this model, the IS manager must understand how users are authenticated. If the service provider controls all user names and passwords, then all additions, deletions and other changes must be made through the service provider.

Robust VPN solutions allow the company to control the authentication of users through solutions such as a proxy RADIUS server in a remote location.

Partner Question List

There are many factors to consider when evaluating a VPN service. The following is a partial list of topics and suggested questions to ask a prospective VPN or VAN partner:

1) Client
   - ♦ What kind of client will you deploy?
   - ♦ What kind of authentication do you do? (Authentication can be based on user name and password, or can use a third party hard or soft token system.)
2) POP infrastructure
   - ♦ What is the port availability/reliability?
   - ♦ What is the server (DNS, RADIUS) availability?
3) Backbone
   - ♦ What is your backbone? (The backbone can be a superhighway or a block on performance. Data throughput will only be as fast as the slowest link.)
   - ♦ What is the end-to-end latency/throughput? What is the effective throughput of the proposed service? (In some cases the actual rate can be very low. For example, even though users connect at 28.8, the effective throughput can be as low as 10 Kbps.)
   - ♦ Can I get a service level agreement (SLA) in which some type of performance guarantee is required?
   - ♦ What is your history of service outages for the last few months?
4) Help and support
   - ♦ Who administers client support and help desk?
   - ♦ How will you support me in addressing user connection issues?
   - ♦ How much visibility is there into the network operations center (NOC)?
   - ♦ Is tiered support available for the help desk?
   - ♦ Do you offer remote diagnostics and monitoring?
   - ♦ What is the nature of your reporting, accounting and billing?
5) CPE equipment
   - ♦ What kind of CPE equipment is deployed?

Integrating VPN with Direct Dial

Optimizing Your Remote Access Solution

To optimize a remote access solution, the IS manager usually must consider integrating VPN connectivity with direct dial access. VPN and direct dial each have their own strengths and weaknesses.

- Cost

  As the example in the following section shows, the major cost in a direct dial solution is the time and distance-based charge to use the telephone system. VPN can greatly reduces long distance charges. However, in some instances, with VPN the fixed monthly charge for ISP access can outweigh savings in telephone charges.

- Performance

  The components limiting performance in a remote access connection are the last mile infrastructure and the technology used at the client site (e.g., 33.6 Kbps modem over a voice circuit). This is common for both VPN and direct dial. However, once past the last mile, a VPN-based connection will never be faster than a direct dial connection. In many cases, because of the latency associated with the Internet, a VPN will be slower than a direct dial connection. Therefore, direct dial is a better option for performance-sensitive connections.

- Implementation

  Setting up a VPN at a central site generally involves less hardware than a direct dial solution. A VPN tunnel terminating device can, for example, be added behind the firewall on the corporation's existing Internet connection. However, VPN products generally require configuration steps that go beyond the set-up requirements of a remote access server. IP addresses, certificates and integration with the firewall must all be considered with the implementation of a VPN solution.

- Technological maturity

  Because VPN products have not been widely tested in production environments, it is the early adopters who are embracing this new remote access technology. Furthermore, VPNs still lack the ease of use that characterizes direct dial products such as Shiva's LanRover® family.

Comparing Costs: An Example

How does the IS manager determine the right remote access solution? Only a detailed assessment of current needs and viable options can result in the most flexible and cost-effective solution. The example below illustrates a hypothetical company and an optimized remote access solution.

Company X has 900 employees. The nature of the company demands that a large number of staff (70%) be enabled for remote access.

The following table suggests a typical split of work patterns and geographical spread:

User type	Number of users	Avg. hours/wk.	Local Use	Long distance
Home Workers	315	2	85%	15%
Travelers	95	2	2%	98%
Telecommuters	113	8	82%	18%

Table 1. User and usage patterns

Extrapolating these patterns of usage gives a total remote access time per month of almost 8,000 hours. However, the critical measure is the split of local to long distance minutes:

	Percentage of remote access
Local users	74%
Long distance users	26%

Table 2. Company X Weekly remote access

Company X has a user port ratio of 10:1. It feeds T1 lines into an access concentrator and uses an 800 service to reduce long distance call charges. Amortizing the capital costs over 36 months and allowing for maintenance and help desk costs, the remote access costs are:

	Cost/month ($)		
Item	Direct Dial	VPN	Hybrid
Monthly hardware costs (capital amortized over 36 months)	875	438	615
Monthly maintenance (technology protection and support plan costs for hardware)	263	131	184
Monthly central telco costs (T1 lines)	2231	669	1511
Monthly long distance costs (800 number)	10765	0	0
Monthly personnel costs (help desk etc)	16,667	16,667	16,667
Monthly ISP costs (staff accounts with local ISPs, $30/account)	0	15750	4007
Total direct dial remote access cost	$30801	$33655	$22984

Table 3. Company X access costs with and without VPN

The hybrid solution uses VPN for long distance and direct dial for local calls. This increases hardware costs but reduces the ISP bill. In addition to the cost benefit, this approach allows the use of direct dial for mission critical applications, or as a backup to Internet problems.

Success Stories

A VPN proved to be the answer for a Boston-based software vendor with a remote programmer in Atlanta who typically spends 50-60 hours per month downloading large files.

By moving from the public switched network to a VPN for this application, the company reduced $360 per month of long distance charges to $35 per month for a flat rate Internet connection.

In another example, a southern power company implemented a VPN link for intermittent communication between its headquarters and a subsidiary in Minneapolis. By eliminating the costly leased line that previously connected the two points, the utility saved over $500 per month.

However, in most cases, corporations will benefit by integrating direct dial and VPN-based remote access. A long-time Fortune 100 Shiva customer has a worldwide audience of internal employees who require local access to many far-flung locales, including corporate headquarters in Massachusetts. In the past, this has led to some very expensive solutions; in many cases, remote users have simply forgone access to the company's worldwide corporate intranet.

A telecommunications consultant for corporate voice and data planning at the company is currently in the process of reversing past remote access shortcomings by investigating products and services for VPN. The company already uses Shiva LanRover Access Switches® for remote access connectivity. Security is integrated by Shiva via a third-party.

The consultant, who views the company's employees as his customers, makes it clear that he is looking to augment, not replace, the current dial-up system. "We don't feel that VPN is a replacement technology," he states. "We think there's still a need for good, local dial-in systems where you get reliable throughput and dedicated access. We feel that VPN is a good partner to that."

How can companies decide what kind of remote access is best for them?

"Listen to the customers' needs," he says. In our case, the company has a mobile sales force that works on the road and at home. We also have a support organization that continues to work after hours at home, and a number of other employees across the corporation who work at home. Between the traditional Shiva remote access solution and a new virtual private networking solution, we feel we can meet all our needs."

Conclusions

The importance of remote access has changed rapidly. Far from being an executive toy or an IS luxury, remote access is now crucial to corporate operations.

Yet, an increase in the scale of remote access deployments, combined with an increase in communication demands, has dramatically increased the costs of owning and maintaining remote access technology. This is predominantly due to the cost of telephone system charges, particularly long distance.

To address these conflicting demands and pressures, some form of VPN will be a standard component in the future of remote access. VPN deployment will expand rapidly when the technology matures and service providers offer cost-efficient data access.

Furthermore, the future is fast upon us. Infonetics Research maintains that the U.S. VPN market – consisting of VPN products, systems integration and service provider services – totaled $205 million in 1997 and will grow over 100% annually through 2001, when it will reach $11.9 billion. Infonetics also states that by March 1998, 92% of large service providers will have plans to offer VPNs.

To maximize the benefits, corporations are integrating VPNs with their current remote access solutions. The result is an optimized remote access solution that reduces the corporation’s overall costs by reducing long distance charges.

About Shiva

Founded in 1985, Shiva Corporation (Nasdaq: SHVA) is a global, full-line provider of remote access solutions, systems and services. Its products enable users and user sites at enterprises of all sizes to connect with corporate information resources, on-line services and the Internet. Shiva offers a full range of award-winning, analog and digital remote access and communications servers. With partnerships that include IBM, Microsoft, Netscape, Motorola and Nortel, Shiva hardware and software are marketed by many of the world's leading computer and communications equipment manufacturers and telecommunications companies. The company is based in Bedford, Massachusetts, with offices worldwide.

Productive and cost-effective remote access solutions are at the heart of Shiva’s business. Shiva continues to provide comprehensive solutions by combining a strong technology focus with a solid commitment to industry-leading remote networking products. In doing so, Shiva helps customers progress to tomorrow’s remote access future while protecting today’s investment.

For more information about Shiva products and services, contact us at **(800) 977-4482** or **(508) 788-3061**. Or visit our Web site at **www.shiva.com**.