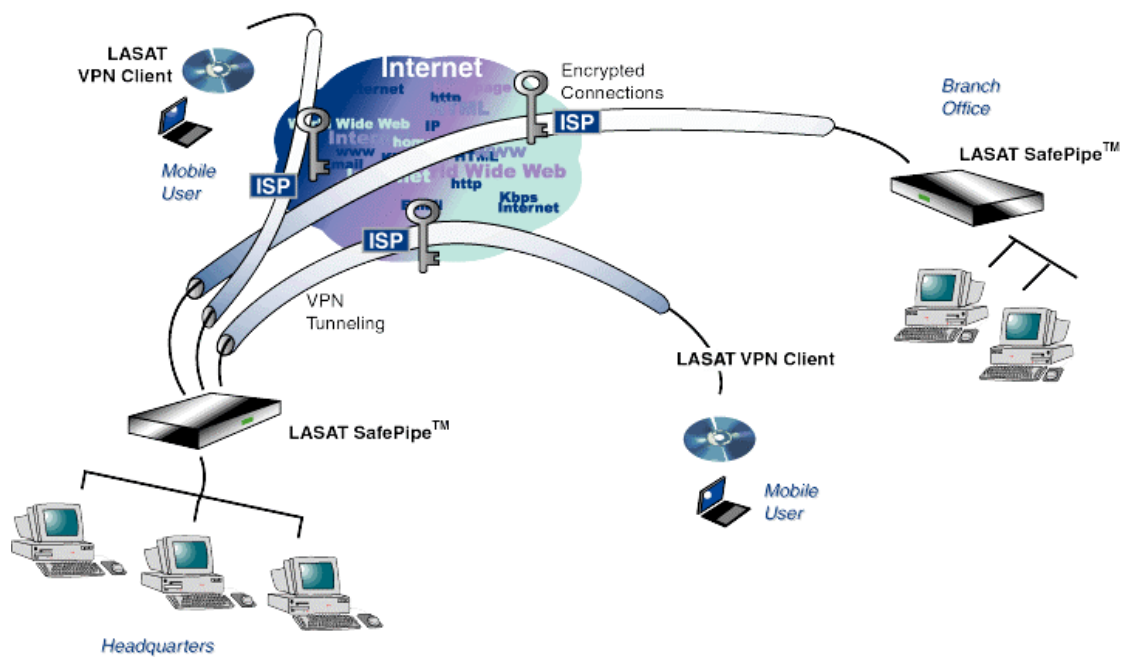


Secure and versatile networking with LASAT SafePipe™ VPN

- a technical white paper about LASAT SafePipe™
 - and the technologies behind it

Revision 1.0



CONTENTS:

1. INTRODUCTION	4
1.1 Virtual Private Networks	4
1.2 LAN-to-LAN VPN	4
1.3 Secure Remote Access	5
1.4 Cost savings	6
2. LASAT SAFEPIPE™ TECHNOLOGY	7
2.1 Introduction	7
2.2 IPsec	7
2.3 Internet Key Exchange (IKE)	8
2.4 Encryption	8
2.4.1 Triple-DES Encryption	8
2.4.2 DES encryption strength	9
2.5 Authentication	10
2.6 Digital certificate and CA's	10
2.7 Firewall	11
2.7.1 Network Address Translation (NAT) and Port mapping	12
2.7.2 Spoofing filtering	12
2.7.3 Packet or IP Filtering	12
2.8 IP routing	13
2.8.1 RIP	13
2.8.2 OSPF	13
2.9 Linux	13
2.10 Leased Line Technologies	14
2.10.1 Point-to-Point Protocol	14
2.10.2 Frame Relay	15
2.11 SNMP Network management	15

3. LASAT SAFEPIPE™ USE	16
3.1 LASAT SafePipe™ set-up scenarios	16
3.1.1 Friendly Networking	16
3.1.2 Friendly Networking with Separate Firewall	16
3.1.3 Business Networking	17
3.2 LASAT SafePipe™ installation procedure	18
3.2.1 Access to the LASAT SafePipe™	18
3.2.2 Configuring basic network information	18
3.2.3 Accessing the Internet	19
3.2.4 Configuring VPNs and remote clients	19
3.3 LASAT SafePipe™ System management	19
3.3.1 Assign an IP address to the public side of LASAT SafePipe™	20
3.3.2 Setting up a secure tunnel	20
3.3.3 CA and X.509 Certificates	20
3.3.4 Setting up remote clients	20
3.3.5 Firewall rules	20
3.4 Service mode	20
3.5 LASAT VPN Client	21
3.5.1 System Requirements	21
3.5.2 Configuration	21
3.5.3 Connection Password	21
3.6 LASAT SafePipe™ Hardware platform	21
3.6.1 MIPS RISC CPU	22
3.6.2 LZS Compression	22
3.6.3 EDHC Technology	23
3.6.4 LAN Interface	23
3.6.5 WAN interface	23
4. LASAT SAFEPIPE™ BENEFITS	24
5. REFERENCES	24

1. Introduction

The computer revolution has evolved into a network revolution. Today, it is a natural part of a company's infrastructure to provide internetworking between local and remote departments, home workers, road warriors as well as Internet access. With the rapidly expanding Internet, the communication infrastructure for the whole world is shifting gears.

The new opportunities mainly concern the use of the Internet as the cheapest, most widely deployed, easiest accessible and most versatile communication service. Using secure protocols, VPNs, it is possible to use the Internet to provide internetworking between enterprises, remote sites and business partners, as well as cheap or even free connectivity for home and travelling users.

Traditionally, remote branches have been connected via leased lines or switched networks like Frame Relay or X.25. This combined network, covering all the connected sites, is called a *private network*. The name "private network" is based on the fact that the same company owns or leases all components.

Many companies also have modem pools that allow the employees to connect to the network from the outside. This is commonly used for working at home or during business travel.

This document describes a way to take advantage of the new internetworking opportunities and the technology required to deploy a secure VPN with LASAT SafePipe™ from LASAT Networks.

1.1 Virtual Private Networks

A virtual private network has the same functionality as a traditional private network, but it uses the Internet, creating "virtual leased lines" that tunnel through the public networks. These virtual lines are protected by strong encryption and authentication mechanisms so that they result in being at least as safe as the traditional connections. Actually, since encryption protocols are always used with virtual private networks, they are in fact better protected than non-encrypted data traffic on leased lines, Frame Relay or X.25, where hackers can access the data traffic if they manage to break into the communications provider.

There is nothing exotic about VPNs. They are based on familiar networking technology and protocols. Basically, VPNs are private network overlays on a public IP network infrastructure such as the Internet.

1.2 LAN-to-LAN VPN

A LAN-to-LAN VPN consists of at least two sites that want to join their networks together. At each site, a VPN device is attached to the network, and these devices create a secure Internet tunnel between them. All data that is transferred from one internal network to the other is passed through the tunnel in encrypted packets.

By creating a VPN over the Internet between corporate offices, organizations can protect sensitive business information from eavesdropping, data tampering and forgery.

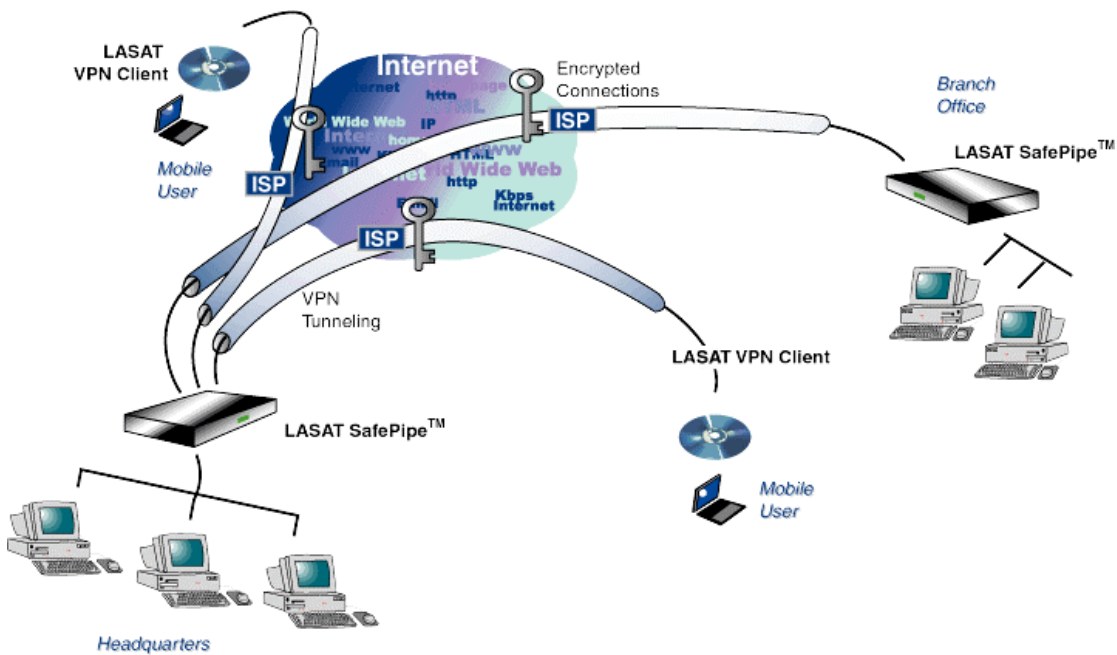
1.3 Secure Remote Access

Many companies have modem pools that allow the employees to connect to the network from the outside. This is commonly used for working at home or during business travel.

Secure remote access gives the same result, using Internet tunnels. Again, the protection is generally better when using the strong encryption protocols in the VPN than by using ordinary, non-encrypted telephone lines.

Secure remote access uses the same VPN device in the company network, and a software driver installed in the remote PC. A tunnel is created between the device and the PC, through which all data is passed. To avoid back doors from the Internet, all Internet access from the remote PC must first pass through the tunnel, and then out via the company firewall.

Tunnels are becoming the most cost-efficient, secure and reliable means of remotely accessing the networks. It becomes the job of the Internet providers to maintain the modem pools.



1.4 Cost savings

The cost savings are achieved by connecting to an Internet Service Provider (ISP) at each network location and by encrypting the data that passes through the public Internet. VPN technology is reasonably priced, and payback periods are often mere months, rather than years. In addition, VPN solutions replace both the direct-dial and leased line solutions, and are subsequently inexpensive enough to enable communication for companies who might never have been able to consider more costly options.

Furthermore, because VPNs use the worldwide Internet, they are able to offer global connectivity of an unparalleled scope. Instead of linking each branch office to corporate headquarters, you can link all your offices together by simply connecting them to the Internet. And as your company grows, your VPN savings grow as well. When corporations scale up to a larger network, VPN savings increase exponentially. The larger the network, the higher the savings. It's really that simple.

A fast VPN connection is less than half the price of a leased line. As an example, a 2Mbit Internet connection from Copenhagen to the USA or Europe is less than 1/3 of the price of a leased line. On a 128 Kbit Internet connection from Scandinavia to Europe you save 35% on your connection fee the first year.

2. LASAT SafePipe™ Technology

2.1 Introduction

LASAT SafePipe™ from LASAT Networks is a hardware-based IPSec VPN device that supports Triple-DES encryption, X.509 digital certificates and is even able to issue certificates via the internal certificate authority (CA). It has a built-in firewall and supports routing updates via standard router update protocols (RIP and OSPF). The box is based on the Linux operating system.

The next section describes briefly the technology behind the LASAT SafePipe™ IPSec VPN box.

2.2 IPSec

IPSec, the security standards developed by the Internet Engineering Task Force (IETF) IP Security Working Group, is the most secure and comprehensive standard available today for **encryption, authentication and key management**.

IPSec ensures confidentiality of your data by encrypting it and providing the information for the decryption. They also safeguard your important communications by guaranteeing user authentication and network access control—no one can tamper with your network, or disguise themselves as a user.

IPSec's interoperability lets LASAT SafePipe™ products exchange keys and encrypted communications with all other IPSec-compliant products, so customers can use different IPSec vendors for multiple scenarios.

A conventional TCP/IP packet contains an IP header, which consists of a source and destination address, control fields, and information about the packet contents (the payload). Because all protocols in the TCP/IP suite have a "next protocol" field as part of the header format, they can be combined in different ways—and that's where IPSec fits in.

IPSec (**I**nternet **P**rotocol **S**ecurity) adds new fields to packet headers, and these fields are what make authentication and encryption possible. Authentication assures the recipient that an IP packet has come from the actual sender. IPSec protocols validate the packets with a secret key that the receiver and sender share. Since most network communications involve packets sent in both directions, this process authenticates both parties. But that's where it ends: Authentication simply confirms the source of the packet—it does nothing to hide its contents.

That's where encryption comes in. IPSec uses an encryption algorithm, e.g. **DES**, to change the data into something completely unreadable for anyone without the proper decryption key. Authentication, for instance **MD5** or **SHA**, ensures that no one is sending fake packets or repeating packets already sent.

But by themselves, neither authentication nor encryption assures a secure connection - it's the combination that counts and that's what makes IPSec the complete toolkit and the most trusted secure communication available today.

The goal of the IPSec group has been the creation of a new secure version of IP, the Internet Protocol. This new version works within the network layer to ensure that your data is protected. This version will be mandatory within the next release of IP. Right now, IPSec is the security standard that VPNs must meet.

2.3 Internet Key Exchange (IKE)

Before a secure session can begin, the communicating parties need to negotiate the communication terms. The automated protocol for this in IPSec is the Internet Key Exchange (IKE). IKE is used each time a VPN tunnel is activated. IKE agrees on which protocol, algorithms, and keys to use, verifies that both ends are who they're supposed to be, and then negotiates a *session key*, which is a temporary key used by IPSec for data encryption and authentication.

This temporary session key in IKE is generated by Diffie-Hellman public key distribution, a cryptosystem that uses a pair of mathematically related keys; a private key that is kept secret within the system, and a public key that can be made known to the public. The Diffie-Hellman key length use in LASAT SafePipe™ is 1024 bit. IKE can use *certificates*, or shared keys to verify that both ends are the correct ones.

2.4 Encryption

Using a secret key, data can be converted to something that is completely unreadable unless you use the key again to convert it back. A private key crypto system uses the same secret key for both encryption and decryption, the key being a fixed-length bit string.

Symmetric encryption schemes are very fast. Each of these differs in bit length, which is more commonly referred to as the “strength” of the encryption algorithm. The strength of the algorithm establishes the amount of effort required to break the system using one of the more common attacks known as a brute force attack.

Many computers can combine forces in an attempt to calculate all of the possible key permutations. Therefore, the longer the key length, the “stronger” the encryption algorithm and the greater the effort required to break the system.

With good algorithms the only attack is to actually try all possible keys until the resulting data becomes readable again. For each extra bit in the key length, there are twice as many possible numbers (keys) to try, so a 57-bit key is twice as hard to break as a 56-bit key. A 112-bit key takes about 70,000,000,000,000,000 times as long to break as a 56 bit one.

2.4.1 Triple-DES Encryption

The most trusted encryption algorithm is Triple-DES, which can encrypt with the strength of 112 bits. Other algorithms are presumed to be as good or better, but they have not been checked and tested to the same degree. To put it simple: they are not as trusted as Triple-DES, since they may have undiscovered weaknesses. The predecessor to Triple-DES, which is simply named DES, is no longer very trusted since the key is so short that it should be breakable without too much trouble. DES is an algorithm developed in the 1970s by IBM, which was made a standard by the US government.

Triple-DES is not three times as strong as single DES. Cryptographers tell us that it is at least 2^{56} times stronger than single-DES and it could be even stronger.

When using Triple-DES, the data is first cut into 64-bit pieces. Each piece is passed through the Triple-DES machinery and changed into another, encrypted, 64-bit piece. All the resulting pieces make up the encrypted data

The receiver of the encrypted data cuts it into 64-bit pieces again, and passes these through the Triple-DES machinery, this time in decryption mode. The encrypted pieces turn back into original pieces and are joined together to become the original data.

For both encryption and decryption, the Triple-DES machinery is initialized with the key.

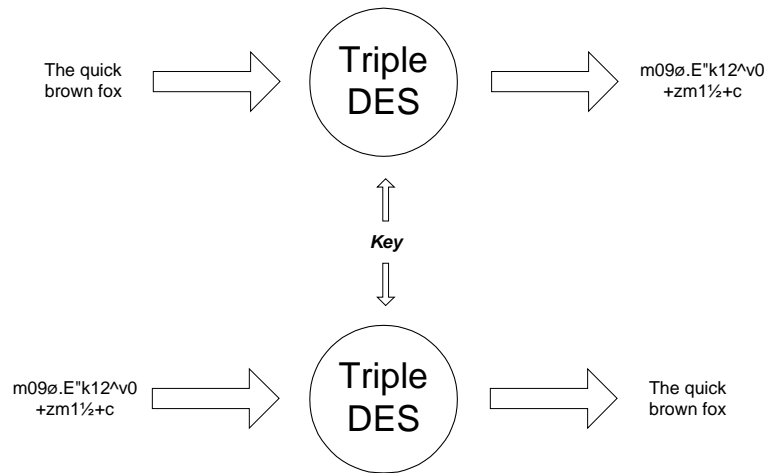
Triple-DES is based on using DES three times (normally in an encrypt-decrypt-encrypt sequence with three different, unrelated keys).

2.4.2 DES encryption strength

There is some confusion regarding the possible key lengths. When using DES, every eighth bit is used as a parity bit, which means it does not add to the security. In other words, DES with a 64-bit key has the strength of 56 bits.

Triple-DES uses 192 bits of key as its input. Every eighth bit is parity, which means that it only corresponds to an effective key size of 168 bits. To complicate matters even further, only two thirds of these bits directly add to the strength of the algorithm—the last third is needed to make the first two thirds secure. In short, Triple-DES has the strength of 112 bits.

Various sources may claim the strength of 56 or 64 bits in single DES, but it should always be understood as 56 bits. They may claim 112, 128, 168 or 192 bits of strength for Triple-DES, but this should always be understood as 112 bits. That's the effective protection you get.



2.5 Authentication

Hash functions can be used to ensure that the message is in its original form (data integrity). When used for data integrity, the hash functions (and resultant digital signatures) are more commonly called data authentication algorithms. Using a digital signature system for data authentication accomplishes the following:

- Proves it was sent by the sender and was not forged (data authentication)
- Proves that the message was not modified or tampered with (data integrity).

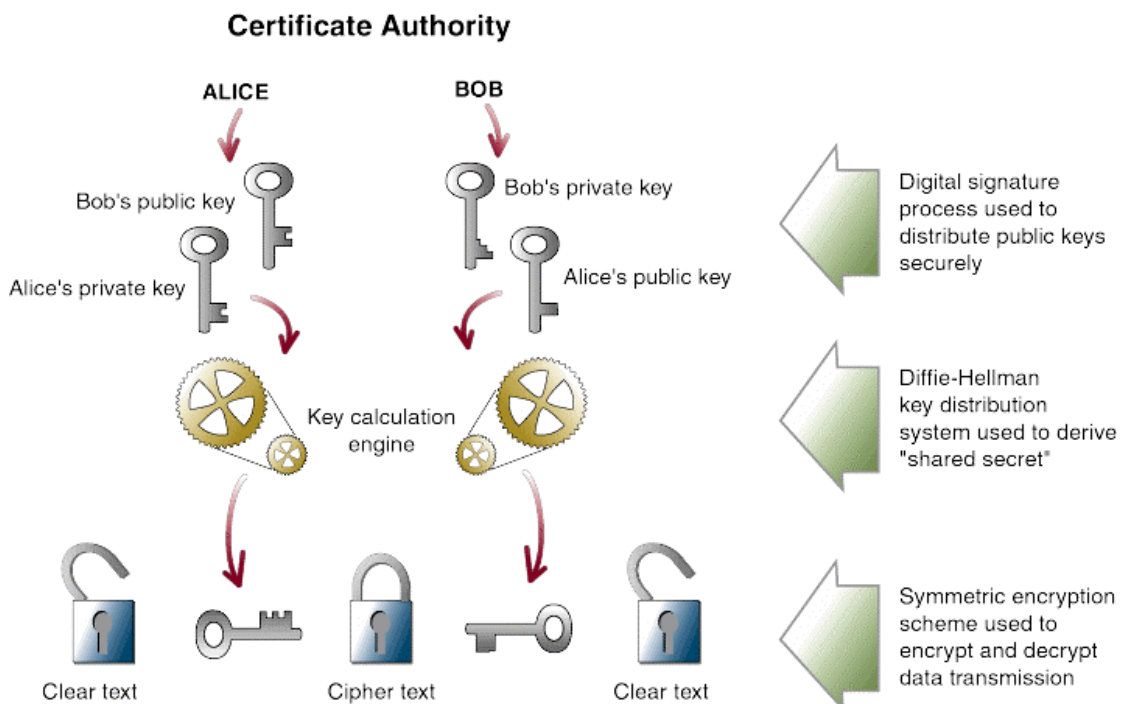
MD5 (Message Digest Algorithm 5) can be used to hash an arbitrary length byte string into a 128-bit value. MD5 is in wide use, and is considered reasonably secure.

SHA (Secure Hash Algorithm) is the other digital signature process in IPSec and this cryptographic hash algorithm produces a 160-bit hash value from an arbitrary length string. Many people consider it quite good. It is a fairly new algorithm.

2.6 Digital certificate and CA's

Digital signatures can also be used to testify (or **certify**) that a public key belongs to a particular person. This is done by signing the combination of the key and the information about its owner by a trusted key. The digital certificate is actually the owner's public key that has been digitally signed by the CA. CA is a trusted third-party organization, company or internal CA that issues digital certificates used to create digital signatures and public-private key pairs. The role of the CA in this process is to guarantee that the individual granted the unique certificate is, in fact, who he or she claims to be.

Digital certificates give the Internet a high level of certainty, much the way a passport or driver's license verifies a person's identity.



Digital certificates are a bit of structured information issued by a certificate authority. The CA system can be run by either a department in a company or a separate company altogether. It verifies a digital certificate user's identity offline—by telephone, postal mail, or in person. Then the information is put into the CA server, which generates a public key and issues a digital certificate to the user, along with a related private key. The user can then encrypt data using the private key and a recipient's public key. The recipient then decrypts the data with his or her private key and the sender's public key—verifying the identity of the sender in the process.

The digital certificate normally contains the following data:

X.509 Certificate
User Name
User Organization
Certificate Start Date
Certificate End Date
User Public Key Parameter
CA Name
CA Signature on Certificate

X.509 is an International Telecommunications Union (ITU) standard, which has become a de facto industry standard for user authentication in an open systems environment.

2.7 Firewall

Firewalls are barricades at the edge of your company's network to keep intruders from entering. Firewalls can be stand-alone devices or fully integrated firewalls built into routers or remote access servers. They can also be implemented at the application level using proxy gateways and servers.

Basic firewalls have rules that only allow certain traffic to pass to and from the Internet. These rules generally prohibit devices on the Internet to use services inside the company, except for instance a Web server. This type of firewall will keep most attackers at bay. Advanced firewalls have comprehensive knowledge about how to avoid various attack methods, and filter out packets that are known to make trouble because of errors in some operative systems (e.g. "Ping of Death"). This type of firewall prevents all but the most resourceful attackers to gain access.

On the market today, there are two different types of firewalls. They are IP firewalls and application firewalls, the so-called proxy servers that serve as a relay between two networks, breaking the connection between the two.

LASAT SafePipe™ has a NAT firewall with IP packet filtering, and as the name implies it includes a **Network Address Translation (NAT)**, based on source IP address and **port mapping**. **Packet filter or IP filter** blocks traffic based on IP address and/or port numbers,

and it works on the Network layer. The firewall also provides an **IP spoofing filter** functionality and a traffic type specific filter.

2.7.1 Network Address Translation (NAT) and Port mapping

The combination of explosive growth in TCP/IP networking and the long-standing practice of assigning globally unique IP addresses to all hosts on TCP/IP networks has resulted in rapid depletion of the available IP address space. Because a unique address is required for each host connected to the global Internet, this presents a serious problem for new enterprise connections.

The solution is NAT (network address translation) defined in RFC 163. A firewall or router using NAT essentially takes all private addresses of outbound traffic (traffic from the internal network to the Internet) and converts the source address to that of the router or firewall's external interface (or to a series of addresses if there are multiple external interfaces). For inbound traffic, the process works in reverse: The NAT box converts destination addresses to those used by the private network.

But address conversion is just one of the advantages of NAT. Security is another: Attackers can't go after machines they can't see—and private addresses aren't visible on the public Internet.

It would be very difficult indeed to take full advantage of reusable addresses on a private network without employing NAT functionality.

All requests originating from the internal network will have their source IP address replaced by the mapped IP address. In a case where a company has only purchased one public IP address, and where port mapping is not used, it will only be possible to have one internal person surfing at a time. This is why port mapping is used.

If the address of the internal host does not appear in the NAT table, a new entry is created for the host. Upon this new entry the packet is also provided with a new available port number. This information then replaces the original packet information. After a time-out period, during which no translated packets for a particular address mapping have been transferred, the entry in the NAT table is removed.

2.7.2 Spoofing filtering

This feature restricts access into an external interface by disallowing a packet to get through if it contains a source address from an internal network. This feature prevents IP spoofing attacks originating from any outside network. Spoofing is faking the sending address of a transmission in order to gain illegal entry into a secure system. This filtering process is performed automatically.

2.7.3 Packet or IP Filtering

The most secure firewall will make sure that no traffic will get from the internal network and on to the external one, being the Internet, and vice versa. Obviously, this limits the functionality too much. The default configuration is that traffic going from the internal network to the Internet is allowed, but traffic from the Internet to the local network is not

allowed. This is where IP filtering comes in, allowing certain rules to be added to the firewall. Packet filtering or IP filtering is discarding unwanted network traffic and allowing other traffic based on its originating address or range of addresses or its type (e-mail, file transfer, etc.).

2.8 IP routing

LASAT SafePipe™ supports IP routing updates via RIP and OSPF, and is very important when you're working with large installations. When you have a large internal network at the "other" site, it's VERY time costly to keep updating the IP addresses. VPN is linking your network together, but without IP routing there is no update on network changes on the network at the other end of the tunnel. With IP routing the two networks act as one virtual network with automatic routing updates on the entire network.

A routing protocol is a formula used by routers to determine the appropriate path onto which data should be forwarded. The routing protocol also specifies how routers report changes and share information with the other routers in the network that they can reach. A routing protocol allows the network to dynamically adjust to changing conditions, otherwise all routing decisions have to be predetermined and remain static.

A routing table is a database in a router that contains the current network topology.

RIP and **OSPF** has become a de facto standard for the exchange of routing information among gateways (routers). It is implemented for this purpose by most commercial vendors of IP gateways.

2.8.1 RIP

RIP (Routing Information Protocol) is a simple routing protocol that is part of the TCP/IP protocol suite. It determines a route based on the smallest hop count between source and destination. RIP is a distance vector protocol that routinely broadcasts routing information to its neighboring routers and is known to waste bandwidth.

2.8.2 OSPF

OSPF (Open Shortest Path First) is a routing protocol that determines the best path for routing IP traffic over a TCP/IP network. OSPF is an interior gateway protocol (IGP) and a link state protocol that provides less router to router update traffic than the RIP protocol (distance vector protocol) that it was designed to replace.

Interior Gateway Protocol is a broad category of routing protocols that support a single, confined geographic area such as a local area network (LAN).

2.9 Linux

Many new Internet products are based on the operating system Linux, an open-source freeware system and therefore a very stable system. The open-source community is a loosely knit group of developers worldwide who write applications and utilities that are usually available for the asking. There is a myriad of Web sites and bulletin boards loaded with

freeware, shareware, testimonials, advice, tips, bug fixes, and endless discussions of the finer points of Linux. In fact, those resources typically take the place of traditional tech support. Open-source software is nothing new, despite all the attention it's getting. BSD Unix, the first open-code OS, has been around for more than 20 years. Linux was developed nine years ago by Linus Torvalds, then a student at the University of Helsinki. Linus Torvalds turned Minix, a popular classroom teaching tool, into Linux, which is closer to the real UNIX. Torvalds created the kernel, and most of the supporting applications and utilities came from the GNU project of the Free Software Foundation. Many programmers have contributed to the Linux/GNU system.

Many Linux lovers first encounter the code when they're in college or grad school and Netware and NT are priced way out of reach. They must like what they see, since Linux's growth rate is nothing short of astounding. It could claim all of six users in 1991. By 1992 that number had hit 1,000; six years later there were 7.5 million. Consider this: Linux grew 212 percent in 1998 alone, grabbing 17.2 percent of the server OS market. That's up there with Unix, which accounts for 17.4 percent of the same market. All this activity helps explain why open-source software is now available from 10 vendors and industry groups.

Linux is a version of UNIX that runs on x86, Alpha, MIPS and PowerPC machines. Linux is essentially freeware; however, the full distribution of Linux along with proprietary add-ons and support are available for a fee from vendors such as Red Hat Software and Caldera. The distribution CD-ROMs include the complete source code as well as hundreds of tools, applets and utilities.

Due to its stability, Linux has gained popularity with ISPs as the OS for hosting Web servers. Its usage is expected to grow as a server OS as well as for the desktop.

How reliable is it? Well, the Linux kernel is much smaller than NT, which means there's much less to go wrong, the Linux kernel runs to some 500,000 lines where the NT kernel comes in at 7 million lines.

The Linux OS is not visible to the user of LASAT SafePipe™, but is managed through a graphic web browser interface.

2.10 Leased Line Technologies

When you talk about the old leased line technology, it could either be a point-to-point circuit switching connection running **synchronous PPP** or a one point to many packet switching running **Frame Relay**.

2.10.1 Point-to-Point Protocol

A data link protocol that provides dial-up access over serial lines. It can run on any full-duplex link from POTS to ISDN to high-speed lines (E1/T1 etc.). PPP encapsulates protocols in specialized Network Control Protocol packets; for example, IPCP (IP over PPP) and IPXCP (IPX over PPP). It can be used to replace a network adapter driver, allowing remote users to log on to the network as if they were in-house.

PPP also provides password protection in clear text using the Password Authentication Protocol (PAP) and the more rigorous Challenge Handshake Authentication Protocol (CHAP).

The protocol can run on any full-duplex link from ISDN to Frame Relay. The protocol is mostly used as a method for carrying higher level protocols.

2.10.2 Frame Relay

A high-speed packet switching protocol used in wide area networks (WANs). Providing a granular service of up to E3 (34Mbit) or DS3 speed (45 Mbps), it has become very popular for LAN to LAN connections across remote distances. Services are offered by all the major carriers. Frame relay is much faster than X.25 networks, the first packet-switching WAN standard, because frame relay was designed for today's reliable circuits and performs less rigorous error detection. Although X.25 was never widely used, frame relay has become a major wide area technology. The name comes from the fact that frame relay does not do any processing of the content of the packets; rather, it relays them from the input port of the switch to the output port.

Frame relay provides permanent and switched logical connections, known as Permanent Virtual Circuits (PVCs) and Switched Virtual Circuits (SVCs). These are logical connections provisioned ahead of time (PVCs) or on demand (SVCs). Committed Information Rate (CIR), provides a certain amount of transmission capacity for the connection. CIRs are adjusted with experience and given in %. Many Frame Relay lines have a CIR of 50 meaning a 50% guaranteed bandwidth.

A customer attaches to the frame relay network via a frame relay access device (FRAD) which resides on the customer's premises. The FRAD may be a separate device or software built into the router. The FRAD connects to a port on a frame relay switch on the service provider's network via an interface known as the User-to-Network Interface (UNI). This line/port is typically some multiple of 64 Kbps, and all traffic for one customer generally travels through the same port.

2.11 *SNMP Network management*

(Simple Network Management Protocol) is a widely-used network monitoring and control protocol. Data is passed from SNMP agents, which are hardware and/or software processes reporting activity in each network device (hub, router, bridge, etc.) to the workstation console used to oversee the network. The agents return information contained in a MIB (Management Information Base), which is a data structure that defines what is obtainable from the device and what can be controlled (turned off, on, etc.). Originating in the UNIX community, SNMP has become widely used on all major platforms.

SNMP 2 provides enhancements including security and an RMON (Remote Monitoring) MIB, which provides continuous feedback without having to be queried by the SNMP console. RMON is proactive and can set alarms on a variety of traffic conditions, including specific types of errors. RMON2 can also monitor the kinds of application traffic that flow through the network. The full RMON capabilities are very comprehensive and generally only portions of it are placed into routers and other network devices.

3. LASAT SafePipe™ use

LASAT SafePipe™ from LASAT Networks is a series of hardware-based IPSec VPN products with Triple-DES hardware encryption (EDHC technology) and compression. Throughput in Triple-DES is from 5- to 45Mbit/s. LASAT SafePipe™ is based on the Linux operating system and a MIPS CPU. It supports IP routing via RIP and OSPF, X.509 digital certificate and has built-in Certificate Authority and NAT firewall. It comes with an optional X.21/V.35 WAN interface. LAN interface is both 10 and 100Mbit Ethernet. The configuration interface to the LASAT SafePipe™ is very simple and is done via a standard web browser like Microsoft's Internet Explorer or Netscape communicator.


3.1 LASAT SafePipe™ set-up scenarios

This text describes the various network setups with LASAT SafePipe™, based on network environment and needs.

3.1.1 Friendly Networking

With friendly networking, the purpose is to join sites and users together with all resources visible at all connected sites, but of course secured against the rest of the Internet. The joined subnets become one single private network. The typical application is a number of branches belonging to a single company, and the remote and traveling workers from this company.


The firewall in the LASAT SafePipe™ is used to secure the internal network, as well as giving the internal devices access to the Internet with a single shared IP address (if needed).



3.1.2 Friendly Networking with Separate Firewall

With this setup, the purpose is the same as with ordinary friendly networking. A separate firewall is to be used (for instance, one has been installed earlier), so the LASAT SafePipe™ is installed in parallel, and direct Internet access through the LASAT SafePipe™ is disabled in the LASAT SafePipe™'s firewall setup. This means that only tunneled traffic to other friendly sites passes through the LASAT SafePipe™; Internet traffic goes through the separate firewall. There is no single point of failure, which means that either type of traffic can continue, even if the other type has been disabled or fails for other reasons.

Each internal device is usually set up with at single gateway. It can be the firewall, which must then be configured to forward all traffic that is to go to remote, friendly networks through the LASAT SafePipe™.



The internal devices can also be set up to use the LASAT SafePipe™ as default gateway (or the LASAT SafePipe™ is installed with the IP address currently used as default gateway for internal devices). The LASAT SafePipe™ must then be set up with the IP address of the secure interface of the firewall as default gateway. In this case, no manually entered route information for remote, friendly networks is needed, since the connected LASAT SafePipe™ will exchange the necessary information and update their routing tables accordingly.

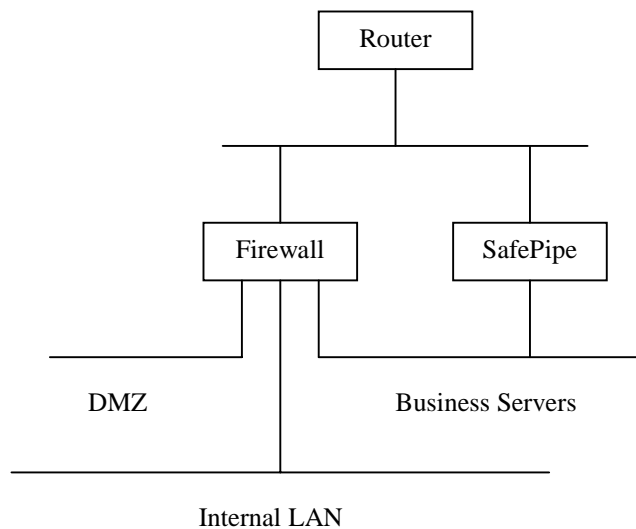
3.1.3 Business Networking

With business networking, other parties are allowed secure access to selected resources—business to business servers. To protect the internal, private network, these resources are placed on a separate subnet, much like a DMZ. The private and the business subnets are connected via a firewall (a router is theoretically enough, if the business servers do not have security holes).

In this scenario, the compatibility tunnels of the LASAT SafePipe™ can be used. These tunnels allow exactly one remote subnet to access exactly one local subnet (the business subnet), in stead of trying to open as much access as possible.

Alternatively, some business partners can be set up with remote clients.

In the example shown, a single firewall is used to connect the Internet, the DMZ, the business subnet and the internal network.



3.2 LASAT SafePipe™ installation procedure

To install LASAT SafePipe™ is a two-phase process and the following information is needed:

- An IP address to assign to the *public interface* on the Internet.
- An IP address to assign to the *private interface* on the LAN.
- The IP address of the remote LASAT SafePipe™ to establish a tunnel to.

The first installation phase takes place using the LASAT SafePipe™ administration tool. This tool does not use a secure protocol, and hence it is strongly suggested that the installation process is carried out on an isolated network environment.

The second installation phase is completed with the encryption of the browser, and should take place in the final network environment in which LASAT SafePipe™ is to work.

3.2.1 Access to the LASAT SafePipe™

As delivered, LASAT SafePipe™ does not have an IP address and no admin password has been selected. This needs to be configured before the browser-based management interface can be accessed.

The IP address of the *private* interface and the admin password of LASAT SafePipe™ are set up using the LASAT SafePipe™ administration tool.

It is strongly suggested that no other devices can listen in on the LAN on this step, in order to keep the password perfectly secret.

This can be accomplished by connecting the crossed cable directly from the PC to the LASAT SafePipe™, or making sure no other cables are plugged into the HUB/switch between the PC and the LASAT SafePipe™.

At the start, the LASAT SafePipe™ administration tool will search the network for LASAT SafePipes™ that are in service mode.

Enter the IP address and subnet mask you have selected for LASAT SafePipe™ and type in the password you want to use.

After the password has been entered, the LASAT SafePipe™ switches to secure mode and can no longer be managed with the installation program. If it is necessary to use the LASAT SafePipe™ administration tool at a later stage, LASAT SafePipe™ must be switched to service mode.

3.2.2 Configuring basic network information

The LASAT SafePipe™ is now ready to be relocated to its final place in the network.

At this point management will be completely secure, so the LASAT SafePipe™ and the management computer may be connected to the rest of the network. It is most secure to use a browser version that uses 128-bit security. 128-bit upgrades for some of the most popular browsers can, if necessary, be found in the “browser” directory of the CD-ROM delivered with the LASAT SafePipe™.

3.2.3 Accessing the Internet

For setting up the public IP address, enter the external IP address and network mask. The IP address can be the same as the internal (default) one, or an address on a different subnet, depending on your preferences. If the LASAT SafePipe™ is installed in the path between the internal users and the Internet, and the internal users have to share one IP address, this must be configured as so.

3.2.4 Configuring VPNs and remote clients

At this point, internal devices should be able to browse the Internet. Use the configuration options in the Web-interface to configure VPNs and remote clients as needed. Each configuration screen will explain the steps.

3.3 LASAT SafePipe™ System management

LASAT SafePipe™ can be monitored from a standard web browser or via an SNMP-based management system. SNMP is a widely-used network monitoring and control protocol that reports activity and sends it to the workstation console used to oversee the network.

Start your web browser, and enter the IP address of the LASAT SafePipe™ as the URL. The login page of the LASAT SafePipe™ will now show up, and you are prompted to enter the management password. After entering the password you are presented with the LASAT SafePipe™ main page. The LASAT SafePipe™ web interface is designed around a main menu with five items on the left side. Some of the menus have submenus, which are presented on the right side. The top of the page always shows you a path to your location in the menu system.

The five top menus are:

1. **System**
For controlling the X.509 certificates, and access to LASAT SafePipe™.
2. **Tunnels**
Creation, deletion and configuration of tunnels.
3. **Remote**
Creation, deletion and configuration of remote clients
4. **Network**
Configuration of network interfaces, firewall and bridging.
5. **Information**
Status, log and statistical information about LASAT SafePipe™

To complete the configuration the following five steps have to be performed:

1. Assign an IP address to the public interface of your LASAT SafePipe™, the interface that is connected to the “insecure” Internet.
2. Set up a secure tunnel between your LASAT SafePipe™ and the LASAT SafePipe™ at the branch office.
3. Optional creation of X.509 certificates in the CA.
4. Create remote clients as needed.

3.3.1 Assign an IP address to the public side of LASAT SafePipe™

During installation of LASAT SafePipe™, an IP address was given to the *private* network interface on the LASAT SafePipe™. The *private* network interface is the one that is connected to the company LAN. The second interface on LASAT SafePipe™, the *public* network interface, is the one that is connected to the “insecure” Internet. The IP address to the *public* network interface can be one given by your ISP, or it can be an IP address on a network that connects LASAT SafePipe™ to your company Internet router.

3.3.2 Setting up a secure tunnel

It’s now possible to set up the secure tunnel between the LASAT SafePipe™ and the LASAT SafePipe™ at the branch office. LASAT SafePipe™ can maintain a number of simultaneous secure tunnels with other LASAT SafePipe™ or IPSec compatible devices. To establish a secure tunnel with another LASAT SafePipe™, enter the IP address of the *public* network interface of the other LASAT SafePipe™, and enter a shared secret or certificate.

3.3.3 CA and X.509 Certificates

Will be added later.

3.3.4 Setting up remote clients

LASAT SafePipe™ not only allows secure tunnels between two LASAT SafePipes™. LASAT SafePipe™ also allows setup of secure tunnels with remote clients running the LASAT SafePipe™ client software. This is ideal for providing travelling workers with secure access to the company LAN. To provide access to a remote worker, create a new client, enter a temporary key, and download the client configuration file to a disk. The remote client uses the temporary key and the configuration file during the installation of the software on the PC.

3.3.5 Firewall rules

LASAT SafePipe™ contains a firewall that allows you to restrict access to and from the company LAN. By entering firewall rules, you can control access to specific services on you company LAN. You can also create rules that restrict access to the Internet from the company LAN, allowing you to prevent specific machines on you LAN from accessing the Internet.

3.4 Service mode

LASAT SafePipe must be in Service mode to allow the LASAT SafePipe™ administration tool do the following:

- Reconfigure IP address and subnet mask
- Set administration password
- Update software
- Restore factory settings or restore a previously saved configuration

The service button on the front panel puts the LASAT SafePipe™ in service mode.

3.5 LASAT VPN Client

LASAT VPN Client, the client software for the distance worker, provides fast, efficient remote access to all the network resources which you are authorized to use.

The next sections explain how easy it is to install and configure LASAT Networks' VPN Client Software for Windows.

3.5.1 System Requirements

A PC equipped with a 486/33 CPU or higher with Windows 95, Windows 98, or Windows NT 4.0, Service Pack 4 or 5. CD-ROM drive and a 3.5" diskette station, if loading the configuration file from a diskette. For the Internet connection you need a functional Internet account using TCP/IP.

3.5.2 Configuration

After LASAT SafePipe™ Client is installed on the PC, the program needs to be configured. The same procedure is used for Windows 95, Windows 98 and Windows NT. The configuration file and a file password are needed from the network administrator. The configuration file is best supplied on a 3.5" diskette. It can also be sent as an e-mail attachment, but this is obviously a less secure option. The configuration file is encrypted, but should be treated confidentially. For maximum security, delete it from the hard disk once LASAT SafePipe™ is up and running. Do not duplicate or distribute it to others. The password should NOT be included on the diskette, but in a separate form, e.g. in printed format or verbally.

3.5.3 Connection Password

During the installation a connection password for activating LASAT SoftPipe Client Software needs to be entered. Only the user knows this password. It must consist of at least 5 letters.

3.6 LASAT SafePipe™ Hardware platform

LASAT SafePipe™ has its own dedicated hardware. This means that the hardware has been designed especially for LASAT SafePipe™. The specifications for the platform are as follows:

- 175 Dhrystone Mips RISC CPU
- 8 MB Flash
- 16 MB SDRAM
- Stac LZS compression chip
- EDHC chip
- Dual 10/100 Ethernet LAN ports
- PCI expansion slot

3.6.1 MIPS RISC CPU

The CPU in LASAT SafePipe™ is **Reduced Instruction Set Computer** - a computer architecture that reduces chip complexity by using simpler instructions. RISC compilers have to generate software routines to perform complex instructions that were previously done in hardware by CISC computers. In RISC, the micro code layer and associated overhead is eliminated.

RISC keeps instruction size constant, bans the indirect addressing mode and retains only those instructions that can be overlapped and made to execute in one machine cycle or less. The RISC chip is faster than its CISC counterpart and is designed and built more economically. The RISC machine executes instructions faster because it doesn't have to go through a micro code conversion layer. The RISC compiler does more work than the CISC compiler. It has to generate routines using simpler instructions that would normally be performed by a single, complex instruction in the CISC computer.

A x86 or a Pentium processor is a CISC processor.

MIPS RISC processors is a widely spread processor standard and it comes from MIPS Technologies, Inc. MIPS is a company that designs and licenses high-performance, cost-effective 32- and 64-bit reduced instruction-set computing (RISC) processor intellectual property and core technology to the leading semiconductor suppliers like IDT and NEC. The amount of computer instructions per second is also referred to as MIPS, or Millions of Instructions Per Second. MIPS rates are not uniform. Some are best-case mixes while others are averages. In addition, it takes more instructions in one machine to do the same thing as another (RISC vs. CISC, mainframe vs. micro). MIPS rate is just one factor in overall performance. Bus and channel speed and bandwidth, memory speed, memory management techniques and system software also determine total throughput.

3.6.2 LZS Compression

On the average, data can be compressed with a factor of 2:1. There are a number of different compression algorithms, and the major difference is not the compression factor, which is in general close to 2:1, but rather the speed of the compression. When a product promises a certain bandwidth, this should include compression speed.

Basically, compression takes advantage of repeated patterns in the data to save space. If the compression machinery encounters a piece of data that is identical to an earlier part, it replaces the data. This means that there are no specific requirements for the compression, the important issues are included in the performance data.

LZS developed by Hi/fn is an enhanced compression algorithm. The LZS algorithm is a general-purpose loss less compression algorithm for use with a wide variety of data types. Its encoding method is very efficient, providing compression for strings as short as two octets in length. The LZS algorithm uses a sliding window of 2,048 bytes. During compression, redundant sequences of data are replaced with tokens that represent those sequences. During decompression, the original sequences are substituted for the tokens in such a way that the original data is exactly recovered. LZS differs from lossy compression algorithms, such as those often used for video compression, that do not exactly reproduce the original data. The details of LZS compression can be found in [ANSI94]. The efficiency of the LZS algorithm depends on the degree of redundancy in the original data.

3.6.3 EDHC Technology

EDHC stands for Encryption, Decryption, Hashing and Compression and is LASAT Networks' own chip technology for encryption. The technology enables LASAT Networks to make products that perform encryption and compression at very high speeds. The chip technology is used in LASAT SafePipe™ VPN products.

3.6.4 LAN Interface

The two on-board fast Ethernet ports provide auto-sensing functionality, which means that when a network cable is connected to the adapter, the auto-sensing feature is used for automatic detection of the right network speed (10 or 100 Mbit/s).

Several different standards are used for the fast Ethernet protocol. The most common one is 100Base-Tx, which requires a UTP cat. 5 cable with 8 wires. This is the standard that LASAT SafePipe™ supports. The supported Ethernet protocol variant is 802.3.

3.6.5 WAN interface

LASAT SafePipe™ can be supplied with an optional WAN interface. This makes it possible to use the LASAT SafePipe™ without the need for an external router.

This interface supports a wide range of synchronous serial interface standards including V.35, X.21bis, EIA449 and EIA530A (RS-422/423 connectors).

It meets the following electrical standards: V.10 (RS423), V.11 (RS422), V.28 (RS232), V.35, and the Synchronous Interface can operate in DTE or DCE mode and supports RTS, CTS, DSR, DTR, DCD, RI, LL, RL and TM signals.

The interface operates at speeds up to 10 Mbps, supports HDLC framing and is ideally suited for Frame relay and PPP applications.

These additional interfaces can either be purchased as an out-of-the-box ready product, or as upgrades. When using the Ethernet WAN option, any third party product for the actual WAN connection is possible, e.g. xDSL or cable modems.

4. LASAT SafePipe™ Benefits

- IPsec compatible hardware VPN device
- strong Triple-DES encryption
- high throughput
- firewall
- X.509 digital certificate
- built-in CA
- automatic update routing tables with bridging of other protocols than IP

All this makes it an ideal choice for medium-sized and large enterprises that want to save cost on their international private lines and on communication with their remote workers.

The installation and configuration of the box is done via a standard browser interface, and all in all it's an easy and cost-saving partner for any IT manager.

The LASAT SafePipe™ VPN series is supported by a LASAT SafePipe™ client, a complete IPsec software package for the remote worker. This software is totally transparent to the remote user, and consequently there is no change for the user if he is sitting at the office or working remotely.

5. References

This paper is only an introduction to the VPN techniques. For more information, the following books are both easy to read and quite comprehensive on their subjects:

TCP/IP:

Inside TCP/IP, by Karanjit S. Siyan (New Riders).

Networking in general:

Computer Networks, by Andrew S. Tanenbaum (Prentice Hall).

Encryption, authentication and key generation:

Applied Cryptography, by Bruce Schneier (John Wiley & Sons, Inc).

Firewalls:

Building Internet Firewalls, by D. Brent Chapman and Elizabeth D. Zwicky, (O'Reilly).

SSH IPsec Express by SSH Communications

TechEncyclopedia, CMP's Tech Web.

How LZS works, by Hifn

Data Communication, by CMP