SANDIA REPORT
SAND2001-2264
Unlimited Release
Printed August 2001

Implementing Virtual Private Networking for Enabling Lower Cost, More Secure Wide Area Communications at Sandia National Laboratories

Marc M. Miller, George A. Yonek

Prepared by
Sandia National Laboratories
Albuquerque, New Mexico 87185 and Livermore, California 94550

Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company, for the United States Department of Energy under Contract DE-AC04-94AL85000.

Approved for public release; further dissemination unlimited.

Issued by Sandia National Laboratories, operated for the United States Department of Energy by Sandia Corporation.

NOTICE: This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government, nor any agency thereof, nor any of their employees, nor any of their contractors, subcontractors, or their employees, make any warranty, express or implied, or assume any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represent that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government, any agency thereof, or any of their contractors or subcontractors. The views and opinions expressed herein do not necessarily state or reflect those of the United States Government, any agency thereof, or any of their contractors.

Printed in the United States of America. This report has been reproduced directly from the best available copy.

Available to DOE and DOE contractors from
U.S. Department of Energy
Office of Scientific and Technical Information
P.O. Box 62
Oak Ridge, TN 37831

Telephone: (865)576-8401
Facsimile: (865)576-5728
E-Mail: reports@adonis.osti.gov
Online ordering: http://www.doe.gov/bridge

Available to the public from
U.S. Department of Commerce
National Technical Information Service
5285 Port Royal Rd
Springfield, VA 22161

Telephone: (800)553-6847
Facsimile: (703)605-6900
E-Mail: orders@ntis.fedworld.gov
Online order: http://www.ntis.gov/ordering.htm

SAND2001-2264
Unlimited Release
Printed August 2001

Implementing Virtual Private Networking for Enabling Lower Cost, More Secure Wide Area Communications at Sandia National Laboratories

Marc M. Miller
Advanced Networking Integration Department

George A. Yonek
Telecommunications Operations Department

Sandia National Laboratories
P.O. Box 5800
Albuquerque, NM 87185-0806

Abstract

Virtual Private Networking is a new communications technology that promises lower cost, more secure wide area communications by leveraging public networks such as the Internet. Sandia National Laboratories has embraced the technology for interconnecting remote sites to Sandia's corporate network, and for enabling remote access users for both dial-up and broadband access.

Acknowledgements

The Computer Security Technology department laid significant groundwork for Sandia's VPN development. In particular, John Long's expertise regarding VPN authentication methods and remote access security concerns proved invaluable in developing a VPN strategy for Sandia.

Grateful assistance was received from the Advanced Networking Department in debugging a Kauai link problem. In particular, Steve Gossage, Richard Hu, Luis Martinez, John Naegle, and Tom Pratt helped resolve a problem in the Kauai network that was prohibiting the VPN from operating successfully.

Telecommunications Operations personnel Roger Adams, Pat Manke, Bruce Whittet, and George Yonek provided expertise, help, and support in getting VPNs off the ground.

Trademarks

Cisco and Cisco Systems are registered trademarks of Cisco Systems, Inc.

DCE is a trademark of the Open Group.

Entrust is a trademark of Entrust Technologies Inc.

KyberPASS is a trademark of KyberPASS Corporation.

Linux is a registered trademark of Linus Torvalds.

Macintosh is a registered trademark of Apple Computer, Inc.

Microsoft, Windows 95, Windows 98, Windows NT, and Windows 2000 are registered trademarks of Microsoft Corporation.

Network Alchemy, CryptoCluster, and CryptoConsole are trademarks of Network Alchemy, Inc.

Ravlin is a registered trademark of RedCreek Communications, Inc

SecurID is a registered trademark of RSA Security Inc.

Smartgate and V-ONE are registered trademarks of V-One Corporation.

SnareWorks is a trademark of IntelliSoft Corporation.

TimeStep is a registered trademark of Timestep Corporation.

UNIX is a registered trademark of Unix System Laboratories, Inc.

VPNet and VPNWare are registered trademarks of VPNet, Inc.

All other trademarks and registered trademarks are the property of their respective owners.

Content

Figures

Tables

Nomenclature

AH	Authentication Header
ATM	Asynchronous Transfer Mode
CA	Certificate Authority
CAST	Carlisle Adams and Stafford Tavares
CCHD	Corporate Computing Help Desk
CMP	Change Management Process
CSU	Computer Support Unit
DCE	Distributed Computing Environment
DES	Data Encryption Standard
DHCP	Dynamic Host Configuration Protocol
DSL	Digital Subscriber Line
ESP	Encapsulating Security Payload
FIPS	Federal Information Processing Standards
FTP	File Transfer Protocol
GUI	Graphical User Interface
HTTP	Hypertext Transfer Protocol
HTTPS	Hypertext Transfer Protocol Secure
IDEA	International Data Encryption Algorithm
IETF	Internet Engineering Task Force
IKE	Internet Key Exchange
IP	Internet Protocol
IPSec	Internet Protocol, Security
ISAKMP	Internet Security Association and Key Management Protocol
ISDN	Integrated Services Digital Network
ISP	Internet Service Provider
KTF	Kauai Test Facility
L2F	Layer 2 Forwarding
L2TP	Layer 2 Tunneling Protocol
LAN	Local Area Network
Mbps, Kbps	Mega-bits per second, Kilo-bits per second
MD5	Message Digest
MDC	Main Distribution Closet
NAT	Network Address Translation
NIST	National Institute of Standards and Technology
OS	Operating System
PPTP	Point to Point Tunneling Protocol
PDA	Personal Digital Assistant
RC2	Rivest's Cipher
RFC	Request for Comments
SHA-1	Secure Hash Algorithm
SNMP	Simple Network Management Protocol
SSH	Secure Shell
SSL	Secure Sockets Layer
UDP	User Datagram Protocol
URL	Uniform Resource Locator
VPN	Virtual Private Network(ing)

Introduction

What is a Virtual Private Network or VPN?

In recent years, a new technology has been developed for wide area networking that connects networks at lower cost by taking advantage of the ubiquitous Internet. Virtual Private Networking typically utilizes a public network such as the Internet to privately transport data by way of encryption. This produces a virtual network since circuits are not hard-wired as in leased lines, and produces a private network since the data is not viewable by others.

This is just one way of defining a VPN. There are many variations. For example, the definition "a Virtual Private Network is a network of virtual circuits for carrying private traffic." (Kosiur, 1998), although succinct, is very ambiguous and can be applied to most any imaginable communications architecture. However, to say "that it is a scheme for using the Internet as a backbone for computer networks." (Wilson and Doak, 2000) is a more direct definition that well describes Sandia's effort to connect remote sites and remote users to the corporate site.

Why a VPN?

The bottom line for most VPN implementations is cost savings. Virtual Private Networks (VPNs) now can provide cost saving of 50 to 75 percent by replacing more costly leased lines and remote access servers and reducing equipment and training costs (Kosiur 1998).

For example, assume a corporation is located in New York, with a major branch office in Hong Kong. The company wishes to connect the networks of the two locations together. Leasing a communications circuit halfway around the world is an expensive proposition. However, a VPN between the sites utilizes the Internet, incurring no distance-dependent transport costs. The only costs are for local Internet access at the two endpoints, a cost most corporations already incur today for general Internet access.

This example can be extended to traveling employees needing access to either the New York or Hong Kong corporate networks. Historically, this would be accomplished using remote access dial-up over the telephone system, typically utilizing 1-800 or long-distance service whose cost is distance and time dependent. A VPN for traveling employees again takes advantage of the Internet by providing local dial-up access to the Internet in most major metropolitan areas of the world, and establishing a VPN to the corporate network. Thus, an employee in London can dial a local Internet access number and connect across the Internet to the corporate network in New York or Hong Kong, paying only the local access charges.

Standardization

Several protocols have been developed over the years for providing VPN functionality. The most common early developments were L2F, PPTP, and the combination of the two, L2TP. These are network layer 2 protocols, designed primarily for remote access solutions and non-IP environments.

A relatively new standard called IPSec was developed to provide network layer 3 protection through data privacy, integrity, authentication, and key management. As the name implies, this protocol is for IP networks only. IPSec provides protection for "the entire network" as opposed to schemes that protect only data within an application, such as the use of SSL in a web browser. Of course, protection of the network only applies where IPSec is operating, such as in a virtual network operating between two sites.

IPSec was developed by an IETF working group. Although many Internet RFCs define IPSec and its components, the three core RFCs are:

1. RFC 2401, Security Architecture for the Internet Protocol, November 1998, defines the implementation of IPSec.
2. RFC 1827, IP Encapsulation Security Payload (ESP), August 1995, defines the encapsulation protocol used to create virtual tunnels through the Internet
3. RFC 2401, IP Authentication Header (AH), November 1998, defines the construction of the authentication header.
   (Wilson and Doak, 2000)

IPSec Data Protection and Authentication

An IPSec VPN protects data from eavesdropping by using encryption, and authenticates data sources by using hash algorithms. IPSec encryption is optional, but is used in almost all cases. There are two types of IPSec packets: ESP and AH. ESP, or Encapsulating Security Protocol, defines a packet format for encapsulating an encrypted IP packet into a new IP packet that incorporates authentication information. AH, or Authentication Header, on the other hand, only incorporates authentication information into a new packet. Therefore, an IPSec VPN can be constructed that authenticates packets but does not encrypt. This is somewhat useful, but the true value of an IPSec VPN comes from applying both encryption and authentication by using ESP.

The IPSec protocol does not specify the type of encryption or authentication to be utilized. This is left to the vendor to implement. However, several standard encryption and authentication algorithms are available. The most commonly used encryption is Triple DES, adopted in 1977 by the National Bureau of Standards. Due to concerns over the capability to crack DES, a more robust form called Triple DES was developed using two keys and three stages of encryption. Other encryption options include RSA, Blowfish, RC2, RC4, and IDEA.

A range of choices also exists for packet authentication, albeit a smaller one. Typical authentication algorithms are MD5 and SHA-1.

Transport or Tunnel Mode

IPSec can operate in one of two modes called transport and tunnel. In transport mode, the IP packet payload is encrypted and/or authenticated, with the original IP packet header left intact, including the original source and destination addresses. This is required for peer-to-peer IPSec connections such as between two hosts. For tunnel mode, required for operation between sites, the entire original IP packet is encrypted and/or authenticated and placed inside a new IP packet whose source and destination addresses are those of the VPN gateways. The original source and destination IP addresses are hidden in the payload of the new packet.

Key Management

To establish a VPN, keys must be exchanged for authentication and encryption. IPSec requires either a manual key process or an automated process called IKE, or Internet Key Exchange. Manual keying is suitable for small sites or for testing purposes, but does not scale well. Automatic keying eases the burden of manual key exchanges. IKE is a combination of the Internet Security Association and Key Management Protocol (ISAKMP, typically pronounced ice-ah-camp), which serves as a framework for authentication and key exchange, with the Oakley protocol, which describes various modes of key exchange.

IKE is designed to provide four capabilities:

1. Provide the means for parties to agree on which protocols, algorithms, and keys to use.
2. Ensure from the beginning of the exchange that you're talking to the right person.
3. Manage those keys after they've been agreed upon.
4. Ensure that key exchanges are handled safely.
   (Kosiur, 1998)

IKE operates in two phases. The first phase establishes a secure channel between two IPSec nodes for performing ISAKMP functions. In the second phase, the two nodes negotiate parameters for the VPN.

For further details of the IPSec protocol, please refer to the book "IPSec: the new security standard for the Internet, intranets, and virtual private networks." (Doraswamy and Harkins, 1999)

For a different spin on IPSec, the paper "A Cryptographic Evaluation of IPSec" points out several shortcomings of the protocol and labels it a "disappointment" due to "its

complexity", but concludes "it is the best IP security protocol available at the moment." (Ferguson and Schneier, 1999)

Performance

One might suppose that a VPN built over the Internet is subject to uncertainty regarding performance. This is certainly true. There are no guarantees of network performance over the Internet. A company that builds a VPN over the Internet must understand the consequences of lost connectivity and ask whether outages can be tolerated. In cases where outages cannot be tolerated, an Internet solution is probably not appropriate. However, this brings the solution back to leased lines or bandwidth allocation from a service provider with service level agreements that guarantee certain network performance levels. In this case, the cost advantage of a VPN running over the Internet is reduced, but performance guarantees may be of greater importance than cost in some cases.

Promise of interoperability using IPSec

In the early days of VPNs, solutions were proprietary. Although VPN implementations within a company do not necessarily suffer from proprietary solutions, establishing VPNs between companies has been difficult due to the likelihood that a partner site has chosen a different VPN vendor. To counter this difficulty, IPSec has been developed as an Internet Standard Protocol, promising interoperability among differing vendor solutions. Even with IPSec defined, interoperability has been hit or miss as VPN technology has matured. To help spur interoperability, the ICSA (International Computer Security Association) has been certifying VPN equipment for IPSec compliance. This has gone a long way in helping get equipment operating across different vendor product lines. However, even with interoperability compliance, there are still issues involving the mixing and matching of VPN options. One vendor's set of options may not mesh with another's set, leaving only a subset of compatible options, and therefore limiting choices when attempting to establish a VPN between differing equipment types.

Building a VPN

Site-to-site

A site-to-site VPN connects one private network to another private network over some public network such as the Internet.

Figure 1 shows a schematic of a typical site-to-site VPN. Site A and Site B are connected by a VPN over the Internet. All traffic passing between sites is protected by encryption and authenticated to ensure each packet is protected from eavesdropping and is originating from the appropriate VPN server. The two sites can be across town, across the

country, or across the world. Each site must have a local connection to the Internet, which most companies already have for general Internet access such as email and World Wide Web access. Anyone intercepting the traffic between sites A and B over the Internet will have access only to encrypted packets protected by strong a encryption algorithm. Any attempts to inject packets between sites will be thwarted by packet authentication.

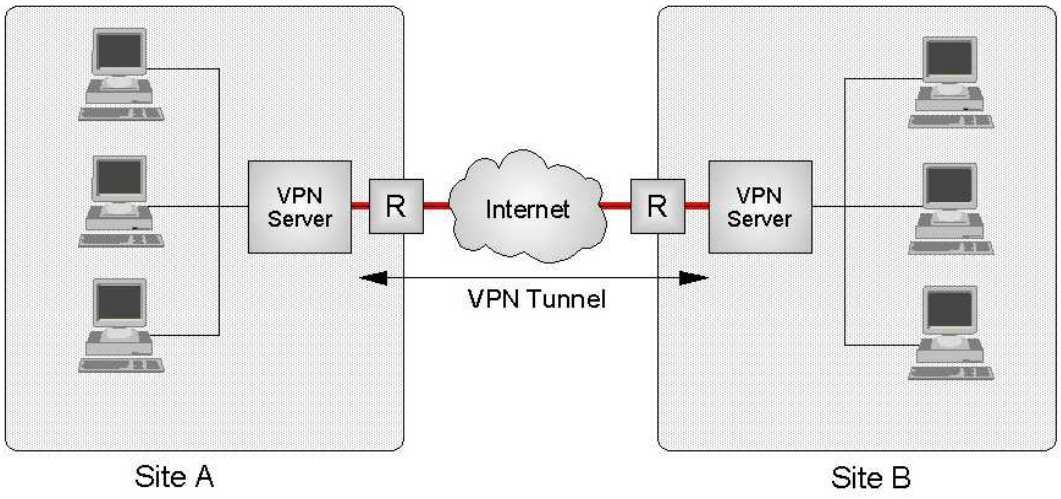

Figure 1. Site-to-Site VPN Schematic

Client-to-site

A client-to-site VPN differs from a site-to-site VPN only in that one end of the tunnel is a user machine and not a network as shown in Figure 2. Typically, a traveling employee will utilize a laptop running VPN software to dial a local Internet access number provided by an ISP, and connect through the Internet back to the corporate network.

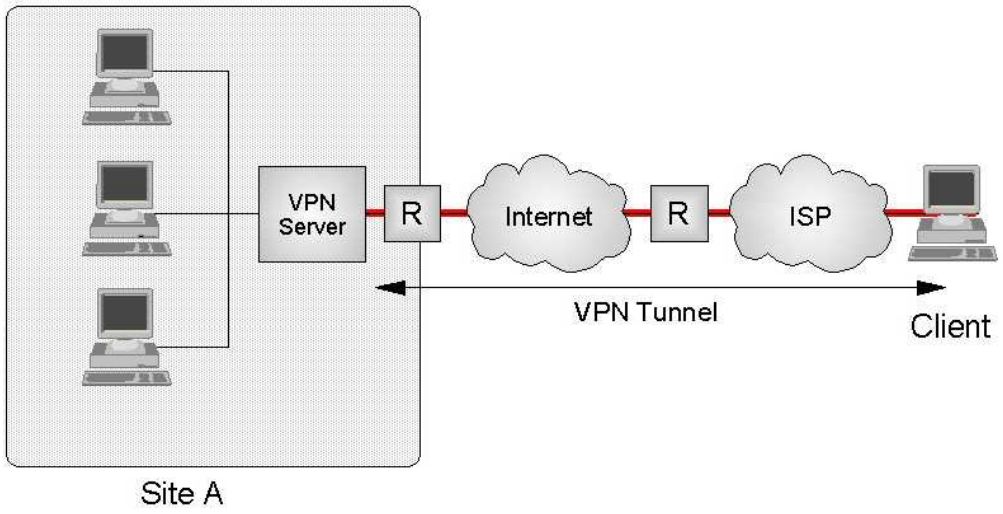

Figure 2. VPN Client-to-Site

Extranet

An Extranet VPN is similar to a site-to-site VPN, but access is limited at one or both ends. It is most often implemented using HTTPS (HTTP with SSL) for web applications. However, for situations where multiple IP protocols are required, a VPN can supply a transparent transport between two dissimilar sites, such as between a manufacturer and a supplier.

Remote Site Protection Policy

For any case in which a corporate protected network is extended to another site, VPN or otherwise, there must be policy in place to govern physical protection of the remote network in order to provide proper protection of the local corporate network. Using a VPN to protect data traversing a public network is of little use if the remote network is compromised internally due to inadequate physical security, thus enabling an intruder to ride the VPN into the corporate network. Put another way, Sandia will not establish a site-to-site VPN to a remote site that is open to unauthorized users.

Sandia's policy for extending the internal, protected network is that the remote site must have the same or similar physical protection as the corporate site. This includes limited physical access, controlled minimally by fences and locked entryways, to keep unauthorized users from accessing the network.

In the case where the Sandia restricted network is extended to an employee's home, such as in creating a site-to-site link over a DSL or cable modem circuit, there must be

equivalent physical access control such as a locked house when unattended. However, locking the unattended house does not address unauthorized access by non-employee family members or visitors. This problem can be overcome by utilizing user authentication as part of the network extension. However, user authentication is not normally a feature of site-to-site VPNs. One possible solution that is in development by Cisco Systems is a hardware VPN client, which operates just as a software VPN client, but in a box that sits between the computer and the network. Such a device can prompt for user authentication before enabling a VPN, appearing as a hybrid between a client-to-site VPN and a site-to-site; the best of both worlds. An added feature of the hardware client is its ability to work with any operating system. Thus, Linux, Mac OS, and other non-Windows platforms can be accommodated. The drawback is that the device is not particularly portable and therefore is of little or no use to the road warrior.

Architecture

There are three methods of connecting a VPN server to a corporate network with respect to the corporate firewall as shown in Figure 3: behind the firewall, in front of the firewall, or parallel to the firewall.

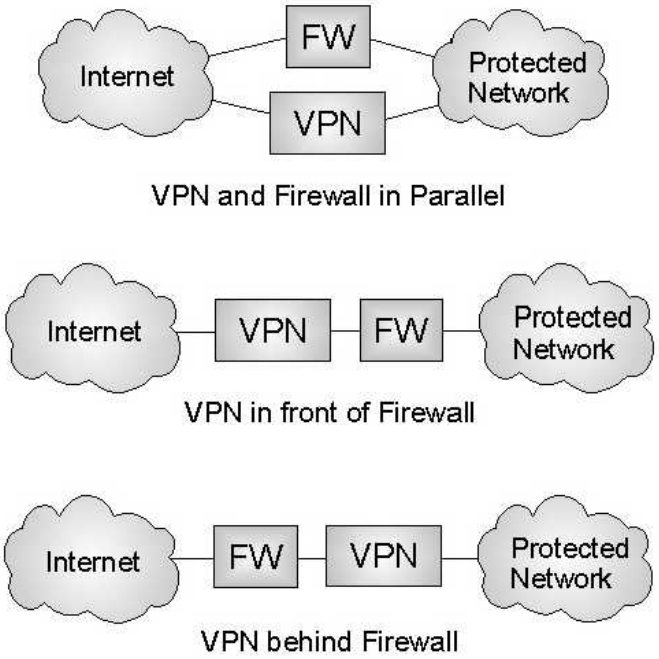

Figure 3. Location of VPN with respect to the Firewall

When locating the VPN server behind or in front of the firewall, all traffic passing into and out of the protected network must also pass through the VPN server. This is problematic for networks that have a high level of traffic passing through the firewall: the

VPN server must be sized to also accommodate this traffic, whether it is VPN traffic or not. In other words, the VPN server cannot be a bottleneck.

However, in-line location simplifies routing. No changes are required to routing to enable VPN activity since all traffic must pass through both the VPN server and the firewall.

Behind the Firewall

A VPN server located behind the firewall is only effective in cases where the remote site addresses are known in advance. This is because the firewall must be configured to allow VPN packets into the protected network in order to reach the VPN server. Only known IP addresses can be utilized for this purpose. In the case of a site-to-site VPN, this is possible. All remote sites will be known and fixed and a firewall can be programmed to allow these VPN packets to pass. However, for a client-to-site VPN, the client address is not normally known in advance since the client address can change with each connection to an ISP. Therefore, a client-to-site VPN cannot utilize a VPN server located behind a firewall, unless the firewall is configured to pass all IPSec packets no matter the origin (not a particularly good policy!).

In Front of the Firewall

Locating the VPN server in front of the firewall is more effective than locating it behind the firewall since the VPN server is free to receive VPN packets from any source. Once the packets are decrypted and authenticated, they can be passed through the firewall, and can be subjected to firewall rules. A disadvantage here is that the decrypted packets are in the unprotected network until they enter the firewall. However, this liability can be mitigated by restricting access to this portion of the open network, either by physical means or by network design, or both.

Parallel to the Firewall

Locating the VPN server parallel to the firewall appears to be the easiest method of connection. However, routing becomes an issue. In this configuration, the VPN server only serves packets that are part of a VPN. All other traffic flows through the firewall. Thus, the VPN server can be sized for the VPN traffic, which in many cases is far less than the traffic flowing through the firewall. However, getting traffic to the VPN server requires specific routing entries. The default route will bypass the VPN server. Therefore, all destination addresses that are part of a VPN must be included in router configurations that direct VPN packets to the VPN server.

Bridging Vs. Routing

Another consideration that can have considerable impact on a VPN implementation is the matter of bridging and routing. Some VPN servers are built as bridge-type devices, meaning they must be placed in the network stream and are invisible to other network devices at the IP layer because they operate at network layer 2. Other VPN devices are built as layer 3 routing-type devices where the server has IP addresses assigned to the inside and outside interfaces, and must span two different IP networks.

The two types have their advantages and disadvantages. Bridge-type VPN servers require fewer IP addresses (one must be assigned for device management) and do not require two networks for operation (one inside and one outside). They can easily be connected into a network by simply plugging into the appropriate network path. On the other hand, access to the network stream may be difficult since a bridging VPN device must be inserted between two environments. This stream may be virtual, as in the case of a router function residing in a network switch, with no physical router interface available to connect the VPN to another router interface.

VPN servers that operate as routing-type devices require at least two IP addresses: one for the inside interface and one for the outside interface. They must be connected across two networks. This can be a disadvantage if two networks are not readily available since a new network would have to be added. However, no matter what architecture exists in a network, a routing VPN device can be inserted where needed for spanning environments.

Some products have the capability of being both bridge and routing types, providing greater flexibility. The trade-off is the additional complexity in device configuration.

VPN Management

Like any networking equipment, VPNs must be managed. Exactly how this is accomplished varies from manufacturer to manufacturer, but most VPN devices rely on a common set of management tools including character-based applications such as serial port and Telnet, and graphical user interfaces utilizing Java and HTTP.

Serial port management involves the old-fashioned serial cable connected to a terminal, but since few terminals exist these days, a cable is normally connected to the serial port of a laptop running a terminal emulation program. This mode of connection is normally only required for initial device configuration such as network address configuration that enables such management tools as Telnet and HTTP. It is also available for emergency connection when network access has been lost.

Most VPN devices can be accessed by Telnet once the network is configured. Since Telnet is a character-based application, a VPN device accessed by Telnet presents a

character-based user interface. Some devices use a command line and some use a command menu.

Virtually all VPN devices offer a graphical user interface or GUI, for device management. Two different methods of running a GUI have been encountered. In one method, an application runs on a user machine and presents parameters that can be configured for a VPN. When the configuration is ready, the data is pushed to the VPN device. Another method utilizes HTTP (or encrypted HTTP using SSL), with the VPN device acting as a web server. Configuration changes are made in a web browser and stored on the VPN device.

Figure 4 is an example of a browser application for managing a Cisco VPN 3000 device. The 3000 is acting as a web server. The user simply starts a web browser, points the browser to the URL of the 3000, and logs in. In this example, the web session is encrypted using SSL. Configuration changes are stored directly on the VPN 3000. The user machine has no local storage of device configuration. This is a versatile method of VPN management since any user machine that is authorized to connect to the VPN server can do so from any authorized location.

Figure 5 shows a typical screen for a Network Alchemy CryptoConsole application for managing CryptoCluster devices. This is a stand-alone application written in Java. Any changes made to a device configuration are stored in a local database on the user's machine and then pushed to the appropriate VPN device. This is a very constrained VPN management technique, since only one user machine is able to communicate management information to the VPN server.

Protection of management information passing between a VPN device and a control station is important to prevent unauthorized VPN access or modification. Telnet access to a VPN device for management is a weak method due to exposure of username and password, as well as all configuration data. Therefore, if enabled at all, Telnet should be limited to basic troubleshooting only. On the other hand, graphic user interface applications normally protect the data using some form of encryption. Cisco's VPN 3000 can use SSL to encrypt browser sessions. Network Alchemy uses a proprietary protocol between the management application and their VPN devices, but does utilize 3DES encryption. SSH has not been widely utilized, and is only recently been made available for the Cisco VPN 3000.

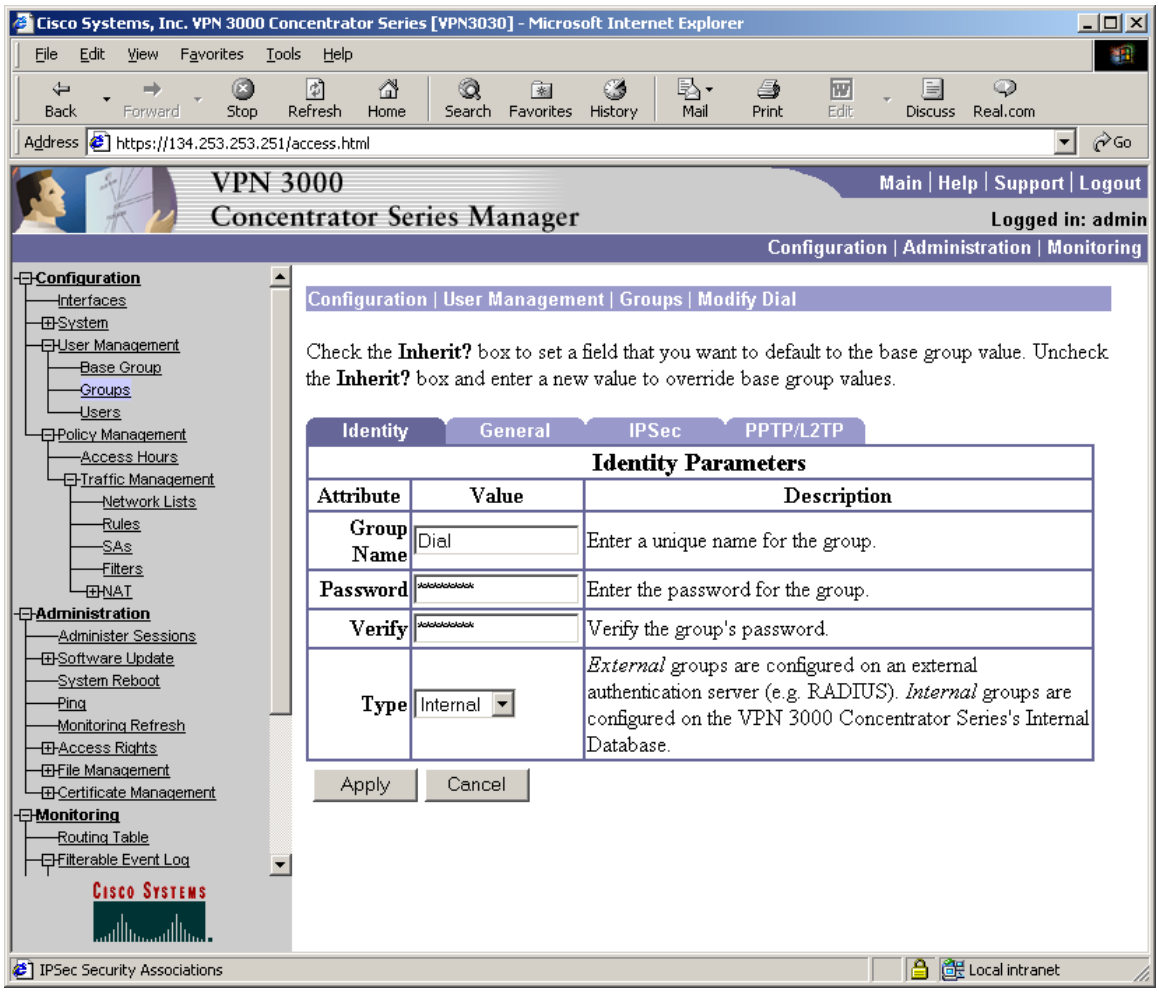

Figure 4. Cisco Systems VPN Management GUI for VPN 3000 Series

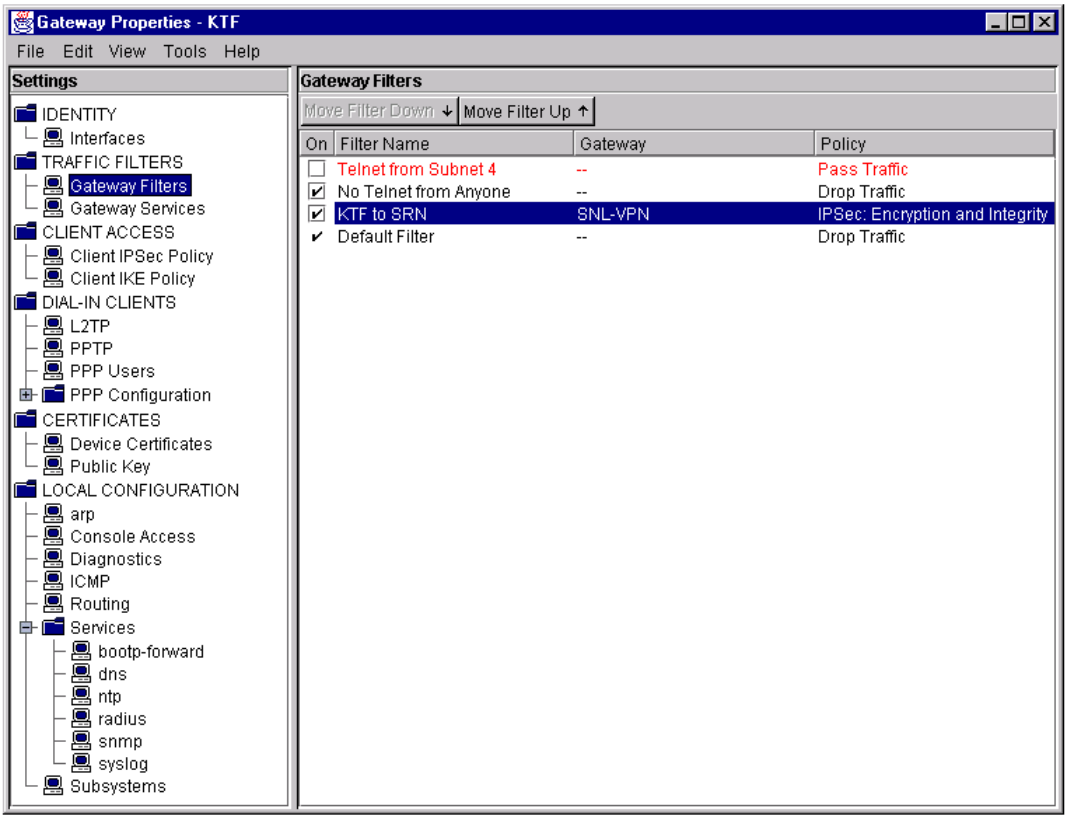

Figure 5. Network Alchemy Management GUI

Project History

LIP project (combining mobility with smart card authentication)

An FY99 project, Location Independent Professional (LIP), was an attempt at identifying and addressing issues involving computing from any location. Rather than being tied to a port on a wall at a particular site, building, and office, a computer should be able to access information from literally anywhere in the world. The LIP project examined technology processes that enable this type of computing freedom. Two areas were examined: movement within a campus environment, and movement outside the campus environment. Within a campus environment, computer mobility can be achieved by either wired or wireless networking. Outside the campus, mobility is currently constrained to wired communications, using the telephone system. Wireless is possible, but not currently practical for the computing “road warrior.” Although advancements are coming quickly in the area of wireless use of PDAs and mobile phones, computers are not yet a big part of these technology advancements.

VPN Technology Explored for Providing Improved Remote Access

Mobility outside the campus utilizes remote access technology. This has been a common mainstay of corporate computing for some number of years. However, it has been poor performing, unreliable, subject to eavesdropping, and costly when utilizing 1-800 or long-distance service.

An improved model involves VPN technology, which can address all the above shortcomings.

The private nature of a VPN results from encryption, which addresses the problem of eavesdropping on the Public Switched Telephone Network. Although it is arguably difficult to eavesdrop on telephone calls (and certainly not impossible), VPN encryption makes this a moot point. The unreliability and typical poor performance of dial-up connections can be largely overcome by dialing a local ISP and then starting the VPN software, rather than having to dial long-distance.

Two VPN products were evaluated during the LIP project: SnareWorks and KyberPASS. Both were early adopters of the emerging VPN technology, with SnareWorks being a very aggressive, feature-rich implementation for large enterprises. Both products were software-based, meaning the product runs on a customer-supplied server computer.

SnareWorks was the first to be evaluated, proving to be a valuable experience towards future VPN efforts. This system is built around DCE (Distributed Computing Environment), a functionality that adds a complex but feature-rich security layer to the VPN product. For a more detailed review of SnareWorks, see the report *Location Independent Professional Project: A Pilot Study,* SAND99-0100 (Miller, Hudson, Long, 1999).

Of all the components of SnareWorks, including security and single sign-on, only the remote access VPN feature was evaluated. The SnareWorks system was built around a Sun Sparc workstation running the Solaris operating system and the SnareWorks application, with DCE an underlying software module for providing security partitions. A laptop was configured with the SnareWorks client module and an external smart card reader used for user authentication. Rather than using a one-time password technology such as a SecurID card, SnareWorks operated with a digital certificate installed on a removable smart card, thus maintaining two-factor authentication (something you have and something you know). The smart card technique promised an improved authentication method over the sometimes difficult-to-use SecurID card, but the smart card reader proved to be disappointing in terms of ergonomics.

The reader was a cigarette pack-sized device, connected to the laptop by a cable that plugged into the serial port. This is a cumbersome device for taking on the road: one more piece of equipment to remember to pack, to lug around, and potentially break.

However, it was easy to use. A smart card, nothing more than a credit card sized piece of plastic with a small memory chip, is simply inserted into the reader before activating the VPN. When the VPN application starts, it prompts for a username and password to unlock the digital certificate stored on the card. It reads the certificate and uses the data to establish a VPN to the end point. Thus, as a two-factor authentication tool, a smart card is easy to use, which was the intended consequence. There are no small buttons to push, a difficult task for large fingers. There are no small numbers to read, a difficult task for poor eyesight or low light levels. However, the drawback is the cumbersome card reader that would undoubtedly be poorly received by most travelers. A better solution, not available at the time of the evaluation, but becoming commonplace, is a card reader that plugs into a laptop PC card slot.

KyberPASS was evaluated after SnareWorks. KyberPASS was a much simpler VPN product, focusing on small to medium enterprises. Like SnareWorks, KyberPASS was a software-based VPN, running on a customer-supplied server. Unlike the complex modules of SnareWorks, KyberPASS was capable of utilizing a simpler authentication system based on PKI (Public Key Infrastructure). KyberPASS was designed primarily as a client-to-site system, but could also connect site-to-site. The system utilized DES, 3DES, or CAST encryption, and SSL or IPSec protocols. IPSec was just becoming available at this time, and therefore was available in KyberPASS as an option. The primary protocol was SSL between a client machine and the VPN server. User authentication was limited to a digital certificate, which was unlocked by a username and password. No two-factor authentication was available at the time for this product.

A KyberPASS class was held to learn as much as possible about the product in the shortest time. The hands-on learning was beneficial.

Two significant results emerged from experimenting with SnareWorks and KyberPASS that laid groundwork for future VPN development: 3DES encryption and two-factor client authentication. While there are several encryption schemes available, 3DES is more secure than the standard and widely used DES, and can operate effectively in a VPN. Two-factor client authentication was demonstrated using smart card technology with digital certificates. While this technology was shown to work, it was not ready for prime time as it was too cumbersome and required added hardware that few travelers would welcome. However, this was a good starting point in demonstrating that a VPN client can utilize two-factor authentication with the promise of simplified and more robust authentication through use of digital certificates.

Follow-up Research

Once the LIP project formally ended with the publication of the final report *SAND99-0100* (Miller, Hudson, Long 1999), additional work was done investigating VPN technology, since it was becoming clear that this technology was important and should be further investigated with a goal toward implementation.

The World Wide Web was used extensively for surveying the VPN market. Several vendors were promoting their VPN efforts. Information was gathered for comparison purposes and to eventually choose one or two products for further investigation. Trade publications were monitored for articles on VPN technology. Magazines were especially helpful in providing product comparisons and evaluations. The VPN '99 conference was attended in March in Atlanta, Georgia, sponsored by Ascend, Compatible Systems, and V-ONE, with participation by other communications heavyweights such as 3COM, Cisco, Sprint, and MCI Worldcom. Corporate presentations included Delta Airlines, Federal Express, McDonalds, and Sears Roebuck. A pre-conference tutorial was attended on advanced VPN security. Although much of the subject matter was new and overwhelming, a good foundation was established.

It was at this conference that the concept of a hardware-based VPN became a reality. Prior to this, only software solutions had been studied, as these were prevalent during the early VPN days. Not one presenter or vendor at the conference presented a software-based product. All were based on some form of dedicated hardware. This is a desirable approach from both performance and security standpoints. A dedicated hardware device can be built to exact requirements, therefore achieving the best possible performance for a given design. Additionally, a dedicated hardware device is likely easier to maintain, and certainly less prone to potential compromise by an attacker.

Several user presentations touted the benefits of VPN applications. These were all large-scale deployments, literally spanning the globe. It was apparent that while the technology was doable, and was showing considerable benefit, there were still considerable concern amongst the presenters as to the development of VPNs, and in particular, VPN standardization.

Lessons learned from the conference:

1. Use a hardware-based solution. Software-based products suffer from performance, maintenance, and security issues.
2. Use a product built with IPSec. This is where the industry is going, so go with the flow. Watch for IPSec compliance; don't settle for IPSec-like as this may have faults and prevent future interoperability.
3. Stick with one vendor. While IPSec promises interoperability, this is not a good thing to try to start with. Getting a VPN working with a single vendor is tough enough at this point. Interoperability will follow.
4. Look for IPSec certification, typically by ICSA (a worldwide leader in security assurance services for Internet-connected companies, http://www.icsa.com).

The LIP project investigation of VPN technology focused on two products that were software-based. It was only during the VPN '99 conference that it became apparent that a hardware-based solution was both superior and available. Early adopters of VPN technology went the easy route by simply developing software, and taking advantage of

existing hardware. A particular limitation of a software solution is performance. A workstation or server configured to run a VPN is not tuned for optimum performance. It must run a general-purpose operating system, carrying with it all the baggage that is included.

Maintenance is another potential headache. A server platform must be maintained separately from the VPN, including OS patches and updates, and hardware.

Additionally, a workstation or server has potential security vulnerabilities, made all the worse by a poorly configured or out-of-date operating system.

These limitations are greatly reduced with a dedicated VPN hardware platform. These devices only operate a VPN. They are not vulnerable to other applications that a workstation or server might be running, either intentionally or unintentionally. They run a dedicated operating system that is efficient and secure.

Formal Project

A formal VPN effort was begun in FY99. Taking account of the wealth of knowledge gained since the original VPN investigations, a plan was developed to implement a VPN.

The plan was built around two phases. The idea was to attack the easier site-to-site implementation first, then go after the more challenging client-to-site implementation.

Phase One – Site-to-Site

Vendor Selection

As VPN technology was in its infancy, changes were certain and often. New developments were coming as quickly as one could comprehend the previous developments. The Web was again used to gather the latest product information and developments in VPNs and IPSec.

One of the best methods of gaining an understanding of a product is to sit down with the vendor and discuss the product. Several vendors were invited to Sandia to pitch their wares. Since the technology was new, this was an opportunity to learn the ins and outs of VPNs as well as what a vendor had to offer.

In the early part of 1999, a VPN company called Network Alchemy was invited to Sandia for a presentation. This company had just formed, and had no released product. However, they gave an impressive presentation, with a product pitch that featured a clustering technology that greatly increases availability and enables easy performance upgrades.

Network Alchemy's clustering technology enables multiple VPN nodes to work together as a single unit. One acts as master, through an election process, and the other nodes work as members of the cluster. If any node drops out of the cluster, the other nodes pick up the sessions of the failed node without dropping any sessions. If a node is added to a cluster, that node is given some of the load. Therefore, node failures do not disrupt sessions and performance can be improved by adding additional nodes as necessary.

Alchemy's VPN devices are router types, meaning they work at network layer 3 and require separate networks on the inside and outside interfaces, passing packets between the networks as appropriate.

Shortly after the Network Alchemy presentation, a company called VPNet was invited to speak about their VPN offerings. Although relatively new, this company already had a record of accomplishment and significant market share over approximately the past year (an eon for VPN startups).

VPNet's devices are bridge types, meaning the interfaces work at network layer 1, and do not require separate networks for the inside and outside interfaces. The devices can be inserted into a network at any point a VPN is desired.

Next in line was a presentation by ODS Networks. ODS has been in the communications business for several years, and the VPN business only a short time (no surprise here). Of particular interest with the ODS product was the lack of IPSec. ODS had not caught up with other VPN vendors in the use of IPSec and were still using proprietary protocols, although they did report that IPSec was to be incorporated in the future.

Site visit

During discussions with VPNet regarding possible equipment purchases, they suggested a visit to a neighboring site where VPNet equipment was being utilized. The site happened to be on Kirtland Air Force Base, at AFOTEC (Air Force Operational Test and Evaluation Center). After making contact with the VPN administrator, a visit was scheduled. AFOTEC was using the VPNet equipment to enable secure dial-in to their protected network. They were very satisfied with the equipment operations, but did point out some minor bugs with the graphical user interface for the administration software.

The visit provided a good user testimonial regarding VPNet. AFOTEC was using the equipment in full production with plans to expand.

Requirements for a Site-to-Site VPN

Having gained some additional insight into VPN technology from research, vendor presentations, and a site visit, a requirements list was developed in order to facilitate product choices. Many of the parameters were well understood and could easily be

identified as important. These were labeled as mandatory. Other parameters were not clear and therefore labeled as desirable. From this list, products could be narrowed down to a final list of candidates.

For example, while it was obvious that a hardware solution was the right solution and therefore mandatory, the need for a tamper-proof box was not so obvious, and was therefore listed as desirable.

Although the initial effort would involve site-to-site, the intent was to find a VPN solution that could address both site-to-site and client-to-site situations. Therefore, the following requirements list addresses VPN aspects such as client capacity, authentication, platform support, and reverse tunneling, that won't come into play until well after site-to-site VPNs have been established.

Requirements for VPN

- Server must be hardware-based with a tamper-proof box desirable.
- Tunneling Protocol must utilize IPSec, with ICSA certification desirable.
- Encryption must utilize 3DES (probably the only one that matters).
- Client Authentication (not significant for the initial site-to-site implementation) must utilize SecurID (this would give us a quick selling point to management since we could leverage our existing authentication system, at least initially) and be Entrust-compatible (to take advantage of our existing Entrust CA and emerging smart card technology).
- Remote management is mandatory. Out-of-band management capability is mandatory, and in-band management is desirable.
- Site-to-Site Tunnels must support at least 10's of sites.
- Client Capacity must support up 2000 connections, with local, national, and international access (majority is local).
- Client Types must include Windows 95/98/NT, with Macintosh/UNIX desirable (Macintosh more than UNIX).
- Communications Interfaces must support 100BaseT (or higher) for main server(s); remote servers may be 10BaseT (or higher).
- Client Reverse-Tunneling must not be allowed.

As was presented earlier, a hardware-based server was deemed necessary to achieve performance and security goals. Tamper-proof (or tamper-resistant) hardware was a concept that some vendors were beginning to embrace, however, none of the products under review considered this.

The tunnel protocol of choice was IPSec. Although it was at the time just beginning to mature, it was clear this was the direction the industry was headed. Additionally, in order to help ensure proper IPSec operation, certification by ICSA was considered a desirable

feature. However, if ICSA certification were considered mandatory, some products would have to be excluded, as they had not sought the certification.

Several encryption algorithms were available, but the strongest was 3DES. This was labeled as mandatory, although virtually all VPN vendors were supporting 3DES.

Client authentication was considered as part of the initial requirements in order to be positioned for future client-to-site work. If a single product could be found to handle both site-to-site and client-to-site applications, this would simplify efforts. SecurID two-factor authentication was chosen as mandatory because it was being used in the production dial-up system. The existing authentication infrastructure could therefore be leveraged for the future client-to-site implementation. Additionally, Entrust compatibility was sought for the same reasons. An Entrust certificate authority was being developed for other needs, and this CA could be potentially used for client-to-site authentication with smart cards carrying digital certificates.

VPN device management comes in several forms, such as Telnet, web-based, and dial-up. Remote management of a VPN device is crucial to efficient operation. How to remotely manage a VPN device is not clear. In-band management, such as Telnet or web-based, may be a security risk, depending on the configuration. Out-of-band management is more secure, but problematic due to the need to establish a separate control channel either by modem or by a separate management network.

A VPN server must support at least tens of sites. It is anticipated that a few dozen could eventually be installed. It is anticipated that a single VPN server will meet this need.

Client tunnels can potentially number 2000 or more. Presently, there are approximately 1500 traditional dial-up users. Although the number of simultaneous dial-up connections is limited to less than 100, it is anticipated that a VPN client-to-site service could be more popular and therefore require more connection points.

Again, considering client-to-site needs, a VPN server must support the full range of Microsoft Windows platforms, 95/98/NT. Macintosh and UNIX platforms are desirable, with Macintosh more so than UNIX[1]. A notable problem with VPN clients is the lack of Macintosh and UNIX products. Most vendors support Windows clients only, due to the large installed base.

Communications interfaces for a VPN server must be 100BaseT (or higher) for the main server. Remote servers can be 10BaseT (or higher) for lower-bandwidth needs at remote sites. Although Sandia's network backbone is currently ATM based, choosing 100BaseT for VPN devices is really a given, since virtually all VPN manufactures are building 100BaseT devices. Only a few higher end VPN devices, such as some VPN-enabled routers, can support higher speed interfaces.

[1] With Linux rapidly gaining popularity, the original statement "Mac more than UNIX" is losing some validity.

Client reverse tunneling, also referred to as split tunneling, is a feature that allows the VPN client to simultaneously access the Internet and a tunnel. This can be a convenience, but is more likely a security concern. A hacker could potentially gain access to a VPN client that is "visible" to the network (ie: not exclusively tunneling traffic). With reverse tunneling disabled, a VPN client cannot be probed, scanned, or attacked directly. The only attack available is on the IPSec packets themselves, which are designed to be extremely difficult to break. Therefore, this feature should be disabled.

Product matrix

A product comparison matrix was then developed, using the requirements as a guide (see Table 1). It was clear that no one product met all the requirements, a frustrating conclusion. Note that this matrix was current at the time of development. Since then, many product changes have taken place but were not incorporated into the matrix in order to indicate the status of the products at the time of evaluation. For example, some of the listed companies have been purchased by other companies, and some products have subsequently changed names. Compatible Systems was purchased by Cisco, with the IntraPort product line name changed to VPN 3000. TimeStep was purchased by Newbridge, then Alcatel, with the product names changing in various ways. VPNet was purchased by Avaya, with the product names maintained. Network Alchemy was purchased by Nokia, with the CryptoCluster product name changed to the CC series.

Eleven products from seven manufacturers are represented in Table 1. These eleven products cover a wide range of VPN features. The first feature, "Server in Hardware", demonstrates that most vendors were embracing the hardware server solution. Only one product in the matrix, V-ONE's SmartGate, was software based and was included here only for comparison since Sandia's requirement was for a hardware-based solution. The next feature column, IPSec, shows that only one product has not yet incorporated the emerging security protocol. As described earlier, ODS had not yet working IPSec into their products. The next feature column, 3DES, demonstrates that this encryption algorithm is widely accepted. For ICSA compliance, most vendors have chosen to be certified, ensuring proper cryptographic operation and interoperability. A notable exception was Network Alchemy, although they did claim IPSec compliance. The next two feature columns, SecurID and Entrust, show that user authentication is far from universal. Only one vendor, VPNet, supported both SecurID and Entrust, with Network Alchemy promising to do so in the next release. The next column, Client Capacity, varied considerably, depending on the power of the particular equipment, ranging from 200 client tunnels to 30,000. The UNIX/Macintosh client column shows that only one vendor could supply a software client that operates on both non-Windows platforms. The LAN Interface column shows that all vendors cater to some form of Ethernet, with most supporting 100BaseT (Fast Ethernet). The Remote Management column shows that most vendors support remote management techniques, ranging from direct serial cable connection (RS-232) to Java-based GUI.

Summarizing the matrix, no product met all the requirements. The Compatible Systems IntraPort 2+ came close, but failed to support Entrust, and fell short with client tunnel capacity. The compromise winner was VPNet's VSU-1010 and VSU-1100, with the only shortcoming being a promise to release a UNIX/Macintosh client.

VPN Product Comparisons

(As of Early 1999)

Table 1 VPN Product Comparisons

	Server in Hardware	IPSec	3DES	ISCA Compliant	SecurID	Entrust	Client Capacity	UNIX/ MAC Client	LAN Interface	Remote Management
Compatible Systems IntraPort 2+	√	√	√	√	√	No	200	Power Mac 8.0+ & Intel Linux	10/100BaseT	GUI/Telnet & out-of-band
ODS CrytoWatch 4210	√	No	√	√	No	No	?	No	10BaseT	?
Red Creek Ravlin 10	√	√	√	√	√	No	?	No	10BaseT	In-band & RS232
TimeStep 2520/4520	√	√	√	√	No	√	500	Mac	10BaseT	In-band & RS-232
TimeStep 7520	√	√	√	√	No	√	2000	Mac	10/100BaseT	In-band & RS-232
V-One SmartGate	No	√	√	No	√	√	?	Mac	Whatever is in user-supplied hardware	?
VPNware VSU-1010	√	√	√	√	√	√	600+	Q3 '99	10BaseT	Java(SSL), SNMP, RS-232
VPNware VSU-1100	√	√	√	√	√	√	5000	Q3 '99	10/100BaseT	Java(SSL), SNMP, RS-232
Network Alchemy VPN Server 2500/5000	√	√	√	No, but claims IPSec com-pliance	No, but coming Q4 99	No, but coming Q4 99	10,000/ 30,000	No	10/100BaseT	Java for 95/98/NT, Solaris, RS-232, Telnet

Product Purchase/Borrow

A philosophy adopted for this project was to not study the market to death, but rather make some purchase or borrow decisions quickly to get the project moving. If the choice later appeared to be a poor one, the product could always be changed.

In order to leverage an anemic budget and lessen the impact of possible false starts, a two-phased acquire and borrow acquisition approach was utilized.

The purchase and borrow decisions were based on the product matrix and additional factors. Network Alchemy had a unique clustering technology that helped them stand apart from the competition. VPNet was well positioned in the marketplace, but seemingly not so eager to make a sale since they would not loan equipment for evaluation. They preferred to make a sale with a money back guarantee.

CryptoCluster 5000 and 2500 Loan

With Network Alchemy eager to make a sale, they offered to lend some equipment for evaluation. In addition, they offered to send a customer engineer to help. They sent three CryptoCluster 5000s. These are their top-of-the-line boxes. Two would be clustered together and tunneled to the third device for a demonstration.

Upon the support engineer's arrival, a small testbed was built using the two clustered 5000s and the single 5000 configured back to back as shown in Figure 6.

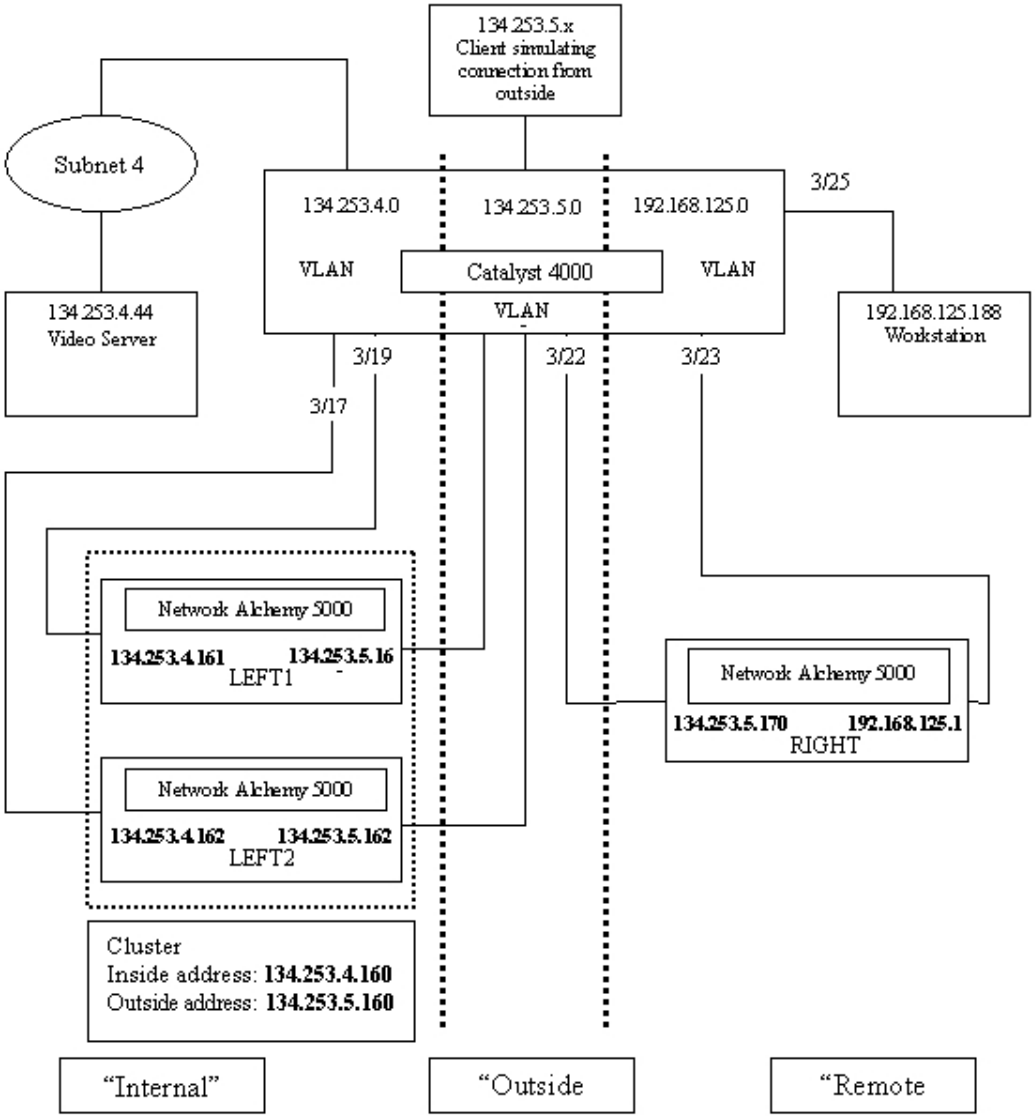

Figure 6. Network Alchemy Demo Testbed

The goal was to simply get the devices talking to learn how they operate. The support engineer did all the work configuring the boxes, while explaining what he was doing. However, due to the short, one-day visit, there was a great deal of information to absorb, and not much time to absorb it.

The Network Alchemy CryptoCluster series can be managed from the console either by way of a serial port or Telnet, or by using the GUI from a workstation. Although the console is fully featured, Network Alchemy recommends that the GUI be used for

management, and the console reserved for troubleshooting only. This is because the GUI stores configuration data in the user workstation. If any changes are made at the console, these changes will not be reflected in the GUI database, and any subsequent updates from the GUI will override changes made at the console.

For this temporary testbed, the Network Alchemy engineer never operated the GUI. All work was done from the console since time was short and he was well versed in the console operation.

Once the two clustered 5000s in the "internal" network were connected to the single 5000 in the "remote site" network, tests were run to determine throughput. A single stream TTCP test resulted in 80 Mbps with the CryptoClusters configured for 3DES. This is a good value for a single stream, considering the maximum possible throughput in one direction is limited to 100 Mbps due to the 100BaseT interfaces. No attempts were made to determine aggregate throughput capability due to limited available time.

Some testing was done with the client VPN capability of the 5000, but only to a very limited extent. The Network Alchemy VPN client software was not capable of two-factor authentication at the time of evaluation (only username/password), and therefore was not seriously considered for use at Sandia. However, some knowledge was gained be trying the software and observing the performance. The client worked well within its limits. It was easy to install, and no appreciable problems were encountered.

CryptoCluster 500 and 2500 Purchase

Having seen the equipment first-hand, it was decided to go ahead with a purchase. The CryptoCluster 5000 was ruled out as being too much hardware for our needs, and too expensive for our budget. The CryptoCluster 2500 was a better fit both operationally and financially. Additionally, the CryptoCluster 500 was appropriate for remote sites. An order was placed for two 2500s and two 500s. If chosen for implementation over other products, the two 2500s would be installed in a cluster at the edge of the protected network, with the 500s available for remote sites.

VPNet VSU-10 and VSU-1010 Purchase

During this same time frame, an order was placed with VPNet for a package containing 1 VSU-1010, 1 VSU-10, the VPNmanager, and 1,200 remote clients. The VSU-1010 was their middle-of-the-line, rack mountable unit capable of 10 Mbps with 3DES encryption. The VSU-10 was the bottom-of-the-line remote site box, a small, tabletop device capable of 8 Mbps throughput with 3DES encryption.

The VSU product line included a Windows application for device management. This application was Java-based and required a particular version of the Netscape Navigator browser that was some versions behind the current Netscape release. This made use of

the VSU manager problematic since a separate version of Netscape was required for managing the VSU devices along with the current version of Netscape for corporate use.

Testbeds

With two sets of VPN devices in hand, testbeds were developed for gaining a better understanding of VPN technology, and exploring the products that were purchased. A great deal of study had been done of the market and product descriptions, but little hands-on had been done up to this point (other than a fast and furious demonstration from Network Alchemy as previously described).

VPNet

VPNet VPN equipment is bridge-based; therefore, the equipment is placed in line with the communications path without requiring any changes or additions to subnet addressing. A simple testbed was developed by putting VPNet devices in line with two workstations separated by a router for exploring VPN operation across separate networks.

A VSU acts as a proxy for the downstream router. When packets are destined for the router, the VSU will step in and answer the ARP query for the router (proxy ARP), giving its own MAC address in place of the router. Therefore, the packets will be sent to the VSU as if going to the router. The VSU will then act on those packets if necessary, (providing VPN services), then forward the packet to the router.

The VPN was configured with 3DES encryption, MD5 authentication, using data compression to boost performance.

After struggling through the initial configuration process and several false starts, a successful VPN was established between the VPNet devices as shown in Figure 7. Throughput tests using TTCP resulted in 3.2 Mbps performance, considerably lower than the advertised 8 Mbps. The difference could not be accounted for.

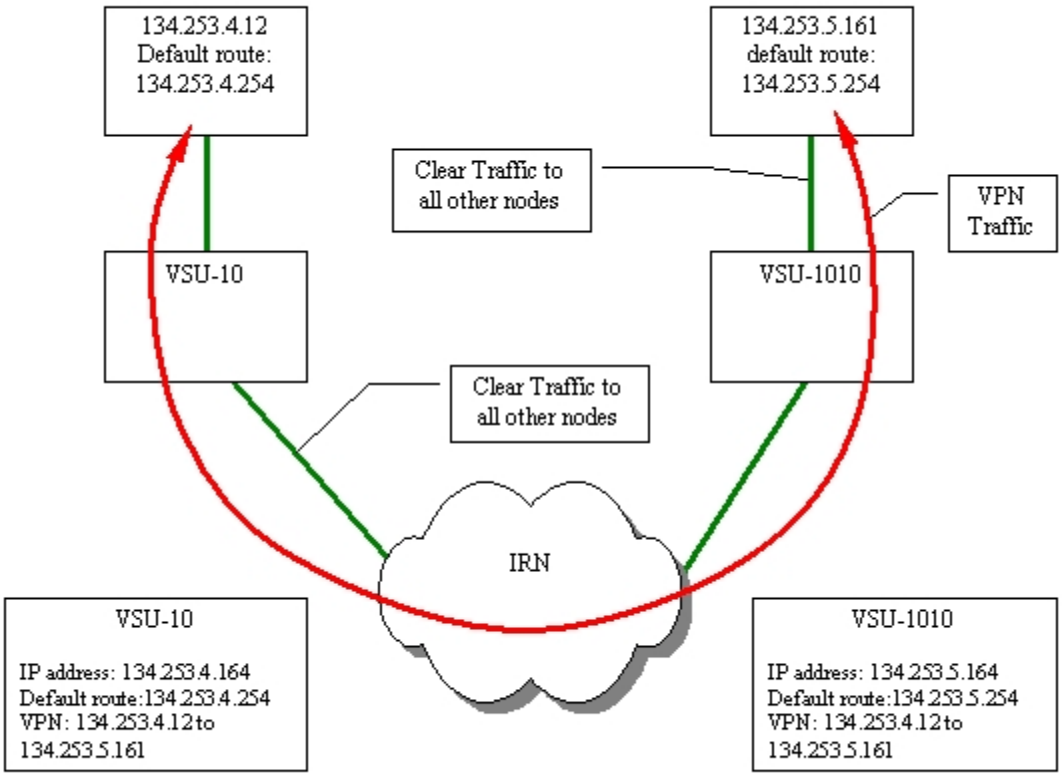

Figure 7. VPNet Testbed

Network Alchemy

A second testbed was developed for the Network Alchemy equipment as shown in Figure 8. For this testbed, a router was placed outside one end of the VPN to explore routing issues outside the VPN. This configuration demonstrates the ability to route to the VPN server using a static route in the routers. The VPN server must also have a static route in order to point traffic in the correct direction (default route to one interface, static route to the other interface). In this configuration, the second VPN server does not need a static route since it does not talk to a router on the private side (in a real installation, this server would require a static route to talk to a router on the private side since the private side would be comprised of many networks).

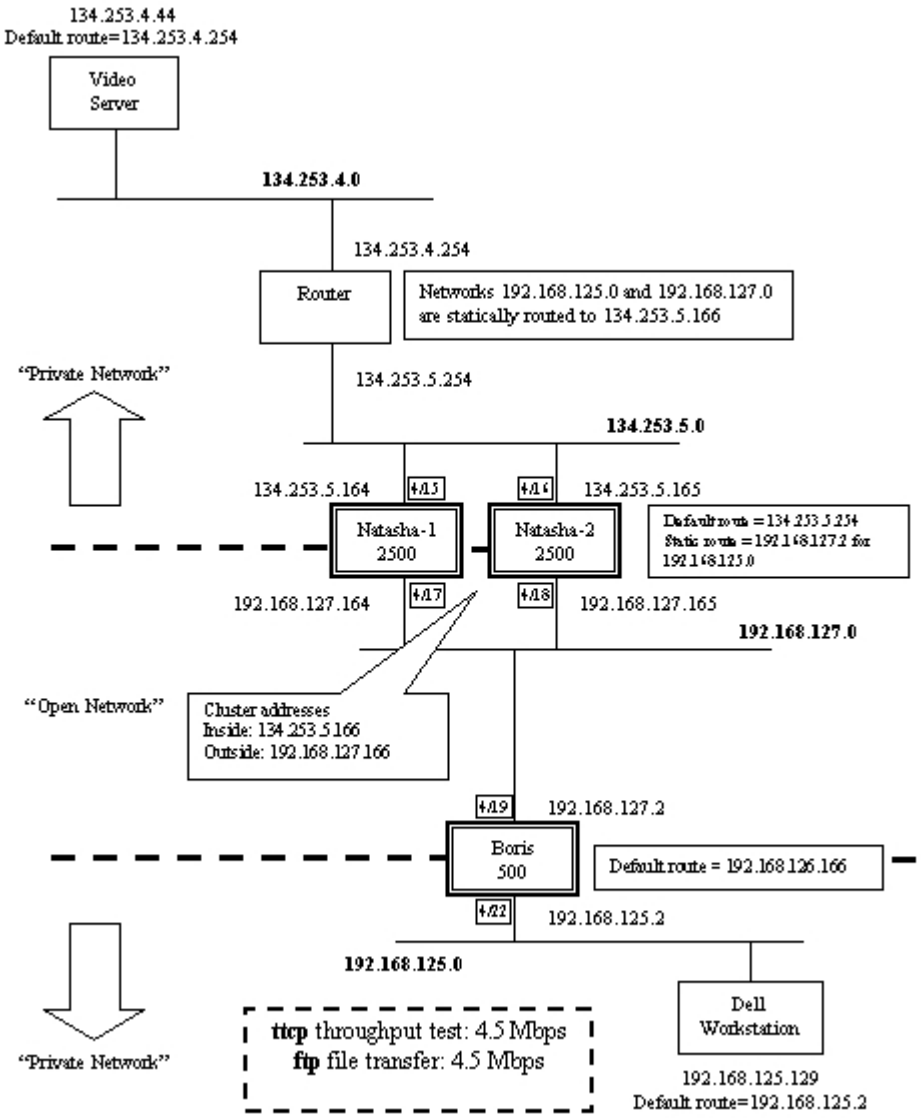

Figure 8. Network Alchemy Cluster Testbed

This testbed utilized two 2500's in a cluster, connected to a single 500. Note that the 2500 cluster has its own inside and outside IP addresses in addition to inside and outside addresses assigned to each individual cluster member. All communication to the cluster involves the cluster inside and outside addresses. The cluster master is elected among cluster members and responds to cluster communications, doling out sessions to the appropriate cluster member, including itself, based on utilization.

Clustering was tested by dropping a 2500 from the cluster, adding it back in, and observing tunnel behavior. Tunnel traffic continued to flow uninterrupted. To ensure that

tunnel traffic would be affected by removal of a cluster member, each 2500 was taken out of the cluster and returned. Removal from the cluster was done by simply turning the unit off.

Testing of the Network Alchemy VPN was done by connecting a workstation on one "private network" and a server on the other "private network." Streaming video was used to load the VPN tunnel and observe performance and stability. Streaming video operated well. No drop-outs occurred.

Performance

Throughput measurements were made using the TTCP test program running between two Dell Precision 410 PCs with 100Mpbs Ethernet interfaces. With a 500 connected to a 2500 cluster, throughput is limited to the maximum available in the 500. This proved to be about 4.5 Mbps with 3DES encryption, which is exactly the value specified by the vendor. FTP file transfers also resulted in 4.5 Mbps. These are reasonable values for a small VPN device about the size of a box of tissues. For small remote sites, 4.5 Mbps is a huge improvement over 56 Kbps dial-up.

In order to baseline the throughput of the 2500s, the testbed was arranged with only two 2500 nodes set back to back as shown in Figure 9. In this configuration, throughput should be much higher with the vendor quoting about 44 Mbps. Again testing with TTCP, the throughput measured about 13 Mbps with 3DES encryption. This is much lower than the vendor specification. Discussions with the vendor revealed that their method of testing utilized UDP packets with a test called NETTEST rather than TCP packets used in the TTCP test. Since UDP packets operate without acknowledgement, this can result in higher throughput.

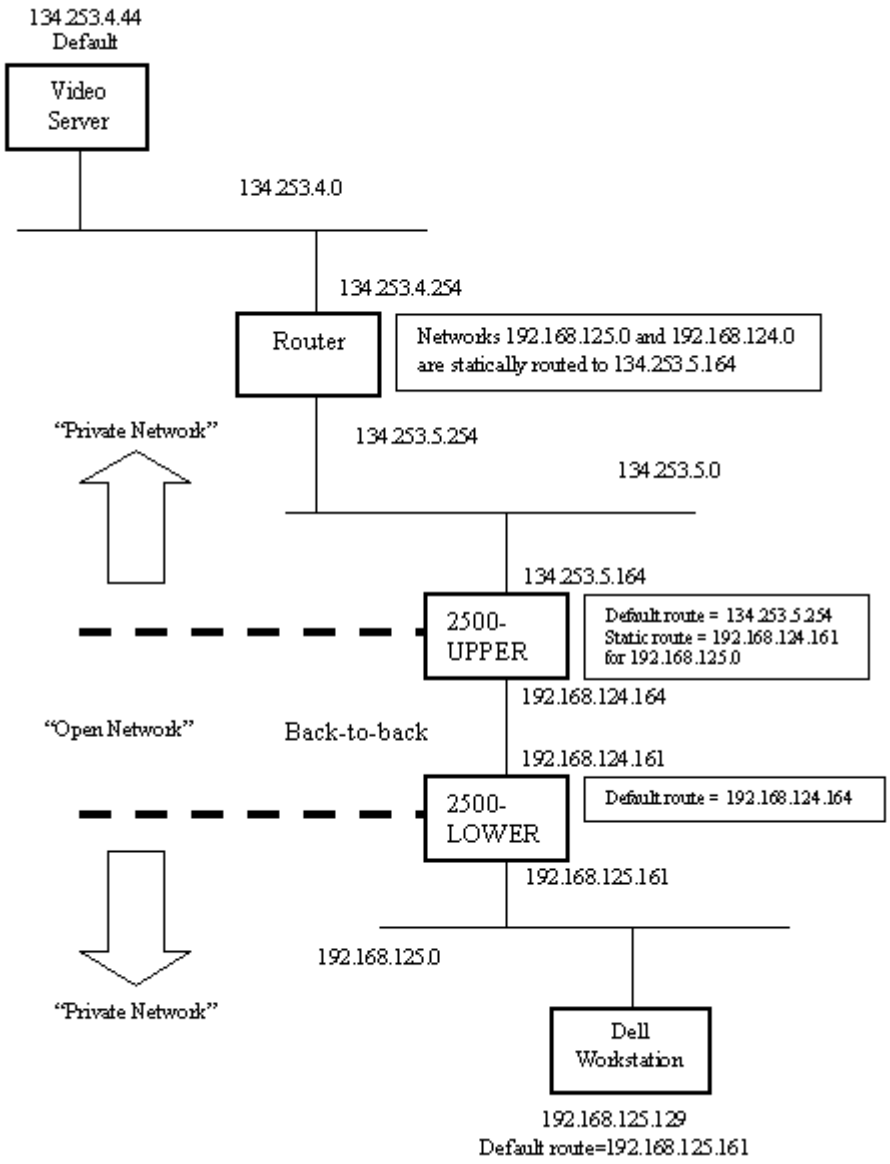

Figure 9. Network Alchemy PerformanceTestbed

Decision Point

In order to move forward, a point was reached where a decision was to be made as to what equipment to use for a VPN. Both the VPNet and the Network Alchemy equipment proved operational, however, the Network Alchemy equipment was chosen over VPNet for a number of reasons. Chief among them was the disadvantage of the bridge operation of the VPNet equipment. In some cases, this can be an advantage, but for Sandia's ATM network backbone, a bridge device is a poor match. Finding a point in the network where

a bridge device can be connected would be difficult. Although there are many Ethernet segments within the network, few of these segments connect to routers using Ethernet. Rather, the Ethernet switches either connect to routers over ATM, or contain virtual routers that connect over the switch backplane.

On the other hand, with the Network Alchemy equipment operating in a router-like mode, it would be much simpler to engineer a connection by simply assigning appropriate IP addresses to the inside and outside interfaces and connecting to the appropriate Ethernet networks. Other features that put Network Alchemy ahead of VPNet were clustering, manageability, and company responsiveness.

Security Proposal

Having made a VPN equipment decision, a proposal for site-to-site VPN implementations was drafted for submittal to computer security for their review and approval. The specifics of Network Alchemy were documented in the proposal. The proposal is attached as Appendix A.

The gist of the proposal is a description of the technology and how it can be applied to Sandia to reduce some communications costs and provide added services. IPSec details were covered, including use of ESP for authentication and encryption, and use of IKE for key negotiation. Testing was done on the equipment to ensure proper operation of various security functions.

SC99

The Supercomputer ’99 show came along in November of 1999. In order to facilitate a particular demonstration, a VPN was requested for connecting the showroom floor to Sandia’s restricted network. This was the first opportunity to get a real VPN operational. As this was a trade show, the VPN would be needed for one week only. After developing a plan, it was presented to the Computer Security department for review and approval.

The proposal was effectively an addendum to the full proposal previously submitted to Computer Security. As such, the technology was understood and the proposal was quickly accepted. The proposal suggested use of two Network Alchemy 2500s with one at SC99 and one in Albuquerque. The configuration of two 2500s was successfully tested in the SC99 testbed in Albuquerque prior to implementation at SC99.

The SC99 VPN was successfully employed as previously described in SAND2000-1812 *(Pratt et al., 2000).* Although it was a considerable challenge to get the VPN running due to lack of experience and a troublesome SC99 network, the end result was satisfactory and the experience invaluable for moving forward with a production implementation.

SC2000

SC2000 presented another opportunity to implement a temporary VPN. However, the requirements where different from those of the previous year. Rather than establishing a VPN from the show floor to a particular host, a VPN was needed from the show floor to an entire subnet within the Sandia Restricted Network. Another change was a desire to utilize different VPN equipment in order to simplify the implementation. This was done by using the Lucent Firewall that was already in place protecting the Sandia open network. The firewall was to be operated in a dual role as a firewall and VPN server. Details of the effort are documented in a yet-to-be released SAND report tentatively titled *The ASCI Network for SC 2000: Gigabyte Networking*.

For SC2000, the VPN was more of a challenge to design than to implement. Several different configurations and equipment choices were considered before settling on a solution using Lucent's VPN firewall. The difficulty lay in getting the VPN into a subnet without affecting the existing production services. The solution was to use the existing firewall/VPN server located at the border between the Internet and Sandia's open network. The VPN would terminate at this point and pass data in the clear to the destination subnet. The implementation was smooth, but a human glitch did interrupt the VPN for a short period of hours. It was restored with some clever effort involving work at both ends of the VPN.

Initial Production Implementation

Finally, the time had come to initiate a production implementation. However, before beginning implementation, the Telecommunications Operations Department suggested that the VPN implementation go through their newly developed Change Management Process, or CMP. This process is designed to review a change to the network to ensure a smooth implementation. Requirements include documentation of the proposed change, testing of equipment involved in the change, and design reviews. This process slowed the implementation considerably, but proved to be an effective method of ensuring success.

An additional outcome of the CMP process was a desire for operations personnel to be trained on the VPN equipment prior to implementation. This is a proper path to success, but entailed another delay in implementation. A class was scheduled and held. Operations personnel were introduced to VPN technology and learned the operation and management of the Network Alchemy VPN equipment.

The CMP design reviews provided a forum for discussion of placement of the VPN server with respect to the corporate firewall. As discussed previously, there are three options for placing a VPN server in the network: parallel to, in front of, or behind the

firewall. The review team agreed that placing the VPN server parallel to the firewall was the best solution since there would be no changes required to the firewall, the VPN server would not be a potential bottleneck for the non-VPN traffic, and the VPN server would adequately protect itself from attack.

Figure 10 shows the chosen VPN server position. Sitting parallel to the firewall requires that the VPN server be able to protect itself. This is indeed the case as the VPN server only responds to authenticated IPSec packets. No general services run on the VPN server since it is a dedicated VPN platform.

Network diagram

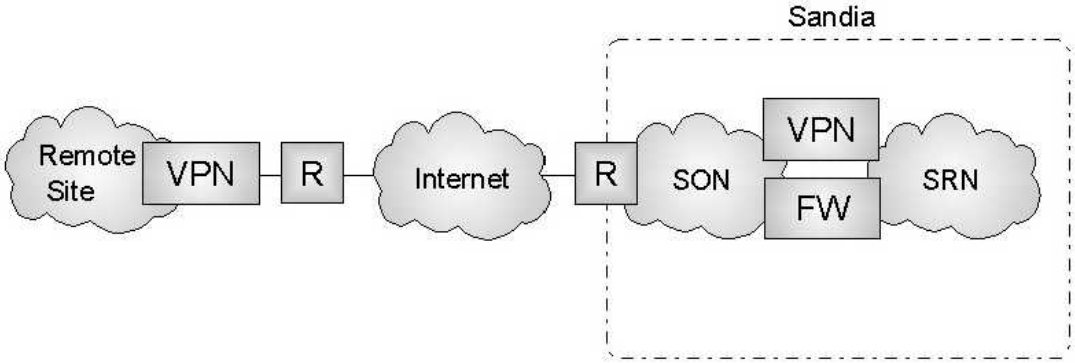

Figure 10. VPN Server Position with Respect to the Firewall

KTF (Kauai Test Facility, Hawaii)

Sandia's KTF facility has long experienced communications problems due to the long distance from the mainland and the remote location. Digital communications back to Sandia have traditionally been done over phone lines using modems, but connection quality of 2.4 Kbps has been poor if not occasionally useless. KTF personnel required digital communications for submitting time cards and expense vouchers as a minimum, and this would typically require up to four hours to complete. There are only six trunk lines available for access off the complex, so when a mission was at the facility, there was always a conflict for voice and data communications.

A VPN solution had been identified as early as 1998 as the only viable solution to KTF digital communications problem. The long distance from KTF to Albuquerque would necessitate a T-1 (1.536 Kbps) leased line at a cost between $15,000 - $18,000 per

month, a cost that neither the Telecommunications Department nor the KTF departments were willing to finance. Therefore, KTF was identified as Sandia's first VPN production site.

A VPN solution requires a connection to the Internet. The only possible connection to the Internet was through the Navy's Defense Research Network (DRN). This is an ATM T-3 network from the Hawaiian Islands to the mainland with a connection to the Internet and DOE's ESNET.

An arrangement was worked out to get an Ethernet port connection into the DRN. This necessitated installing an extension using media converters across single mode fiber for a distance of about 3 miles from the KTF facility to the PMRF facility (Pacific Missile Range Facility). The KTF personnel purchased some devices and installed them. The link was initially unstable. Network testing showed dropouts and intermittent long delays for data transfers. It was decided to still proceed with the VPN equipment installation and perform troubleshooting of the circuit with the technician on site.

Once the CryptoCluster 500 was installed and initially configured at KTF, it was necessary to make contact with the 500 from Albuquerque using the CryptoConsole in order to complete the installation. This proved to be challenging because the link to KTF was still operating poorly, with afternoons almost impossible. A morning attempt at contact was finally successful, and the configuration was completed. A VPN was now established between the SRN and KTF, however, it did not function well due to the intermittent behavior of the KTF link.

With time running out, several members of the Advanced Networking Integration and the Telecommunications Operations departments put their collective experience together in a marathon session in an attempt to resolve the networking problem at KTF prior to return of the KTF personnel to Albuquerque.

It was determined through a process of elimination that a problem existed in the fiber optic link between KTF and PMRF. The media converters that change the network signal from electrical to optical for transport over the fiber cable were experiencing a problem interacting with the Ethernet port on the router at PMRF. The media converters operate in full duplex mode, but the PMRF router Ethernet port was operating in half duplex mode. After contacting the media converter manufacturer to verify the duplex mode operation, the DRN network administrator was contacted and asked to change the PMRF router Ethernet port configuration to properly match the media converter. After the change was made, the link performance improved dramatically, but was still not operating satisfactorily. Further troubleshooting determined that one of the media converters was not working properly and was replaced with a spare, providing a solid link and a good VPN connection.

With time running out for the KTF crew to return to Albuquerque, only some quick, informal throughput testing was done on the VPN. A value of 400 Kbps was obtained, which is a vast improvement over the original dial-up speed of 2.4 Kbps.

Since its inception, the KTF VPN has required virtually no maintenance or troubleshooting. The circuit has remained continuously in operation. It has been closely monitored by reviewing log records. The logs are routinely checked for anomalies. A more formal monitoring function using SNMP is yet to be established. SNMP provides an automated process for monitoring communications equipment for the purpose of generating alarms and statistics.

Oops – FIPS 140

After implementing the Phase One site-to-site, attention was brought to a Federal Government requirement for cryptographic systems, which applies to VPN use at Sandia. FIPS is a set of standards governing computer security for the Federal Government and its contractors. FIPS is developed and maintained by NIST, the National Institute for Standards and Technology.

Below is the description of the FIPS 140-1 requirement as specified at the NIST web site http://csrc.nist.gov/cryptval/140-1/1401val.htm:

The NIST Cryptographic Module Validation (CMV) Program was announced on July 17, 1995. This program validates cryptographic modules for conformance to FIPS 140-1, *Security Requirements for Cryptographic Modules*. In the "Applicability" section of FIPS 140-1, it states that:

> "This standard is applicable to all Federal agencies that use cryptographic-based security systems to protect unclassified information within computer and telecommunication systems (including voice systems) that are not subject to Section 2315 of Title 10, U.S. Code, or Section 3502(2) of Title 44, U.S. Code. This standard shall be used in designing, acquiring and implementing cryptographic-based security systems within computer and telecommunication systems (including voice systems), operated by a Federal agency or by a contractor of a Federal agency or other organization that processes information (using a computer or telecommunications system) on behalf of the Federal Government to accomplish a Federal function. Federal agencies which use cryptographic-based security systems for protecting classified information may use those systems for protecting unclassified information in lieu of systems that comply with this standard. Non-Federal government organizations are encouraged to adopt and use this standard when it provides the desired security for protecting valuable or sensitive information."

Because initial VPN investigations at Sandia did not consider FIPS certification, and with the help of Murphy's Law, all VPN equipment purchased to date did not have FIPS 140-1 certification. It would not be prudent to ignore the FIPS requirement, and therefore the position adopted was to add FIPS 140-1 certification to the Sandia VPN requirements list. This meant going back to the drawing board.

Because FIPS 140-1 certification applies to cryptographic modules and not to systems, it is important to look carefully at a particular vendor solution to determine if all cryptographic modules of a system meet FIPS 140-1. It is recommended here that any module of a system that does not have FIPS 140-1 certification should invalidate the system. For example, a remote access VPN product that has FIPS 140-1 certification for its hardware, but not for its client software, cannot be considered for implementation at SNL.

Searching through the NIST FIPS 140-1 certification list revealed a relatively small set of equipment to choose from, compared to the non-FIPS approved set of VPN equipment. For site-to-site applications, the Cisco 7100 VPN Router series was a good balance between higher-bandwidth requirements and cost. The RedCreek Ravlin 10 was a possibility for lower-bandwidth situations, and was already in use at Sandia for encrypting a wireless network extension. For client-to-site applications, taking into consideration the requirement for a FIPS 140-1 approved 'system', the NIST list dropped to zero! However, a timely missive from a vendor representative indicated that one of their products was currently undergoing FIPS 140-1 certification, including the client software. Assuming this true, and considering all remaining requirements, the Cisco VPN 3000 Series product emerged as a possible saving grace.

Kodiak

After the successful VPN implementation for KTF, a request was received for a VPN to Sandia's launch facility on Kodiak Island in Alaska. Communications is just as problematic at Kodiak as at KTF due to the long distance and poor telephone service. Access to Albuquerque has been by modem, but further complicated by a satellite link from Kodiak to Anchorage, which adds considerable latency. Measurements of latency between Kodiak and Albuquerque showed that the latency jumps from 110 ms to 1342 ms when the satellite link is encountered (contrast this with KTF latency of approximately 180 ms). High latency should not have any effect on a VPN, but can affect applications that run over the VPN. A typical latency problem encountered by an application is a time-out that occurs when packets are not acknowledged within a set period because of long latency time from a satellite system. These time-outs can often be adjusted within an application to accommodate satellite latency.

Due to the enlightenment of the FIPS 140-1 requirement, a decision was made to not utilize Network Alchemy equipment for the Kodiak installation, despite the fact that the equipment was already in use and ready for expansion. Accommodating Kodiak with Network Alchemy would require only installing a unit at Kodiak and connecting it to the

existing Network Alchemy VPN server in Albuquerque. However, a new set of equipment was chosen, based on previous experience with a wireless link established to Sandia's remote site in Albuquerque. This wireless link uses a RedCreek Ravlin 10 unit to provide 3DES encryption. Since this device was already in use and had been through the CMP process, it was chosen for Kodiak.

A RedCreek Ravlin 10 unit was installed between Sandia's Open Network and the Restricted Network at the building 880 MDC, alongside the Network Alchemy VPN node. The remote unit was installed at the Kodiak Launch Center as a temporary connection for initial testing prior to Sandia personnel arriving at Kodiak for a launch. Because this circuit had long satellite delays and a low data rate of 128Kbps, testing prior to the arrival of the launch crew seemed appropriate, and proved to be the right choice.

Just as with the KTF VPN installation, the Kodiak VPN installation encountered a network problem. Although the VPN worked across the link, with some minor complications, the link was causing packet loses when the packet sizes were above 412 bytes. This caused the circuit to be mostly useless. Troubleshooting from Albuquerque narrowed the problem down to the satellite link into Kodiak. The contract personnel that maintain communications for the Kodiak facility were left to work with their Internet service provider to correct the problem.

In the fall, when the launch crew returns to Kodiak for the next launch, it is anticipated that the Kodiak satellite link will be operational and the VPN will be reestablished.

Phase Two – Client-to-Site

The initial plan for a client-to-site VPN was to develop a proposal only. This was changed to doing a pilot to demonstrate the feasibility of the technology. More and more employees have been requesting DSL and cable modem access to the SRN. Only a VPN can accommodate these requests. Additionally, with corporate emphasis being placed on telecommuting, a VPN could help here as well.

The client-to-site phase began in earnest once the site-to-site implementations were well established. However, as discussed above, FIPS 140-1 certification threw a wrench into the works and temporarily derailed the operation.

Vendor Selection

Prior to understanding the applicability of FIPS 140-1 certification to VPNs, a near-perfect product had been identified for use in the client-to-site configuration. For a separate but related project intended to interconnect the NWC (Nuclear Weapons Complex), a Cisco VPN 5001 was selected by mutual agreement of the members of the

complex. As this device was being evaluated at SNL for the NWC application, it became clear that the 5001 was well suited for VPN client use as well as the intended site-to-site. The 5001 was originally called the Intraport 2+ before Cisco purchased Compatible Systems. The Intraport was evaluated on paper as part of the original product screening process, but was not selected. Hindsight being what it is, this was unfortunate. The 5001 (and previously the Intraport) is well suited to client VPN use for several reasons, chief among them the fact that clients are available for Macintosh and Unix as well as Windows. This is rare if not unique among VPN vendors.

After some testing of the 5001, the process of implementing a pilot program was about to begin, until the FIPS hammer fell. The 5001 was immediately out of the running, a scramble began for a replacement, and the Cisco VPN 3000 Series was identified as a likely replacement.

Since no VPN 3000 was on hand, the pilot implementation was again on hold while a 3000 was obtained. For the sake of prudence, a Cisco 3030 was acquired by loan from Cisco for evaluation and initial pilot implementation. The VPN 3000 series is comprised of five units, ranging from the low-end 3005 unit, up to the high-end 3080. The units are basically identical, differentiated only by redundancy, memory, and simultaneous client capacity. The 3030 was chosen because it is designed for medium- to large-sized organizations with bandwidth requirements from full T1/E1 through fractional T3 (50 Mbps maximum performance) and up to 1500 simultaneous sessions, which is a good fit for Sandia's current remote access needs.

Security Proposal

Having made a client VPN equipment decision, a proposal was drafted for submittal to computer security for their review and approval. The specifics of the VPN 3030 were documented in the proposal. The proposal is attached as Appendix B.

As with the site-to-site proposal, the gist of this proposal is a description of the technology and how it can be applied to Sandia to reduce some communications costs and provide added services. IPSec details were covered, including use of ESP for packet authentication and encryption, and use of IKE for key negotiation. One significant difference in this proposal is the discussion of user authentication and how SecurID is used to verify user identify. Testing was done on the equipment to ensure proper operation of various security functions.

Implementation

As with the site-to-site implementation, the client-to-site required a pass through the CMP process, including documented system tests and a design review. The 3030 was run

through a set of security tests to determine if it operated as advertised. This included such tests as checking password operation, console timeout, protection of system parameters by password, device behavior during software and power resets, and encryption of tunnel data. The test results were presented at a design review along with the proposed network connection. With CMP approval, the 3030 was connected to the network in order to begin initial user testing.

A schematic of the test configuration is shown in Figure 11. A client machine, typically a laptop, running the VPN client software, connects to an ISP at some location in the world. This could be a local connection at the user's home, or somewhere on the road. Once the ISP connection is established, the user invokes the VPN software to make a connection to the VPN server located at the boundary of the SRN. The user is challenged for a SecurID authentication. Once the user enters the username and pass code, this data is passed to the VPN server, then forwarded to a Radius server, then finally forwarded to the SecurID server. Once the user has authenticated, a tunnel is established from the client machine to the SRN over the Internet.

Note that the client connection to the Internet is not restricted to dial-up. Any ISP connection will work, whether it is dial-up, ISDN, DSL, or cable modem. ISDN, DSL, and cable are often referred to as broadband, due to the higher communications rates. Dial-up in only rarely referred to as narrowband. Also, note that the client is not limited to a laptop. Desktop systems can also employ a VPN client for fixed locations that utilize ISDN, DSL, or cable modem connections.

The Cisco VPN 3000 Series has the capability to compress data to help increase the performance of dial-up connections. The higher speed connections, ISDN, DSL, and cable, are not recommended for compression by Cisco since the VPN server CPU requirement might overdrive the server. Therefore, two connection profiles were developed for distinguishing compression needs.

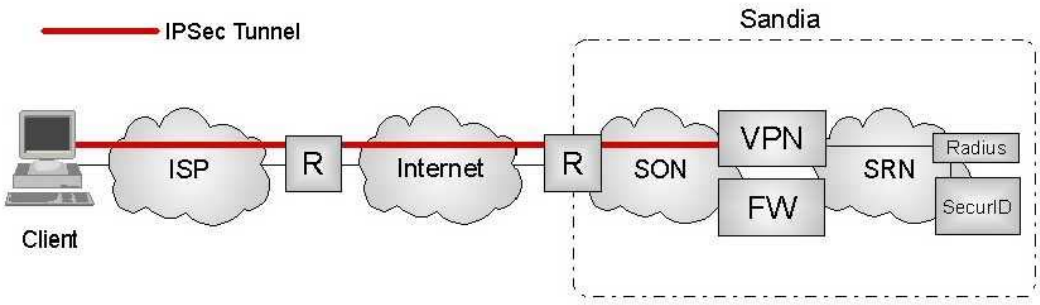

Figure 11. Remote Access VPN

Shakedown

An initial test program was begun with a very small set of users for shaking down the system. Both dial-up and broadband access was tested, with the majority of initial connections being the broadband type. Several minor problems were encountered, but otherwise the system has performed as expected.

Initial problems have involved improper configuration, lack of understanding of the system, and authentication failures. Most new systems will have some bugs to work out and this implementation is no exception. The initial configuration problems were tracked down and corrected. Authentication problems were the result of using a development authentication system that is not kept up to date with the production system. Some of the VPN 3030 test users had new SecurID cards that were not initialized in the development system. As a result, these users were unable to successfully authenticate until their SecurID accounts were corrected in the development system. No other SecurID authentication problems were encountered during the shakedown test.

The traveling employee wishing to connect to the SRN using the VPN must access the Internet first. This requires an ISP account of some sort. Initial testing involved personal ISP accounts, which is altogether appropriate for broadband access in which the account is at a fixed site, but not particularly suited to dial-up situations in which the account might originate from anywhere in the world.

In order to possibly accommodate the need for ubiquitous ISP access, iPass was invited to Sandia to discuss their product. They are one of two such services, the other being Gric Communications, Inc. iPass is a sort of clearing house for Internet access. They contract with ISPs around the world to provide a single point of contact between a user and the Internet. An iPass customer is supplied with a dialer application that simplifies the Internet access process by dialing the correct local access telephone number based on the city or area code of the user's current location. When the dialer makes a successful connection to the Internet, the VPN software is invoked for connecting the client back to the corporate site. Usage is billed per hour with the rate based on the location. Most domestic locations have approximately a 4-cent/minute rate, with international rates higher, depending on the location. Billing is centralized and invoiced monthly.

Why use iPass when 1-800 service is available domestically to Sandia, and direct dial is possible from anywhere in the world, with rates comparable to iPass? From a cost standpoint, there may be no advantage to iPass for Sandia. In fact, directly dialing Sandia would be easier than using a dialing application that requires entry of a country, city, or area code.

Remote Access VPN Performance

Throughput of a remote access VPN is going to vary considerably between dial-up and broadband. Since dial-up connections are limited to 56 Kbps, throughput is relatively low to start. The addition of a VPN running over a 56 Kbps connection can only degrade performance due to the addition of IPSec. This is overcome largely using data compression to squeeze packets down in size and thus make them transmit more quickly. Tests run on a dial-up VPN show that this compression does indeed make a difference. Using the TTCP test program over a 56 Kbps dial-up produced the data shown in Table 2:

Table 2 Remote Access VPN Throughput

56K modem, actual connect speed 49.2 kbps, with 3des encryption, sha1 authentication	Throughput
Server compression on	78.4 Kbps
Server compression off	18.9 Kbps

The 78 Kbps throughput was the average of 11 test runs, with a high value of 105 Kbps and a low of 46 Kbps. The 18.9 Kbps throughput was the average of six test runs, with a high value of 24.4 Kbps, and a low value of 14 Kbps.

The TTCP test determines the maximum throughput a communications system is capable of achieving. Actual throughput will be less depending on the situation. For example, TTCP uses text as packet payload, making these packets highly compressible. In actual use, such as browsing a web site, some of the data such as graphics is already compressed, and will not be appreciably accelerated by additional compression. In addition, the TTCP test does not take in account disk transfer speed, which can have a negative impact on data transfer when the data must be written to or read from disk during the transfer. Therefore, the standard disclaimer applies here: “your mileage may vary.” However, it is clear that compression does make a big difference and is in fact necessary for getting any useful throughput out of a 56K VPN.

Broadband VPN throughput is very much higher than a dial-up VPN, but limited to the maximum throughput of the broadband connection. For example, DSL connection speed varies considerably depending on the type of DSL, ranging from hundreds of kilobits for low-cost service, to megabits for higher-cost service. A typical consumer DSL service might provide a maximum of 640 Kbps download speed, and 256 Kbps upload speed. A typical cable modem service may provide a maximum of 1.5 Mbps download speed, and 125 Kbps upload speed.

In order to generate a baseline throughput for broadband access, a computer was placed on the LAN on the outside of the VPN server, and tested through the VPN server to a

computer on the inside of the VPN server. The test results represent the maximum possible throughput available for a broadband connection for the particular setup. Higher speed computers would most likely achieve higher throughput.

The test results show that a typical laptop computer running the VPN client software can achieve 2 Mbps throughput, which will fill most DSL or cable modem pipes. This is shown in Table 3 for compression on and compression off. Although compression doubles the throughput, Cisco discourages use of compression with broadband access due to the high burden on the server CPU. For a lightly loaded server, this may be acceptable.

Table 3 Broadband Baseline Throughput

LAN connection to VPN server, IBM T20 700 MHz laptop to Toshiba Portege 333 MHz laptop	Throughput
Server compression on	2.03 Mbps
Server compression off	0.98 Mbps

VPN Client Usage Policy

Besides the obvious speed difference, the feature that distinguishes dial-up and broadband access types with respect to a VPN is the security vulnerability inherent in high-speed connections due to their persistent nature. “Always on” DSL or cable modem connections are opportunities for hackers compared to dial-up connections where the user machine’s IP address changes with each connection and the connection speed is relatively slow. Therefore, there are two different policies proposed for handling these connection types. These policies are proposals only, subject to review and approval. They are being adopted as recommendations for the trial:

High-speed access (persistent)

A client machine connecting to SNL over a high-speed link such as DSL or cable modem must operate a firewall product primarily for protecting the computer while not operating a VPN tunnel. The computer used for SNL access may be personally owned.

A firewall for this scenario can be either in software or hardware form. A software firewall, commonly called a personal firewall, runs directly on the computer that is initiating VPN connections. The firewall can block all connection attempts to the computer to prevent hacking. This is particularly important when the computer is not running the VPN. Hardware firewalls are often used for home networks that connect to a broadband Internet service. The hardware firewall will protect all computers on the home

network, rather than just the VPN machine. Some products are now available that provide a home or small business network hub or switch in combination with a simple firewall for connecting to a DSL or cable modem service.

An increasingly common approach for connecting multiple home machines to broadband services automatically provides a firewall function. The technique utilizes a small, relatively inexpensive router that sits between the DSL or cable modem and the user's machines. This router may contain an Ethernet hub or switch, or may have a hub or switch connected to it to provide a LAN in the home. The router's outside interface, which connects to the DSL or cable modem, is assigned the IP address given to the user by the service provider. Typically, in this configuration, network address translation, or NAT[2], is then used to provide IP addresses for the home machines. An example configuration is shown in Figure 12.

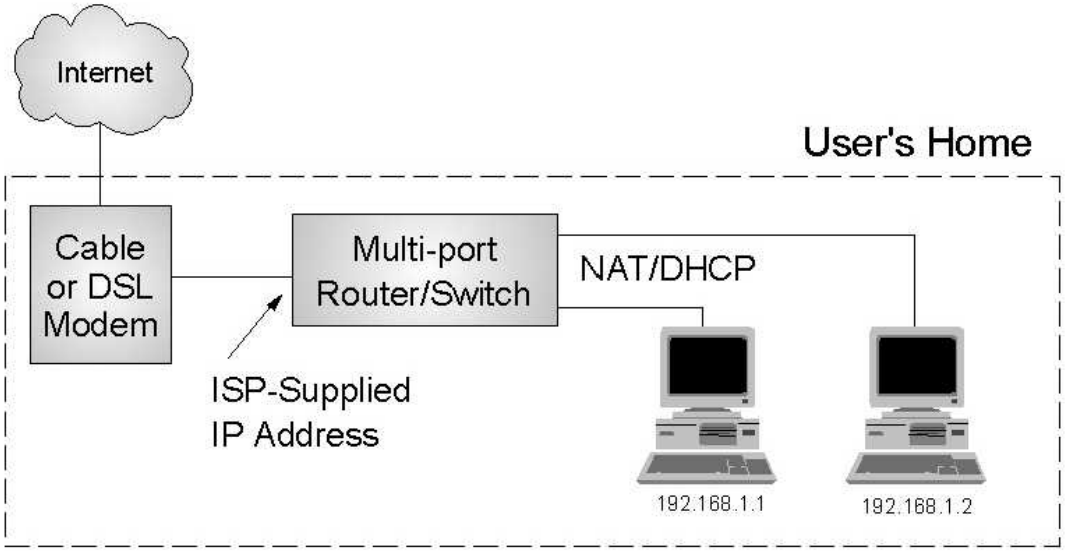

Figure 12. An Example of a Home LAN Connected to the Internet Through a Cable or DSL Modem

In this example, the home machines are assigned private network addresses 192.168.1.1 and 192.168.1.2 from the router using the DHCP protocol (Dynamic Host Configuration Protocol, a technique for dynamically assigning IP addresses). The two machines can communicate directly with each other over the LAN using these private addresses. When either machine wishes to communicate with the Internet, the router translates the source machine's private IP address to the router's public outside address and preserves or translates any port numbers involved in the particular protocol in use in order to create a unique mapping between the source machine address and the router address. The packets

[2] Although commonly called NAT, this particular application is more appropriately called PAT, or Port-Address Translation. NAT typically translates one address to another, while PAT translates multiple addresses to a single address, using protocol port numbers to create unique mappings.

are then sent to the destination with the appearance of originating from the router and not the user's home machine. The return packets arrive at the router, which must translate them back to the user machine's private address, and then forwards them on to the appropriate machine.

Although this technique is primarily designed to share a single Internet IP address among multiple user machines, the result is to make the home machines invisible to the Internet. Under normal circumstances, no machine from the Internet can access the home machines, thus providing a firewall function without an explicit firewall.

A VPN can be safely operated from one or more of the home machines in this example. However, a complication can prevent the VPN from successfully operating. The NAT function must operate on port numbers that are typically contained in an IP packet, but an IPSec VPN does not utilize port numbers and must therefore use a different technique. In this case, a VPN client can wrap an IPSec packet in a UDP packet, providing port numbers that can be manipulated as needed by the NAT device. Unfortunately, this does reduce performance somewhat due to additional overhead.

Low-speed access (non-persistent)

A client machine connecting to SNL over dial-up Internet access may operate a personal firewall product. The computer used for SNL access may be personally owned.

Although a dial-up connection can be hacked, it is of less interest to hackers than persistent connections due to the ever-changing IP address of the dial-up machine, and the slow-speed connection. Therefore, use of a firewall is considered optional.

Note, however, that a low-speed connection can utilize NAT in a similar manner as described above for a high-speed connection with the same firewall-like result. In this case, one machine dials the Internet using a normal analog modem and acts as the NAT device. Other locally connected machines then access the Internet through the NAT machine. Although performance is considerably reduced compared to a high-speed connection, protection of machines behind the NAT device is accomplished.

Next Steps

The initial testing of the client-to-site VPN has been successful. Testers have been happy with the technology, especially broadband access. Consensus among project members was reached to move forward with implementation of a remote access VPN.

Getting to full production mode will require several steps. The next step will be to purchase a pair of VPN 3030 servers and put one into production. Testing can continue with the original set of users, plus a few additional users. The network operations group

will be trained in the operation of the VPN server and client. CSU members and the Corporate Computing Help Desk, CCHD, will be trained in the operation and support of the VPN client. When all training is complete, the general user population will be invited to utilize this new service.

Results

Over the course of two years of research and experimentation with VPN technology, both site-to-site and client-to-site VPN implementations have been successfully established.

Requirements were developed to aid in selecting appropriate equipment. These requirements were developed over a two-year period of research, experimentation, and incorporation of existing computer security requirements. For site-to-site VPNs, two vendors were selected and eventually narrowed to one. Network Alchemy's CryptoCluster 500 and 2500 were chosen. For the client-to-site VPN, the awareness of FIPS 140-1 certification resulted in no equipment matching the requirements. However, the Cisco VPN 3000 Series was identified as a potential solution as it is undergoing FIPS 140-1 certification for both the hardware server and client software.

It has been demonstrated that a remote site can be successfully connected to the corporate headquarters by operating a VPN over the Internet, eliminating the need for a costly leased line. Sandia's KTF was limited to poor dial-up access to the corporate intranet due to the prohibitively high cost of a leased line from Albuquerque to Kauai. A VPN not only enabled KTF to communicate reliably with the corporate intranet, but to communicate at considerably higher speed. Dial-up access is physically limited to 53 Kbps and routinely limited to much less. VPN access improved this by at least a factor of 10, enabling KTF personnel to efficiently complete corporate tasks such as time card and expense report submissions as well as access department resources in a timely fashion.

It has also been demonstrated that remote users can be connected to the corporate site by operating a client VPN over the Internet by either dial-up or broadband access.

The SC99 and SC2000 supercomputer shows both demonstrated the usefulness of VPN technology for quickly and easily establishing secure connectivity from a remote site to the corporate network over a public network.

Conclusion

VPN technology has rapidly evolved, is reaching a level of maturity, and is becoming a standard tool in the network engineer's bag of tricks.

The IPSec protocol has the ability to become a ubiquitous component of networking, enabling secure communications that go beyond today's site-to-site or client-to-site implementations.

Interoperability of IPSec devices has been slow to become a reality, but is gaining ground. The IPSec protocol was designed with interoperability at its heart, but early implementations did not necessarily interpret the specifications uniformly. Additionally,

extensions to the protocol, such as user authentication, made for compatibility problems. Over time, the protocol has matured, albeit in a relatively short time.

IPSec only protects data in transit. A complete VPN solution requires diligence in maintaining secure systems to protect data at the source and destination.

Site-to-site VPNs are relatively straightforward to implement and can achieve significant cost savings over leased lines, while taking into account the lack of guaranteed service from Internet use.

Client-to-Site VPNs are more complex to develop due to user authentication and the need for client software. Client VPNs may utilize a personal firewall for dial-up access, and must use a personal firewall for broadband VPN access due to the increased risk from hackers attempting to exploit machines on persistent connections. Broadband access is of particular interest due to its significantly higher speed versus dial-up and its growing availability for both local and remote connections.

VPN management is relatively straightforward. Once a VPN is established, it tends to operate autonomously, requiring little manual intervention or oversight. However, monitoring is essential to ensure the system is operating properly. Logs should be routinely reviewed for system errors or invalid connection attempts. Additionally, the SNMP protocol can be used to automatically monitor for system problems.

References

Dave Kosiur, 1998, *Building and Managing Virtual Private Networks*. John Wiley & Sons, Inc., New York.

Casey Wilson, Peter Doak, 2000, *Creating and Implementing Virtual Private Networks*. The Coriolis Group, Scottsdale, Arizona.

Naganand Doraswamy, Dan Harkins, 1999, *IPSec: the new security standard for the Internet, intranets, and virtual private networks*. Prentice Hall Inc., New Jersey.

Niels Ferguson, Bruce Schneier, 1999, *A Cryptographic Evaluation of IPSec*, http://www.counterpane.com.

Marc M. Miller, Jack A. Hudson, John P. Long, 1999, *Location Independent Professional Project: A Pilot Study*, SAND99-0100. Sandia National Laboratories, Albuquerque, NM, February 1999.

Thomas J. Pratt, Thomas D. Tarman, Luis G. Martinez, Marc M. Miller, Advanced Networking Integration; Roger L. Adams, Telecommunications Operations; Helen Y. Chen, James M. Brandt, Peter S. Wyckoff, Security & Networking Research, *The ASCI Network for SC '99 : A Step on the Path to a 100 Gigabit Per Second Supercomputing Network*, SAND2000-1812. Sandia National Laboratories, Albuquerque, NM, July 2000.

FIPS 140-1, *Security Requirements for Cryptographic Modules*, 1995, National Institute for Standards and Technology, http://csrc.nist.gov/cryptval/140-1/1401val.htm.

Appendix A - Proposal for Providing Enterprise VPN Site-to-Site Service for SNL-NM

Introduction

Concurrence is sought from the Computer Security and Passwords and Computer Security Technology departments for deploying VPN technology as described in the following proposal.

VPN technology is proposed for SNL-NM for providing enterprise-wide site-to-site access to the SRN from remote locations (this proposal excludes remote client access (dial-up VPN) – this will be proposed at a later date). The goal is to provide lower cost access with greater privacy. Lower cost comes from utilizing the Internet or other networks that are either free or the cost is covered elsewhere. Greater privacy comes from Triple-DES encryption, and authentication that guarantees the integrity of the data.

Data will be protected using the IPSec protocol with ESP (Encapsulating Security Payload) for providing data encryption and authentication.

Several remote sites have been identified as potential VPN implementations. The following two are highest priority based on customer requirements:

- A VPN is proposed for linking Kansas City to the SNL-NM SRN over the Internet to initially supplement and eventually replace an existing dedicated low-speed link.

- A VPN is proposed for linking Sandia's Kauai Test Facility (KTF) to the SNL-NM SRN over the DREN (Defense Research and Engineering Network) network. KTF currently has SON access, but has no SRN access.

Site to site VPN service will be constructed with dedicated hardware devices from Network Alchemy, Inc. These devices have 10/100 Ethernet interfaces capable of Triple-DES throughputs at 4.5 Mbps for the CryptoCluster 500 and 45 Mbps for the CryptoCluster 2500.

Network Alchemy VPN devices are capable of "clustering", which enables multiple devices to be logically combined as a single device for improving reliability and performance. One node of the cluster acts as the master and has the responsibility for distributing the workload across all the other members of the cluster. Clusters can be expanded or shrunk while in service without interruption to current sessions because every node is fully aware of all of the session states and security associations being handled by every other cluster node.

All devices, whether stand-alone or clustered, can be managed from a single Java-based application called CryptoConsole. This application will be run on a stand-alone PC with limited physical access. Management accounts will be granted to a limited set of people, such as the administrator and a backup.

Figure A-1 shows the configuration of VPN devices linking remote sites to the SRN. The VPN server sits parallel to the SRN firewall in order to separate VPN traffic from non-VPN traffic.

Figure A-2 shows the physical connection of the VPN device to the SRN and SON network devices.

The VPN team will be composed of Marc Miller (4616), George Yonek (4614), and Pat Manke (4614). Fran Current (4812) may be utilized to facilitate interaction with Kansas City.

Project Acceleration funds this project.

Network Diagrams

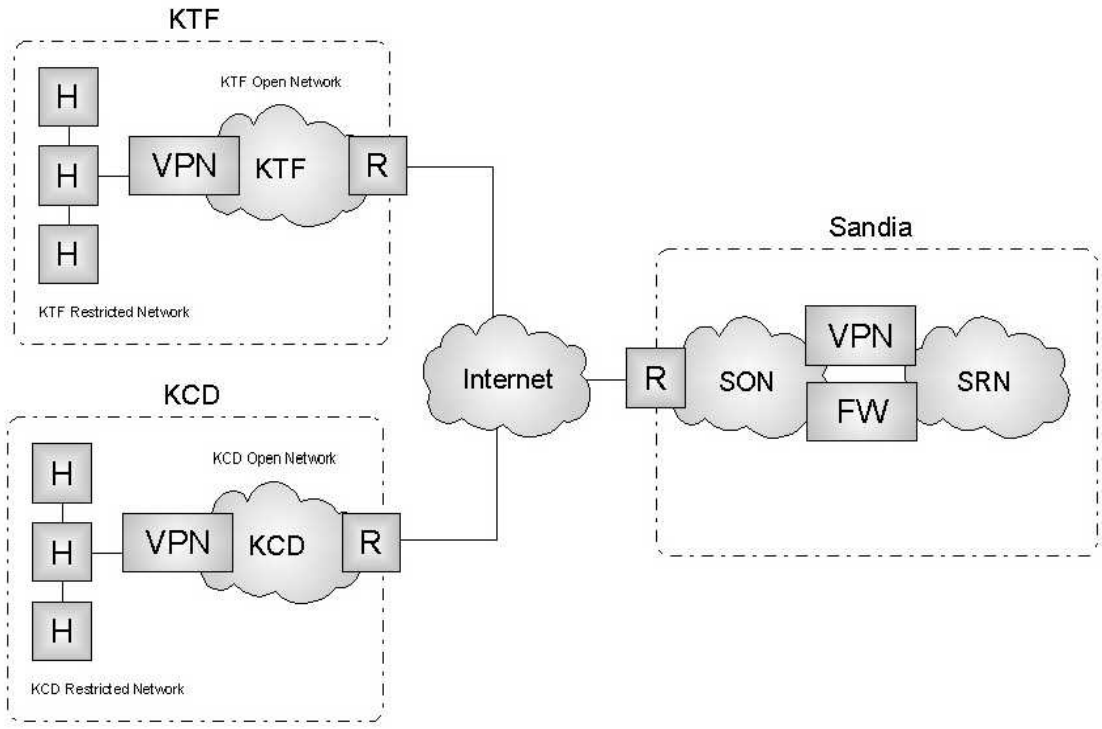

Figure A - 1 VPN Devices Linking Remote Sites to the SRN

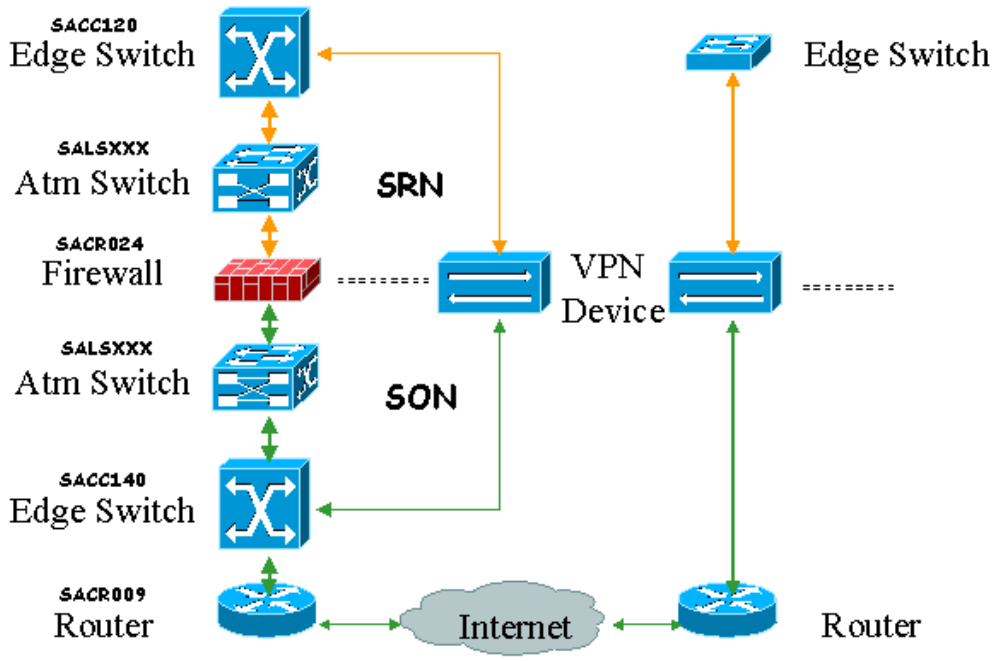

Figure A - 2 Physical Connection of the VPN Device

Definitions

- **CA**: Certificate Authority.
- **CryptoCluster**: any of the family of VPN devices manufactured by Network Alchemy (e.g. CryptoCluster 500 (4.5 Mbps Triple-DES), CryptoCluster 2500 (45 Mbps Triple-DES)).
- **CryptoConsole**: a Java-based stand-alone application that manages a set of CryptoCluster devices. It runs on Windows 95/98/NT.
- **Gateway**: another name for a VPN server.
- **HMAC SHA-1**: Hash Message Authentication Code Secure Hash Algorithm 1 (SHA-1 was created by the National Institute of Standards and Technology (NIST) along with the National Security Agency (NSA). It produces a 160-bit hash value.
- **HMAC MD5**: MD5 was created by Ron Rivest as an improvement to MD4. It produces a 128-bit hash value.
- **SA**: Security Association, all objects needed to securely communicate with another entity.

Details

The IPSec protocol will be used for communication between VPN devices.

IPSec policies:

- Encapsulating Security Payload (ESP) protocol for privacy and integrity.
- Privacy provided by Triple-DES encryption.
- Integrity and replay prevention provided by HMAC SHA-1 message authentication code.
- Perfect Forward Secrecy (PFS) for maximum protection of session keys.
- Internet Key Exchange protocol (IKE) for key management.

All VPN traffic will operate in tunnel mode as opposed to transport mode in which the original IP header remains intact. In tunnel mode, the entire IP packet including address headers is compressed, encrypted, and authenticated. This secure packet is then placed into the payload section of a new IP packet with a new, unencrypted address header for originating and destination VPN devices. Thus, originating and destination IP addresses will be hidden, with only the VPN device public addresses exposed.

At the SRN, the CryptoCluster will be established in parallel to the firewall and operate in blocking mode: all non-member traffic is blocked. Therefore, only IPSec packets, with proper authentication, will be processed by the CryptoCluster.

PPTP and L2TP subsystems will be disabled. No PPTP or L2TP sessions will be serviced by the VPN.

Key Management - IKE

The Internet Key Exchange (IKE) keying method will be used for agreeing on authentication methods, encryption methods, keys to use, how long to use the keys before changing them, and for secure key exchange. The authentication mechanism will be digital certificates. Certificates will be generated by Sandia's CA. Triple-DES will be used for encryption.

IKE Phase 1 and 2 negotiation options:

- PFS enabled to protect session keys (for phase 1, PFS will hide the endpoints for which the IKE exchange; for phase 2, PFS will ensure that a Diffie-Hellman key exchange is done each time a new SA is established, rather than relying on the keying material obtained from the keying material negotiated in phase 1).

- Default Diffie-Helman Group #2 (MODP 1024-bit) will be utilized.
- New IKE Security Associations will be generated when time elapsed reaches 8 hours.

Traffic Filters (Tunnels)

Network Alchemy refers to tunnels as traffic filters. These filters act as access control lists or ACLs, defining how to process a tunnel. A filter can pass data in the clear, drop data, or perform IPSec functions on the data.

Tunnels (IPSec SAs) will be re-keyed every hour (default) with the capability of re-keying at some interval of data (gigabytes or megabytes).

Tunnels will be authenticated using digital certificates generated by the Sandia CA.

For efficiency, traffic can share an SA and tunnel rather than creating separate tunnels based on local host or port, remote host or port, or IP protocol. The default is to have all traffic between two gateways share a single tunnel.

Access Control

Each VPN tunnel will grant specific remote subnets or hosts/servers access to specific SRN subnets or hosts/servers based on IP address. If necessary or appropriate, additional control can be placed on services and/or ports.

For example, a remote subnet with address X.Y.Z.0 could be granted access to SRN subnet 134.253.4.0 by building a tunnel with these subnets as the endpoints. All hosts in X.Y.Z.0 could then tunnel to any host on 134.253.4.0. Any other hosts on any other subnets at the remote site would not have access to the SRN since they are not associated with any VPN to the SRN.

Alternatively, a tunnel could be built for remote hosts X.Y.Z.1, X.Y.Z.2, and X.Y.Z.3 only to an SRN subnet such as 134.235.4.0, the entire SRN (134.253.0.0), or specific SRN hosts.

System Management

The CryptoCluster VPN server's kernel software is cryptographically signed by Network Alchemy. This signature ensures that the image was built at Network Alchemy and has not been altered in any way. If the CryptoCluster VPN server cannot verify the signature on the kernel, it will not load the kernel into RAM and run it.

To ensure that an unauthorized node cannot be booted into the CryptoCluster VPN server, the CryptoConsole automatically generates a Security Token whenever nodes are added to a configuration. The Security Token is entered into the CryptoCluster VPN server during bootstrap configuration, and it is used as a pre-shared session key to authenticate an anonymous SSL communication between the CryptoConsole and the CryptoCluster VPN server. The first 32 bits of the token are a hash of the bootstrap configuration information (this includes the inside and outside IP addresses and netmasks as well as the default route). The last 32 bits are a randomly generated value used to help ensure security. Therefore, an unauthorized node would need to utilize a security token generated from the same CryptoConsole as all other authorized nodes, and this would not be possible without authorized access to the CryptoConsole. In addition, for initial configuration, the security token expires within a specified period of time (1 to 30 days, with 7 days as the default).

All intra-cluster traffic is authenticated using keyed MD5, and all keying material that is passed between clusters is encrypted using Triple-DES.

In-band system management will be implemented with a stand-alone Java application running on a Windows platform (the CryptoConsole). The application is password protected. The application communicates with the VPN servers using SSL with Triple-DES encryption. During the installation, each gateway, either standalone or clustered, is configured with a special SSL-only Certification Authority. As part of this installation process, digital certificates are generated for the CryptoConsole client and for servers that run on the CryptoCluster VPN server.

Console access is protected by username and password. This same username/password combination is used with the CryptoConsole management application. The default account is ‘Administrator’. This account cannot be deleted. Other accounts can be created and deleted. These secondary accounts can have either full administrator privileges, or only the ability to monitor server activity. In the event that all usernames and/or passwords are forgotten, it is then necessary to re-install the entire VPN deployment, which requires physical access to the hardware to initiate the installation.

The CryptoConsole will run on a Windows NT system located on the SRN, with controlled physical access (computer annex) and limited accounts.

Out-of-band management will be implemented on a local basis only, utilizing a terminal or PC with terminal-emulation software connecting through a serial port to the console. No dial-up capability will be provided into the console port. This console is normally utilized for diagnostics only; configuration is controlled through the CryptoConsole.

No Telnet capability will be allowed into the management system. This capability will be turned off at each VPN server.

SNL will maintain control of remote VPN devices installed by SNL with regard to ownership and in-band management. However, remote device console expertise will be

necessary for cases where a device is inoperable or improperly configured, and in-band management is not possible.

The version of each node's configuration is displayed in the management station for configuration control. Any changes cause the version number to increment.

Sniffing Results

Packets passed through a CryptoCluster tunnel were "sniffed" using a Network Associates sniffer with the intent of showing that these packets are converted to IPSec.

Packets in the Clear

The following two packets are the result of a single ping (ICMP echo and ICMP echo reply) from a workstation on one end of a VPN testbed (192.168.125.129) to a workstation at the other end (134.253.4.44).

```
Sniffer Network Analyzer data from 2-Sep-99 at 09:49:54, file
C:\ENCAP\ping.ENC, Page 1
- - - - - - - - - - - - - - - - Frame 9 - - - - - - - - - - - - - - - - -
SUMMARY  Delta T     Destination   Source        Summary
     9            [134.253.4.44]  [192.168.125...  ICMP Echo

ICMP:  ----- ICMP header -----
ICMP:
ICMP:  Type = 8 (Echo)
ICMP:  Code = 0
ICMP:  Checksum  = E85B (correct)
ICMP:  Identifier = 256
ICMP:  Sequence number = 25600
ICMP:  [32 bytes of data]
ICMP:
ICMP:  [Normal end of "ICMP header".]
ICMP:

ADDR  HEX                                                 ASCII
0000  00 50 5A 01 A0 1A 00 C0  4F 8E D4 64 08 00 45 00  .PZ.....O..d..E.
0010  00 3C 8C 43 00 00 20 01  45 2B C0 A8 7D 81 86 FD  .<.C.. .E+..}...
0020  04 2C 08 00 E8 5B 01 00  64 00 61 62 63 64 65 66  .,...[..d.abcdef
0030  67 68 69 6A 6B 6C 6D 6E  6F 70 71 72 73 74 75 76  ghijklmnopqrstuv
0040  77 61 62 63 64 65 66 67  68 69                   wabcdefghi


- - - - - - - - - - - - - - - - Frame 10 - - - - - - - - - - - - - - - - -

SUMMARY  Delta T     Destination   Source        Summary
    10    0.0053  [192.168.125... [134.253.4.44]   ICMP Echo reply

ICMP:  ----- ICMP header -----
ICMP:
ICMP:  Type = 0 (Echo reply)
ICMP:  Code = 0
ICMP:  Checksum  = F05B (correct)
```

```
ICMP:  Identifier = 256
ICMP:  Sequence number = 25600
ICMP:  [32 bytes of data]
ICMP:
ICMP:  [Normal end of "ICMP header".]
ICMP:

ADDR  HEX                                                ASCII
0000  00 C0 4F 8E D4 64 00 50  5A 01 A0 1A 08 00 45 00  ..O..d.PZ.....E.
0010  00 3C F1 E4 00 00 7C 01  83 89 86 FD 04 2C C0 A8  .<....|......,..
0020  7D 81 00 00 F0 5B 01 00  64 00 61 62 63 64 65 66  }....[..d.abcdef
0030  67 68 69 6A 6B 6C 6D 6E  6F 70 71 72 73 74 75 76  ghijklmnopqrstuv
0040  77 61 62 63 64 65 66 67  68 69                   wabcdefghi

Sniffer Network Analyzer data from 2-Sep-99 at 09:49:54, file
C:\ENCAP\ping.ENC, Page 2
```

Packets Encrypted

The following two packets are the same ping packets shown above, but after conversion to IPSec. The ICMP echo packets have been converted to IP packets by encapsulating the original ping packet into an IP packet after encrypting the original packet with Triple-DES. The packets are now addressed to the VPN nodes (192.168.127.166 and 192.168.126.2).

```
Sniffer Network Analyzer data from 2-Sep-99 at 09:54:42, file
C:\ENCAP\na_ping.ENC, Page 1


- - - - - - - - - - - - - - - - - Frame 24 - - - - - - - - - - - - - - - - -

SUMMARY  Delta T     Destination   Source        Summary
    24            [192.168.127... [192.168.126.2]  IP  D=[192.168.127.166]
S=[192.168.126.2] LEN=92 ID=44545

IP:   ----- IP Header -----
IP:
IP:   Version = 4, header length = 20 bytes
IP:   Type of service = 00
IP:         000. .... = routine
IP:         ...0 .... = normal delay
IP:         .... 0... = normal throughput
IP:         .... .0.. = normal reliability
IP:   Total length    = 112 bytes
IP:   Identification  = 44545
IP:   Flags           = 0X
IP:         .0.. .... = may fragment
IP:         ..0. .... = last fragment
IP:   Fragment offset = 0 bytes
IP:   Time to live    = 63 seconds/hops
IP:   Protocol        = 50 (SIPP-ESP)
IP:   Header checksum = 4E61 (correct)
IP:   Source address      = [192.168.126.2]
IP:   Destination address = [192.168.127.166]
IP:   No options
IP:   [92 byte(s) of data]
```

Note that the sniffer correctly identified the packet as ESP.

```
ADDR  HEX                                                ASCII
0000  00 50 5A 01 A0 5B 00 60  1D 05 48 85 08 00 45 00  .PZ..[.`..H...E.
0010  00 70 AE 01 00 00 3F 32  4E 61 C0 A8 7E 02 C0 A8  .p....?2Na..~...
0020  7F A6 BF D6 04 00 00 00  00 B3 83 1D B2 5B 7D B2  .............[}.
0030  AF 0A 29 3B 9D EE 7B D1  F8 0C A0 84 B6 67 DD BD  ..);..{......g..
0040  20 FC D8 A3 25 08 9D 61  4D BF 80 4B 28 B1 88 3E   ...%..aM..K(..>
0050  48 47 0F E1 97 D7 C3 1C  66 E5 67 62 03 E9 09 25  HG......f.gb...%
0060  78 AB AE C2 F8 40 08 FA  77 F5 3F 8A 2C 0A 8E 73  x....@..w.?.,..s
0070  5C E7 DF E2 34 19 02 47  90 09 43 C6 D9 01        \...4..G..C...



- - - - - - - - - - - - - - - - - Frame 25 - - - - - - - - - - - - - - - - -

SUMMARY  Delta T     Destination   Source        Summary
    25    0.0064  [192.168.126.2] [192.168.127...  IP  D=[192.168.126.2]
S=[192.168.127.166] LEN=92 ID=29873

IP:   ----- IP Header -----
IP:
IP:   Version = 4, header length = 20 bytes
IP:   Type of service = 00
IP:         000. .... = routine
IP:         ...0 .... = normal delay
IP:         .... 0... = normal throughput
IP:         .... .0.. = normal reliability
IP:   Total length    = 112 bytes
IP:   Identification  = 29873
IP:   Flags           = 0X
IP:         .0.. .... = may fragment
IP:         ..0. .... = last fragment
IP:   Fragment offset = 0 bytes
IP:   Time to live    = 64 seconds/hops
IP:   Protocol        = 50 (SIPP-ESP)
IP:   Header checksum = 86B1 (correct)
IP:   Source address      = [192.168.127.166]
IP:   Destination address = [192.168.126.2]
IP:   No options
IP:   [92 byte(s) of data]

ADDR  HEX                                                ASCII
0000  00 60 1D 05 48 85 00 50  5A 01 A0 5B 08 00 45 00  .`..H..PZ..[..E.
0010  00 70 74 B1 00 00 40 32  86 B1 C0 A8 7F A6 C0 A8  .pt...@2........
0020  7E 02 15 A3 08 00 00 00  00 1E C1 8B 14 38 AA 48  ~.............8.H
0030  E3 4C 59 70 82 A3 34 F0  36 40 24 62 A9 70 76 1A  .LYp..4.6@$b.pv.
0040  77 E2 E0 50 04 B0 69 66  1B 43 77 C5 99 E8 9F 6B  w..P..if.Cw....k
0050  60 38 F2 39 08 CC A1 91  6F 61 76 D7 2D AA B6 CA  `8.9....oav.-...
0060  BF E6 32 82 D8 58 36 00  0F 69 1E 27 29 AF EE C4  ..2..X6..i.')...
0070  AA 56 C4 7D 97 89 00 0A  09 16 EF F6 4D 14        .V.}........M.
```

Note that the sniffer correctly identified the packet as ESP.

Implementation Steps

Install a CryptoCluster (two CryptoCluster 2500 nodes) at the edge of the SRN. No communication will be possible until the communications configuration is completed.

Generate configurations for the CryptoCluster nodes in CryptoConsole (security token, IP addresses, node names, and routes).

Connect a terminal to each CryptoCluster node and manually enter the configuration data generated by CryptoConsole (security token, IP addresses, node names, and routes).

Complete the configuration process by downloading remaining configuration information from the CryptoConsole to each of the CryptoCluster nodes.

Once the communications configuration process is completed, the CryptoCluster will respond to pings and communicate with the CryptoConsole, but will not pass any VPN traffic as no filters have been established.

Establish a CryptoCluster node in the SON for testing connectivity to the SRN. The "inside" of the CryptoCluster node would be on an isolated subnet to prevent any unexpected access.

Once all services are verified (routing, DNS, WINS, etc.) for the SRN/SON test, a CryptoCluster node will be connected at a remote site to be determined (such as Kansas City or Kauai). Filters will be built to establish the appropriate tunnels.

After the initial site is established, it will be monitored closely for security, reliability, and performance. If any significant problems occur, the link may be dropped or restored to a previous configuration (pre-VPN) until the problem is resolved. If no significant problems occur, additional VPNs will be established.

Summary

How is user traffic protected?

All user traffic is protected by running the IPSec protocol in tunnel mode, providing Triple-DES encryption and SHA-1 authentication/replay prevention. Tunnel mode protects the originating and destination IP addresses by compressing, encrypting, and authenticating the original IP packet and placing this into the payload of a new IP packet whose originating and destination IP addresses are those of the VPN devices.

How is CryptoCluster software authenticated?

Network Alchemy cryptographically signs the CryptoCluster kernel software. If the CryptoCluster VPN server cannot verify the signature on the kernel, it will not load.

How is an unauthorized node prevented from joining a cluster?

In order for a node to join a cluster, the node configuration *must* first be created in the CryptoConsole, using the cluster PIN. This configuration is then entered manually into

the node joining the cluster. The remaining configuration information is then downloaded to the node from the CryptoConsole. Therefore, a node that has been configured from a different CryptoConsole will be rejected when attempting to join a cluster.

How is an unauthorized CryptoConsole prevented from managing CryptoCluster devices?

CryptoCluster nodes must initially be configured from a CryptoConsole. The node configuration information (inside and outside IP addresses, netmasks, and default route) is hashed in 32 bits and combined with 32 random bits to generate a security token that is used to authenticate a node to the CryptoConsole during initial configuration. An unauthorized CryptoConsole would attempt to communicate with a CryptoCluster using the wrong security token and PIN. The CryptoCluster device would therefore ignore the CryptoConsole.

How is an unauthorized VPN device prevented from joining a VPN?

The VPN must be configured in the CryptoConsole. Therefore, without access to the CryptoConsole, no additional VPN can be established. Also, digital certificates will authenticate any devices that are authorized to join a VPN.

How is unauthorized access to the CryptoConsole prevented?

The CryptoConsole is password protected. The password is case-sensitive and must be between 8 and 32 characters in length. The password may contain any combination of the characters a-z, A-Z, 0-9, "-", or "_" (except the username itself). Of course, the CryptoConsole will be physically protected from unauthorized access.

How is traffic protected between the CryptoConsole and CryptoClusters?

Management traffic is protected with SSL, using Triple-DES encryption and special SSL-only digital certificates.

How is Intra-Cluster traffic protected?

Intra-cluster traffic is encrypted using Triple-DES, and authenticated with MD5.

How often are SAs re-negotiated during their lifetimes?

IKE SAs will be re-negotiated every eight hours. IPSEC SAs will be re-negotiated every hour.

Appendix B - Proposal for Providing Enterprise VPN Client-to-Site Service for SNL-NM

Marc M. Miller
Advanced Networking Integration

March 14, 2000

Acronyms

DSL – Digital Subscriber Line
DES – Data Encryption Standard
ESP – Encapsulating Security Protocol
FIPS 140-1 – Federal Information Processing Standards, *Security Requirements for Cryptographic Modules*.
GUI – Graphical User Interface
SHA-1 – Secure Hash Algorithm
HTTPS – Hyptertext Transfer Protocol with SSL
ICMP – Internet Control Message Protocol
IKE – Internet Key Exchange
IP – Internet Protocol
IPSec – Internet Protocol Security
L2TP – Layer 2 Tunneling Protocol
PFS – Perfect Forward Secrecy
PPTP – Point to Point Tunneling Protocol
RSA ACE – RSA Security, Inc. SecurID authentication server
SA – Security Association
SRN – Sandia Restricted Network
SSL – Secure Socket Layer
VPN – Virtual Private Network

Proposal for Providing Enterprise VPN Client-to-Site Service for SNL-NM

Proposed by Telecommunications Operations
and Advanced Networking Integration Departments

Introduction

Concurrence is sought from the Computer Security and Computer Security Technology departments for deploying VPN remote access technology as described in the following proposal.

VPN technology is proposed for SNL-NM for providing enterprise-wide client-to-site access to the SRN from remote locations. The goal is to provide lower cost access, greater privacy, and in some cases, improved performance. Lower cost comes from utilizing the Internet rather than expensive 1-800 access. Greater privacy comes from Triple-DES encryption, and authentication that guarantees the integrity of the data. Improved performance can come from utilizing a VPN over a DSL or cable modem connection where available.

Data will be protected using the IPSec protocol with ESP (Encapsulating Security Payload) for providing data encryption and authentication.

Client to site VPN service will be constructed with dedicated hardware devices from Cisco Systems, Inc. The VPN 3000 Series concentrator, model 3030, is capable of 1500 simultaneous connections.

The 3030 is an integrated remote access VPN solution. The client software is tightly coupled to the hardware server. Client configurations are set in the server and uploaded to the client at connect time. The user has little control over the client configuration. The only significant parameters that can be modified are the group name and password and the destination server. The only ramification of changing these parameters is an inability to connect. There are no options for modifying IPSec parameters. When a user initiates a connection and successfully authenticates, the tunnel is transparently created and the user continues on with normal communications actions.

Users will be authenticated using SecurID through a Radius server that will provide additional logging capability. No VPN connection can be made into the 3030 without successfully passing SecurID authentication. The SecurID system will be the same used for traditional dial-up. No special requirements are needed to use SecurID with the VPN. Users that already hold a SecurID card can used the same card with the VPN.

Split tunneling will not be enabled. Split tunneling is a feature that, when enabled, allows a client machine to access the Internet while simultaneously operating a VPN tunnel. Disabling split tunneling eliminates any Internet-spawned hacking possibilities while a

tunnel is operating. The client machine will not respond to any communications outside the tunnel.

The 3030 is managed from a Java-based browser GUI, Telnet, or direct connect console. Access is protected by username and password. Browser access will be secured using HTTPS. Telnet access can be secured by using a special Telnet/SSL application. All management access will be initiated from within the SRN only. This will be a trial period for analyzing the viability and manageability of the technology.

Encryption of SRN data falls under the requirement for FIPS 140-1 certification. The 3030 is currently undergoing FIPS 140-1 certification, for both the client software and the server hardware. During the period of time the 3030 is operational prior to certification, no UCNI data will be transmitted over a tunnel.

There are two general types of VPN remote access: slow speed (dial-up) and high speed or broadband (DSL and cable modem). Besides the obvious speed difference, the feature that distinguishes these access types with respect to a VPN is the security vulnerability inherent in high-speed connections due to their persistent nature. “Always on” DSL or cable modem connections are easy prey for hackers. Therefore, there are two different policies proposed for handling these connection types:

High-speed access (persistent)

A client machine connecting to SNL over a high-speed link such as DSL or cable modem must operate a personal firewall product primarily for protecting the computer while not operating a VPN tunnel. The computer used for SNL access must be Sandia-supplied and utilized strictly for business purposes.

Low-speed access (non-persistent)

A client machine connecting to SNL over dial-up Internet access may operate a personal firewall product. The computer used for SNL access may be personally owned.

Figure B-1 shows the configuration of a VPN client linking to the SRN. The VPN server sits parallel to the SRN firewall in order to separate VPN traffic from non-VPN traffic.

The VPN team will be composed of Marc Miller and Steve Gossage (9336) and Pat Manke (9334) and others to be determined (Computer Security, CSU Tech Dev, etc.)

Network Diagram

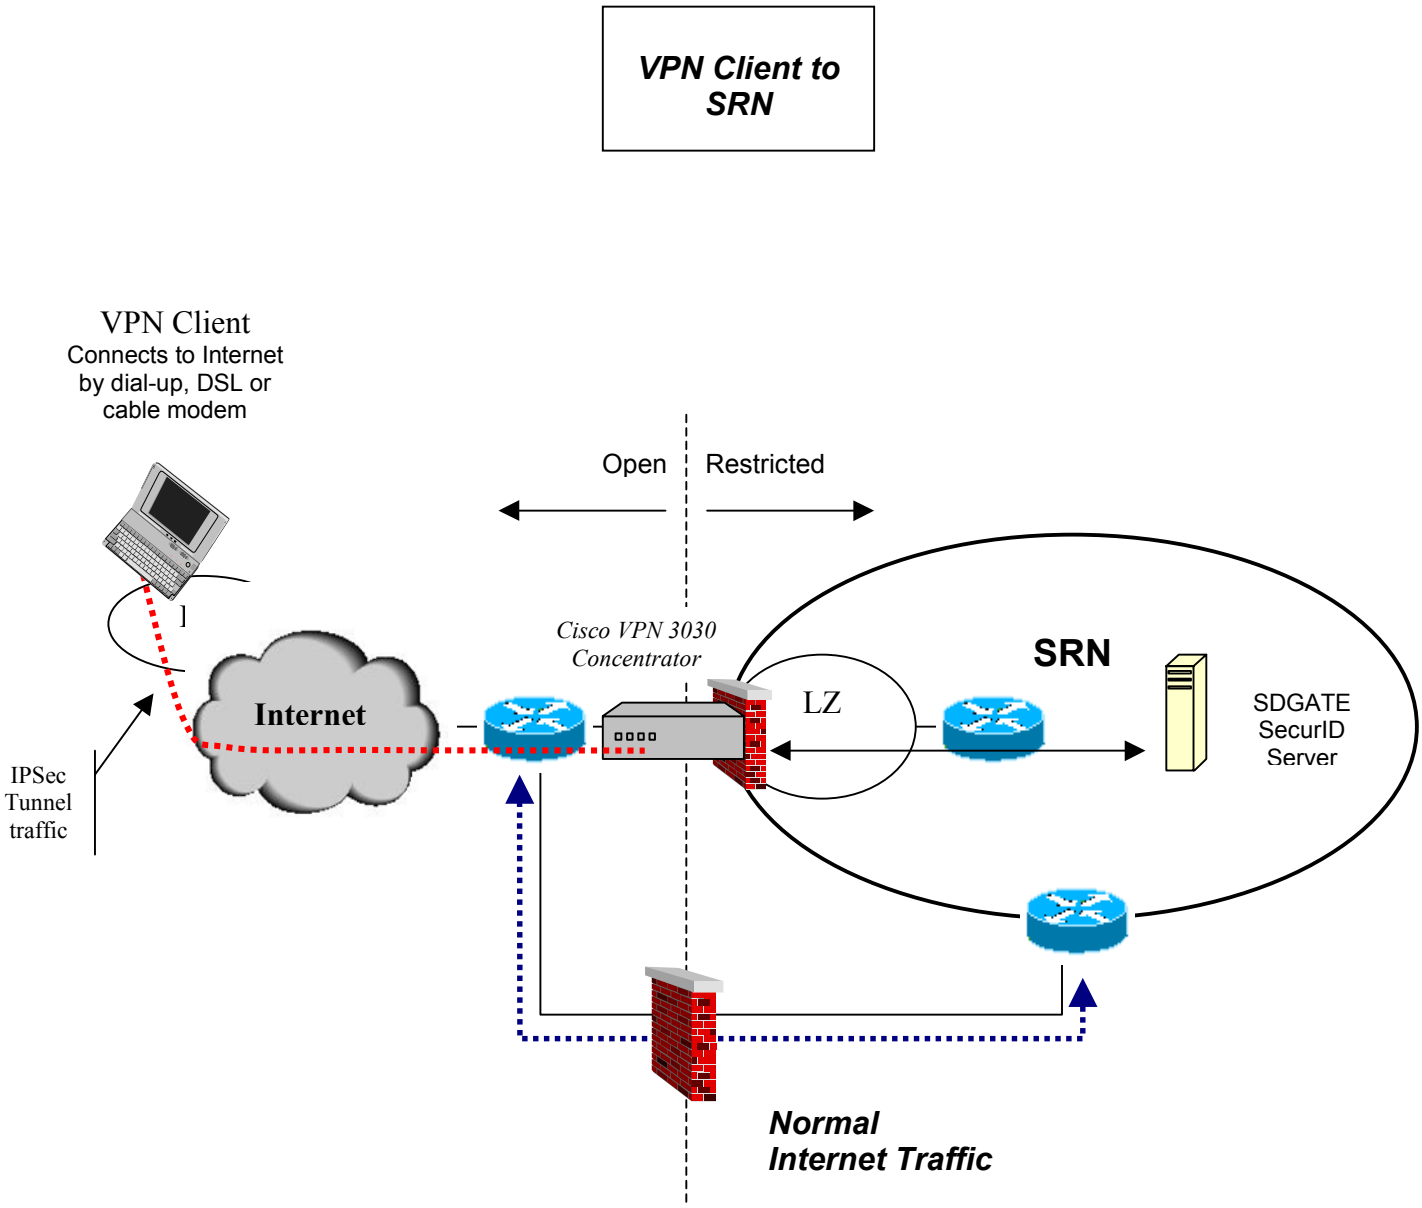

Figure B - 1 VPN Client Connection to SRN

Details

The IPSec protocol will be used for communication between a VPN client and the VPN server.

IPSec policies:

- Encapsulating Security Payload (ESP) protocol for privacy and integrity.
- Privacy provided by Triple-DES encryption.
- Integrity and replay prevention provided by SHA-1 message authentication code.
- Perfect Forward Secrecy (PFS) for maximum protection of session keys.
- Internet Key Exchange protocol (IKE) for key management.

All VPN traffic will operate in tunnel mode, in which the entire IP packet including address headers is compressed, encrypted, and authenticated. This secure packet is then placed into the payload section of a new IP packet with a new, unencrypted address header for originating and destination VPN devices. However, unlike a site-to-site tunnel in which both the original source and destination IP addresses are hidden, the client IP address must be maintained as it acts as its own VPN server.

At the SRN, the VPN 3030 will be established in parallel to the firewall. The VPN 3030 will process only IPSec packets, involving a properly authenticated client.

PPTP and L2TP protocols will be disabled. No PPTP or L2TP sessions will be serviced by the VPN 3030.

Key Management - IKE

The Internet Key Exchange (IKE) keying method will be used for agreeing on authentication methods, encryption methods, keys to use, how long to use the keys before changing them, and for secure key exchange. A group name and password will be used to authenticate the VPN client with the VPN 3030. SecurID will be used to authenticate the user (by way of a Radius server).

IKE negotiation options:

- PFS is enabled to protect session keys. Perfect Forward Secrecy is a cryptographic concept where each new key is unrelated to any previous key (for phase 2, PFS will ensure that a Diffie-Hellman key exchange is done each time a new SA is established, rather than relying on the keying material obtained from the keying material negotiated in phase 1).

- Default Diffie-Helman Group #2 (MODP 1024-bit) will be utilized (where the prime and generator numbers are 1024 bits. This option is more secure than Group 1, 768 bits, but requires more processing overhead).
- New IKE Security Associations will be generated when elapsed time reaches 8 hours (Network Alchemy's default value of 8 hours was carried over to the 3030; the rekey time is a balance between performance and security).

Tunnels

Each client machine establishing a VPN to the SRN will create an IPSec tunnel. Up to 1500 tunnels can be supported in a 3030.

Tunnels (IPSec SAs) will be rekeyed each hour of existence (again, this is the default value for Network Alchemy and is carried over to the 3030).

Authentication

There are two stages to authentication of a client VPN as shown in Figure B-2. The first stage of authentication is between the client and the 3030. This process uses a shared group name and password. Once this information is entered into the client, typically only during configuration, it is stored in the client and the user does not need to reenter the information at each connection. The client application automatically sends the group name and password when the connection is initiated.

Once 3030 authentication passes (group name and password), the user is authenticated using SecurID. The SecurID username and passcode are passed securely to the 3030 over the Internet, then in clear text to a Radius server, which in turn communicates with the RSA ACE server to complete the user authentication. The Radius server is being used for its superior reporting capabilities rather than communicating directly with the ACE server.

Once user authentication is successfully completed, the 3030 supplies the client with IPSec parameters for establishing an IPSec tunnel.

All authentication data is passed in encrypted form during the key exchange process, unlike a traditional dial-up SecurID session in which the username and passcode are passed in clear text.

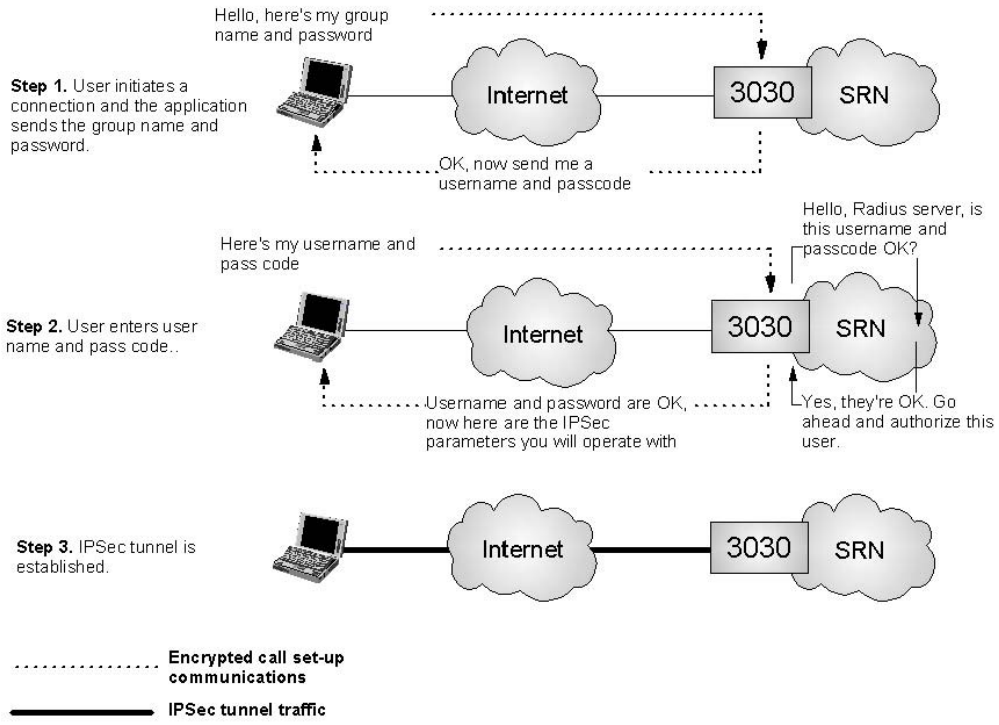

Figure B - 2 VPN Client Connection Process

Access Control

Typical VPN clients will have unrestricted access to the SRN, just as in normal dial-up access.

System Management

Console access times out after a set period of time, not to exceed 30 minutes (limited by the management application).

The 3030 will be configured to limit console access based on the IP address of the source. Only authorized IP addresses or IP subnets will be allowed console access.

Telnet access to the 3030 console will be disabled.

In-band system management will be implemented with a Java-based browser application. The management application is username and password protected. The application communicates with the VPN server using SSL with RC4-128 encryption.

Console access is protected by username and password.

Out-of-band management will be implemented on a local basis only, utilizing a terminal or PC with terminal-emulation software connecting through a serial port to the console. No dial-up capability will be provided into the console port. This console is normally utilized for diagnostics only; the preferred configuration method is by use of a secured browser.

No Telnet capability will be allowed into the management system. This capability will be turned off at the VPN server.

System Tests

Packets passed through a VPN tunnel were "sniffed" using a Network Associates sniffer with the intent of showing that these packets are converted to IPSec.

Packets in the Clear

The following two packets are the result of a single ping (ICMP echo and ICMP echo reply) from a VPN client machine on one end of a VPN testbed (134.253.5.187) to a workstation at the other end (134.253.4.75) with no tunnel. Note in the IP header that the protocol is identified as 1, ICMP.

```
- - - - - - - - - - - - - - - - - - - - Frame 1 - - - - - - - - - - - - - - - - -

DLC:  ----- DLC Header -----
      DLC:
      DLC:  Frame 1 arrived at  14:47:07.3129; frame size is 74 (004A hex)
bytes.
      DLC:  Destination = Station 00027DF32080
      DLC:  Source      = Station 3com  73A69D
      DLC:  Ethertype   = 0800 (IP)
      DLC:
IP: ----- IP Header -----
      IP:
      IP: Version = 4, header length = 20 bytes
      IP: Type of service = 00
      IP:       000. ....   = routine
      IP:       ...0 .... = normal delay
```

```
      IP:       .... 0... = normal throughput
      IP:       .... .0.. = normal reliability
      IP:       .... ..0. = ECT bit - transport protocol will ignore the CE bit
      IP:       .... ...0 = CE bit - no congestion
      IP: Total length    = 60 bytes
      IP: Identification  = 51497
      IP: Flags           = 0X
      IP:       .0.. .... = may fragment
      IP:       ..0. .... = last fragment
      IP: Fragment offset = 0 bytes
      IP: Time to live    = 32 seconds/hops
      IP: Protocol        = 1 (ICMP)
      IP: Header checksum = B997 (correct)
      IP: Source address      = [134.253.5.187]
      IP: Destination address = [134.253.4.75]
      IP: No options
      IP:
ICMP: ----- ICMP header -----
      ICMP:
      ICMP: Type = 8 (Echo)
      ICMP: Code = 0
      ICMP: Checksum = 1B5C (correct)
      ICMP: Identifier = 256
      ICMP: Sequence number = 12544
      ICMP: [32 bytes of data]
      ICMP:
      ICMP: [Normal end of "ICMP header".]
      ICMP:

- - - - - - - - - - - - - - - - - - - - Frame 2 - - - - - - - - - - - - - - - - -

DLC:  ----- DLC Header -----
      DLC:
      DLC:  Frame 2 arrived at  14:47:07.3137; frame size is 74 (004A hex)
bytes.
      DLC:  Destination = Station 3com  73A69D
      DLC:  Source      = Station 00027DF32080
      DLC:  Ethertype   = 0800 (IP)
      DLC:
IP: ----- IP Header -----
      IP:
      IP: Version = 4, header length = 20 bytes
      IP: Type of service = 00
      IP:       000. ....   = routine
      IP:       ...0 .... = normal delay
      IP:       .... 0... = normal throughput
      IP:       .... .0.. = normal reliability
      IP:       .... ..0. = ECT bit - transport protocol will ignore the CE bit
      IP:       .... ...0 = CE bit - no congestion
      IP: Total length    = 60 bytes
      IP: Identification  = 22233
      IP: Flags           = 0X
      IP:       .0.. .... = may fragment
      IP:       ..0. .... = last fragment
      IP: Fragment offset = 0 bytes
      IP: Time to live    = 126 seconds/hops
      IP: Protocol        = 1 (ICMP)
      IP: Header checksum = CDE7 (correct)
      IP: Source address      = [134.253.4.75]
      IP: Destination address = [134.253.5.187]
      IP: No options
      IP:
ICMP: ----- ICMP header -----
```

```
      ICMP:
      ICMP: Type = 0 (Echo reply)
      ICMP: Code = 0
      ICMP: Checksum = 235C (correct)
      ICMP: Identifier = 256
      ICMP: Sequence number = 12544
      ICMP: [32 bytes of data]
      ICMP:
      ICMP: [Normal end of "ICMP header".]
```

Packets Encrypted

The following two packets are the same ping packets shown above, but after conversion to IPSec. The ICMP echo packets have been converted to IP packets by encapsulating the original ping packet into an IP packet after encrypting the original packet with Triple-DES. The source packet maintains the original IP address of the client, 134.253.5.187. However, the packets are now addressed to the VPN server, 134.253.5.180. Note that the IP header now identifies the protocol as 50, ESP.

```
- - - - - - - - - - - - - - - - - - - - Frame 16 - - - - - - - - - - - - - - - -
-
DLC:  ----- DLC Header -----
      DLC:
      DLC:  Frame 16 arrived at  14:50:05.3751; frame size is 126 (007E hex)
bytes.
      DLC:  Destination = Station 0090A4003C39
      DLC:  Source      = Station 00027DF32080
      DLC:  Ethertype   = 0800 (IP)
      DLC:
IP: ----- IP Header -----
      IP:
      IP: Version = 4, header length = 20 bytes
      IP: Type of service = 00
      IP:       000. ....   = routine
      IP:       ...0 .... = normal delay
      IP:       .... 0... = normal throughput
      IP:       .... .0.. = normal reliability
      IP:       .... ..0. = ECT bit - transport protocol will ignore the CE bit
      IP:       .... ...0 = CE bit - no congestion
      IP: Total length    = 112 bytes
      IP: Identification  = 1
      IP: Flags           = 0X
      IP:       .0.. .... = may fragment
      IP:       ..0. .... = last fragment
      IP: Fragment offset = 0 bytes
      IP: Time to live    = 63 seconds/hops
      IP: Protocol        = 50 (SIPP-ESP)
      IP: Header checksum = 61F2 (correct)
      IP: Source address      = [134.253.5.187]
      IP: Destination address = [134.253.5.180]
      IP: No options
      IP:
ESP: ----- IP ESP -----
      ESP:
      ESP: Security Parameters Index = 19930080
      ESP: Sequence Number           = 1
```

```
      ESP: Payload Data             =
1E4037E04C7DB50D4FC5A152A6FD6BF5F79FD147041D6A470AA5BEB1DED5BB463DEB7AC008D2D0D
D0EAE2CC2E8310466D2...

- - - - - - - - - - - - - - - - - - - - Frame 17 - - - - - - - - - - - - - - - -
-
DLC:  ----- DLC Header -----
      DLC:
      DLC:  Frame 17 arrived at  14:50:05.3764; frame size is 126 (007E hex)
bytes.
      DLC:  Destination = Station 3com  73A69D
      DLC:  Source      = Station 0090A4003C39
      DLC:  Ethertype   = 0800 (IP)
      DLC:
IP: ----- IP Header -----
      IP:
      IP: Version = 4, header length = 20 bytes
      IP: Type of service = 00
      IP:       000. ....   = routine
      IP:       ...0 .... = normal delay
      IP:       .... 0... = normal throughput
      IP:       .... .0.. = normal reliability
      IP:       .... ..0. = ECT bit - transport protocol will ignore the CE bit
      IP:       .... ...0 = CE bit - no congestion
      IP: Total length    = 112 bytes
      IP: Identification  = 22574
      IP: Flags           = 0X
      IP:       .0.. .... = may fragment
      IP:       ..0. .... = last fragment
      IP: Fragment offset = 0 bytes
      IP: Time to live    = 64 seconds/hops
      IP: Protocol        = 50 (SIPP-ESP)
      IP: Header checksum = 08C5 (correct)
      IP: Source address      = [134.253.5.180]
      IP: Destination address = [134.253.5.187]
      IP: No options
      IP:
ESP: ----- IP ESP -----
      ESP:
      ESP: Security Parameters Index = 3620713157
      ESP: Sequence Number           = 1
      ESP: Payload Data              =
F52A28A36157B61A52211713B250B4C4EB04B3D9879E31342FF375F336EF2CC43E1543EBB038150
81FEC84A500A918CFBD...
```

Implementation Steps

Install a VPN 3030 at the edge of the SRN. No communication will be possible until the communications configuration is completed.

Configure the VPN 3030 to enable communications: set the IP addresses of the private and public interfaces, the network masks, default and tunnel routes.

Complete any necessary configuration changes such as entering the Radius server parameters.

Once the communications configuration process is completed, the VPN 3030 will respond to pings. Management is possible through the GUI from selected SRN nodes (not possible from the Internet). It will not pass any VPN traffic until an authenticated client establishes a tunnel.

Initial test of the 3030 will be done with a small set of users. Additional users will be added if the 3030 proves reliable. Full production will not occur until the 3030 receives FIPS 140-1 certification.

After the installation, the 3030 will be monitored closely for security, reliability, and performance. If any significant problems occur, the 3030 will be disabled.

Summary

How is user traffic protected?

All user traffic is protected by running the IPSec protocol, providing Triple-DES encryption and SHA-1 authentication/replay prevention.

How are users authenticated?

All users must authenticate using SecurID. The SecurID username and passcode are encrypted during the key exchange process.

How is an unauthorized person prevented from managing a VPN 3030?

The VPN 3030 console (browser, Telnet, direct cable connection) is protected by username and password. Browser access to the 3030 can and should be limited to specific IP addresses. Additionally, for maximum protection, browser access will be protected using SSL. Telnet access can also be protected using a special Telnet application that utilizes SSL. Otherwise, Telnet access should be disabled.

How is an unauthorized VPN device prevented from establishing a site-to-site tunnel to a 3030?

The 3030 will not be configured for site-to-site VPN tunnels. Unauthorized attempts to establish a site-to-site connection to the 3030 will fail. Attempts will be logged.

How often are SAs re-negotiated during their lifetimes?

IKE SAs will be re-negotiated every eight hours. IPSEC SAs will be re-negotiated every hour.

DISTRIBUTION:

1	MS	0801	W. F. Mason, 9320
1		0801	M. O. Vahle, 9300
1		0806	J. P. Brenkosh, 9336
1		0806	C. D. Brown, 9332
1		0806	L. B. Dean, 9336
1		0806	J. M. Eldridge, 9336
1		0806	M. J. Ernest, 9336
1		0806	S. A. Gossage, 9336
1		0806	R. L. Hartley, 9336
1		0806	T. C. Hu, 9336
1		0806	J. A. Hudson, 9336
1		0806	P. C. Romero Jones, 9332
1		0806	B. R. Kellogg, 9336
1		0806	J. P. Long, 9332
1		0806	L.G. Martinez, 9336
10		0806	M. M. Miller, 9336
1		0806	J. H. Naegle, 9336
1		0806	T. J. Pratt, 9336
1		0806	J. A. Schutt, 9336
1		0806	L. Stans, 9336
1		0806	T. D. Tarman, 9336
1		0806	L. F. Tolendino, 9336
1		0806	E. L. Witzke, 9336
1		0812	R. L. Adams, 9334
1		0812	M. D. Gomez, 9334
1		0812	C. M. Keliiaa, 9334
1		0812	E. J. Klaus, 9334
1		0812	J. H. Maestas, 9334
1		0812	P. L. Manke, 9334
1		0812	M. R. Sjulin, 9330
1		0812	T. J. Spears, 9334
1		0812	P. M. Torrez, 9334
1		0812	B.C. Whittet, 9334
1		0812	V. K. Williams, 9334
5		0812	G. A. Yonek, 9334
1		0813	R. M. Cahoon, 9327
1		0813	T. R. McMullen, 9327
1		0813	D. N. Packwood, 9327
1		0813	A. A. Quintana, 9327
1		0813	J. D. Stratton, 9327
1		0813	R. A. Suppona, 9327
1		0813	J. L. Taylor, 9327
1		9011	G. J. Blair, 8910
1		9011	J. A. Hutchins, 8910

1	9012	R. D. Gay, 8930
1	9012	S. C. Gray, 8930
1	9018	Central Technical Files, 8945-1
2	0899	Technical Library, 9616
1	0612	Review & Approval Desk, 9612 For DOE/OSTI
1	0161	Patent and Licensing office, 11500