A Closer Look at Remote Access:

Can Your Organization Benefit from VPNs?

INDUS
RIVER™

Indus River Networks, Inc. 31 Nagog Park, Acton, MA 01720 phone: 978.266.8100 fax: 978.266.8111 web: www.indusriver.com

Table of Contents

Executive Summary

Virtual Private Networks offer exciting new opportunities for companies to improve their internal and external communications, business systems, and overall responsiveness. By leveraging the tremendous scope and ubiquitous nature of the Internet, organizations are able to enhance communication and sharing of information among employees, customers and partners while benefiting from significantly lower transmission costs.

VPNs come in three distinct application types: Remote Access, Site-to-Site and Extranets. Each application has unique features, benefits, and requirements that are important to understand and evaluate properly. This paper will analyze Remote Access VPNs in detail while briefly covering some aspects of Site-to-Site and Extranet applications.

Remote Access VPNs can potentially provide enormous cost savings over traditional RAS approaches in place today, driven mainly by lower transmission costs. The amount of savings depends on many variables, including the nature of a company's business, its existing connectivity methods, the type of VPN system used, and the size and capability of its support staff. Actual costs savings models are outlined in the paper. Conservatively, 60% savings or greater can be achieved with VPN solutions that pay for themselves in 5 months or less.

Care must be taken to ensure that these cost savings are realized. Though line transmission costs are much lower, new challenges present themselves: increased connectivity steps, greater complexity, the need to monitor Internet usage efficiency, remote troubleshooting problems, increased usage time, and more. These are all solvable problems as long as proper planning is done and good, specialized VPN solutions are chosen.

Network administrators should carefully evaluate VPN products and deployment plans according to some general planning principles, each covered in this paper. The planning factors to be considered by the network planner are:

- ❒ What types of network access policies are needed for my enterprise network?
- ❒ How will security be implemented? Should I be worried?
- ❒ How will my company's remote access capabilities scale with the business and with usage?
- ❒ How will my staff manage the process?
- ❒ How can a VPN system simplify the remote user's connectivity experience, leading to increased productivity?

VPNs are an exciting new area that will have an enormously positive impact on a company's overall performance. The key is to keep the implementation strategy productive and manageable.

The Role of Virtual Private Networks

What is a Virtual Private Network (VPN)?

A Virtual Private Network, or VPN, is a sophisticated enterprise network application designed to run over the Internet – a ubiquitous transport between remote users, branch offices, worldwide partners, and corporate application systems. By leveraging the tremendous scope of the Internet, organizations enable immediate communication and sharing of information among employees, customers and partners while benefiting from significantly lower transmission costs.

VPNs are often discussed in fairly generic terms; however, most systems vendors and industry analysts categorize them into three distinct applications: Remote Access, Site-to-Site and Extranets. Each application has unique features, benefits, and requirements. For example:

- *Remote Access VPNs* provide connectivity for a common interest group of mobile remote users and telecommuters into a central corporate site through the Internet (many-to-few connection relationship).
- *Site-to-Site VPNs* provide connectivity for a common interest group among multi-user branch offices and the central sites through the Internet, replacing more expensive existing leased line or multiple dial-up line methods (few-to-few connection relationship).
- *Extranet VPNs* provide business-to-business (partners, resellers, affiliates, etc.) connectivity for multiple interest groups through the Internet for electronic commerce, product information, business support systems, and other aspects of day to day commerce (a mix of connection relationships).

While all three VPN applications represent a radically different, albeit greatly improved, way of doing business, most industry experts agree that the remote access application is leading in market demand for the following reasons:

- Today's remote access line costs are huge, and network implementations often warrant a radical system redesign as the number of remote users/telecommuters/off-hour workers grows exponentially.
- Although today’s Internet is often labeled “slow and overburdened”, most access facilities provide more than enough bandwidth for the average remote dial-in client (28/33 Kbps modem speeds still dominate).
- Dependable service level agreements (SLAs), along with bulletproof and standardized security offerings, must be in place before any widely deployed site-to-site or extranet VPN will be used.

Analysts agree. Infonetics Research predicts that by 2001 there will be over 8,000,000 Remote Access VPN users in the United States alone.[1]

[1] Study from Infonetics Research Inc. *User Plans for VPN Products and Services, 1998.*

What Technologies are Used to Run a Remote Access VPN?

VPN systems, at a minimum, are designed to provide secure point-to-point (user-to-resource) “tunnels” through the Internet, in effect, emulating a private point-to-point circuit at a much lower cost. However, this technique does present a few significant challenges. Because the Internet is a public transport mechanism, sophisticated security techniques are required for the implementation of secure private networks. In addition, the Internet may not provide the same reliable performance as that of dedicated point-to-point links. Congestion is a common occurrence on the Internet. It results in packet loss and it causes packets to be retransmitted, often in very inefficient ways that adversely affect overall performance. Finally, companies using Internet-based VPNs rely on one or more network service providers to operate their internal data network – at times a discomforting proposition.

There are many technologies available to solve these problems.

Tunnel Protocols are used to establish secure point-to-point transport through the Internet. Most Remote Access VPN systems today use PPTP (extensions to Point-to-Point protocol – PPP), L2TP, or the emerging IPSEC standard. There are two important aspects of the tunnel protocols: performance and security. No one standard approach has been followed by vendors, resulting in a wide variation of technologies and limited interoperability.

Performance through a tunnel can be improved through a variety of techniques, including reducing packet loss, implementing data compression prior to encryption and adapting the transmission window to use the actual Internet Round Trip Time (RTT) as the basis for session flow. Although these techniques can pay huge dividends and greatly increase the overall efficiency of a VPN system, they are not widely implemented.

Tunnel Servers (TS) are used to aggregate all tunnel connections back at the central site or at any other destination point in the enterprise. The tunnel servers must be high performance, highly scalable devices that are capable of handling hundreds or thousands of user connections simultaneously. While they often include some access control, authentication and encryption capabilities, they are designed to be a dedicated tunnel endpoint “processor.”

Authentication techniques are used to validate users prior to establishing the VPN connection. There are a variety of techniques used today to authenticate users, ranging from the operating system username / password, RADIUS authentication, to hardware-based token authentication cards. Commonly based on a challenge-response method, the authentication process originates at the VPN points of access.

Encryption algorithms are used to scramble (disguise) packets as they are transported. Even with user authentication established (remote user Bob is, in fact, remote user Bob), the VPN is not totally secure if packets are transmitted through an unsecured Internet in original plain text. Common sniffing techniques can be used to capture, even alter, information flow if encryption is not implemented. Each end of the tunnel encrypts data prior to transmission and decrypts data once received by the authenticated user.

Although security is an important aspect of the tunnel server, most tunneling protocols don't specify which encryption or authentication technique to use. As a result, many different algorithms are used. The choice of one algorithm versus another (RC-4 versus DES) has far less an effect on security than the actual implementation differences. Key exchange, for example, can be implemented a variety of ways and can greatly affect a system's ability to interoperate with one another.

Isn't My Existing Router/Firewall a VPN System?

A common misconception is that "conventional" routers, firewalls, or server-based software solutions can provide all the capabilities required for a reliable enterprise VPN solution. However, these devices lack the features and scalability needed for most production VPN designs. For example, conventional routers typically do not monitor or control end-user traffic flows. They are designed to process network-level flows and often lack techniques to monitor or troubleshoot a remote user. Also, because multi-purpose routers are not designed to be high performance tunnel aggregation servers, they often lack the throughput required for encryption and compression processing.

Firewalls provide effective end-user authentication and security enforcement but they cannot scale tunnel processing performance across an enterprise network of thousands of users. In addition, they often lack critical features such as multiprotocol encapsulation, necessary for supporting IPX and NetBEUI applications. Server-based software solutions are not scalable and often lack the comprehensive security, management and fault tolerance features needed to maintain a large-scale data-networking infrastructure. To address these critical issues of performance and functionality, dedicated VPN systems have been developed to provide the full complement of system capabilities that are needed to build enterprise-class remote access networks.

A Closer Look at Remote Access VPNs

Remote access VPNs use the Internet for ubiquitous connectivity for remote users. By utilizing local, low cost Internet point-of-presence (POP) access for hundreds or thousands of users, companies can gain enormous cost savings by eliminating the monthly line access charges for traditional dial-up schemes. Ideal applications for VPN solutions include sales force automation, telecommuting and mobile decision-makers.

In many cases, Remote Access VPNs replace a significant portion of an existing RAS (remote access server) solution. The major advantage of replacing a company's RAS solution with a VPN is line cost savings. Dedicated dial-in numbers (usually 800 numbers) and the associated cost of remote access server equipment are the two largest cost items contributing to high remote access costs. With VPNs, companies have a lower cost alternative.

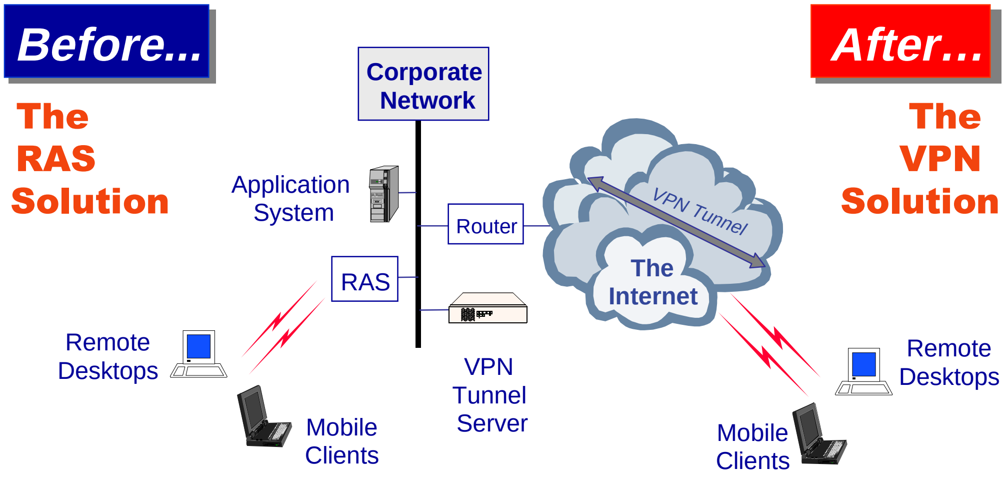

Figure 1 RAS vs. VPN Solutions

Even with the huge cost savings associated with migrating from a RAS solution to a VPN, many companies choose to deploy this technology in stages or operate them in parallel. This approach allows some groups of users to be moved over as appropriate and it allows users who don't have Internet access to maintain connectivity to the central site. In fact, many companies settle on a hybrid approach that provides the best of both worlds.

Cost Savings

What are the real cost differences between a RAS and VPN? A simple example illustrates the potential savings[2].

Suppose a company has 500 remote users (25% of the total employee population of 2000). The remote users include 250 mobile workers (travelers), 50 telecommuters, and 200 after-hour workers (connected from home). Both mobile workers and night workers average 30 minutes connection time per day, 20 days a month, while telecommuters average 240 minutes (4 hours) per day, also for 20 days / month. Ignoring equipment start-up costs and maintenance charges, assume that line costs average $.08 / minute domestically (toll call or 800-service) and $0.60 / min internationally. Usage breakdown is as follows: mobile / after-hour workers use domestic service 90% and international service 10%. Additionally telecommuters use a dedicated $30 / month line. The line costs alone for this scenario are $40,590 per month. Annual setup and maintenance costs could easily add another 10% to this total.

Now consider a Remote Access VPN. Assume the same usage breakdown and the same usage time, a per-user Internet access fee of $20 / month, international ISP roaming fees of about $10 / hour, and a VPN system cost per user of $200 (this includes all central site hardware, client software and a network management system). Overall, the Remote Access VPN costs about $203,000 per year, a savings of more than 60% on the operating costs of the traditional dial-in system. To examine the details of this cost comparison, refer to the cost of ownership model provided in Appendix A.

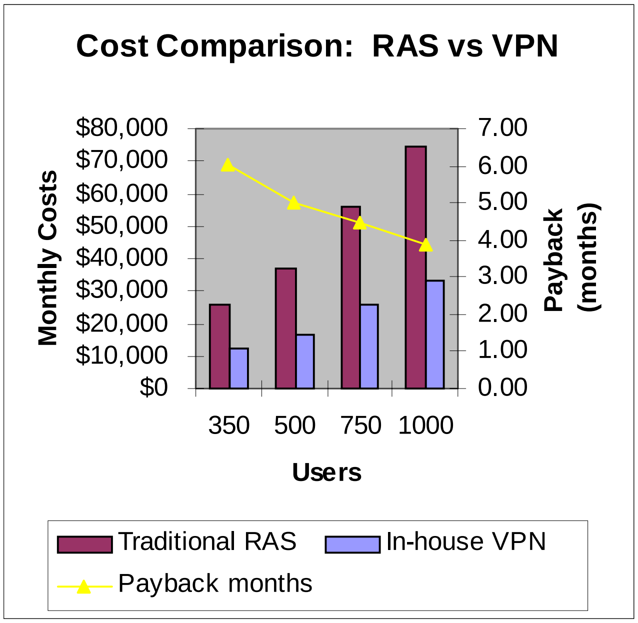

Assumptions

Line Costs	
800 numbers	$.08
US toll average	$.08
Intl average	$.60
Remote user profile	
Mobile worker	50%
Telecommuter	10%
Day extender	40%
Usage by profile	
Mobile worker	30 min/day
Telecommuter	240 min/day
Day extender	30 min/day

Figure 2 Comparing the Costs of RAS and VPNs

[2] Refer to Appendix A in this paper for a complete cost analysis.

It is important to note how quickly the Remote Access VPN system pays for itself through line costs savings – **less than 5 months** for a 500 user VPN! This cost savings will become even more compelling in subsequent years as the number of remote users grows significantly. How significantly? According to Infonetics Research, the number of remote access users per organization is expected to more than double over the next 2 years.[3]

Cost savings aren't the only benefit realized when moving from a RAS solution to a VPN. Other potential benefits include more connectivity flexibility (56K, ISDN, cable modems, DSL, etc.), better global connectivity (number of Internet POPs versus private company access), better availability (SLAs that offer 99.7% uptime), and a reduced management and support staff since much of this is now out-sourced.

So What's the Catch?

There are some significant challenges to overcome.

- ❐ The previous cost models assume that Internet (POP) access is a local call (nearly free), but in many areas, including most countries outside North America, it is not. Line costs must be managed to prevent potentially severe impacts on the savings outlined earlier. Unfortunately, this is a problem that is very difficult to solve! How does a user know if a call is local? Rate schemes for most calls do not follow geographic convention (e.g., state / county lines), therefore, it is a major challenge to predict call rates in many areas. Worse yet, many remote users (50% in the cost model) dial in while traveling and have no idea what the call rates are away from home. The ideal solution to this problem is for the VPN's client software to automatically select the appropriate call-in number based on source and destination matching and the preferred ISP.
- ❐ User complexity increases with the additional steps required for remote access (connection through the Internet). Unless this challenge is overcome, network administrators trade toll costs for productivity costs.
- ❐ Some ISPs and third party companies offer an Internet "roaming" capability within areas where there is no local POP for ubiquitous connectivity. Roaming charges can average from $0.05 / minute ($3.00 / hr) up to $0.25 / min ($15 / hr) or higher depending on the location. These charges are often necessary for some of the remote user population and must be factored into the overall VPN costs.

In addition to the high line costs associated with remote access, the cost of end-user support and fault management is also significant. Infonetics Research found that although network management and help desk costs typically average about 15% of the total remote access costs, the "hidden" costs of supporting remote access users themselves is even higher. In fact, Infonetics found that remote users, on average, are forced to spend 6-10 hours / month installing, reconfiguring, and troubleshooting software and / or connection parameters.

[3] Study from Infonetics Research Inc. *User Plans for VPN Products and Services*, 1998.

Imagine the opportunity cost of the remote users' time coupled with the additional burden back at the help desk (typically after hours support). In Company XYZ the opportunity cost would be close to $165,000: 500 users x 6hrs / month x $50 / hr cost = $150,000 / year, plus approximately $15,000 for the additional help desk support costs. Suddenly these costs represent 30–40% of the overall remote access costs. And, this doesn't take into account the lost opportunity costs of downtime incurred when a sales rep is NOT working on revenue generation tasks for the company!

The Road Warrior's Nightmare

Remote access has traditionally posed significant challenges for users. To put these problems in human perspective, consider the beleaguered "road warriors" It's late at night, he/she is alone in a hotel room after an exhausting workday, wanting simply to connect into the central office to download the day's E-mail and the information needed for tomorrow's big meeting. Many things go wrong … PC/modem configuration problems (what does error 13 mean?), can't get a dial tone (what's the hotel's outside line access code?), forgotten passwords, noisy line (session disconnected – do you want to retry?), huge call rates and surcharges (wait until he gets his bill!), etc.

Now substitute a Remote Access VPN and bingo – problems go away, right? Unfortunately, no. All of the above problems and others still exist with conventional VPNs. New challenges include required knowledge of local ISP POPs (where's my dial-up number list?), potential IP addressing problems (what does "error - can't resolve network address mean?"), Internet congestion (really slow downloads late into the night), and, of course, busy signals.

In the midst of all of these problems, the beleaguered road warrior – perhaps one of the top salespeople in the company – gives up, picks up the phone, and calls the emergency help desk (an expensive resource late at night). While on the phone with the support personnel, they are burning valuable work time needed to prepare for tomorrow's revenue-producing activity.

There has to be a better way – and there is … ***insist on a VPN system that provides user simplification.***

From a manageability standpoint, Remote Access VPNs must be viewed differently. Unlike traditional RAS or even conventional LAN environments where the network "edge" is usually an onsite intelligent wire closet switch, the network endpoint of a VPN is the remote user's laptop itself. Because the transport network is a third party service, the remote user's desktop must be managed as if it is a network node. For this reason, it is vital that any Remote Access VPN system include intelligent client software that can monitor the local ISP connection and manage the performance, connection policies, and scalability of the VPN system as usage increases.

The conclusion: Remote Access VPNs offer significant cost savings potential over traditional RAS approaches, conservatively estimated at 60% or greater. Care must be taken, however, to keep user connectivity simple and easy-to-use, and to make intelligent Internet access decisions to ensure that the savings and benefits are realized. A good VPN solution will provide these important features.

Establishing VPN Design Requirements

Where Do I Begin?

There are some basic planning tasks that must be done prior to VPN design and implementation. For example, what are the overall network policies? What is the security strategy, and how will it be implemented? How do I integrate my VPN with my existing network infrastructure? How do I select an ISP?

Gartner Group recently published a five-tier model outlining the Enterprise and Application Needs for enterprise networks. Shown in Figure 3, it outlines five “variables” to consider when considering a VPN solution: Security, Scalability, Manageability, Simplicity, and Quality of Service.

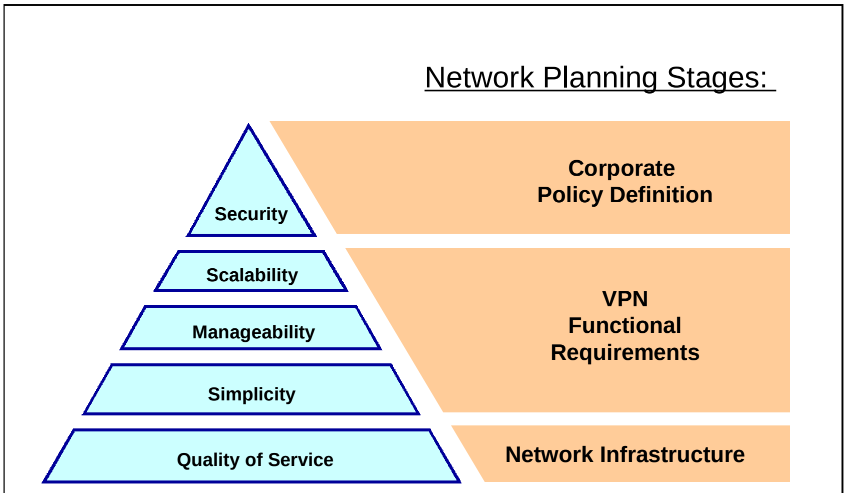

Source: Gartner Group, “Enterprise and Application Needs”, Information Security Conference, June10-12, 1998

Figure 3 The Three Essential Requirements of a Remote Access VPN

Using Gartner's model as a starting point, the overall task of designing an enterprise VPN solution can be broken down into three fundamental planning components:

- ❐ Defining the overall *policy* rules for all users and resource of the network.
- ❐ Defining the *functional requirements* of the VPN to provide appropriate service to users.
- ❐ Planning the overall *network infrastructure* needed to integrate the VPN solution and Internet services into the existing enterprise network without reducing remote user productivity or increasing costs.

Policy Definition

Network policies are rules established by network administrators to govern the usage of the network and all corresponding resources. Many network policies need to be established prior to designing and installing the VPN. Fortunately, many of these policy definitions may already be defined and in place for the existing network infrastructure (pre-VPN) and need not be changed.

Many policies exist for an enterprise network; however, most fall within 2 or 3 main groups: security, cost thresholds, and class-of-service usage privileges. Some security policies, such as access rights, are usually straightforward rules and are often defined by organizational boundaries (Marketing, Engineering, Finance, Sales, etc.) and enforced at the application server. Others, such as network-wide policy components like bandwidth commitments (class of service) require specialized VPN systems that are capable of handling policy-based network access controls end-to-end. Most systems and services do not provide this today but likely will in the future.

Policy definition must be done up-front in order to make optimal implementation decisions as the VPN is designed. Though cost thresholds and class-of-service policies are important to consider, proper implementation of network-wide security often ranks highest on the network planner's list of priorities. In its VPN study, Infonetics Research reports that of all the concerns users have about VPN implementations, security is, by far, number one. Below is a closer look at the choices available once the objectives are set.

Security Strategy: Implementation Choices

One of the most commonly used methods to provide security from unauthorized network access is the use of password-based authentication and encryption keys. Great care must be taken when using passwords to avoid some simple mistakes that can significantly weaken security. Simple, easy-to-remember passwords are also easy-to break. Longer and more complex passwords are much more resistant to attacks but often meet with more resistance from the user community. Network administrators must examine the risk / cost / complexity trade-off for VPNs before deciding what level of security is appropriate for each application.

Before reviewing the various techniques used to implement a secure VPN, it's worth noting that physical infrastructures and human factors play a critical role in any security strategy. The first line of defense for any network is physical security; i.e., restricting access to network resources such as important servers, internetworking devices, access hubs and VPN systems. Another critical element is controlled "human factors" – confidentiality of passwords, restricted sharing of PCs and network equipment, employee-only access points, etc. Regardless of how advanced the authentication, authorization, and encryption techniques, failure to implement uniform security policies will leave any enterprise network open to unwelcome access at its weakest points.

Once the first-line physical infrastructures and human factors are in place, a Remote Access VPN requires authentication and encryption technology. Authentication ensures that the remote user is, in fact, an authorized user and encryption scrambles data traffic as it transverses the VPN (Internet). Following is a closer look at each.

Authentication

The *native database* password, also known as a "shared secret", is the most common authentication approach in use today. While simple to create and manage, password-based authentication techniques such as PAP, which send the password through the network, can be vulnerable to a variety of attacks such as network monitoring, man-in-the middle attacks and password detection. For this reason, they are often coupled with a challenge-response protocol such as CHAP (Challenge Handshake Authentication Protocol).

Coupled with CHAP, the native database password (shared secret) authentication technique works as follows:

1. Network administrator enters usernames/passwords into a database (VPN management server).
2. User requests a connection (a login request).
3. Server sends back a random challenge.
4. User calculates a hash based on a one-way function using the random challenge and the password and sends the hash back to the server.
5. Server calculates the hash based on the same one-way function, the random challenge and the user's password and compares the result with the user's response. If a match results, the authentication is successful.

Increasingly, disparate remote access and VPN systems are managed by a central user account and policy database. The Remote Access Dial-In User Services (RADIUS) protocol was the first step in unifying devices from multiple vendors into one authentication scheme. Any VPN or RAS device that supports RADIUS can authenticate against a central RADIUS server which not only defines the user names and passwords, but can also maintain remote user policies such as IP addresses, allowed length of call and number of concurrent logins. The RADIUS server authentication process, slightly different from the one outlined above, is shown in Figure 4.

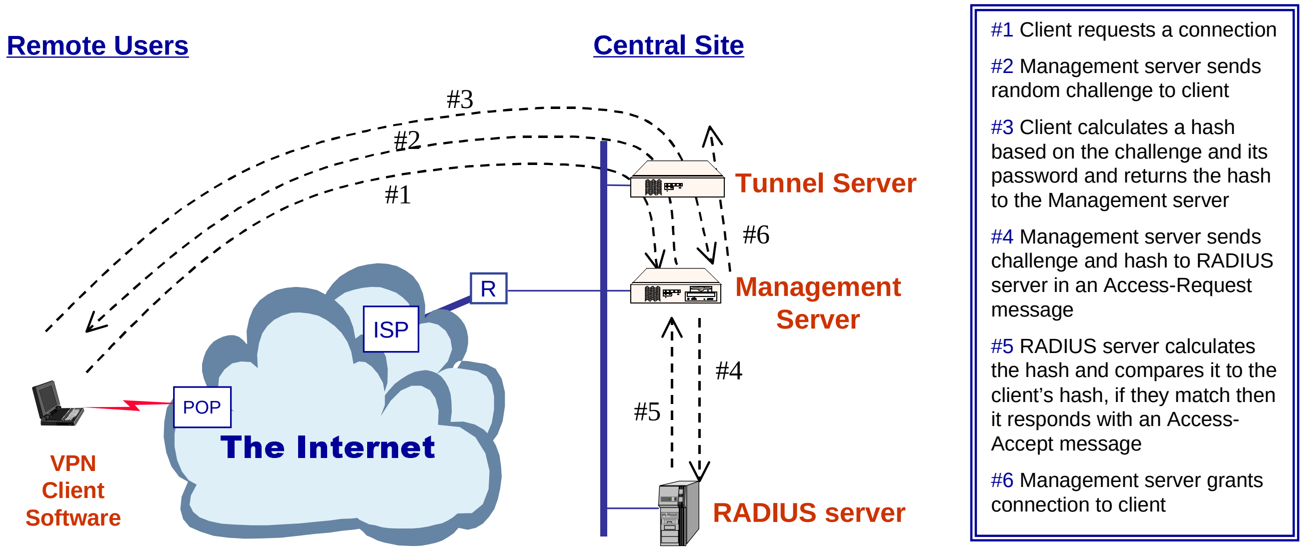

#1 Client requests a connection

#2 Management server sends random challenge to client

#3 Client calculates a hash based on the challenge and its password and returns the hash to the Management server

#4 Management server sends challenge and hash to RADIUS server in an Access-Request message

#5 RADIUS server calculates the hash and compares it to the client's hash, if they match then it responds with an Access-Accept message

#6 Management server grants connection to client

Figure 4 User Authentication Using a RADIUS Server

There are initiatives underway to consolidate these central access lists with the existing authentication frameworks for the enterprise LAN, such as Microsoft NT Domains and Novell NetWare Directories. This allows directory-based management and control of all aspects of the enterprise network, including remote user access. In fact, Microsoft and Novell have released RADIUS support for their respective network operating systems in order to provide remote user access management.

"Stronger" Authentication Methods

Two-factor authentication provides additional security improvements over the password protection schemes described earlier. To authenticate using the two-factor method, the user must have a unique password (something he knows) and the authentication token (something he has). The authentication token may be as rigid as a hardware-based token card, such as SecurIDR from Security Dynamics, or a software-based algorithm specific to an individual user's machine. Although some two-factor authentication solutions are proprietary, some, such as SecurID, support the RADIUS protocol and can be easily integrated into many VPNs.

Encryption

Common encryption algorithms in use today are based on shared secret techniques that use the same “key” (or fixed-length bit string) for both encryption and decryption. Encryption is performed by running an algorithm (based on the value of the key) against the plain text to produce an encrypted output. Decryption is done at the other end using the same key to recover the plain text.

The bit length of the key determines the “resistance” of the encryption algorithm to brute force attacks; a 32-bit length yields 4.3 x 10^9 alternative keys, while a 128-bit key yields 3.4 x 10^{38} alternative keys. The 128-bit key would take thousands of years of computational time to perform an exhaustive search; therefore, it’s not susceptible to brute force attacks. RC-4 (usually 40 or 128 bit), DES (56 bit), and Triple-DES (128 bit) are the common shared secret encryption algorithms used in most VPNs today.

A public key algorithm, on the other hand, uses a pair of keys; one is public and the other is kept private. For example, if Alice is to transmit data to Bob, she encrypts the data with Bob’s public key (B2) so that only Bob can decrypt the data with his private key (B1). Popular public key encryption algorithms include Diffie-Hellman (DH) and Rivest Shamir Adleman (RSA).

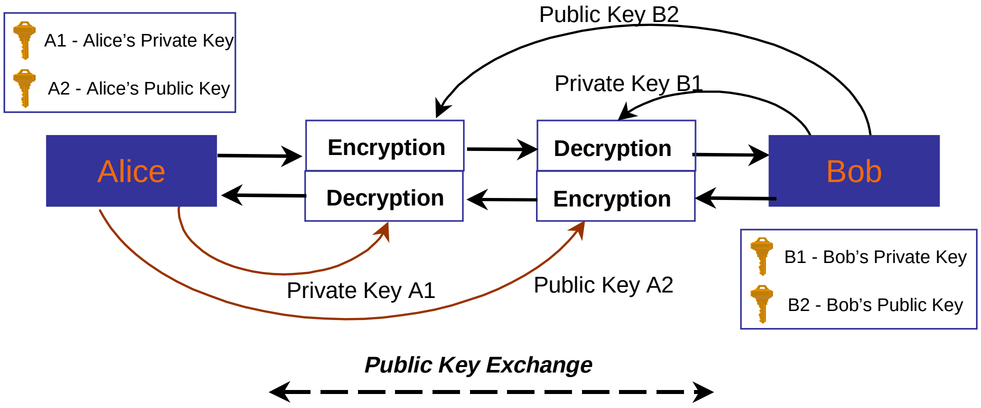

Figure 5 Basic Public Key Encryption Methodology

Key management is also an evolving technology category. There are many proposals for standardization under consideration; ISAKMP / Oakley is the most popular. For two implementations of PPTP or IPSEC to be compatible, they must both use the same key agreement technique. While many vendors claim to fully support one method of key agreement or another, different implementations have led to incompatibility among the various devices.

Digital certificates and certificate authorities combine to provide a public key infrastructure (PKI), where the user's public key is shared with other users and only the user knows the private key. Corporations can implement their own certificate authorities using Entrust Technology's Entrust® or Verisign's OnSite®, for example. Verisign also provides a service whereby they act as a trusted certificate authority on your behalf.

IPSEC

According to Infonetics Research, the two tunneling protocols that are targeted for deployment over the next 12-24 months are PPTP and IPSEC. PPTP is a layer 2 tunnel protocol with inherent advantages such as multi-protocol encapsulation of IP and non-IP traffic. It is widely deployed as a VPN protocol today and is available on Microsoft Windows 95 desktops and NT servers.

IPSEC is now viewed by the marketplace as the next generation tunnel protocol for wide spread deployment. It offers a more comprehensive security framework with robust authentication and encryption schemes being defined by the IETF. Specifically, IPSEC's Encapsulating Security Payload (ESP) provides authentication and encryption for IP datagrams, with the user defining which encryption algorithm is used. The IPSEC standard, nearing completion, broadly defines how key management, encryption and digital certificates fit together.

IPSEC, likely to become the most commonly used protocol in Remote Access VPNs, is a significant step forward towards standardizing VPN solutions; however on its own, it cannot provide for multi-vendor interoperability. For example, IPSEC does not define specific techniques for key management (PKI), network address assignment techniques, compression algorithms, route forwarding protocols, multiprotocol techniques, and other important factors which are key to realizing truly interoperable Remote Access VPN systems. Since IPSEC is an evolving protocol, it is important to select a VPN solution that allows easy migration to the emerging standards, both at the central site and the remote user's workstation.

✓ **RECOMMENDATION**

At a minimum, use CHAP and proven encryption techniques such as MPPE. For greater security, use approaches such as token-based authentication and an IPSEC-based encryption. Choose a VPN solution that offers a smooth migration to evolving IPSEC interoperability. Longer term, certificates are on the horizon – watch for maturity.

VPN Functional Requirements

Once the basic policy definitions and implementation methods are chosen, network administrators must understand the functional characteristics of the VPN system to ensure that VPN features match the demands of remote users and network administrators. For VPN products, the functional requirements fall into three categories: *Scalability* attributes, *Manageability* characteristics, and *Simplicity* of the solution.

Scalability Requirements

There are a number of factors to consider when defining the scalability of a VPN system. How many users must be supported? What is the projected growth of the user population? How will all the users be "installed" on the network? Who will support the users as the population grows? How will authentication scale? What types of reports are required?

One of the most difficult VPN design steps is to estimate the "scale" of the system, specifically the usage requirements: number of remote users, amount of time logged in and performance needs. Infonetics Research found in their 1997 study that 12% of users, on average, connect remotely; however, there is no doubt that this percentage will increase dramatically. Infonetics expects it to double, as a percentage, in two years. Driving factors for this continued growth include:

- ❐ Lower cost connectivity (unlimited Internet access for $19.95 / month or lower, to local POPs)
- ❐ Higher speed modems that allow for more practical application use when dialing in (it's fast enough to use!)
- ❐ More powerful laptop machines that allow users to take their work with them
- ❐ Further proliferation of the Internet (expanding the electronic business environment).

Considering these factors, one planning scenario is the overall number of remote users will grow significantly, certainly at a rate much greater than the corporation's overall headcount growth. Another likely scenario is that connection times will increase because users will have lower cost connectivity from virtually anywhere and improved performance as POPs upgrade their modem banks to higher speed connections. Furthermore, continued improvement in tools and applications will allow employees to easily work while away from the office.

User Population Planning

User population is a critical design element for any VPN. First, it drives tunnel capacity requirements. The number of users is directly proportional to the number of tunnels (every user requires a unique tunnel from laptop to corporate network – no exceptions). It is unlikely that every user will be logged in simultaneously; therefore, 1500 remote users may need only 500 – 1000 tunnel capacity (2:1 – 3:1 oversubscription), though increasing usage time will push the oversubscription number lower. Tunnel capacity must also be supported in a distributed fashion. Some corporations do not have a central site environment and often require purely distributed VPNs to handle this capacity across any / all individual tunnel servers.

Second, user population, combined with growing usage characteristics, increases the load on the dedicated Internet / VPN connection for the central site. It also increases the load on authentication systems, servers, firewalls, routers, and other resources used by the remote user community. To preserve performance levels (response time) these devices must scale along with the loading. See the central site bandwidth-planning model in Appendix B for a more detailed discussion.

Third, as usage and user population increase, the management and overall support burdens grow significantly and the likelihood of problems increase.

Clearly, user population planning has several significant side effects on size, performance, and overall availability of a VPN. For this reason, overall employee growth and remote access usage must be carefully considered.

✓ **RECOMMENDATION**

Plan for growth aggressively, not conservatively. Aggressive usage planning in the initial design will avoid capacity constraints later.

Maximizing Link Performance

Like all router-based internetworks, the Internet's actual performance versus theoretical maximum (wire speed) is adversely affected by two performance "killers": latency and packet loss. The severity of the performance impact caused by latency and packet loss is defined by the characteristics of the VPN's path through the Internet, congestion on that path and VPN protocol implementation.

Latency

The path a packet takes through the Internet determines the latency (or round trip delay). Greater latency slows down interactive applications. More importantly, it negatively impacts protocols such as TCP that don't effectively manage window size (the number of packets in simultaneous transit). Unfortunately, most VPN products use TCP to guarantee data channel delivery and therefore fall victim to this window-size management inefficiency.

Even though the Internet carries huge amounts of bursty traffic and is often oversubscribed (congestion), latency is fairly predictable and often averages about 200-400msec (one way). The latency value is largely a function of the router hops (more routers = higher latency) and end-station processing power. Very little latency is contributed from the links themselves.

Packet Loss

Packet loss accounts for a much higher performance degradation than latency. Lost packets require source-end retransmission after the destination has discovered packet loss. This process takes at least three times the average 200 msec latency because it requires the loss recognition, retransmit request, and the retransmission itself to recover. Worse yet, many transport protocols (TCP) will discard all packets in the window if a single packet is lost. Consequently, a 20% packet loss rate can amplify to over 50% during peak congestion periods, resulting in huge VPN performance degradation.

Packet loss may have adverse effects on some encryption and compression algorithms. They require ordered delivery of packets, therefore, packet loss causes additional performance problems.

Can Packet Loss and Latency Effects Be Avoided?

No, but there are techniques that can be implemented to minimize the effects. One obvious way to improve the performance of any network link is to use compression techniques. The effect is to reduce the overall traffic volume thereby reducing the overall number transmitted packets. Fewer packets will yield less transmission time and better overall throughput.

A general problem with many VPN products is that encryption is done prior to compression. Unfortunately, good encryption creates random looking data that cannot be efficiently compressed by modems. Unless the VPN software compresses data prior to encryption, encrypted packets must be sent in the expanded form. For E-mail and text based data this can reduce performance by 40-70% compared to the performance of a dialed up remote access connection that doesn’t require end-to-end encryption. To solve this problem, data should be compressed prior to encryption.

A second way to minimize overall latency is to adjust the transmission window size (the number of packets sent prior to receiving an acknowledgment of receipt). This approach maintains a heavy channel load during normal transmission while minimizing the recovery time when packets are dropped. During normal data transmission, it's advantageous to increase the transmit window size as much as possible since acknowledgments are kept at a minimum and the channel is loaded to capacity. However, when packets are dropped (during congestion periods), TCP requests retransmission of all packets that come after the lost packet. Since they have already arrived at the destination, this is a very inefficient method for recovering packets.

The window size needs to be balanced. It should be large enough to accommodate increased Internet load without incurring packet loss and small enough to allow for efficient packet recovery when packets are lost. For example, Microsoft’s PPTP[4] uses a window size of only 3 packets (acknowledgments are required after every 3 packets sent), resulting in good packet loss recovery attributes but yielding higher than necessary idle time (latency). A better approach is to “learn” the optimum window size by measuring the packet loss level and adjusting window size up or down accordingly.

Finally, selective retransmission of lost packets at the transport layer greatly improves overall throughput. A transport that synchronizes packets and retransmits only lost packets is much more efficient than basic TCP. By performing these functions at the transport layer, compression and encryption can view the network as reliable, which allows them to perform optimally, greatly accelerating performance

Scalable Tunnel Server Architecture

A VPN’s tunnel server is the point at which user access to the corporation’s enterprise network is terminated. The tunnel server’s performance characteristics must scale to accommodate the bandwidth requirements driven by the remote user community (as outlined in Appendix B). In many cases, bandwidth is defined by the high-speed Internet connection to the central site. Usually this connection is a T1 line (or greater) and requires a tunnel server architecture capable of “filling the pipe”. As multiple links are added to increase the central site capacity, the tunnel server’s performance must scale proportionately.

Because no server architecture can scale without limit, tunnel capacity, at some point, becomes limiting. For this reason, it is advantageous to “cluster” tunnel servers to distribute the tunnel processing, using round robin techniques to balance the user load across multiple engines. Shown in Figure 6, clustered tunnel servers can be configured as a pure cluster (with a single Internet access point) or as completely parallel servers (each carries a unique Internet attachment).

[4] Dial-Up Networking 1.2 used with Windows 95

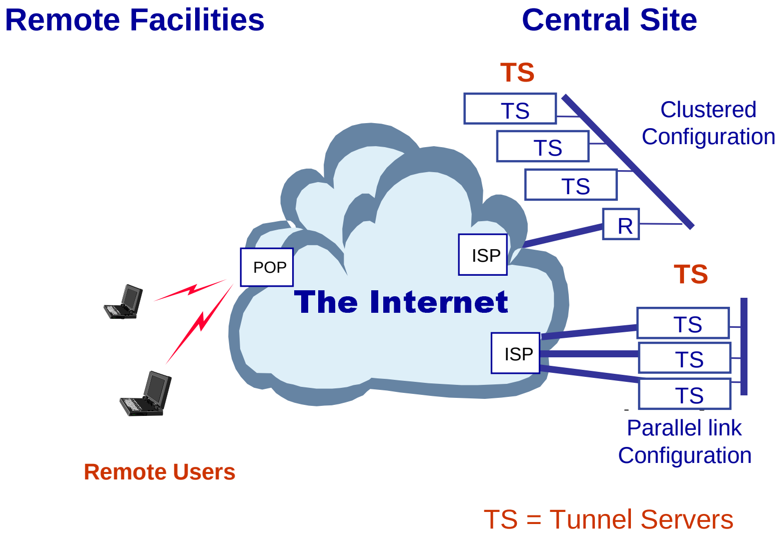

Figure 6 Clustered Tunnel Server Configurations

Multiple tunnel servers do not necessarily have to be clustered; they can be used in a distributed fashion to allow network administrators to locate tunnel concentration physically close to the application servers. A distributed tunnel server design still provides load sharing characteristics (though not as comprehensive as the clustering approach) while providing direct access to decentralized resources. This approach will further reduce transmission costs by connecting LAN-attached remote office workers to the VPN using a single dedicated Internet access line.

VPNs designed with distributed servers provide better overall fault reliability since users are no longer dependent on a single tunnel server. Though redundant power supplies, watchdog timers, and other fault tolerant features are frequently designed into server hardware, it is impossible to completely eliminate the possibility of a system failing. Instead, network designers often prefer a distributed design that spreads the user load across multiple systems, increasing their uptime potential.

✓ **RECOMMENDATION**

Choose a VPN system with proven performance claims and quantifiable scalability attributes. Important scalability attributes to look for include performance measurements (using data compression), system capacity numbers, protocol “extensions” that can boost transmission efficiency, and clustering techniques to distribute overall processing load.

End-to-End Manageability

Like any enterprise network, a VPN must be manageable from a central control point. Single views of user population, login status, health of the VPN system, performance and traffic loading are all important network characteristics. VPNs, as opposed to traditional private networks, add additional management challenges such as managing the remote desktops and managing the Internet itself.

Getting Users Connected

VPN management starts with configuration management and often consists of easy to use client-access lists that commonly include some aspects of usage privileges and authentication techniques. Centralizing this function yields a much more secure operating environment that simplifies adds, moves, and changes. IP addresses can be managed within this same system as well.

Even though the management system is centralized, it must provide the ability to extend out to the remote user’s laptop to ensure that remote login procedures are followed. Ideally, a robust VPN system will include a client application that not only autodials the local POP location for the user, but also reports back to the central site on connection status.

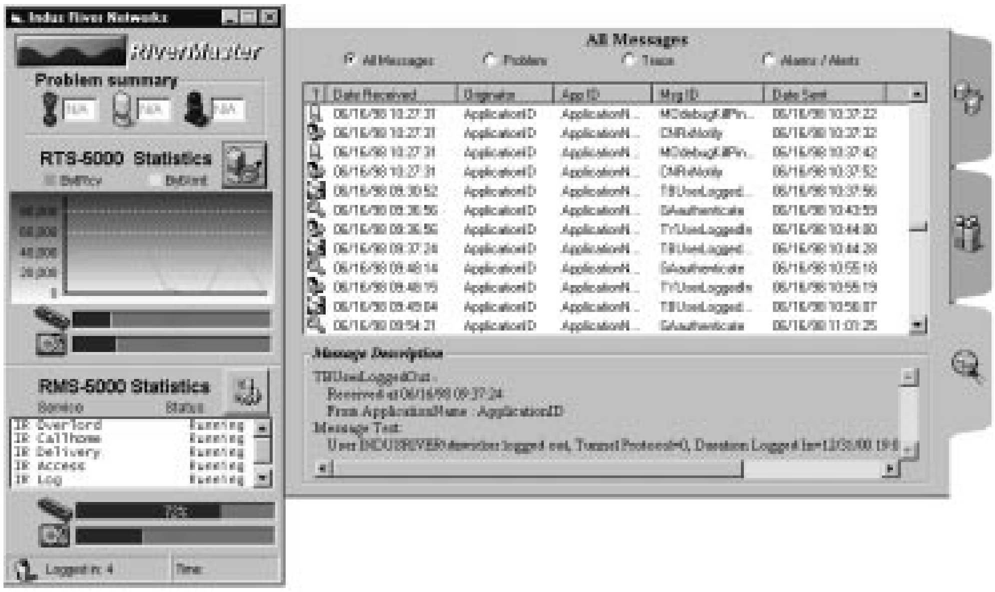

Service	Status
IR Overlord	Running
IR Callhome	Running
IR Delivery	Running
IR Access	Running
IR Log	Running

!	Date Received	Originator	App ID	Msg ID	Date Sent
	06/16/98 10:27:31	ApplicationID	ApplicationN...	MCdebugKillPin...	06/16/98 10:37:22
	06/16/98 10:27:31	ApplicationID	ApplicationN...	CMRxNotify	06/16/98 10:37:32
	06/16/98 10:27:31	ApplicationID	ApplicationN...	MCdebugKillPin...	06/16/98 10:37:42
	06/16/98 10:27:31	ApplicationID	ApplicationN...	CMRxNotify	06/16/98 10:37:52
	06/16/98 09:30:52	ApplicationID	ApplicationN...	TBUserLogged...	06/16/98 10:37:56
	06/16/98 09:36:56	ApplicationID	ApplicationN...	GAauthenticate	06/16/98 10:43:59
	06/16/98 09:36:56	ApplicationID	ApplicationN...	TYUserLoggedIn	06/16/98 10:44:00
	06/16/98 09:37:24	ApplicationID	ApplicationN...	TBUserLogged...	06/16/98 10:44:28
	06/16/98 09:48:14	ApplicationID	ApplicationN...	GAauthenticate	06/16/98 10:55:18
	06/16/98 09:48:15	ApplicationID	ApplicationN...	TYUserLoggedIn	06/16/98 10:55:19
	06/16/98 09:49:04	ApplicationID	ApplicationN...	TBUserLogged...	06/16/98 10:56:07
	06/16/98 09:54:21	ApplicationID	ApplicationN...	GAauthenticate	06/16/98 11:01:25

Message Description

TBUserLoggedOut -
  Received at 06/16/98 09:37:24
  From ApplicationName : ApplicationID
Message Text:
  User INDUSRIVER\[illegible] logged out, Tunnel Protocol=0, Duration Logged In=12/31/00 19:0

Figure 7 Example of a VPN-Wide Login Status View

Cost Control, Security Rules, Connection Management

Runtime monitoring is often easier with a centralized management function. Though centralized, remote desktops serve very important management roles here as well. All connection policies must be enforced at the network edge (the desktop) in accordance with centrally defined policies. Examples of some of these policies include:

- ❒ Auto-dialing the least expensive, speed-matched POP from any location in the world
- ❒ Terminating and reestablishing an alternate Internet connection if availability or performance is poor
- ❒ Monitoring user login time and restricting (or terminating) connectivity if cost thresholds are exceeded

As a result, real-time monitoring is done by the client application in the background while the remote user is using the VPN. Policy control elements (disconnect due to exceedingly high login costs, for example) can be initiated if needed.

High Availability – Keep the VPN Running

Overall network availability and fault monitoring should also be managed centrally. It would quickly become overwhelming if all network faults were reported directly to the network administrator from the various VPN components. The remote desktops serve an important management role here. Intelligent client applications can monitor the connection process and automatically fix configuration problems (e.g., modem string error), download current revisions of drivers or utilities and maintain the most recent ISP POP information to ensure best-possible POP connections. This level of VPN intelligence will enable quick troubleshooting, resulting in less downtime and higher overall VPN availability.

Can the Internet Be Managed?

The tools described above can also be used to manage the Internet connection. VPN systems with feature-rich client applications, tunnel servers, and management systems can effectively monitor basic connectivity parameters (calls answered vs. busy signals, connection rates, number of disconnects) of the ISP, providing a quick comparison of real usage versus committed service level agreements. More advanced usage parameters can be reported as well, including Internet performance (round trip delay characteristics), average authentication time for each ISP and Internet availability provided by each ISP.

Are Any Management Operations (Other Than the Desktops) Distributed?

Yes. Since all support personnel are not centrally located and don't work 24 hours a day in the office, there needs to be a mechanism for providing distributed management operations. The best approach is to centralize the management server and distribute management applications throughout the VPN, including the ISP facility. A management server should collect information from all VPN components while presenting summary data to the management application. The management application can then be distributed anywhere within the VPN to allow network administrators to "manage from anywhere."

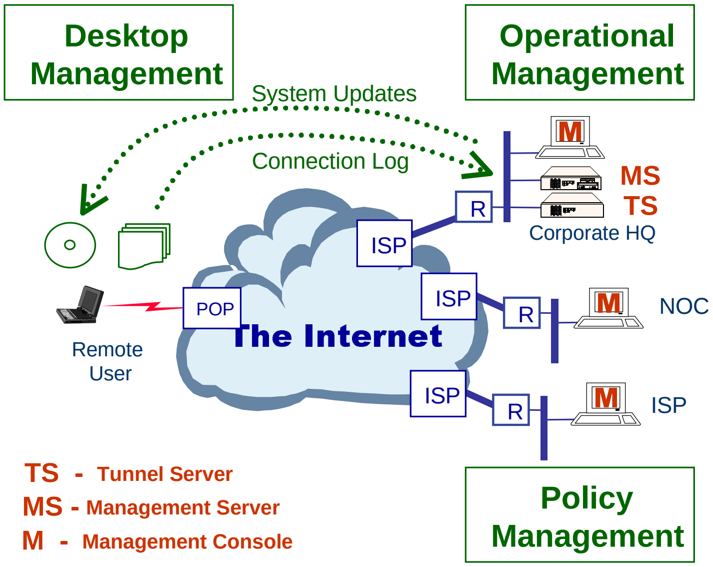

Figure 8 Distributed Management System with Centralized Control

A distributed management architecture for VPNs has another significant advantage – the ability to easily split management operations between customer site and ISP. As part of an overall service agreement, many network administrators will out-source some management elements to the service provider. A distributed management architecture allows them to use common, integrated management components to obtain a single view of their enterprise VPN while relying on the service provider to trouble-shoot a segment of the network.

RECOMMENDATION

Choose a VPN system that provides complete end-to-end management views that include the remote users. Look for techniques to quickly resolve user-connection problems and to monitor user log time. The ability to distribute the management console and its applications is a valuable one.

Simplicity of Deployment and Operation

A Remote Access VPN is a very simple concept; however, it can become quite complex as it's rolled out to remote users who are unfamiliar with its internal “workings” (ISPs, POPs, policies, costs, etc.). Simplicity of both deployment and operation of the VPN must be accomplished by providing the users with two important VPN “features”: (1) a simple, consistent graphical user interface, and (2) intelligence in the client that can simplify connectivity rules as well as isolate and repair fault conditions within the VPN.

What Should the Remote User See?

All users should have the same view of the remote access VPN through an easy to use GUI, similar to the one shown in Figure 9.

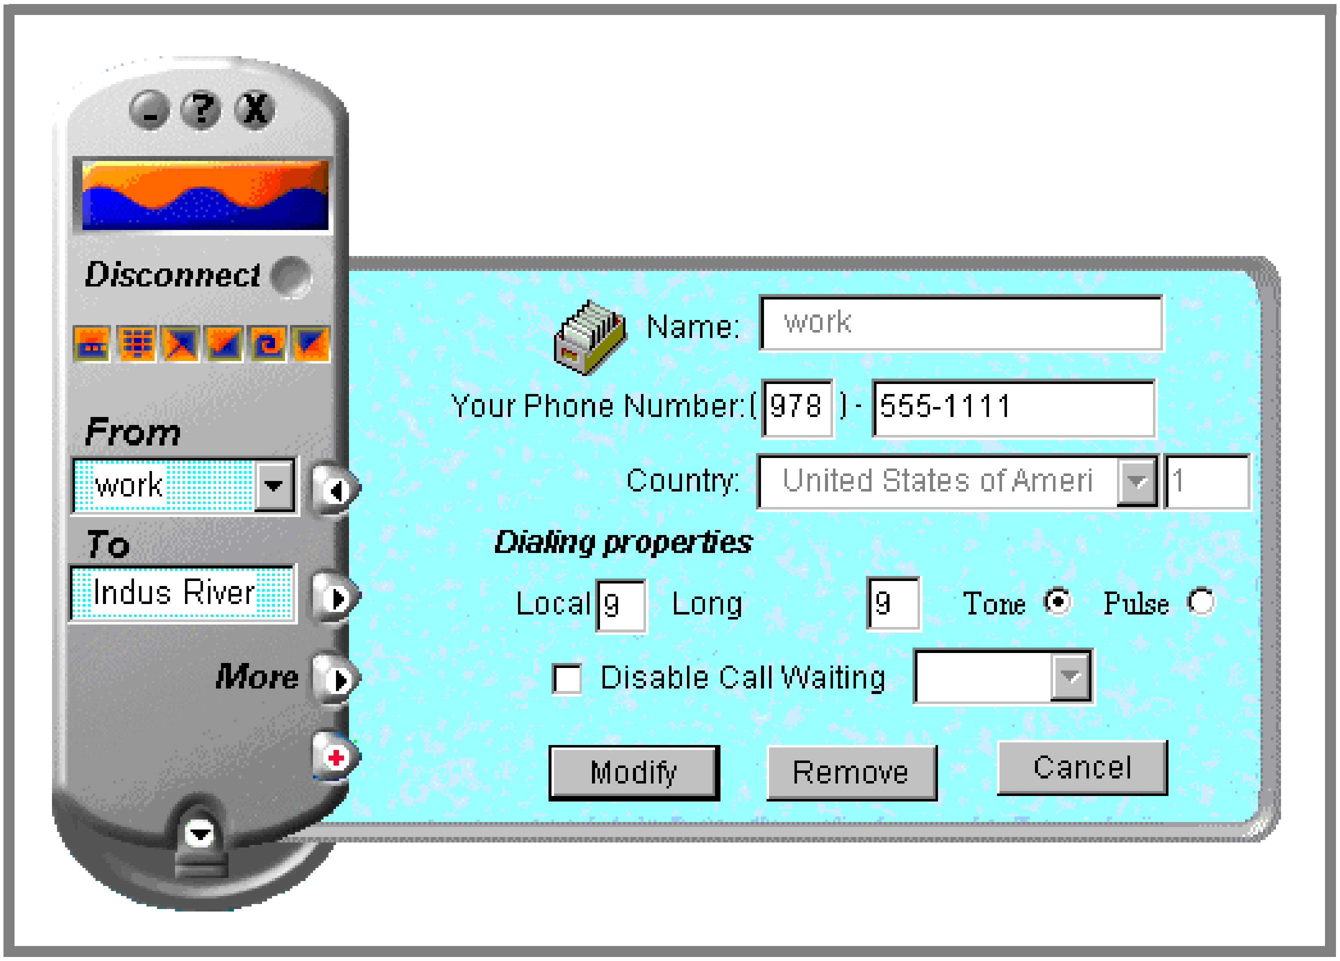

From: work
To: Indus River

Name: work
Your Phone Number: (978) - 555-1111
Country: United States of Ameri 1

Dialing properties

Local 9 Long 9 Tone ● Pulse ○

□ Disable Call Waiting

Figure 9 Remote User’s View of the VPN

The user interface must be consistent across all types of users, whether they are:

- ❒ Connected through LANs, dial-up or cable modems, or other technologies
- ❒ Dialing from the US, or from anywhere in the world
- ❒ Using "conventional" PPTP-based dial-up techniques, or proprietary protocol implementations
- ❒ Using a single ISP or multiple service providers, with multiple service agreements

In addition to consistency, the remote user interface needs to be a controlled environment. Since many usage parameters are established and controlled centrally (they are enforced remotely), hiding these parameters greatly simplifies the remote user's experience.

Behind-The-Scenes Features That Simplify Connectivity

The software tasks that are needed to effectively run the VPN should be transparent to the remote user. Examples include choosing the best local call for Internet access (can easily be done through a local database lookup), choosing the appropriate type of access technology (is a LAN-attach method used to get to the Internet?) and logging in to the local POP (how many passwords must be remembered?). All of these confusing variables can be easily handled by the client application without the user's involvement, simplifying the connection process for the remote user. It's a single button process.

The connectivity process can be simplified even further as more and more intelligence is pushed out to the client. The client application, without the user's knowledge, is capable of managing the ISP access information (not just access numbers, but call rate tables, and POP connect speeds) and monitoring the connection itself. In doing so, it can report back to the network administrator if frequent busy signals occur, modem speeds are lower than advertised, link performance is slow, and other characteristics that are critically important to the health and performance of the VPN. As outlined in the previous section, this information is required to more effectively manage the VPN.

Dynamic Fault Recovery

Fault monitoring within the client application itself is not enough. The application must provide the mechanisms to auto-recover from problems that are encountered throughout the installation, connection, or runtime processes. Auto-recovery is an important element as both the remote user and the network administrator benefit from this capability.

Consider the dial-up process itself. The application should self test and diagnose problems with the modem / network setup (is the modem working? is there a dial tone?). Once a problem is encountered, the application should be able to automatically fix it. Most dial-up clients will inform the user that the dial-up isn't working; however, they won't fix the problem.

What about the VPN (ISP) connection? The client application should monitor the connection process (authentication procedure results, login time, etc.) and the real-time connection (packet loss, overall performance, number of disconnects) to obtain a detailed profile of the remote connection. In addition, the client application should have the ability to recover from a failure (disconnected, for example) by logging-in again or by choosing an alternate access number or ISP.

What about software revisions? What happens if a driver is out of rev and, as a result, the remote user is unable to connect? What if POP numbers change? The VPN should have the capability to automatically download current versions of software, including up-to-date ISP tables, to all clients whenever they log in. The user should not have to request this information.

These features will simplify the experience of using the VPN and drastically reduce the number of help calls.

✓ **RECOMMENDATION**

Don't underestimate the importance of these features! Choose a VPN system with observable user-simplicity features – this will ensure that the cost savings goals of the VPN are realized. Important features to look for are dial assistance capabilities, diagnostics and dynamic fault recovery.

Network Infrastructure

The last step in the planning process is to consider the integration aspects of the overall network infrastructure. In particular, there are two issues to consider: (1) the integration of the VPN system with other "legacy" networking resources, and (2) choice and integration of Internet services into the enterprise network.

How Do I Integrate the VPN with "Legacy" Systems?

The VPN system is not autonomous, so careful consideration must be given to the transition of the VPN system into the existing remote access and network infrastructure. This infrastructure may include a wide variety of products (routers, switches, firewalls, authentication servers, RAS equipment, and application servers); however, most VPN deployment issues arise with existing routers and firewalls.

What About the Routers?

Routers, in nearly all cases, provide Internet access routing and, in many cases, routing for the corporate enterprise network. Internet access routing, usually based on BGP-4, directs all traffic to / from the Internet and, in many circumstances, provides access control (layer 3 firewall). Because dedicated VPN systems do not provide this capability, Internet access routers should remain in their existing location on the network when the VPN is deployed.

Similarly, VPNs do not route inside the Corporate LAN. Most LAN environments require routing features and LAN interfaces that VPN systems do not typically provide. Features such as multi-protocol routing of IP, IPX, AP, DECnet operate over LAN interfaces such as 100M / 1000M Ethernet, FDDI and ATM. For this reason, these "Corporate LAN" routers should not be redeployed either.

The VPN System does need to provide some elements of routing. Consider the diagram in Figure 10, where a virtual private network is used to connect remote users to the corporate network and the connections (tunnels) are terminated across the tunnel server shown in the central site.

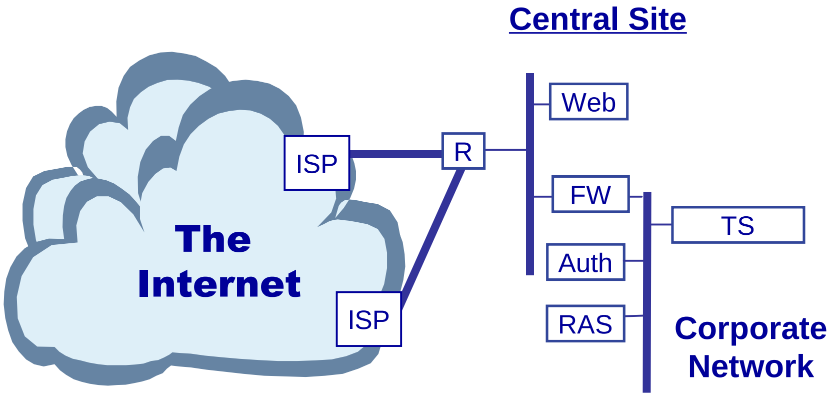

Figure 10 Enterprise VPN with Legacy Router

Remote users will probably have the Internet access router (R) as their default gateway. Since the tunnel server is reachable (known route via auto learning or static configuration), it will direct traffic accordingly. However, resources on the corporate network, such as an application server, do not know that the tunnel server is the proper access point to respond to the users. Therefore, the VPN tunnel servers must:

- ❐ Learn the physical topology of the attached network(s) to be able to forward packets arriving from remote users via the virtual network;
- ❐ Advertise the virtual network topology to the physical network (resources) so that the resources are able to reach remote users via the VPN.

What if the tunnel server didn't provide this routing capability (some do not)? Gateway router configurations in each remote user's system would have to be changed. Or, the IP addressing schemes would have to be changed. Or, a fair amount of addresses would have to be manually configured into the tunnel servers. These are all very unattractive propositions for most network administrators.

What About Firewalls?

VPN systems may provide basic access control capabilities, but nothing that is sufficient to replace the sophistication of a dedicated firewall. Like the router, the firewall usually stays in place, though its location relative to the VPN's tunnel server can provide different results.

The two scenarios shown in Figure 11 illustrate the most popular options for interconnecting a firewall and a VPN tunnel server.

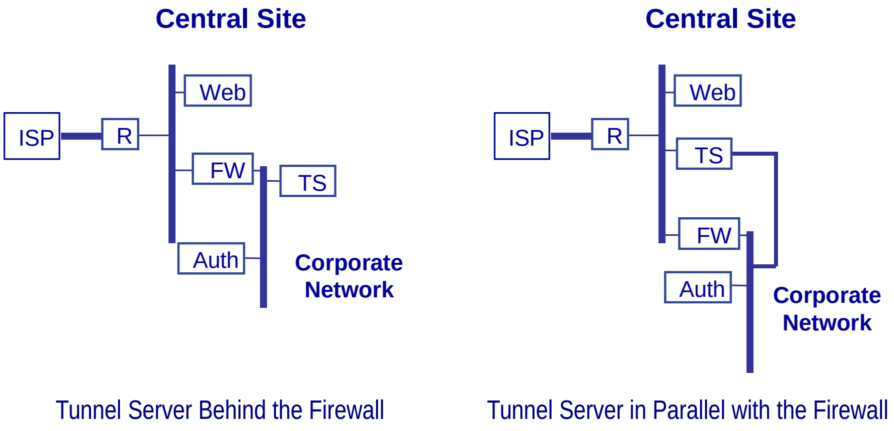

Figure 11 Firewall and Tunnel Server Placement Scenarios

The left diagram shows the VPN tunnel server (single LAN interface only) behind the firewall. This is a very common initial VPN implementation since few changes have to be made to the firewall. Generally, this is the lowest risk approach since the VPN tunnels are secured (terminated, authenticated, encrypted and decrypted) only if traffic is allowed through the firewall. Some reconfiguration of the firewall is required to allow the tunnel protocols through. The principle drawback of this approach is that the firewall must scale its capacity and performance as the amount of VPN traffic flowing through it increases. It is probably unnecessary to filter VPN traffic in a firewall given the security measures already inherent in the VPN solution.

Another variation of this model (not shown) is to connect the tunnel server off of a separate LAN interface in the firewall and not directly to other devices. In this model, the tunnel server is located in the DMZ. This is also a very safe approach although it is subject to the same drawbacks stated above.

The right diagram shows the tunnel server (with dual LAN interfaces) connected in parallel with the firewall. In this configuration, all VPN-based remote access traffic flows through the tunnel server (assuming that authentication and encryption checks are passed) while all non-VPN traffic is subjected to the firewall. Because the VPN tunnel server can act as a "closed stance" firewall, this approach is the most scalable one since it off-loads the firewall from handling all of the traffic to / from the Internet.

✓ **RECOMMENDATION**

Existing network infrastructure components – routers, firewalls, and route topology - require only minor changes with the installation of a VPN. Don't make things more complicated by redesigning. Better yet, choose a VPN system that can peacefully coexist without change. Longer term, look to improve performance and manageability by removing bottlenecks or by adjusting traffic control.

How Do I Choose an Internet Service Provider?

A common question for many network administrators is: now that I understand how to design and implement a Remote Access VPN, how do I choose the right ISP? Unfortunately, there is no simple (or right) formula. Service providers differ widely on the products they provide and the service options they offer (there are literally thousands of choices). They also differ in size and scope, ranging from the large, global service providers (WorldCom, GTE, Sprint, IBM, etc.) to national providers to regional providers. The choices depend on the user environment, the geographic location, the usage level and the billing requirements.

Each corporation is different, and it's up to the network administrator to decide which ISP decision criteria are most important to his business. Common questions to consider include:

- ❒ What kind of global / national coverage do my users need?
- ❒ What is the operating budget for the VPN and for Internet access?
- ❒ How many ISP accounts are needed?
- ❒ What are the service level agreements needed? Are there any?
- ❒ Are there specific features required? Security? RADIUS proxy? Distributed Management?
- ❒ Are there specific billing requirements? Who can provide them?
- ❒ Do I want my VPN provided (and managed) from an outside source (Out-Sourced), or is it an In-House design?
- ❒ What level of outside support is needed for the remote users?

Global Coverage

Global coverage for remote VPN users located anywhere in the world is available from many of the larger service providers. VPN size and scope is not limited as long as a service provider with extensive geographic coverage is chosen. Many ISPs offer “roaming services” with partners such as iPass to provide global connectivity for an additional roaming fee. These fees may need to be closely monitored since they can greatly impact the overall VPN operational costs. Managing access fees is a good application of VPN-wide policies that may be set by the network administrator and managed by the remote user’s VPN client software.

Service Pricing

With a few exceptions, basic Internet access with unlimited access time is available from multiple sources throughout the world for less than $20 / month per client[5]. More restricted access time options are available for less than $7 / month. This option allows network administrators to lower the VPN costs if some groups of users log less time than others. These services seldom have service level agreement commitments or custom management services, but they do provide reliable Internet connections with adequate performance. These “basic” service provider options are excellent choices for many corporations, particularly when they are used with feature rich, high performance VPN systems.

For example, a comparison of basic Internet access offered at a flat monthly rate and a premium service based on usage-sensitive pricing ($3 / hour) is provided in the table below. Also shown is the cost of “800 service” at 8 cents / minute to gauge the relative cost savings of Internet service levels against a dedicated dial access network.

Min/Day	Hrs/Month	Basic Service	Premium Service	800 Service
30 min	10 hr	$20/mo	$30/mo	$48/mo
60 min	20 hr	$20/mo	$60/mo	$96/mo
240 min	80 hr	$20/mo	$240/mo	$384/mo

The table shows that low usage levels of 10 hours / month generate similar cost levels for basic and premium Internet services and a moderate (35%) savings over “800 service”. However, as usage increases, the flat rate structure of the basic service delivers substantial cost savings over networks based on either premium rate Internet service or “800 service” dial networks. The cost difference between the basic-rate and premium-rate Internet services highlights the role of the “service level agreement” to define the higher quality of service that users are entitled to expect from these premium services

[5] Based on June 1998 pricing from several ISPs.

Service Level Agreements

Service providers offer some level of service level agreement to customers that guarantee minimum service delivery, a characteristic that can boost the overall reliability and performance of an enterprise VPN. SLAs often guarantee performance (round trip latency maximum) and network availability.

SLA guarantees from service providers must be considered carefully since the rates are higher than basic rate access and the benefits may already be offered by the basic rate service! When considering SLA premiums, keep in mind that:

- ❒ Many service providers outsource some, perhaps all, of their services to a larger ISP with broader coverage. Since your ISP can't guarantee these services, you may be better off using the same carrier.
- ❒ Some SLA commitments may be for the ISP's own “backbone” only. Users seldom have the opportunity to know or to control this.
- ❒ Sometimes, SLAs may not be better than standard services (performance and availability). In this case, the service provider is doing nothing more than quantifying their own service at a premium.

Over time, service agreements will enhance the VPN system overall because clients experience on-going improvement in the performance and reliability of their tunnel connections. SLA offerings will continue to advance to include higher levels of manageability and better performance. These guarantees will increase the costs of the VPN overall, in some cases, substantially. Network administrators need to consider these options carefully, weighing the cost of the services and the potential benefit.

Out-Sourced vs. In-House VPNs

An out-sourced VPN is a relatively new service provided by some ISPs. The ISP manages most, if not all, of the company's remote access / site-to-site / business-to-business networks. Although out-sourced VPN services are fairly new, many promise premium level access – or service level agreements - that offer backbone latency maximums (performance commitments), overall network availability (uptime per month), and more. Some ISPs offer premium service for dial up access like guaranteed minimum modem speeds and minimum busy signals for remote access users. These VPN services typically cost $2 - $5 / hr, a slight-to-significant premium over basic Internet access charges depending on the usage time. A few of today's offerings provide some level of end-to-end security and management; however, most plans are POP-to-POP based services that do not extend directly to the remote user.

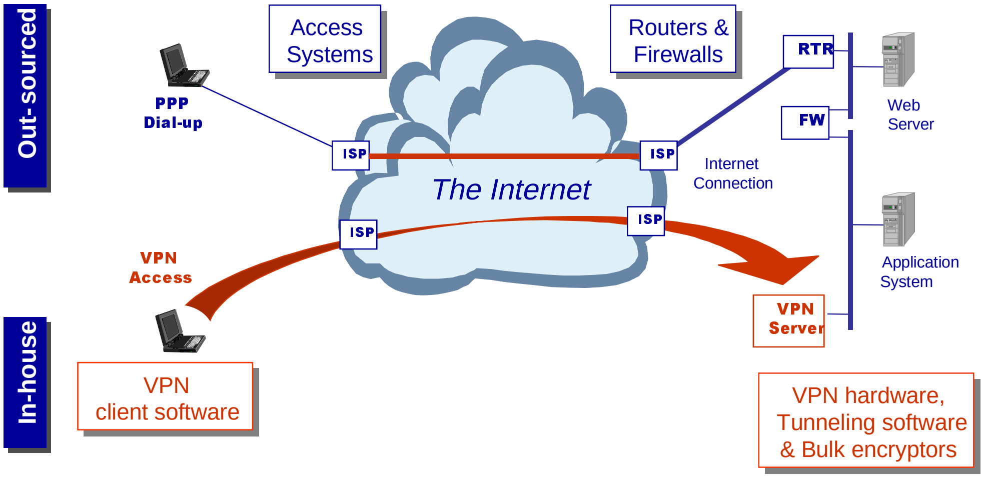

Figure 12 Out-Sourced vs. In-House VPNs

This white paper has focused on the requirements for building an enterprise-class VPN using in-house technology to scale, manage and simplify the network. An In-House VPN is a system that is designed and managed by a company for its own use; the VPN equipment, Internet service, and network management products and services are purchased from external suppliers. In-House systems usually support end-to-end (user-to-user) security mechanisms and manageability, therefore, they are able to support a richer set of end-to-end features. Because In-House VPNs use the same Internet access methods as out-sourced solutions (same ISPs) they can also benefit from premium access with guaranteed service level agreements.

The In-House versus out-sourced decision should be based on the nature of the company’s support staff size and skill level. A company with an ample in-house support staff and expertise is unlikely to choose an out-sourced VPN solution, especially if its business is dependent on an operational VPN. Conversely, a business that lacks the in-house support staff necessary to design and maintain a VPN may have little choice but to out-source.

So How Do I Choose Then?

In many cases it may be advantageous to choose a multi-tiered approach in which a combination of global, national, and local ISPs are selected for a wide variety of services. An effective approach may be to:

- ❒ Purchase low cost, flat-rate services from a local or regional ISP as the primary provider for a large percentage of the remote user population.
- ❒ Purchase broader coverage from a national provider for traveling users, or service level agreements for “power users”.
- ❒ Supplement the above services with high-end premium services from a global ISP for enhanced management, guaranteed services (security, performance, etc.) and worldwide coverage.

Over time, this approach provides network administrators with the ability to easily react to changes in user profiles, business conditions, service features and pricing from ISPs.

Indus River Network's Remote Access VPN Solution

RiverWorks™ is a new generation Remote Access VPN system designed to provide unprecedented levels of performance and ease-of-use to Internet-based remote users. By combining powerful central site tunnel concentration, tunnel acceleration, intelligent VPN desktop access software and a central VPN management authority, RiverWorks meets the needs of the largest corporate remote access networks.

System Components

The RiverWorks Remote Access VPN System, shown in Figure 13, consists of five main components:

- ❒ *RiverWorks Tunnel Server (RTS-5000)*, providing high speed, scalable tunnel processing with advanced security for up to 2,000 simultaneous connections per server.
- ❒ *RiverWorks Management Server (RMS-5000)*, providing complete, end-to-end control and management of the VPN tunnels and all of the remote users (clients).
- ❒ *RiverMaster™ Management Application*, a PC-based software application used in conjunction with the RMS to perform all network management activities on the RiverWorks VPN.
- ❒ *RiverPilot™ Universal Access Manager*, providing intelligent, easy-to-use client access that combines Internet and traditional dial access methods to insure that the user is connected through the lowest cost, most efficient ISP POP.
- ❒ *RiverWay™ Subscription Service*, providing an up-to-date database (for use in the Management server and all clients) of current ISP access numbers, ISP and carrier rate tables, system diagnostics files, and other new utilities as they become available.

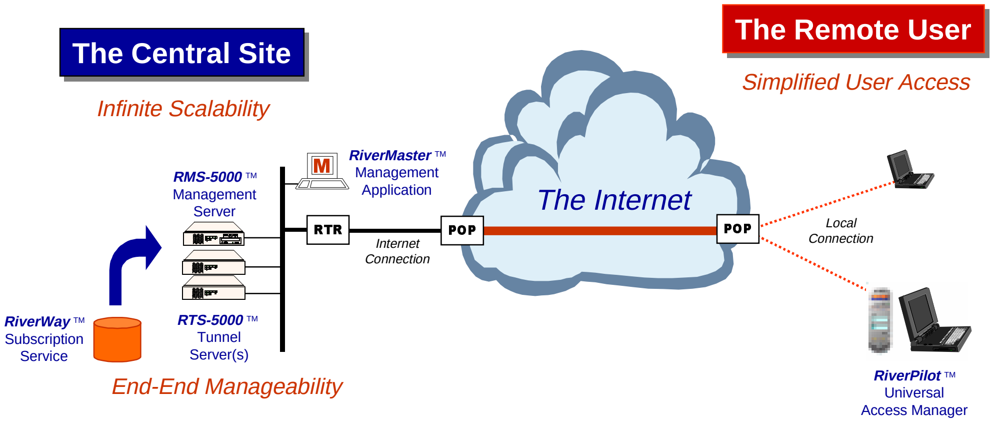
The Central Site
Infinite Scalability
The Remote User
Simplified User Access
RiverMaster ™ Management Application
RMS-5000 ™ Management Server
RTR
Internet Connection
POP
The Internet
POP
Local Connection
RiverWay ™ Subscription Service
RTS-5000 ™ Tunnel Server(s)
End-End Manageability
RiverPilot ™ Universal Access Manager

Figure 13 RiverWorks Remote Access VPN System

Key Features

Three significant attributes differentiate RiverWorks from current product offerings: scalability, manageability and simplicity.

Scalability

RiverWorks is designed for scalability with its dedicated VPN processing capacity and support of stacked and clustered tunnel servers for unlimited user connection requirements and traffic loads. RiverWorks provides leading performance through the use of protocol "extensions" such as channel "learning" to properly size protocol windows and selective re-transmission that retransmits only lost packets. This technique is much more efficient than basic TCP. Tunnel servers may be added to expand central site capacity or decentralized to provide distributed tunnel concentration points. A dedicated management server supports high-volume user authentication, the distribution of new software, database updates, configuration files, and the collection of all management statistics.

Manageability

RiverWorks provides true end-to-end management of the VPN infrastructure with automated diagnostic and recovery tools, remote user policy administration, and Internet performance and availability tracking. Access policies are centrally defined and remotely enforced by the VPN access manager. Connection statistics are logged and reported to the management system for centralized monitoring of VPN resources including remote desktops, Internet availability and responsiveness and multiple VPN tunnel servers.

Simplicity

RiverWorks' VPN desktop software uniquely simplifies the remote user environment by automating network access and problem resolution. The simplicity of the remote user environment is enhanced by a highly intuitive interface that is easier to operate than a cell-phone. More importantly, two innovative software services are operating transparently on the user's behalf. One is called the TollSaver™ connection manager and it performs a connection decision on behalf of the user to find the lowest cost and fastest link available. TollSaver consults an on-board database of ISP access numbers and rate tables to perform this function. The second transparent service is the Prescriptive Diagnostics Engine™ that detects connection faults and automatically corrects them by executing a diagnostic script chosen from an on-board Prescriber database.

What are the Benefits of the RiverWorks Solution?

RiverWorks advances a new model for remote access in several ways. With RiverWorks:

❐ Leading performance is delivered with a highly scalable architecture that allows enterprise networks to grow with usage.

❐ Remote connectivity is no longer limited to the dial-up networking environment; it is now technology-independent and includes all forms of access technologies from direct dial to Internet dial to cable modem to remote LAN-based desktops.

❐ The remote desktop is no longer viewed as a personal computer with a modem; it is now a managed node on the wide-area network to be configured, monitored, diagnosed and optimized like other network resources.

❐ The support burden for remote access can now shift from the user of the network to a virtual help desk embedded in a network that is self-diagnosing and self-correcting.

❐ The Internet can be used for business networking in a high-performance and reliable manner.

APPENDIX A – Model for Remote Access Ownership Costs

This appendix compares traditional remote access costs against the cost of a Remote Access VPN.

Traditional Remote Access Costs

To generate the traditional remote access cost model, the following assumptions were used:

- ❒ Remote user profile is 50% mobile users (travelers), 40% after-hours workers (from home), 10% telecommuters.
- ❒ Minute per day usage is 30 for mobile users and after-hours workers, 240 (4 hours) for telecommuters. All three average 20 days/month usage.
- ❒ Line access charges, based on a survey in May 1998, average $0.08/minute for 800/US long distance service and $0.60 for international calls.
- ❒ Dial-in line density is equal to the sum of: # telecommuters + (10% of sum (mobile + after-hours workers)). Monthly lines are assumed to cost $30.
- ❒ Installation, setup and maintenance costs are assumed to be about 10% of line costs. Breakdown includes $500/port access equipment, 3-4% start-up fees, and 15-20% maintenance fee on installed equipment.

A. Fixed Costs	**Cost/Port**	**Users**	**Users/Port**	**Total RAS Ports**	**Initial Capital Cost**
Equipment Telecommuter	$500	50	2	25	$12,500
Equipment Mobile Users/After-Hours	$500	450	10	45	$22,500
Network Management Software					$3,000
Set-Up Fees (3% HW Cost)					$1,050
TOTALS					**$39,050**
B. Monthly Line Costs	**Usage**	**Rate**	**Total Mins**	**Min/Month**	**Cost/Month**
US Long Distance	10%	$0.08	270,000	27,000	$2,160
International Long DIstance	10%	$0.60	270,000	27,000	$16,200
800# Service	80%	$0.08	270,000	216,000	$17,280
TOTALS					**$35,640**
C. Telecommuter Line Costs	**Total Users**	**Telecommuters**		**Monthly Cost**	**Cost/Month**
	500	50 (10% of total)		$30	**$1,500**
D. Dial-In Site Line Costs	**Total Users**	**Lines**		**Monthly Cost**	**Cost/Month**
	500	115		$30	**$3,450**
E. Total Monthly Line Costs (B+C+D)					**$40,590**
F. Monthly Maintenance (Estimated)					**$417**
TOTAL ANNUAL COSTS				Initial Capital Cost (A)	**$39, 050**
				Monthly Costs (E+F)	**$41,007**
				Resulting Annual Costs	**$492, 080**

Remote Access VPN Costs

To generate the Remote Access VPN cost model, the following assumptions were used:

- ❐ Internet access costs per month for each user is $20 with no usage limits or additional surcharges for domestic use. Dial-in access to ISP Point-of-Presence (POP) is assumed to be a local call.
- ❐ ISP roaming charge of $10/hour for international usage.
- ❐ VPN system for 500 users priced at approximately $100,000 (about $200 per user). This is a one-time charge, not a recurring cost.

A. Fixed Costs - Hardware/Software	**Cost/User**	**Users**	**Server Cost**	**Setup Charge**	**One-Time Setup Cost**
Client software and Tunnel/Management Server	$500	500	$25,000	$3,000	**$103,000**
B. User Internet Access Fee	**Users**	**Access/User**			**Internet Fees/Month**
	500	$20			**$10,000**
C. International Roaming Premium	**Users**	**Intl Mins**	**Total Hours**	**Cost/Hour**	**Roaming/Month**
	500	27,000	450	$30	**$4,500**
D. HQ Internet Access	**Users**	**Bandwidth**			**Access/Month**
	500	T1 (1.536 Mbps)			**$2,000**
E. Monthly Maintenance (Estimated)					**$417**
TOTAL ANNUAL COSTS				Initial Capital Cost (A)	**$103,000**
				Monthly Costs (B+C+D+E)	**$16,917**
				Resulting Annual Costs	**$203,000**

APPENDIX B – VPN Capacity Planning for the Central Site

This appendix provides a brief overview of VPN capacity planning. The most important element of capacity planning is providing the right amount of central site bandwidth for the growing size and connectivity demands of the remote user population. VPN remote access differs from conventional remote access in the following ways:

- ❒ Connections are more economical; more users can be supported by a single high-speed connection eliminating port contention problems.
- ❒ Users will engage in a different connection profile, performing more browsing and casual access over longer connection periods.
- ❒ As users realize they can do much more than read their E-mail on their VPN connection (for example, run client-server applications, browse corporate web sites, access the public Internet, etc.) usage will increase.
- ❒ This new connection paradigm will create large swings in traffic patterns making it difficult to match bandwidth demand with supply.

Network planners will need to make some assumptions about connection rates and network usage in order to plan effectively for central site VPN capacity. Additionally, the network capacity plan should anticipate both normal and heavy usage patterns.

Three Key Variables

The tables provided below are designed to provide basic guidelines for central site capacity by examining three key variables: users, connections and usage.

- ❒ *User Populations* - This variable is an estimate of the total number of VPN users within a given remote access profile. The four profiles modeled include a low speed analog user, a high-speed analog user, an ISDN telecommuter, and a LAN-attached user.
- ❒ *Connection Rate* - This variable replaces the notion of a contention ratio in the traditional remote access network. It reflects the percentage of users that connect to the VPN (authenticate, establish tunnels and log-on) from the total remote user population. Notice that an almost unlimited number of users can log on but with a variable connection quality versus the limited number of users having dedicated circuits in the traditional dial access network approach.
- ❒ *Usage Rate* - This variable is an estimate of the amount of available access bandwidth the user will consume as a result of traffic generated during the network connection.

The Tables

The tables on the next page can be used to calculate the total bandwidth required for a user population of various sizes and profiles. The resulting total bandwidth estimate can be used to size the capacity of the high-speed Internet connection for the central site (or the required portion of a shared Internet connection). The bandwidth estimate can also be used to determine the processing bandwidth needed from one or several VPN tunnel servers used to support the resulting tunnel traffic.

Separate tables are used to measure the bandwidth consumption resulting from the four typical access bandwidth line speeds. Within each of these four remote access profiles, multiple user population sizes are modeled down the rows of the table. Each row has a separate line for normal usage (15% consumption of the available access bandwidth) and heavy usage (50% of the available access bandwidth).

Across the column heading are three different assumptions about connection rates. They are 10%, 20% and 30% of the user population simultaneously connected to the central site network. Pick the connection rate that most closely matches the expected frequency of VPN access for the user population. Notice how the total bandwidth numbers can vary dramatically when connection rates increase from 10% to 30% and as usage fluctuates from normal to heavy usage (a ten-fold variation in total required bandwidth).

The formula used to calculate the total bandwidth shown in the tables is as follows:

Total Bandwidth = User Population x Connection Rate x Access Bandwidth x Usage Rate

Obviously, there are other factors that are equally important when planning a properly designed VPN such as insuring network availability and minimizing end-to-end response time. While these factors are not considered in this capacity planner, they should be included in a total system design effort.

Here are some additional points that may help when using these tables.

- ❒ To obtain a composite bandwidth estimate for a mixed user population, add the results from each of the four tables to represent different remote access user bandwidth requirements.
- ❒ If the reader is unsure about the connection rate or the usage rate that will be experienced on the VPN, try one of the high/low combinations (such as a 30% connection rate combined with normal usage or a 10% connection rate combined with heavy usage). This will provide a conservative estimate.
- ❒ To see the effect of other traffic loading scenarios like 5% usage for very casual connections, divide the “normal usage” numbers by three. Similarly, if theoretical maximum loading is desired, then multiply the “heavy usage” numbers by two to calculate the required bandwidth at 100% usage. Similar adjustments can be made to model other usage scenarios. Remember that all bandwidth numbers are given in Kbps, so convert them to Mbps values as necessary.

Low Speed Analog User			Access Bandwidth: 33.6 kbps		
			Total Bandwidth Required (kbps)		
Users Pop.	Traffic Load	Usage Rate	Connection Rate		
			10%	20%	30%
250	Normal	15%	126	252	378
	Heavy	50%	420	840	1260
500	Normal	15%	252	504	756
	Heavy	50%	840	1680	2520
100	Normal	15%	504	1008	1512
	Heavy	50%	1680	3360	5040
2000	Normal	15%	1008	2016	3024
	Heavy	50%	3360	6720	10080
4000	Normal	15%	2016	4032	6048
	Heavy	50%	6720	13440	20160
8000	Normal	15%	4032	8064	12096
	Heavy	50%	13440	26880	40320

High Speed Analog User			Access Bandwidth: 56 kbps		
			Total Bandwidth Required (kbps)		
Users Pop.	Traffic Load	Usage Rate	Connection Rate		
			10%	20%	30%
250	Normal	15%	210	420	630
	Heavy	50%	700	1400	2100
500	Normal	15%	420	840	1260
	Heavy	50%	1400	2800	4200
100	Normal	15%	840	1680	2520
	Heavy	50%	2800	5600	8400
2000	Normal	15%	1680	3360	5040
	Heavy	50%	5600	11200	16800
4000	Normal	15%	3360	6720	10080
	Heavy	50%	11200	22400	33600
8000	Normal	15%	6720	13440	20160
	Heavy	50%	22400	44800	67200

ISDN Telecommuter			Access Bandwidth: 128 kbps		
			Total Bandwidth Required (kbps)		
Users Pop.	Traffic Load	Usage Rate	Connection Rate		
			10%	20%	30%
25	Normal	15%	48	96	144
	Heavy	50%	160	320	480
50	Normal	15%	96	192	288
	Heavy	50%	320	640	960
100	Normal	15%	192	384	576
	Heavy	50%	640	1280	1920
150	Normal	15%	288	576	864
	Heavy	50%	960	1920	2880
200	Normal	15%	384	768	1152
	Heavy	50%	1280	2560	3840
250	Normal	15%	480	960	1440
	Heavy	50%	1600	3200	4800

LAN-Attached User			Access Bandwidth: 384 kbps		
			Total Bandwidth Required (kbps)		
Users Pop.	Traffic Load	Usage Rate	Connection Rate		
			10%	20%	30%
25	Normal	15%	144	288	432
	Heavy	50%	480	960	1440
50	Normal	15%	288	576	864
	Heavy	50%	960	1920	2880
100	Normal	15%	576	1152	1728
	Heavy	50%	1920	3840	5760
150	Normal	15%	864	1728	2592
	Heavy	50%	2880	5760	8640
200	Normal	15%	1152	2304	3456
	Heavy	50%	3840	7680	11520
250	Normal	15%	1440	2880	4320
	Heavy	50%	4800	9600	14400

APPENDIX C – Glossary

Call Home

A dial-up connection from a **RiverPilot** PC into a **Management Server** to receive an updated **TollSaver database** or bypass a problem that is preventing a tunnel connection.

Firewall

A combination of hardware and software which limits the exposure of a corporate network to outside attack by enforcing a boundary between the network and the Internet. Firewalls normally fall into one of two categories: application-level or network-level (often referred to as a packet filter). An application-level firewall examines traffic at the application level, and only passes packets that are sent by approved applications (such as FTP, E-mail, or Telnet). This type of firewall often readdresses outgoing traffic so that it appears to have originated at the firewall rather than an internal host, thereby concealing the address of the internal host. A network-level firewall examines traffic at the network packet level, and filters packets based on the destination and/or source address.

Generic Routing Encapsulation (GRE)

Tunneling protocol developed by Cisco that can encapsulate a wide variety of protocol packet types inside IP tunnels, creating a virtual point-to-point link over the Internet. For **PPTP**, GRE is used to encapsulate **PPP** data packets within an IP packet (IP packet headers contain address information necessary for routing, while PPP packets do not).

Indus River Tunneling Protocol (IRTP)

A proprietary tunneling protocol developed by Indus River Networks that offers better performance than standard **PPTP**. To take advantage of this protocol, the remote user tunneling into the corporate network must be using Indus River Networks' **RiverPilot software**.

Internet Service Provider (ISP)

A vendor who provides direct access to the Internet. ISPs bill users for the amount of time they are connected, and may also offer additional services such as Web site hosting, E-mail, or news group readers. Remote users reach the ISP by dialing into an ISP **POP** with a computer, modem, and phone line, or over a dedicated circuit (such as a cable modem connection).

Management Channel

A portion of the tunnel connection that is used to download an updated **TollSaver database** from the **Management Server** to the **RiverPilot** PC. When a remote user establishes a tunnel connection to the corporate network, RiverPilot sends a message to the Management Server asking if the TollSaver database has changed. If the RiverPilot's database is out-of-date, the Management Server downloads a new database during low-traffic periods, so that the download does not interfere with regular traffic between the remote user and the network.

Management Server

An Indus River Networks device that manages **Tunnel Servers**. **Network administrators** configure Management Servers from a **RiverMaster Management Station**. The network administrator can create a remote user database on the Management Server or instruct the Management Server to authenticate remote users against an external authentication server (such as a RADIUS or SecurID server). When the network administrator changes tunnel connection parameters, the Management Server provide updated configuration files to Tunnel Servers on request.

Network Administrator

The person responsible for installing and maintaining a company's network equipment, and also insuring that network resources (such as servers and the applications running on them) are consistently available and performing well. In terms of Indus River Networks products, this person physically installs **Management Servers** and **Tunnel Servers**, distributes **RiverPilot software** to **remote users**, and runs **RiverMaster** software on his/her PC to manage the entire **VPN**.

Point of Presence (POP)

In Internet terms, the physical site that contains an **ISP**'s network equipment. Remote users dial into the POP, authenticate against the ISP's customer database, and then gain access to the Internet. ISPs typically have POPs scattered throughout their service area, so that customers can dial a local phone call and avoid paying long-distance charges when accessing the Internet.

Point-to-Point Protocol (PPP)

The Internet standard for sending network traffic over serial lines, such as dial-up phone lines. Unlike its predecessor SLIP (Serial Line Internet Protocol), PPP provides error detection and compression capabilities.

Point-to-Point Tunneling Protocol (PPTP)

A network protocol for linking remote locations over the Internet rather than over costly long-distance or leased lines. To accomplish this, PPTP encapsulates other network protocols (such as TCP/IP, IPX, and NetBEUI) and uses encryption to secure the data sent over the Internet. PPTP was developed jointly by Microsoft and U.S. Robotics.

Policy

A set of rules that governs how **remote users** log onto the corporate network. Corporate policies are defined by the **network administrator** and maintained on the **Management Server**. Policies fall into two general categories: Internet access and user/group administration. For Internet access, the network administrator determines which **ISP**s and telephone carriers the remote user can select, what rates are acceptable for phone calls and Internet connection periods, and which regions of the country the remote user may connect from. These policies are reflected in the customized **TollSaver database** that is distributed as part of the **RiverPilot software**. For user/group administration, the network administrator establishes the log in methods for both ISP access and corporate network access; specifies the use of protocols, encryption, and compression; and determines the user's right to change his or her username or password.

Prescriptive Diagnostics Engine

A feature of Indus River Networks products that diagnoses why a tunnel connection failed and attempts to correct the problem, either on its own or with user assistance. On Indus River Networks **RiverPilot software**, the Prescriptive Diagnostics Engine performs a step-by-step check of each tunnel connection element, including the COM port or serial driver used, modem or terminal adapter, line to a PBX or the telephone network, local or long distance phone service, connection to the **ISP POP**, ISP authentication settings, and so forth. On the Indus River Networks **Management Server**, the Prescriptive Diagnostics Engine uses the call home feature to provide an alternate route that tests end-to-end operation and isolates tunnel problems, and also allows the remote user to download missing or updated files.

Professional User

A person with little or no experience with software and modems, and who is not interested in the underlaying technology of a product or employing all of its features. Instead, this person wants a computer product to operate on the "point and shoot" principle, making it perform basic functions with a minimum of thought and effort.

Remote User

A computer user who wants to access data on a corporate network from a remote location, such as a field office, home office, or temporary lodging. Remote users working from a fixed location are often referred to as telecommuters or day-extenders (if they access the network after regular work hours). Remote users who travel frequently and attempt to access the network from different locations are often called mobile users. *See also* **Professional User** *and* **Technical User**.

RiverMaster Management Station

A computer running the Indus River Networks management application which communicates with **Management Servers** and **Tunnel Servers**. Using the RiverMaster PC, a **network administrator** creates user databases, sets policies for user groups, views activity logs, and generates usage reports.

RiverPilot Software

Indus River Networks client software that runs on a Windows 95 PC and allows a remote user to create a secure tunneling connection to a corporate network. This software features the **TollSaver database** that automatically presents a list of ISP **POP**'s to allow the user to select the lowest-cost connection, and the **Prescriptive Diagnostics Engine** that automatically diagnoses connection problems and either corrects the problem itself or directs the remote user on how to solve the problem.

Routers

Devices which direct network traffic among LANs or WANs until the data reaches its destination. To do this, routers communicate with one another using dedicated protocols such as IGRP (Interior Gateway Routing Protocol) and BGP (Border Gateway Protocol) to transfer information on network addressing, status, and configuration.

Technical User

A person experienced using computer software and modems and who understands general remote access concepts. Not intimidated by technology, this person customizes his or her computer environment and pushes any product's functionality envelope.

TollSaver Database

A feature of Indus River Networks products that provides remote users with a list of ISPs, phone numbers of available POPs, and connection rates. The master TollSaver database is maintained on the **Management Server** and downloaded to the **RiverPilot software** over the **management channel** portion of the tunnel connection.

Tunnel Server

An Indus River Networks device that creates a secure virtual private circuit over the Internet between itself and a remote user's computer. The Tunnel Server encapsulates data packets using **PPTP** and encrypts data to prevent third-parties from intercepting and examining it. Multiple Tunnel Servers can be managed by one **Management Server**. Tunnel Servers receive their configuration settings from Management Servers and pass login information to the Management Server when a remote user attempts to authenticate a tunnel connection.

Tunneling

Technology that lets a network transport protocol carry information for other protocols within its own packets. For example, by encapsulating NetBEUI packets, IP can route them across the Internet, which is not normally possible.

Virtual Private Network (VPN)

An extension of a company's private network that uses the resources of the public Internet. While most private networks use dedicated lines and equipment that are company property, a virtual private network "borrows" resources from the Internet on an as-needed basis.