PROJECT NAME:	Architecture.
TITLE:	DSL Security
SOURCE:	Eicon Networks, BellSouth, Arescom
Contact:	Hugues de Fombelle (hugues.defombelle@eicon.com)

ABSTRACT:

This contribution provides an overview of some of the security aspects of DSL-based corporate networks. With the expansion of the Internet and the increasing use of Internet technology inside the organization, more and more computing resources have become connected to networks that can potentially be reached from outside – and inside - the enterprise or individual PC, and from inside the enterprise as well. As connectivity increases, so does the risk of attack on the network or stand-alone PC. In this environment, two factors drive the need for a network security system: the need to maintain the integrity of data communications and the need to protect intellectual property and information assets.

DISTRIBUTION:

DSL Forum Members and all interested parties

NOTICE:

This document has been prepared to provide a Security positioning statement for the DSL Forum. This document is offered as a basis for discussion and is not a binding proposal on EICON Networks, Bell South or ARESCOM. The content of the document may be subject to change in form and/or numerical value after further study. Eicon Networks, Bell South or ARESCOM specifically reserve the right to add to, amend, or withdraw the statements contained herein.

Table of content:

INTRODUCTION

The recent spate of Denial of Service (DoS) attacks on popular web sites have raised concerns about Internet security. This white paper will give an overview about security issues and DSL related solutions.

DSL offers many benefits such as high-speed connections from 10-to-100 times faster than dial-up, simultaneous voice and data over the same phone line and choice of ISP. DSL also provides with an "always-on" connection, which means consumers can maintain their DSL Internet connections 24 hours a day, seven days a week.

DSL is inherently more secure than other broadband communications, namely the majority of the currently deployed cable modem services. DSL is a point-to-point connection between a consumer's home or office location and the telephone company switching office. Cable, on the other hand, is a point-to-multipoint connection that shares network connectivity among homes in a neighborhood, much like a shared LAN. In addition, with DSL each customer has a separate "Private Virtual Circuit," a unique connection that authenticates and secures the communication between the customer's PC and the Internet.

Despite these obvious advantages of DSL, anybody who establishes a dial-up or "always-on" Internet connection incurs some security risk stemming from the duration of the network connection rather than the access method. A number of standard measures are available that users can apply to protect themselves which will be addressed in this paper as well in the "DSL Security for Home and Telecommuter" Whitepaper.

SECURITY - Aspects of Baseline DSL Network Architecture

DSL Forum has an architectural recommendation of Point-to-Point Protocol (PPP) over asynchronous transfer mode (ATM). Many Telco's have deployed this "*standard*" approach where the PPP over ATM connection is from the customer to a service selection gateway. From this gateway, the customer's PPP session is carried over a "*tunnel*" to the ISPs gateway router.

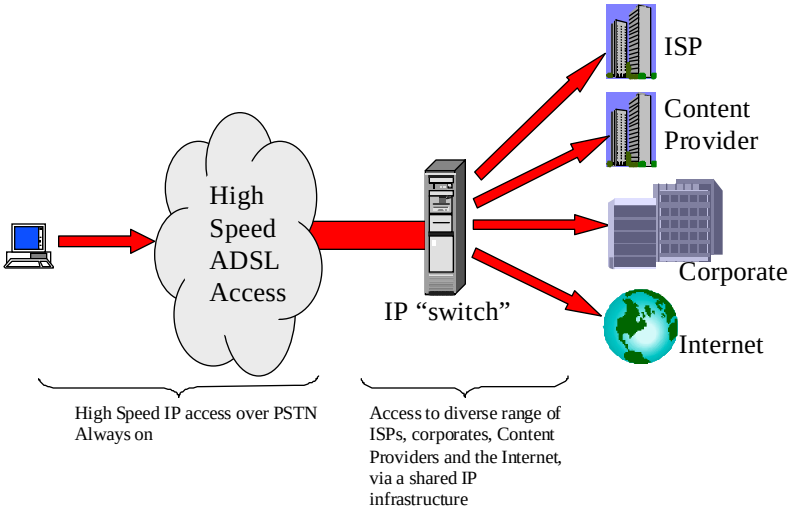

PPP provides authentication using PAP and CHAP. Radius or "AAA" servers at the ISP and/or linked to the gateway can also use such standard ISP authentication procedures. An IP address can be dynamically assigned to the end user from a pool of addresses during a PPP session initiation. The PPP over ATM and tunneling aspects are generally good for security, in the sense that these can provide virtual partitioning of the overall broadband platform.

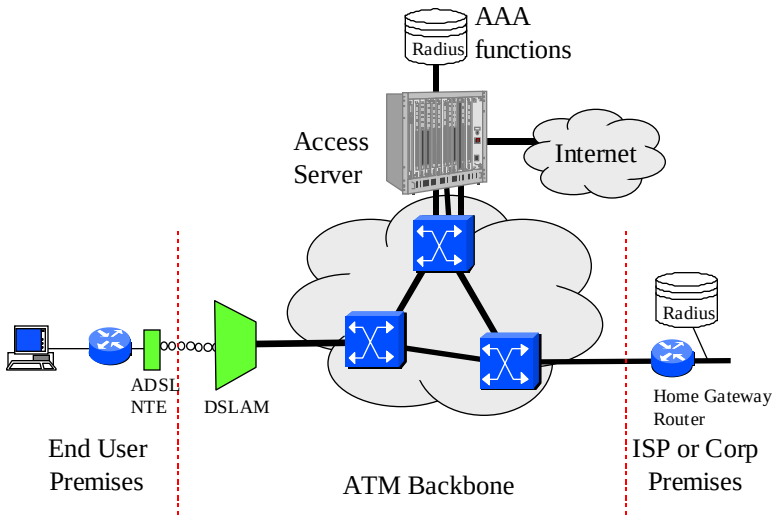

A different architecture being deployed by some service providers is PPP over Ethernet (PPPoE). At first glance this may appear to be less secure because of the use of Ethernet, which implies broadcast capability across a community of users (such as a cable modem provides) and which could result in user addresses being broadcast over the network (e.g. ARP requests) which could then be intercepted. However, in practice no additional real security risk is probable over PPPoA. Since most actual implementations of PPPoE will terminate the bridging session coincident with the PVC termination, no "leakage" occurs from one PVC into another. PPP provides for transient connection so that it is possible to use PPP/DSL without being "always on".

BACKGROUND – The Hacker: Attack Process and Tools

Varied levels of technical expertise make it difficult to defend against a typical hacker.

Generally, a hacker's first step in a break-in attempt is researching the attack target. The hacker's goal is to construct a database of the targeted network and gather information about the hosts attached to the network. The hacker has a number of applications available that can be used to gather information.

- **DNS Servers**. these servers can access a list of host IP addresses and their corresponding hosts
- **Finger Protocol**: this protocol reveals detailed information about the users (login names, phone numbers, etc.) of a specifically targeted host
- **Ping Program:** this is employed to locate a particular host to determine its accessibility
- **SNMP Protocol**: this can be used to examine the routing table of an unsecured router to learn intimate details about the network's topology
- **Trace Route**: this program reveals intermediate network numbers and routers in the attack path to a specifically targeted host
- **WHOIS Protocol:** his information service provides data about all DNS domains and the system administrators responsible for each domain

These tools and this research together will arm the hacker with the information necessary to accurately devise a path to the targeted information.

After all of the pertinent information about the targeted network is gathered, the hacker attempts the first attack on the network's security system to probe for weaknesses. At this stage of the attack, utilizing a series of tools to allow for automatic scanning of the host under attack is common. A typical attack is as follows:

- Establishing vulnerability: Since the list of unknown network vulnerabilities is short, a successful hacker will design a small computer program that will try to connect to specific network access ports of a targeted host. This program will tell the hacker which hosts on the network are vulnerable to attack.
- Tools of the trade: The hacker can now use one or more software tools that are readily available such as the Internet Security Scanner (ISS) or the Security Analysis Tool for Auditing (SATAN) that can scan the entire domain or sub-network looking for holes. These programs will determine the weakness of each system with respect to several vulnerabilities. The hacker will then use this information to gain access.

Security Policies for Businesses

Detecting Signs of Intrusion

Policies: Specify the level of verification that is required when examining each class of data and service provided by the organization.

Look for unexpected changes to directories and files.

Policies: Define the responsibilities and authority of systems administrators and security personnel to examine file systems on a regular basis for unexpected changes. Users should be told about such authority and examination. Require users to report any unexpected changes to their software and data files to system administrators or your organization's designated security point of contact.

Inspect your system and network logs.

Policies: Specify that log files be inspected on a regular basis by authorized personnel, and that anomalies be recorded and reported to your organization's designated security point of contact.

Review notifications from system and network monitoring mechanisms.

Policies: Require notification of users regarding the process inspections that will be done.
Specify the responsibilities and authority of designated systems administrators and security personnel to examine processes for unexpected behavior.
Specify what forms of unexpected behavior users should watch for. Require users to report any such behavior to their designated security officials and system administrators.
Specify what software and data users and administrators are permitted to install, collect, and use, with explicit procedures and conditions for doing so.
Specify what programs users and administrators are permitted to execute and under which conditions.

Look for signs of unauthorized access to physical resources.

Policies: Require the tagging and inventory of all physical computing resources. Specify how to respond when a physical intrusion has been detected.

Configure computers for user authentication.

Policies: Your organization's policy for networked systems should describe under what conditions an account is created and deleted. This should include what account actions are taken (disabled, deleted, transferred) and how files are handled when an employee, contractor, or vendor who has an account on your systems no longer works for your organization.
Require appropriate authentication of all users on all computers that can access information resources; this includes authenticating users of network services hosted by your servers include an appropriate password policy prohibit users from recording and storing passwords in places that could be discovered by intruders.
Your organization's acceptable use policy for workstations should require that users shut down or lock their unattended workstations.
When writing a password policy, remember that requiring users to have complex passwords may have the undesired result of the user's writing the passwords on paper that they keep near the computer (often stuck to the machine) or with personal papers (in a wallet, purse, or briefcase).
If that paper is observed, lost, or stolen, it creates a potential vulnerability.
If a password policy is especially difficult to follow, it creates in users a desire to find ways around it. This attitude can negatively influence users' compliance with other aspects of security policies.

Investigate unauthorized hardware attached to your organization's network.

Policies: Require the maintenance of documented hardware inventories.
Require the maintenance of a documented network topology.
Specify the authority and responsibility of designated security personnel to perform physical audits of installed hardware and software.

Specify what kinds of hardware and software users are permitted to install themselves on their desktop machines.

Review reports by users and external contacts about suspicious system and network events and behavior.

Policies: Require users to report any unexpected or suspicious system behavior immediately to their designated security officials and system administrators.
Require users to report any physical intrusions to networked systems or offline data storage facilities immediately to their designated security officials and system administrators.
Require system administrators to investigate each reported suspicious activity to determine whether it represents an intrusion.
Require system administrators to notify users in advance of any changes that will be made to the systems they use, including software configurations, data storage and access, and revised procedures for using systems as a result of the changes.

Offer only essential network services and operating system services on the server host machine.

Policy: Your organization's security policy for networked systems should require that network servers be configured to offer only essential services.

Eliminate all means of intruder access.

Policy: We recommend that your organization's purchasing guidelines mandate the specification of security requirements for all information technologies.

Develop a computer deployment plan that includes security issues.

Policies: Require that a detailed computer deployment plan be developed and followed whenever computers are being deployed (or redeployed).
Require that access to your deployment plan be permitted to only those who require the information to perform their jobs.

Keep operating systems and applications software up to date.

Policy: Your organization's security policy for networked systems should require that systems administrators install necessary security-related software updates in a timely manner.

Protect computers from viruses and similar programmed threats.

Policies: Your organization's workstation acceptable use policy or security policy for networked systems should define users' authority (or lack thereof) to download and/or install software on the computer.
Specify who has the responsibility to scan for viruses and eradicate them -- users or system administrators.
Prohibit users from running executable files that they have received as email attachments or downloaded from un-trusted sites.

Design the firewall system.

Policies: Your organization's networked systems security policy should include the risks you intend to manage with the firewall.
The services you intend to offer to untrusted networks from your protected network: These could be offerings to the Internet or to other internal networks.
The services you intend to request from untrusted networks via your protected network: These could be requests to the Internet or to other internal networks.
The objective that all incoming and outgoing network traffic must go through the firewall (i.e., that no traffic which bypasses the firewall is permitted, for example, by using modems) — or conversely, that specific loopholes are permitted and under what conditions (e.g., modems, tunnels, connections to ISPs).
In the offering and requesting of services, your policy should ensure that you only allow network traffic that is determined to be safe and in your interests that minimizes the exposure of information about your protected network's information infrastructure.

Acquire firewall hardware and software.

Policies: Ensure that you have all hardware and software components available before attempting firewall system deployment. Conduct a preliminary installation of the firewall software and operating system on the target hardware to ensure that nothing is missing. It is particularly important that you do this upon receipt of hardware and software if your deployment is delayed. If something is missing, you have time to correct the omission before deployment deadlines. If you skip this step, you may not realize the omission until much later. Plan to do this type of preliminary installation many times. The more comfortable you are with the installation process, the more quickly you can perform major reconfigurations or recoveries. If your firewall operating system (OS) resides in nonvolatile memory (e.g., flash memory), make sure that you can erase its contents completely and rewrite the OS image onto the hardware. Do this for both your primary and all spare OS hardware. This will ensure that your OS hardware works correctly and that you can load a new version of the operating system once the firewall system is deployed. If you have limited experience with the target hardware or operating system, bring in a knowledgeable consultant or vendor. Document your understanding, the actions that they take, and the recommendations that they make. Have the consultant/vendor sign this document in the event you encounter problems in the future. It may give you some leverage to have them return without incurring additional expense.

Configure IP routing.

Policies: Require that configuring IP routing for your firewall system is performed in an environment isolated from your operational networks. Specify what connectivity is to be permitted with the specific statement that all other connectivity is denied

Configure firewall packet filtering.

Policies: It is rarely the case that packet filtering can implement an organization's security policy exactly. Your managers must decide what level of accuracy and precision is required in implementing your security policy. It will almost certainly be the case that you will have to implement rules that are either more permissive or more restrictive than your policy. You need to determine how to handle this discrepancy. That all network traffic that is not explicitly permitted should, by default, be denied. That configuring packet filtering for your firewall system is performed in an environment isolated from your operational networks.

SECURITY REQUIREMENTS

Maintaining security is a continuous challenge. Just when a user thinks an airtight system is in place, a new hacker technology or an especially diabolical adversary enters the picture. Regardless of the type or location of a perceived threat, an effective system for securing the integrity of information while maintaining availability of information assets must:

- **Allow access to information by authorized parties only**
- **Implement policies determining who is authorized for what level of access to which information**
- **Employ a strong user authentication system**
- **Deny malicious or destructive access to any information assets**
- **Protect data from end to end**

THREATS TO SECURITY

A computer networking system can be attacked in a number of ways, resulting in differing degrees of damage. These attacks can take several forms:

- **Denial of service**: the attacker disrupts the smooth flow of information by crashing or overloading a critical device such as a server, router or firewall. This is an attack on the availability of information.

- **Theft of information:** the attacker acquires information that is proprietary to the organization. This can be done by eavesdropping, by masquerading as an authorized entity or by a brute-force attack such as the use of a computer program that guesses passwords. This is an attack on the ownership of information and intellectual property.

- **Corruption of data**: the attacker either destroys or corrupts data stored on disk or corrupts data as it is transmitted across the network. This is an attack on the integrity of information.

Establishing adequate or even impenetrable security at one point of attack while leaving one or more of these other points uncovered is like posting a guard at the front desk and leaving the company's doors and windows wide open. An employee seeking revenge or a serious thief will try every avenue of entry, particularly if the value of the information is great and the access is relatively easy.

ELEMENTS OF A SECURITY SOLUTION

A security solution that maximizes the benefits of networked data communications must contain these elements:

- User authentication
- Access control
- Encryption
- Management

An enterprise may employ any or all of these elements to achieve integrity and access control. The best strategy depends on the risk involved, the cost of the deployment and the cost of a security breach or lost data. The following sections more closely scrutinize each element in a total security solution.

USER AUTHENTICATION

Proof of identity is an essential component of any security system. It's the only way to differentiate authorized users from intruders. User authentication to the network is a necessity for any enterprise that is serious about protecting information assets and knowing who or what is attempting to gain access to the network. Authentication becomes particularly important when some of the more sophisticated communication methods are used.

In addition to proving identity, authentication systems are used to determine what information the requestor can access—for example, a human resources database or corporate financial database. True authentication generally in Corporations two or three of the following elements:

- Possessions (smart card, certificate)
- Information (password)
- Physical attributes (fingerprint or other biometric information)

Authentication is most often achieved through challenge and response, digital certificates, or message digests and digital signatures.

- **Challenge and response:** this authentication method uses a software agent within a database system or a workgroup server, that presents the person requesting access to a resource with a challenge, most often a username and password. This is the most common form of security and one that is easily broken when passwords are not carefully chosen (at least 8 characters) and maintained. Does changing passwords frequently also help here?

- **Digital certificates**: One of the earliest uses of digital certificate technology was Privacy Enhanced Mail, the predecessor to S/MIME (Secure/Multipurpose Internet Mail Extensions), a widely used specification that brought a higher level of security to e-mail through encryption and digital signature-based authentication. Since the introduction, the use of digital certificates has continued to grow steadily.

Digital certificates are essential components of a public key infrastructure (PKI), which can be generally defined as a security system that consists of protocols, services and standards that support applications of public-key cryptography.

- **Digital Signatures**: public key cryptography is used to validate messages that have been digitally signed. Such messages can be simple e-mail or part of a protocol for establishing a secure communication's session. The sender of the message to be authenticated digitally signs the message using a private key. The signature can be validated using the sender's corresponding public key, which is contained in the sender's certificate and can either be sent along with the message or retrieved from a certificate repository.

  The association between the sender's identity and the sender's public key can be authenticated through a digital certificate issued by a trusted certificate authority (CA). The CA certificate is issued in advance to all parties, and its public key can be used to authenticate the public key in the sender's certificate. When the sender's public key has been validated, it can be used to authenticate the digital signature of the message itself. Since the CA certificate is already available to both the sender and receiver, this method can be used to authenticate messages in either direction without contacting a third party.

  To implement a secure certificate or signature system, the following conditions must be met:

  A certificate authority service provider or software package must issue a certificate to all potential senders and receivers. The receiver must be able to use the CA certificate to verify the sender's public key. The sender's authenticated public key must then be used to verify the digital signature of the message itself.

  Although a digital certificate system can affect the performance of heavily used servers, this is usually not the case. Typically, the certificate itself is provided by the client, in which case the authenticator does not need to perform any server access. Moreover, the value of preventing a security breach often far out-weighs the inconvenience of slightly delayed access.

- **Message digests (and digital signatures)**: applying a one-way hash function such as MD5 or SHA-1 to a message creates message digests. "One-way" means that the original message cannot be recreated from the digest. A digital signature uses the private key of an individual to encrypt the message digest. At the receiving end, the digest is recreated from the message text, the public key is used to decrypt the digest from the digital signature, and the two message digests are compared. If they match, the messages are in all probability the same. Comparison of the message digests provides both a means of authenticating the signature and checking message integrity.

- **Biometrics:** Biometric security products can't guarantee 100 percent authentication -- nothing can, experts say. But biometric tool kits are being offered on the market today that will tighten precautions a few notches further. A layered biometrics verification Internet tool kit can allow software providers to add biometric voice and fingerprint authentication to traditional security applications that protect Internet servers.

  Value-added resellers (VARs) who deal with e-commerce applications requiring high levels of security, should consider including a public-key infrastructure (PKI) to increase the security. Layered security levels can provide the most value for companies needing high levels of security.

  One area where a tool kit could be used is for enhancing security for Internet banking. A bank, contracting with an application service provider (ASP), could require biometric verification for a high-value transaction over the Internet.

  A vendor seeking to wire money using the Internet would go to the bank's Web page, fill out the required information and submit the transaction. If the transaction is for a high value, the bank would decide it needs biometric verification and automatically send a message to a security server requesting that the vendor speak to a customer service agent.

Biometrics can provide solutions in the following areas:

- Encryption systems training at customer facilities
- E-mail encryption and authentication software
- Web security and E-commerce
- Secure EDI systems
- Automated tools for cryptanalytic testing of client products and traffic
- Digital signatures and certificate authority systems
- Virtual Private Network solutions (VPNs)
- Telephone end-to-end encryption systems
- Radio and wireless encryption technologies
- Telephone counter-fraud measures
- Electronic countermeasures
- Audio interception countermeasures
- Cryptography product(s) certification
- Network security solutions
- Steganographic countermeasures
- Fingerprint, face, voice, signature biometrics

ACCESS CONTROL

Access control governs a user's ability to make a connection to a particular network, computer or application, or to a specific kind of data traffic. Access control systems are generally implemented using firewalls, which provide a centralized point from which to permit or deny access.

- **Internet firewalls:** an Internet firewall is a system or group of systems that enforces a security policy between an organization's network and the Internet. The firewall determines which inside services may be accessed from the outside, which outsiders are permitted access to the inside services.

  For a firewall to be effective, all traffic to and from the Internet must pass through the firewall, where it can be inspected. The firewall must permit only authorized traffic to pass and the firewall itself must be immune to penetration. Unfortunately, a firewall system cannot offer any protection once a hacker has gotten through or around the firewall.

  It is important to note that an Internet firewall is not just a router, a bastion host or a combination of devices that provides security for the network. The firewall is part of an overall security policy that creates a perimeter defense designed to protect the network.

  The benefits of an Internet firewall stem from its manages access between the Internet and an organization's private network. Without a firewall, each host system on the private network is exposed to attacks from other hosts on the Internet.

  An Internet firewall cannot protect against attacks that do not go through the firewall. For example, if an unrestricted dial-out is allowed from inside the protected network, internal users can make a direct SLIP or PPP connection to the Internet.

- **Firewalls**: these can be physical devices or software agents filter packets heading into or out of an organization based on a set of policy rules. They can allow access to an enterprise network by username/password, type of service requested (ftp, http, telnet), location of destination (network or computer), or location of requestor (network address). A firewall can request authentication before allowing any traffic to pass at all, and in so doing can take advantage of the various authentication schemes available. There are two distinct types of firewalls as well as some hybrids that don't fit neatly into any category. The difference between firewall types is primarily related to how they handle external traffic.

- **Packet-filtering firewalls**: this type of firewall controls access and data into and out of the network. Packet-filtering firewalls can simply be routers or switches that are configured with access lists. They can permit or deny access based on the protocol, source or destination port, and source and destination of IP addresses. Moreover, for a higher level of security, they can be configured to allow TCP

communications only when initiated from the internal network. Packet-filtering firewalls typically do not employ any kind of user authentication, because the environments in which they are usually deployed handle levels of traffic that are too high to allow for it.

- **Application/proxy firewalls**: these devices or software agents handle requests in place of the network application or server they are safeguarding. They provide network resources inside the firewall with a layer of protection, ensuring that secure resources are never accessed directly. Application/proxy firewalls typically support local caching of Web content and address translation, thereby hiding internal IP addresses from Internet surfers. Proxy firewalls can allow access based on source address, destination address or an identity (authentication).

- The benefits of a firewall: the key firewall feature is IP address management (i.e., IP filtering) that provides a detailed level of control over inbound and outbound traffic on all interfaces. The firewall can be logically placed either inside or outside of NAT. With IP filtering control on each interface you can specify exactly which traffic you want to allow or deny. Traffic can be controlled based on IP address, port value, syn and ack flags and protocol type of both source and destination IP addresses.

  You can block specific users from accessing the Internet while allowing others to browse freely. Allow access to a web server on your LAN without exposing your mail server or other services running on the same machine. Allow locally initiated connections to be replied to while blocking remotely initiated connections

Numerous vendors offer many firewall products at a wide range of prices. In many cases these products also include a Web server, thus creating a turnkey system. Any organization that has an open connection to the Internet should deploy one or more of these devices to ensure adequate security.

In addition to firewalls, an enterprise may implement a network-level security protocol, such as Internet Protocol Security (IPSec), to protect information as it moves through a network. IPSec works at the packet level. Every packet is protected to provide authentication, integrity and (optionally) confidentiality. Any of the firewall devices described above, as well as each of the individual computers or servers on a network, can implement IPSec. IPSec is discussed in greater detail later in this paper.

ENCRYPTION

Even if both access control and authentication security systems are completely effective, the enterprise can still be at risk when data communications travel over a third-party network such as the Internet. Indeed, the low cost and ease of connecting to the Internet have made it an extremely attractive medium for communication within and between enterprises.

Encryption is used to protect against eavesdropping. It renders information private by making it unreadable to all except those who have the key needed to decrypt the data. It does not matter whether a third party intercepts packets over the Internet; the data still cannot be read. This approach can be used throughout the enterprise network, including within the enterprise (intranet), between enterprises (extranet) or over the public Internet to carry private data in a virtual private network (VPN).

The degree of protection afforded by encryption depends upon the strength of the encryption algorithm. Against brute-force attacks, that strength is determined by the number of possible keys, which in turn is defined by the key size.

A recent brute-force attack was able to try 245 billion keys per second. With this type of computing power, an intruder could try all possible 56-bit keys in 81 hours, finding the key in an average of 40 hours. However, with 112-bit key and the ability to try 245 billion keys per second, it would take an average of 336 trillion years to discover the key. A Triple Data Encryption Standard system (3DES) uses either 112-bit or 168-bit keys.

Encryption systems in common use today include the following:

- **Shared key encryption**: both or all parties possess a previously distributed key that locks and unlocks the data. The sender provides the key to a shared symmetrical encryption algorithm to encode the data before placing it in a packet bound for the remote site; the remote site then provides the key to the same encryption algorithm to decode the data. Shared key encryption systems use DES, 3DES, RC5, IDEA,

and other algorithms that are extremely fast. The system's strength lies in the length of the key and resistance to analyzing encrypted data. The system's weakness is that if the key becomes known, anyone can decrypt the cipher text.

- **Public key encryption**: one party possesses a private unlocking key and makes a public locking key. Any sender can use the public key to encrypt the communication; the receiver then uses its corresponding private key to decrypt the data. Directory servers from Novell, Netscape and others can store a digital certificate that contains a user's public key. This system can also be combined with data exchanged at the time of communication to arrive at a shared, session-specific secret key.

  The public key/private key system can also be used to create a digital signature, which is a digest of a plain message encrypted using a key and appended to the plain message. This digest makes it possible to authenticate the sender of the message and verify that the message has not been altered. Public key encryption is very CPU-intensive. It is typically used for small amounts of data where strong security is required.

- **Secure key exchange**: Both parties first authenticate themselves (often using digital certificates) during a session-specific encryption key distribution process. The session key is created based on data generated by both parties at the time of communication. This key can then be used to encrypt and decrypt all other communications.

All encryption systems place an additional load on the network because one or more round trips are needed to authenticate the parties. The machines involved in the communication must also perform large mathematical operations to encrypt and decrypt data, and this can amount to a noticeable increase of CPU cycles on systems that pass many packets. To free the CPU from this task, the conventional burden of encryption systems can be moved to firmware or hardware, such as a coprocessor on the network interface card (NIC) or elsewhere in an embedded system.

SECURITY MANAGEMENT

A security system should allow for oversight and control by a human authority. Any system that uses authentication requires some central authority to verify those identities, whether it be the /etc/password file on a UNIX host, a Windows NT domain controller, or a Novell Directory Services (NDS) server. The ability to see histories, such as repeated failed attempts to breach a firewall, can provide invaluable information to those charged with protecting information assets. Some of the more recent security specifications, such as IPSec, require the presence of a database containing policy rules. All these elements must be managed for the system to work correctly. However, management consoles or functions themselves represent another potential point of failure of a security system. It is therefore important to ensure that these systems are physically secured and that authentication is in place for any logon to a management console.

INTERNET PROTOCOL SECURITY

As the Internet becomes more critical to organizations and enterprises of all sizes, the need has grown to protect intellectual property and at the same time conduct business has grown. To promote security for business communications, the Internet Engineering Task Force (IETF) developed Internet Protocol Security. IPSec offers standards-based, consistent security for IP networks.

In an IPSec communication, the two communicating entities (which can be individual hosts or intervening devices, such as routers or firewalls) first establish a Security Association (SA). During negotiation of the SA, the two entities agree on what kind of security will be employed.

The Internet Security Architecture document (RFC 2401) specifies two major traffic security protocols: Authentication Header (AH) and Encapsulating Security Payload (ESP).

- **Authentication Header:** every IP packet sent during the life of the SA contains an Authentication Header in addition to standard IP headers. This AH verifies that the packet received is identical to the one the sender sent. It contains a Message Authentication Code (MAC), which authenticates the sender of the packet and ensures that the contents of the packet have not been altered. Two techniques

are used to create the MAC, both involving message digests and a shared secret (established in a separate security session):

(1) ***Keyed MAC (KMAC)***: the sender encrypts the digest with the shared secret, and the receiver decrypts it to check for validity; or
(2) ***Hashed MAC (HMAC)***: the sender creates the digest over the combination of the shared secret and the message text; the receiver does the same and compares the two digests for validity.

- **Encapsulating Security Payload (ESP):** in the ESP method, the actual payload data of a packet is encrypted using a session key. This key is often derived for the session itself using a public key encryption system. First, the entity on the other end of the key exchange is authenticated, then a key derived from the agreed-upon parameters is transmitted. In addition to encrypting the data portion of the message, ESP can also be used to provide authentication and integrity—features provided by AH. However, ESP security is less comprehensive than that provided by AH. The AH header protects entire messages, whereas ESP only covers the portion of the message after the ESP header.

  A single SA specifies one of the two types: AH or ESP. A single communication may employ both AH and ESP by creating two SAs. The calculations required of the computers involved in an IPSec communication depend on the system of security and the level of encryption employed.

- **A Security Policy Database (SPD)**: keeps track of the kinds of security, encryption, and authentication that a particular enterprise can implement, and also keeps track of the active Security Associations. This makes it possible to monitor IPSec activity across the network, and to manage the security systems employed at any given site.

- **IPSec:**

IPSec offers a complete and integrated system for securing data networks. IPSec can be used within the organization or on the Internet because it is based on a set of open specifications including the entire TCP/IP protocol suite. Also, like TCP/IP (and unlike proprietary schemes), IPSec is designed for interoperability between enterprise systems.

Network system designers can integrate IPSec into an existing system in three ways:

By integrating IPSec processing into the TCP/IP network stack of the host or other device. This requires the host CPU to do security processing.

By performing IPSec processing in software before the data packets are processed by the existing TCP/IP networking stack (known as a bump in the stack, or BITS). This approach also requires the host CPU to do security processing.

By performing IPSec processing before the data packets are processed by the host computer (known as bump in the wire, or BITW). This system offloads security processing to a processor on a network component, such as a NIC with an on-board encryption chip, and leaves the CPU free.

The use of an additional processor (or BITW) to handle IPSec security tasks promises the greatest throughput while still delivering the full benefit of a comprehensive security system.

- Viruses and active code:

to be supplied

- Virus protection policies

to be supplied

- CPE solutions

to be supplied

NETWORK DESIGN OPTIONS – Enhanced Security Options

1. One step to improve security is to encrypt the passwords when using PPP authentication via CHAP to avoid clear text being snooped on by others running “LAN Watch” type software on the same local LAN.

2. An additional feature that many DSL service providers are using is Network Address Translation (NAT) on the customer's DSL router. This hides the customer PC's IP address from the outside world to prevents the customer from receiving any unsolicited incoming communication such as a PC Operating System bombs. Unfortunately this also means that NAT has a number of downsides. The customer’s router can’t now forward an unsolicited IP traffic stream to the end user PC. Hence users would then have problems with certain applications such as VoIP, NetMeeting, games, ICQ (community chat). etc. In addition, applications which use IP ephemeral ports could clash with the customer’s router. Applications that bury the IP source address in the IP packet payload could also be problematic. NAT also can complicate service provisioning and trouble shooting.

3. With the service selection gateway architecture, a customer for a broadband DSL ISP service has to establish a session before passing traffic. It is possible to configure the service selection gateway such that as a default the customer is only permitted to access pre-specified ISP s. A customer can then only complete session establishment if they are authorized by the Radius server of the ISP/Corporate in question. The Telco can also record all session establishment and clear down activity, and has the capability to rapidly bar service to any end user indulging in unwarranted behavior.

4. Another possibility is for the telco to offer ISPs/Corporations the option of absolute session length timeouts or idle timeouts. This will (optionally) change the connectivity model from "always on" for those ISPs/Corporations who are particularly worried about security.

5. In addition to the features of the DSL network, the customer could also use encryption of their data or the new IPSec protocol both of which impact throughput performance.

6. An appropriate solution may well be to put firewall technology in place to prevent unwanted intrusions, either in the CPE (which will impact ease of provisioning and troubleshooting but could be the customer’s choice) or in the network (as a value add service). Since each PVC is a separate and distinct connection to the individual end-user, service providers could provide a firewall as a value-added service to the customer based on this individual and distinct connection. Filtering of certain IP traffic (e.g., excessive IP pings/ICMP echo-requests to the customer) could take place in the service selection gateway router and/or the ISPs Home Gateway router. This approach would be a little more difficult to employ on a cable modem network since the firewall would have to be at the head-end and in each customers cable modem (to prevent attack by others on the same cable which is like a large LAN).

7. There are a number of CPE firewall boxes coming onto the market now and being used in conjunction with DSL services in some regions. This may be appropriate for the business market but may not be taken up so rapidly in the more cost sensitive residential market (apart from by the more expert or paranoid users). An alternative for a more technically competent user would be to configure firewall type functionality into their own DSL router. Residential gateways with firewall capability are now available with “pin holes” to overcome some of the application problems that could arise as described for NAT above. Many recent developments in customer DSL routers are focused on auto-provisioning and policy management in the customers box by the service provider. This could enable service providers to sell and manage security options in this equipment for those customers willing to allow their service provider access their box.

CONCLUSIONS

As enterprise and individual resources are connected to a larger global network, the implementation of an effective security system becomes imperative. The security system should provide the following functionality:

- Authenticate users and messages
- Control access to resources based on identity
- Provide protection at all points of entry
- Protect the integrity of data and intellectual property
- Employ encryption systems that cannot easily be broken
- Impose the least possible burden on existing systems
- Interoperate with business partner systems

The IPSec security system offers all of these protections in one integrated, standards-based package. By using coprocessors or other firmware solutions to process security information, the IPSec system can leave existing systems unburdened while still protected.

Other sources:

a. www.icsa.net - The International Computer Security Association
b. www.cert.org - The Computer Emergency Response Team
c. www.sans.org - System Administration, Networking, and Security Institute
d. packetstorm.securify.com Largest Internet security tool database

ACRONYMS & ABBRVIATIONS

3DES	Triple Data Encryption Standard
AAA Server	
ADSL	Asymmetric Digital Subscriber line
AH	Authentication Header
ATM	Asynchronous transfer mode
BITS	Bump in the stack
BITW	Bump in the wire
CA	Certificate authority
CHAP	Challenge authentication protocol
CPU	Central processing unit
DES	Data Encryption Standard
DSL	Digital Subscriber line
ESP	Encapsulating Security Payload
HMAC	Hashed Message Authentication Code
IDEA	Internet Development and Exchange Association
IETF	Internet Engineering Task Force
IPSec	Internet Protocol Security
ISP	Internet service provider
KMAC	Keyed Message Authentication Code
MAC	Message Authentication Code
NIC	Network interface card
PKI	Public key infrastructure
PAP	Password authentication protocol
PPP	Point-to-Point Protocol
SA	Security Association
SLIP	Serial Line Internet Protocol
SPD	Security Policy Database
SNMP	Simple Network Management Protocol
TCP/IP	Transmission Control Protocol/Internet Protocol
VPN	Virtual private network
xDSL	[any] Digital Subscriber line