

Infrastructure of Electronic Information Management

Gregory D. Twitchell and Michael T. Frame

U.S. Geological Survey
12201 Sunrise Valley Drive / MS 302, Reston, Virginia 22092
USA

Gregory_Twitchell@usgs.gov, mike_frame@usgs.gov

ABSTRACT

The information technology infrastructure of an organization, whether it is a private, non-profit, federal, or academic institution, is key to delivering timely and high-quality products and services to its customers and stakeholders. With the evolution of the Internet and the World Wide Web, resources that were once “centralized” in nature are now distributed across the organization in various locations and often remote regions of the country. This presents tremendous challenges to the information technology managers, users, and CEOs of large world-wide corporations who wish to exchange information or get access to resources in today’s global marketplace. Several tools and technologies have been developed over recent years that play critical roles in ensuring that the proper information infrastructure exists within the organization to facilitate this global information marketplace. Such tools and technologies as JAVA, Proxy Servers, Virtual Private Networks (VPN), multi-platform database management solutions, high-speed telecommunication technologies (ATM, ISDN, etc.), mass storage devices, and firewall technologies most often determine the organization’s success through effective and efficient information infrastructure practices. This session will address several of these technologies and provide options related to those that may exist and can be readily applied within Eastern Europe.

1.0 NETWORK INFRASTRUCTURE

There is a need in today’s environment to provide high performance, scalable, and robust systems to almost all businesses, regardless of size. A strong network infrastructure is the key to providing reliable systems with minimal downtime for mission-critical applications. A major component of insuring data integrity is the provision of network security. Gone are the days of restricting physical access and implementing password protection to insure data integrity. That is just a small component in the overall scheme. With the explosion of the Internet, access attacks on networks can occur from anywhere and at anytime. With that type of exposure it is critical that networks are protected not only from hardware and software failures, but also from cyber attacks. The following will provide a general overview of network infrastructure.

2.0 MASS STORAGE DEVICES SYSTEM FAULT TOLERANCE

2.1 Random Array of Independent Drives disk system (RAID)

The Random Array of Independent Drives disk (RAID) system has the capability to protect data and remain on-line with data access despite a single disk failure (RAID storage systems with two concurrent disk failures can continue to operate with a hot disk standby). Once the failed drive(s) is replaced, the system will rebuild the hard drive while remaining online. RAID system can be supported in both hardware or software configuration.

Paper presented at the RTO IMC Lecture Series on “Electronic Information Management”, held in Sofia, Bulgaria, 8-10 September 2004, and published in RTO-EN-IMC-002.

2.2 Disk Mirroring

This is the most elementary type of disk system that provides for system fault tolerance. It requires two hard disk drives. These two drives are duplicates of each other. If there is a failure of one of the drives, the disk system will continue to operate and provide on-line data access. When the failed drive is replaced the system will provide for an on-line reconstruction of the new drive by copying the entire contents of the operational drive to the newly installed replacement drive.

2.3 Extended Data Availability and Protection (EDAP)

Extended Data Availability and Protection (EDAP) is a storage system that provides data protection and access to data, despite failures within the disk system, or, any attached systems, or environmental failures. RAID is considered the lowest level of EDAP.

EDAP techniques for disk are designated as RAID Levels 1-5. Mirroring and Parity RAID are two types that provide various levels of EDAP for disks. The disadvantage of mirroring is that it requires 100% redundancy, while Parity RAID read performance is enhanced. Impact on write performance for Parity RAID is modest. Additionally when failure occurs a higher percentage of a mirrored group may fail in relation to Parity RAID.

RAID levels 3, 4, and 5 are usually identified as Parity RAID. In the RAID Level 3 array, all the disks operate in parallel. RAID Level 3 is useful for high bandwidth applications while RAID Levels 4 and 5 are more suitable for high transaction rate applications. In comparison to the 100% redundancy required by mirroring, RAID levels 3-5 requires only 10% to 33%. The differences in the RAID levels 3-5 are related to how the data and the redundant data are mapped to the disk. EDAP operates in three states:

- Normal (Protected) – EDAP capability is not being employed to counteract a failure.
- Reduced (Vulnerable) – EDAP is in use to counteract a failure of one disk.
- Down – A state in which data cannot be stored or retrieved.

In a normal state, EDAP is running at peak performance, providing on-line storage and retrieval of data. Performance is affected during the reduced state due to regeneration of data from the failed disk and rebuilding of the failed disk. This regeneration of data requires additional processing time and impacts the overall I/O performance. EDAP design is critical to minimize impact for the performance level during the storage system's reduced state.

In addition to maintaining on-line access to reliable data during a failure, EDAP can minimize the time period that the storage system is in a reduced state by supporting on-line sparing of disks. This will facilitate the reconstruction of the reliable data immediately upon a disk failure. EDAP can also minimize the period in which a storage system is in a down state by providing hot swapping of failed disks. This would make it possible to replace disks without powering down the system, therefore putting them in a down state.

2.4 RAID System Fault Tolerance Systems

Protection of computer data from anomalies, such as human errors, hardware failures, or environmental conditions has been a top priority since the development of the first computer system. A strong data backup strategy is a primary and essential component to insure data protection. Advances in hardware reliability have minimized loss of data due to hardware failures and environmental conditions, but minimizing human failures has been more difficult.

Historically, “mirroring” the data was first level of protection, in addition to maintaining the immediate access to data in the event of a failure of one of the disks in a mirrored pair. Implementation of a mirrored system was twice as expensive as a non-mirrored system. The need to develop a less expensive system to protect data and on-line access led to the development of RAID (Redundant Array of Independent Disk Drives).

Early Parity RAID disk systems suffered from poor performance due primarily to the need to generate and write parity during a write operation and to regenerate the data on-line in response to I/O requests for data from a failed disk. This performance problem was addressed by the use of caching and write-assist disks. Different RAID levels basically describe how data and redundant data are mapped across the disks of an array.

Parity RAID requires a minimum of three physical disks. One disk is dedicated to parity, a second to the first data set, and a third disk to the second data set. Data and parity can be configured to map to disks in a manner so that one disk is not totally dedicated to parity.

2.5 Redundancy Performance

Mirroring or Parity RAID that provides redundancy against a disk failure also provides performance advantages. Multiple disks in an array can be used in parallel for applications requiring high transfer rates (bandwidth) or independently for applications requiring high transaction rates (asynchronous I/O requests involving relatively small amounts of data for each I/O task).

Controller, device channeling and redundancy in disk may enhance disk system performance. Performance can be improved by having all components in a redundancy group actively on-line to share I/O tasks. If any one component fails, the performance will degrade until the defective component is replaced.

Table 1: RAID Level Summary

RAID Level	Description	Data Reliability (protection)	Data Transfer Rates	I/O Rates
0	Striping of data across multiple drives in an array. This is a high performance solution, however there is no data protection.	No data protection	Very High	Very high for both reads and writes
1	Raid 1 is known as mirroring. Mirroring is the 100% duplication of data from one disk to another. This is a high availability solution, but due to the 100% duplication, it is a costly solution.	Excellent reliability	Reads are higher than a single drive. Writes are about the same as a single drive.	Reads are up to two times faster than a single drive. Writes are about the same as a single drive.
5	This is the most widely deployed RAID level. This level provides a balance between performance and cost. Striping with parity. Data and parity information is spread among each drive in the drive group. Parity is equal to the total number of disks in the volume minus one drive.	Good reliability	Reads are similar to RAID 0. Writes are slower than a single drive due to penalty of writing parity.	Reads are similar to RAID 0. Writes are usually slower than a single drive.

3.0 JAVA

Security features provided by Java™ are intended for a variety of audiences, including end users and developers.

For users there is built-in security that prevents malicious programs such as viruses from running, while maintaining privacy about their files and any information about them. In Java™ 1.2 security controls can be invoked when desired for applications, similar to those for applets in previous versions.

Developers can use application-programming interfaces (API) to invoke security for programs. The framework for API enables administrators to define, and then integrate, security to control access to resources. These include authentication and authorization services, cryptography service, security manager service, and policy implementations. Java™ Authentication and authorization services (JAAS) provide support that administrators can define by users, groups or roles. Java™ Cryptology Extension (JCE) allow for encryption, key generation and agreement, and message authentication and code. JCE also allows for the addition of other qualified cryptography libraries. To allow secure Internet connections Java™ Secure Socket Extension (JSSE) packages include Java™ versions Secure Sockets Layer (SSL) and Transport Layer Security (TLS) protocols. Use of JSSE enables HTTP, Telnet, NNTP, and FTP and ensures the security of the data passing between the client and the server.

Users can also manage public/private key and public key certificates from people they trust. Java™ tools allow management of database of keys and certificates, digital signatures for Java ARchive (JAR) files, authentication of signatures and integrity of content. Users can create and modify policies files that will allow them to define and control their environment.

These policies have evolved from the original Java™ security model, known as the “sandbox” model. In this model local code is given full access to system resources (i.e. a file system). Applets that are downloaded would have only limited access to system resources with the “sandbox.” Permissions would include checking for file existence, read rights, write rights, renaming rights, directory rights creation, listing of files, file type, file size and file timestamp.

4.0 PROXY SERVERS

Browsers such as Internet Explorer and Netscape allow the user to browse the web without restrictions. The client is free to request and access any web page without regard to its content. Due to this flexibility the user may be able to access web information that is inappropriate for certain situations.

Proxy servers allow administrators to limit access to certain content within a network environment. A proxy server resides between the client or end user and an external server on the World Wide Web. The placement of this server within that environment determines its control over an entire domain or a group of individuals. The proxy server resides strategically on the firewall and will intercept all of an end user's requests at the firewall. If a requested web page is not restricted by access control list (ACL) the proxy server processes the request and the web page is sent to the client. However, if the requested page is on the ACL, the client will receive a message that the web page is not valid or not accessible.

A typical network configuration places the proxy server as the Internet gateway for users whose access to the web is restricted.

Internet performance can be enhanced with the use of a proxy server if it functions as a caching server. The proxy server can cache web pages previously requested by users without the need to go the Internet. For example if a number of users request the same web site, the first user's request will go to the web, the page will be downloaded to the user and the proxy server will write the page – cache that web page data – to its hard drive. Any subsequent request for that web page is served from the proxy server's cache. This will avoid unnecessary duplicate requests and delays that might occur from the Internet. However, with the explosive growth of the web, maintaining the ACL for proxy servers is an administrative function that can be very labor-intensive.

Proxy servers cannot be a total solution for controlling inappropriate or objectionable material from getting to the end user. They will not prevent inappropriate material in an email attachment, nor will they filter transmission of objectionable material in a chat session. Most proxy servers can accept domain names; however, controlling inappropriate pages from downloading is a difficult task.

Proxy servers strengths are their provision of a higher level of control than exists with end users using unrestricted browsers, and their ability to process access to web pages more efficiently.




Figure 1a: Proxy Server.

Workstation A initiates a request for www.nbii.gov. The request is routed through Proxy Server A, which processes the request and returns the web page to Workstation A (i.e. assuming it is on the approved ACL). The web page is also written to Proxy Server's A cache (i.e. hard drive). Any subsequent request for www.nbii.gov from a workstation on the internal network, the proxy server will deliver the information from its local cache. The end result is network performance is enhanced by not requiring a search for www.nbii.gov outside the internal LAN, which reduces overall bandwidth usage.

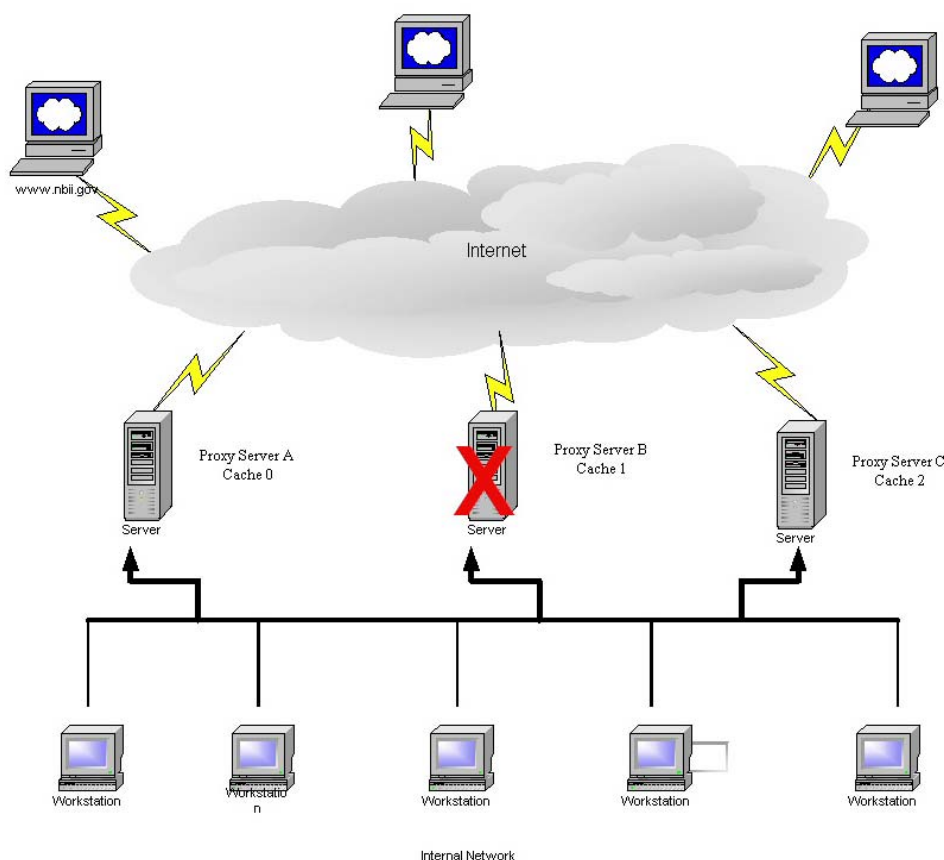


Figure 1b: Multiple Proxy Server Environment.

The internal network workstations will request information from the Internet. On the initial request the proxy server will retrieve that information and serve it to the internal workstation and cache it locally. For a subsequent request for the same information by another internal workstation the proxy server will serve that information from its local cache rather than retrieve it additional times from the Internet. The loss of Proxy server B does not effect network performance, because of server redundancy. The cache on all three servers are identical and any request handled by server B is simply rerouted to server A or C. Proxy servers can also use Network Address Translation (NAT).

5.0 NETWORK ADDRESS TRANSLATION (NAT)

The limits on the availability of IP network addresses in IP version 4.0, coupled with the explosive growth of the Internet, have resulted in an inadequate number of unique IP addresses to meet demand. One solution to this problem is the use of Network Address Translation (NAT).

NAT is able to translate IP addresses by setting up a transition table of all internal IP addresses that will send data packets through the NAT router. With the use of NAT the external world will see only a limited number of valid registered IP addresses. Internally to the organization a private addressing scheme could support a large number of network hosts. Each interface that leads to an outside interface would have a valid registered

address. This is an advantage from a security perspective. Any internal host will be unable to receive an incoming IP connection from an external system unless the external interface (i.e. gateway) is specifically configured to allow the connection. In order to establish integrity in the internal network all external interfaces must be configured with NAT.

NAT exists in two modes; static NAT and dynamic NAT. Static NAT maps internal IP addresses to external IP addresses on a one-to-one basis. Dynamic NAT maps all internal IP addresses to use one external IP address.

The recommended operation to NAT is detailed in Request for Comments (RFC) RFC 1918, which describes the recommended private internal addressing schemes. Standards for NAT are established in RFC 1631.

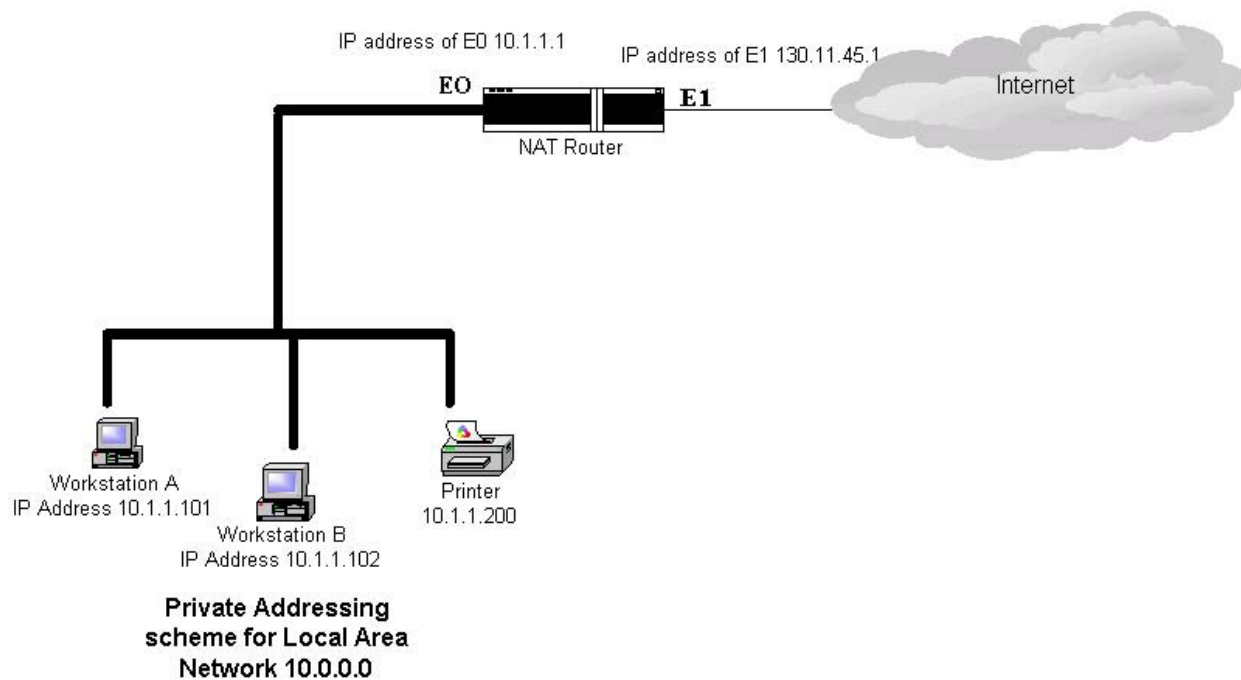


Figure 2: Network Address Translation (NAT).

This figure shows a NAT enabled router with a network address of 10.1.1.1. When any host on the 10.0.0.0 internal LAN makes a request to an external host (i.e. the Internet) NAT will translate the 10.0.0.0 addresses to 130.11.45.1. The internal hosts can access any host on the external network. For the external hosts looking in it has the appearance that all inbound and outbound data are originating from the single IP address of 130.11.45.1 (i.e. E1 the route). Figure 2.3 is an example of NAT in the dynamic mode.

6.0 FIREWALL

A firewall is a security host sitting in between a company's internal network and the external world (Internet). Firewalls are usually the company's first line of defense to prevent attacks from the outside. There are different types of firewalls in use today. The two most popular are Packet Filtering and Proxy Servers. These types of firewalls have different capabilities. Firewalls that filter data packets will allow or deny the entry of

traffic based on the IP address or source and destination ports. Proxy firewalls are based on specific applications used, including http, telnet, ftp and ssl traffic. The firewall will check this type of traffic according to the specific rules that are defined for those applications. The organization's security profile will determine the type of firewall used.

An organization that has an Internet connection should install a firewall for two primary reasons. First, company data must be protected by a malicious attack. An attack could be in the form of malicious code (i.e. virus attack) or an individual hacking into the network. Downtime from this attack could result in loss of productivity, increased expenses and revenue loss. While backups can restore the data, there is potential for stolen data. The value of the data that could be compromised is the second reason to install a firewall. The data that is stolen could contain confidential company data. In many cases, stolen data maybe used against a company. For these reasons, firewall implementation is just the first step to protect data.

6.1 Intrusion Detection

Without the proper monitoring tools in place an intruder may bypass a company's firewall. The damage that can be done could be severe if the organization is not prepared. Some skilled hackers can get through certain firewalls without notice. Deploying intrusion detection tools at the perimeter of a network is just as important as deploying a firewall. These two technologies must be used in tandem at network boundaries; a company would want to know if a hacker penetrates the firewall and gain access to the network. The use of these tools enables real-time monitoring and alerts to the appropriate individuals. With these alerts, necessary actions can be taken to protect valuable data. As an extra precaution companies should deploy a firewall and intrusion detection tools between most departments in the organization.

6.2 The Insider Attack

Hackers are not the only threats to a company's valuable data; many security problems can be directly attributed to their employees. Training the end user is a very important action for a company to take; after all, the end users have access to some or all of the company's data. Investment in training the users is necessary so accidents do not occur and company data is not compromised. If internal attacks do occur, intrusion detection software can monitor and track the electronic activities of the employees.

6.3 Firewall Conclusion

A combination of tools is necessary in order to protect data internally and externally; one tool is usually not enough. Combining various technologies such as firewall and intrusion detection software will maximize the ability to monitor and protect valuable data.




Figure 3: Firewall.

An example of a LAN with a firewall that provides filtering of content for inbound and outbound data. A DMZ on a protected leg allow external users to have read access, but prevents initiation of any requests to and fro internal LAN access.

7.0 VIRTUAL PRIVATE NETWORKS

Virtual Private Networks (VPNs) are becoming a popular way to deploy private networks across a wide geographic area. A VPN device can involve either hardware or software; installation on the server (sender of information) or the client (recipient of information) may be necessary. These devices are used to establish a secured session between the server and the client. Virtual private networks use a public network to link one or more endpoints. An endpoint can be a network device, such as a router or a user, such as a personal computer. If the endpoints are two network devices, that communication is considered a LAN-to-LAN VPN connection. For example, an organization's Virginia office network could connect to its Denver office network over a LAN-to-LAN VPN connection. If the endpoints are an end user's personal computer and a network device,

the communication is a LAN-to-network client VPN connection. A traveling user needing to connect to his LAN to read his email is an example of a LAN-to-network client VPN connection. This communication link exists only when a session is active; once completed, the session is terminated and the link is destroyed.

The usual mode of connection for the traveler is to connect in a dial up mode to an Internet Service Provider (ISP); the software would then create a secure VPN session with the mail server (See Figure 4 for a graphical representation of a virtual private network).




Figure 4: Virtual Private Network.

With rapidly changing technology there are a variety of ways and options for implementing VPNs.

In this example, a remote laptop user dials into an Internet Service Provider (ISP). Once the connection is established, the remote user runs the local VPN software. The software will build a secure VPN tunnel through the Internet and the user is authenticated by the LAN's VPN server. Once authenticated by the VPN server, the remote user has access to the LAN resources. Once the user session ends the virtual path is destroyed.

7.1 Point-to-Point Tunneling Protocol (PPTP)

One of the most widely used VPN protocols in use is Point-to-Point Tunneling Protocol (PPTP). PPTP encapsulates a Point-to-Point (PPP) frame. PPP was first widely used to dial into a remote network

using a modem. PPP transmits a network specific packet by encapsulating it into an IP (Internet Protocol) packet for transmission over the Internet. This enables the transfer of non-routable protocols such as IPX (Internetwork Packet Exchange), NetBeui, and AppleTalk, in addition to TCP/IP (Transmission Control Protocol/Internet Protocol).

PPTP has many of the characteristics of PPP, because it is very flexible in its ability to be used for non-TCP/IP environments. PPTP is widely used due to the popularity of the Windows operating system. The Microsoft Corporation developed PPTP; it is a widely deployed protocol used in Windows 9x/ME, Windows NT, Windows 2000, and Windows XP. This deployment of client software allows for the use of voluntary VPNs. PPTP authentication uses Password Authentication Protocol (PAP) and Challenge-Handshake Authentication Protocol (CHAP). Microsoft has developed its own PPTP authentication called MS-CHAP; this uses NT Domain information for authentication, which is based on RC4, based 40-bit or 128-bit encryption. MPPE is Microsoft's implementation of its PPTP client software and is the solution for voluntary mode access.

7.2 Layer 2 Forwarding Protocol (L2F)

Layer 2 Forwarding (L2F) has many similarities to PPTP. Like PPTP, L2F was designed to work with PPP and support non-routable protocols. Additional benefits of L2F over PPTP are its ability to support more authentication standards, differing types of networks and multiple threads on a single connection. The additional authentication standards include Terminal Access Controller Access Control System (TACACS) and Remote Authentication Dial-in User Service (RADIUS). Both of these standards authenticate at the beginning of the session transmission. Frame Relay and Asynchronous Transfer Mode (ATM) are examples of different types of networks that L2F supports. PPTP allows only one client to connect over a single connection, whereas L2F allows for multiple connections over a single tunnel.

7.3 Layer 2 Tunneling Protocol (L2TP)

Layer 2 Tunneling Protocol (L2TP) is used primarily in mandatory mode to access VPNs. It has the same capabilities as L2F and PPTP. Like L2F, L2TP supports multiple connections through one tunnel, works with various types of networks, and supports non-routable protocols (i.e. both IP and non-IP traffic). Unlike the others however, L2TP is IPsec compliant (see Section 2.4), which provides for stronger encryption standards, authentication and key management. The sender of any packet is able to encrypt and/or authenticate each packet. Encryption and authentication of packets leads to the use of two modes, transport and tunnel mode. In the transport mode only the transport layer is encrypted or authenticated. In tunnel mode, encryption and authentication are applied to the entire packet, not just one segment. The tunnel mode method provides for the greatest protection against attacks due to the increase in security.

7.4 IPsec Encryption

IP Security protocol (IPsec) encryption is an Internet Engineering Task Force standard that supports 56-bit and 168-bit encryption algorithms in client software. IPsec uses two protocols: AH (Authentication header) and ESP (Encapsulated Security Payload). With these protocols, IPsec ensures that transmitted data is delivered to the intended party, without augmentation or interception by unauthorized individuals. IPsec supports certificate authorities and Internet Key Exchange (IKE), and GRE is an optional configuration. IPsec encryption can be deployed in several environments and in various operating systems platforms.

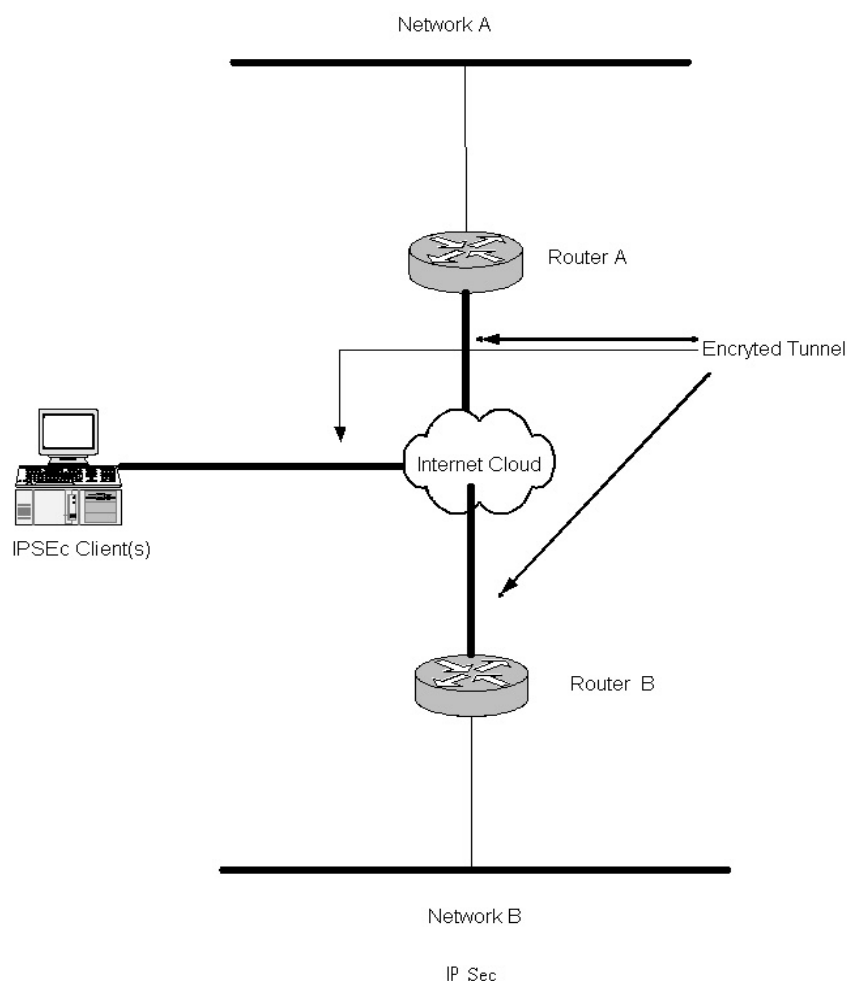


Figure 5: IPsec.

The IPsec client establishes a VPN session through the Internet with Network A and/or Network B. This session can be initiated via a direct or dialup connection. The data passing to those networks are encrypted with either 56-bit or 168-bit encryption.

7.5 Generic Routing Encapsulation Tunneling

Generic Routing Encapsulation (GRE) allows the encapsulation of IP and non-IP traffic for transmission over the Internet and/or an IP network to a specific destination. A measure of security is provided, because the packet can only enter at a specific interface. However it does not provide true security, because the packet is not encrypted.

In Fig. 6, both networks A and B run a combination of non-routable IP, IPX, and Appletalk protocols. GRE can be used to encapsulate data packets for Network A to communicate to Network B across the Internet.

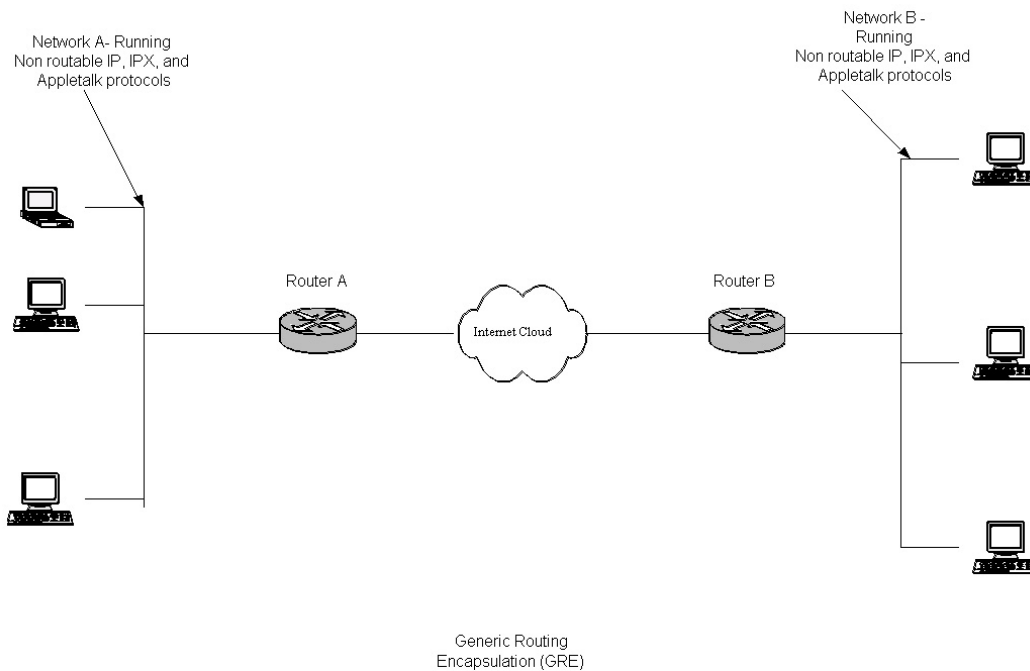


Figure 6: Generic Routing Encapsulation Tunneling.

7.6 Virtual Private Dialup Network (VPDN)

Virtual Private Dialup Network (VPDN) was developed by the Cisco Corporation and allows private network dial-in service. It may allow access to many remote servers.

The user dials into a server, often referred to as a Network Access Server (NAS; the user's destination server is known as the Home Gateway (HGW).

When a user dials into a local access server or NAS using a Point-to-Point protocol client, the NAS will forward the PPP session to the user's HGW, which will authenticate the user and initiate the session. After the user's PPP session is authenticated, then all the frames are sent through the HGW gateway router for that client. The NAS is used for access to the Internet, whereas the HGW is used for authentication to the user's home network. Authentication to NAS and HGW are not necessarily the same, but both must remain active for the remote user when she is accessing the home network.




Figure 7: Virtual Private Dialup Network (VPDN).

The remote PPP client dials into the Network Access Server (NAS) using the Point-to-Point protocol (PPP). NAS will authenticate the user and pass him/her to the home gateway, which will give the user access to his/her network.

7.7 Point-to-Point Protocol over Ethernet (PPPoE) and Multiprotocol Label Switching (MPLS)

Point-to-Point Protocol over Ethernet (PPPoE) and Multiprotocol Label Switching (MPLS) are two emerging technologies. PPPoE allows Layer 3, the Network layer of Open Systems Interface (OSI), which is mainly deployed using existing Ethernet infrastructure, (e.g. DSL service), and allows multiple users access through a single access point. MPLS is an emerging technology that provides for rapid rollout and scalability.

The primary use of the VPN protocol will determine whether PPPoE or MPLS would be best for a specific application. If the main usage is to dial up and establish a VPN session, then PPTP, L2F and L2FP is the most appropriate protocol (i.e. Network-client-to-LAN connection). If the primary application is to connect to LAN or networks devices (i.e. LAN-to-LAN), then an IPSec would be the more appropriate solution. PPTP, L2F and L2FP work at the data link layer, Layer 2, of the Open Systems Interface (OSI) model. IPSec works at the Transport layer, Layer 3, of the OSI model. The advantage of working at the Data link layer is that it offers the capability to transmit non-IP traffic through tunnels. Since IPSec works at the Transport layer it is limited to using IP traffic only.

8.0 CONCLUSIONS

The key to a robust, scalable, flexible, secure, and usable network system is to establish a strong infrastructure. Consideration of all hardware systems and software applications needs to be made at the earliest possible stages. Network hardware and applications are co-dependent. It is useless to build a strong network system with built-in redundant hardware if it is not usable to the internal and external users. With the technology constantly changing it is necessary to continually review the Network Architecture. Organizations must take a proactive role in the planning and review process, which should be a standard component of good network management. Security, reliability, usability and scalability should all be a part of the normal review process.

9.0 REFERENCES

Ciolek, A. (January 4, 2001). Virtual Private Network (VPN) Security. SANS Institute. [Online]. Available: <http://rr.sans.org/encryption/VPN-sec.php> [1 May 2002].

Cisco Systems, Inc. (February 24, 2002). Overview of How IPSec Works. Cisco Documentation. [Online]. Available: http://www.cisco.com/univercd/cc/td/doc/product/software/ios121/121cgcr/secur_c/scprt4/scdipsec.htm#xtocid10 [5 July 2004].

Egevang, K. and Francis, P. (May 1994). Network Working Group RFC 1631: The IP Network Address Translator (NAT). Andrews System Group, Carnegie Mellon. [Online]. Available: <http://asg.web.cmu.edu/rfc/rfc1631.html> [5 July 2004].

Krywaniuk, A. (November 21, 2001). Security Properties of the IPSec Protocol Suite. Internet Engineering Security Taskforce. [Online]. Available: <http://www.ietf.org/internet-drafts/draft-ietf-ipsec-properties-01.txt> [3 May 2002].

Lee, D. (1999). Enhanced IP Services for CISCO Networks: A Practical Resource for Deploying Quality of Service, Security, IP Routing, and VPN Services. Indianapolis: Cisco Press.

USByte.com (n.d.). RAID systems. USByte.com. [Online]. Available: http://www.usbyte.com/common/raid_systems.htm [5 July 2004].

Virtual Private Network Consortium. (n.d.). VPN Standards. Virtual Private Network Consortium Homepage. [Online]. Available: <http://www.vpnc.org> [5 July 2004].

