Overview of Intranet Technologies

Working Paper #35

Louisiana Forest Products Laboratory

Louisiana State University Agricultural Center

Baton Rouge, LA

Olivian T. Pitis[1]

July 21, 1999

[1] Research Associate Specialist, Forest Products Marketing Program, Louisiana Forest Products Laboratory Louisiana State University Agricultural Center, Baton Rouge, LA

TABLE OF CONTENTS

STUDY OBJECTIVE

The objective of the study is to analyze and describe the technologies used to implement an intranet. The main hardware and software components are concerned, with a focus on network physical structure and the TCP/IP suite of protocols that facilitate an intranet. The study may be useful as a starting point for management in understanding the terminology and fundamentals of intranet technology, in order to better communicate with the Information Technology department or outside agencies in the process of intranet implementation.

INTRANET FUNDAMENTALS

1. The Internet

The Internet can be technically defined as a "network of networks". In this respect, it allows for communication among networks using different hardware and software configurations, on a worldwide scale. The foundation on which the Internet is build is a suite of protocols that are usually know by the combined names of two of the better known of them: TCP (Transmission Control Protocol) and IP (Internet Protocol) (9). The TCP/IP protocol allows information to travel using different routes between computers occupying different locations on the network.

While the Internet is global in reach, an intranet groups and facilitates the exchange of information between a limited number of computers, belonging to one organization. The similarities between an intranet and the Internet are great, as the former was developed on the same foundation as the latter: the TCP/IP protocol.

2. Physical Network

As in the case of any network, a physical layer (composed of cabling, switches, connectors, bridges, hubs, etc.) is the transfer medium for the information between computers forming an intranet. In this respect, intranets are typically supported by existing LAN (Local Area Network) infrastructure. A LAN is a group of computers and associated peripheral devices connected by a communication channel, capable of sharing files and other resources between several users; a LAN typically connects computers across small geographic distances (same building or neighboring office buildings). For organizations in need of covering large geographic areas, WANs supplement and interconnect LANs to provide a means of connection across large distances, often crossing the geographical boundaries of cities or states (11).

3. TCP/IP

Because the TCP/IP protocol governs the way information flows over an intranet, it directly influences both hardware and software used to build it. And although protocols are part of the software, the importance of TCP/IP justifies its being described in this introductory section.

3.1 IP Addressing Scheme (9, 11)

Each computer on the Internet is identified by a unique IP (Internet Protocol) address. An IP address consists always of four numbers (ranging from 0 to 255) separated by a period, therefore resulting in a 32-bit number. The IP encodes the identification of the host computer (host identity or *hostid*) and of the network (network identity or *netid*) to which a host attaches. An example of an IP address is 192.149.89.61. Although this address is expressed in decimal form, the classification of IP addresses is based on their binary form, containing 32 bits. By using a different number of bits to represent either the *hostid* or *netid*, IP addresses can be classified according to the number of hosts that can be connected to a network (Figure 1).

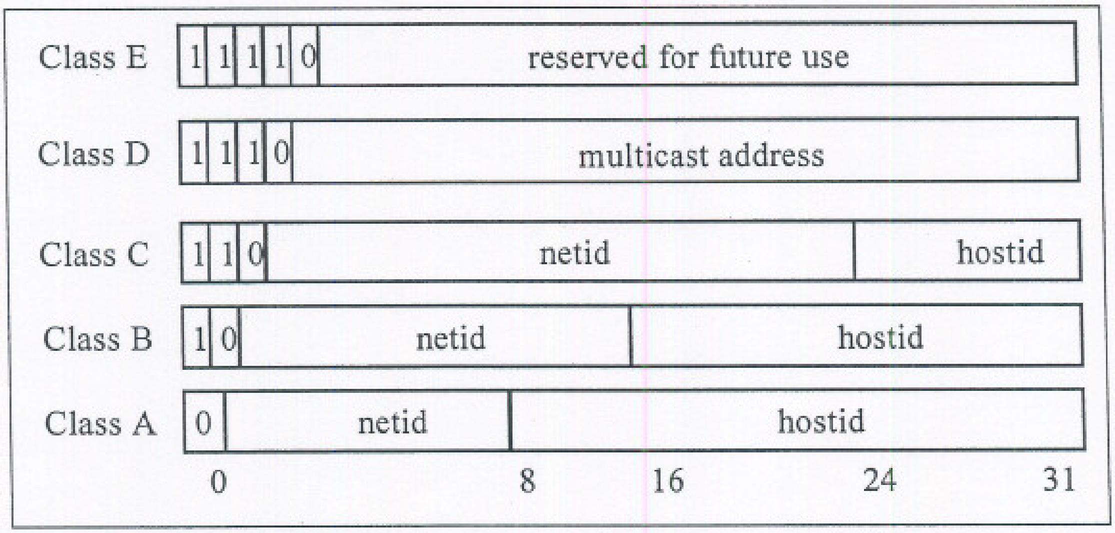

Figure 1. The five forms of IP addresses (after Minoli 1997)

Class A addresses (127 networks): the first byte of the IP address ranges from 0 to 127. This class address is reserved for large networks (up to 16,777,216 hosts).

netid	hostid	IP address
74	103.14.138	74.103.14.138

Class B addresses (16,384 networks): the first byte of the IP address ranges from 128 to 191. This class address is reserved for medium-sized networks (up to 65,535 hosts).

netid	hostid	IP address
134.64	143.24	134.64.143.24

Class C addresses (2,097,152 networks): the first byte of the IP address ranges from 192 to 223. This class address is reserved for small networks (up to 254 hosts).

netid	hostid	IP address
134.64	143.24	134.64.143.24

Classes D (multicast addresses) and E (reserved for experimental purposes) are less encountered in general use of the Internet and therefore are not discussed here.

It is important to note that IP addresses are assigned by a central authority: the Internet Network Information Center (InterNIC). For organizations implementing intranets that are not connected to the Internet, *netid* and *hostid* numbers can be freely assigned. However, if the intranet is subsequently connected to the Internet, a possible conflict of addresses can occur. It is therefore recommended that even private networks apply for IP addresses from InterNIC (11).

3.2 TCP/IP Layering Model

The suite of protocols known as TCP/IP is responsible for various tasks along the communication process. This results in the TCP/IP being a four layered protocol, with each layer responsible for specific processes (Figure 2) (11):

1. The **link layer** (or network interface layer) includes the device driver in the operating system and the network card in the computer, handling the hardware details of physically interfacing with the network cabling.
2. The **network layer** (or Internet layer) handles the movement of information in the form of packets in the network, with respect to the route used to transfer information between two IP addresses.
3. The **transport layer** provides the flow of data between the two communicating systems, to facilitate the application layer above it. The transport layer establishes the basic requirements for reliable data transfer. Any additional requirements concerning data transfer reliability are satisfied by the application layer.
4. The **application layer** handles a particular application, such as TELNET (remote login), FTP (File Transfer Protocol) or SMTP (Simple Mail Transfer Protocol - for transfer of electronic mail).

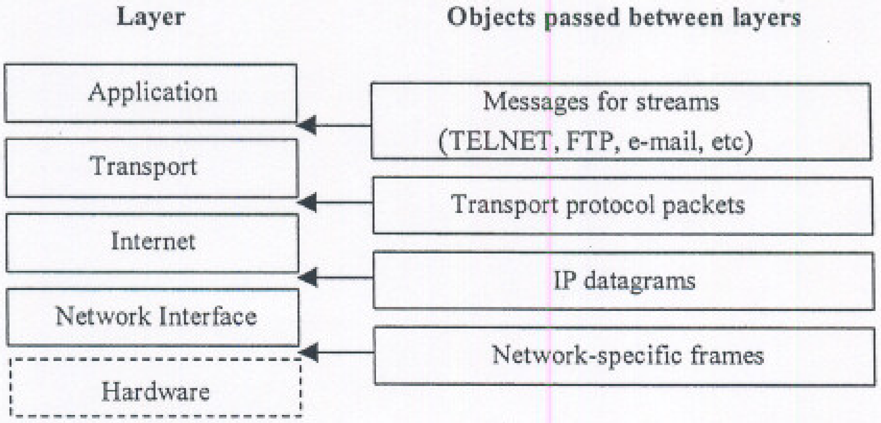
Layer
Objects passed between layers
Application
Messages for streams (TELNET, FTP, e-mail, etc)
Transport
Transport protocol packets
Internet
IP datagrams
Network Interface
Network-specific frames
Hardware

Figure 2. The four layers of TCP/IP protocol (after Minoli 1997)

3.3 TCP/IP Characteristics

As a result of the properties described above, the TCP/IP protocol ensures a reliable flow of information between any two computers connected to the network, even in the event of a part of the network becoming disabled for any reason (military attack, natural disasters, etc.). The information that needs to be transmitted from A to B is first broken into packets at point A. The packets are subsequently transmitted over the network (not necessarily all packets following the same route) and then reassembled at point B. The TCP/IP has the following characteristics (9):

- it is a routable protocol, meaning the information can be directed on a specific route, therefore reducing network traffic in more solicited areas of the network;
- it has reliable and efficient data-delivery mechanisms;
- it is a widely published standard and is completely independent on any hardware or software manufacturer;
- it allows for communication between computers running on different hardware or software architectures (i.e. UNIX to Windows NT, PC to mainframe, etc.)
- it uses a common addressing scheme, which allows any two machines connected to the network to communicate with each other, even in a network as large as the Internet.

3.4 Routing

The process of directing the packets of information from point A to point B is called routing. Routers are an essential element in this process. A router is a hardware device specialized in receiving and transmitting TCP/IP information packets. A packet traveling from address A to address B is sent from one router to the next until it reaches the final destination. Each router has information about its connections with neighboring routers; this information is updated on an ongoing basis. Therefore, each router is able to derive the whole network topology and therefore identify the optimum route for a packet based on its destination address. In the case of an intranet, the router may serve as the connection point to the Internet. Any information travelling to an address inside the network is directed to the correct address. Packets intended for an address outside the intranet are submitted to one of the neighboring routers, according to the optimum route calculated by the router. The whole routing process has been evolving continuously since its inception in 1978. However, the basic principle described above is expected to continue to satisfy the requirements of the Internet traffic during the following ten years, while new principles are being tested and implemented (such as switch based routing or routing in software - performed by a general purpose PC, or multicasting) (11).

3.5 Domain Name Service (4)

The IP address system ensures unique allocation of address for all computers connected to the Internet. However, the system is cumbersome to use for humans. This is the reason why the Domain Name Service (DNS) has been implemented. A domain name allows for use of strings instead of numbers, therefore providing for an easier way to remember and use Internet addresses. A domain name consists of a top level domain (**.edu** - educational, **.org** - nonprofit organization, **.com** - commercial, **.net** - network support, **.gov** - government and **.mil** - military) and at least one (possibly more) subdomains (4). For instance, the domain name www.lsu.edu specifies the Web server machine in the domain lsu.edu. The www.lsu.edu is translated in the IP address 130.39.152.3. "www" can be replaced by other combinations indicating the type of server: "ftp" (File Transfer Protocol server), "mail" (e-mail server), "chat" (chat server), etc. At the same time, if one server plays more than one role, different domain names can be translated in the same IP address (4).

The translation from domain names to IP addresses is done by DNS servers, which are servers permanently connected to the Internet and which maintain DNS tables that indicate the correspondence between domain names and IP addresses. When a computer is soliciting information from a specific address (i.e. www.lsu.edu), it sends first a request to a DNS server, which subsequently returns a IP address. The host computer then sends its request towards this IP address (4).

At the beginning of a domain name address it is possible to add a combination of letters that indicates the application protocol that is used to transfer the information. A URL (Universal Resource Locator) is thus obtained. An example is http://www.lsu.edu. "http" indicates that the protocol HTTP (Hypertext Transfer Protocol) is used to transfer information to and from the computer with the name www.lsu.edu, which after a contact with the DNS server is further translated into the IP address 130.39.152.3.

The above description of the TCP/IP protocol is only intended to draw the frame that allows for an understanding of terminology and functionality of an Intranet as they are described in this paper. The protocol, as the key factor in the success of the Internet, is much more complex and under further development. One of the most significant changes that is expected to occur until the year 2005 is the use of the 16-byte IP address (instead of the 4-byte address used today), to allow for more computers and other devices (such as home appliances) to be connected to the Internet.

INTRANET Hardware

1. Local Area Networks

A typical intranet is developed, at least in its first stages, on the existing LAN within the organization. Further expansion of the intranet can be achieved by using WANs. Alternatively, if the intranet is the first network that the organization implements, the physical network that is built to support the traffic of information still resembles that of a LAN. For this reason, a description of the typical structure of a LAN follows, with a focus on the hardware architecture.

On a typical LAN, organizations implement proprietary (customized) software solutions. In addition to the TCP/IP protocol, an intranet is based on other software that is described in the section "Intranet Software". For this reasons, the paper does not focus on the software characteristics of a LAN, but instead it describes the hardware features that are relevant to the functioning of an intranet.

1.1 LAN Topology (10)

LAN topology refers to the geometric layout of the cable used to interconnect workstations on the network. There are five basic network topologies, described below (Figure 3).

A **loop** topology results in workstations communicating through the use of a common controller on a cable closed into a loop. Although the loop topology reduces the cost of terminals connected to the loop because of the minimum intelligence required from them, it can be limited by the processing power of the controller. Possible lengthy exchange between two terminals, as well as the inoperability of the network in the event of controller failure restrict loop topology to several niche areas.

In a **bus** topology, each workstation is connected directly to the main data highway. Although this structure allows each two stations to talk to each other, special procedures are required to handle situations such as two stations trying to communicate at the same time.

In a **ring** topology, the cable that forms the data highway is connected into a loop. The difference from the loop topology is that the stations do not require a controller to govern the communication process.

In a **star** topology, each station is connected to a central network controller, which is the common point of communication. The network controller is somewhat analogue to a telephone switchboard, in that access from one station to another can only occur only through the central device.

A **tree** network can be considered to represent a complex bus, with the common point of communication at the top of the structure known as the headend.

In practice, it is common to encounter networks that are a mixture of two or more of the above topologies. The choice of any of the available topologies is dependent on performance requirements, available resources, existing configuration (i.e. wiring already in place). Currently, the most commonly types of networks in use by private or government organizations are Ethernet Networks (a mixture of star and bus topology) and Token-Ring Networks (a mixture of star and ring topology).

1.2 Wiring

All the cables that support the network traffic are commonly known as "the cable plant". The cable plant has two major components: backbone wiring and horizontal wiring. Horizontal wiring connects workstations on the same floor or in the same area to the wire center (usually located in a wiring closet). Backbone wiring connects the wiring closets together. Because often wiring closets are usually placed above each other on different floors, backbone wiring is also known as vertical wiring. However, in the case of adjacent wiring closets, the same media is used for connection as in the case of vertical wiring, even though, physically, the cables are laid horizontally (17).

In the case of horizontal wiring, twisted pair cables are most commonly used. Twisted pair wires are in essence copper telephone wires and many times they provide the cost saving and convenience of existing wiring already in place. For LANs where signal noise interference can be a problem, shielded cabling or coaxial cable must be used for horizontal wiring. All these technologies provide a transport medium that satisfies economic costs and traffic requirements from workstations to the wire center (17).

Backbone wiring has to satisfy more stringent requirements with respect to interferences and signal integrity over larger distances. Therefore, a different medium is typically used: fiber optic wiring. Fiber optic wiring is more expensive than copper wiring and requires more specialized installation labor. However, it is more reliable in preserving the integrity of the signal and is immune to electrical interference from elevator shafts, electric facilities, etc (17).

Lately, as a consequence of decreasing prices for fiber optic cables, there is a tendency to use fiber optics for horizontal wiring as well ("fiber to the desktop"). Although the installation cost is higher than for copper cable, fiber optic can accommodate for future increased bandwidth requirements. However, the choice between the two media can be a difficult one, and it is normally made after all factors affecting installation and use of the LAN are considered: budget, time and labor constraints, expected life and future growth of the network, performance requirements, etc. (17).

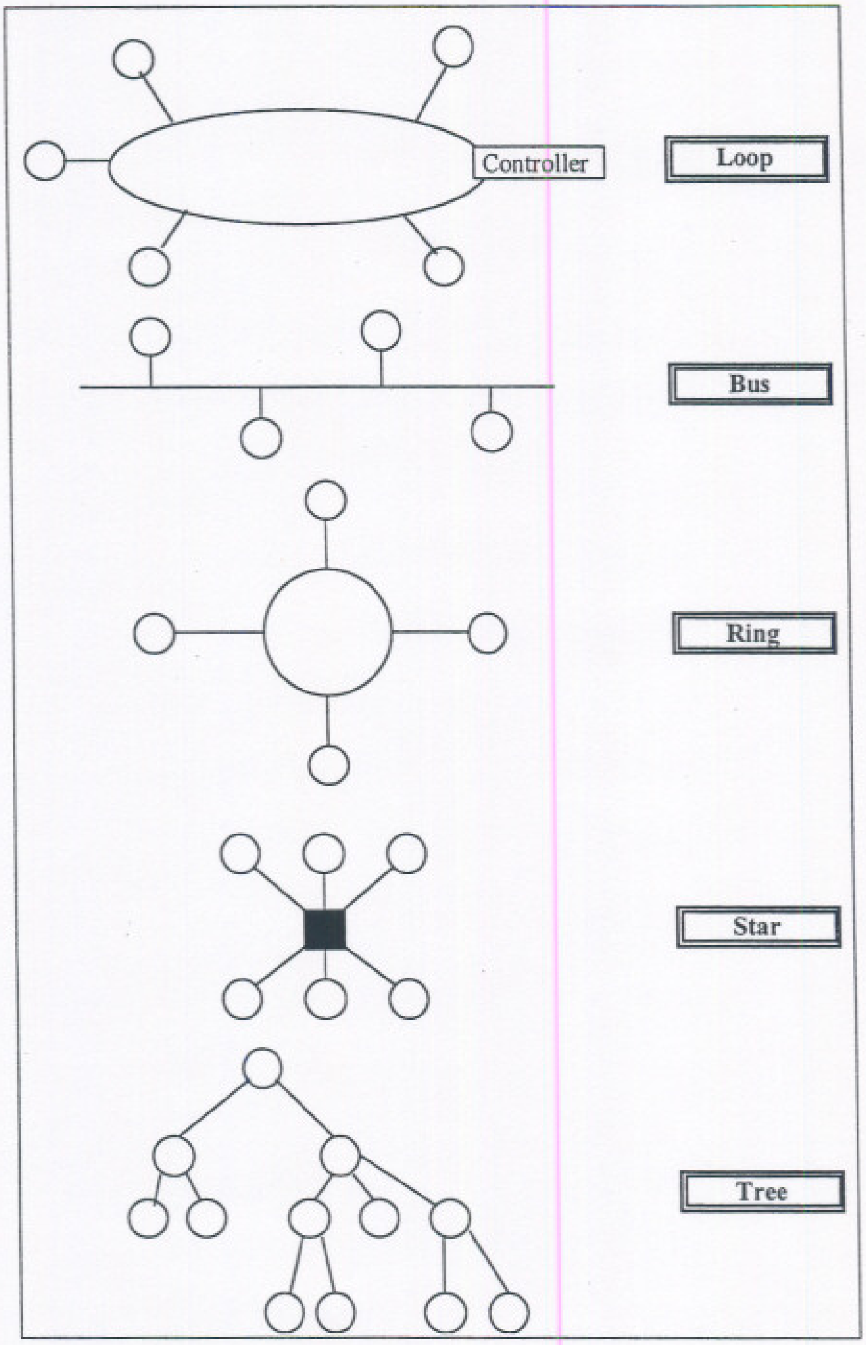

Figure 3. LAN topologies (after Held 1994)

1.3 Key Hub Components

A major step in the LAN evolution has been achieved by the advent of the intelligent hub or concentrator. Hubs provide two main attributes: 1) central connectivity and management of network nodes via diverse cabling media and communication protocols and 2) monitoring and management of the network topology design, segmentation and performance (7). Hubs achieve segmentation of the network, which allows grouping together of devices that exchange data frequently in order to optimize network traffic. Segmentation "keeps local traffic local", allowing only communication that must occur outside the segment to go out. Across the network enterprise, segmentation reduces network traffic with direct positive effects on cost and performance (7). Key hub components in a LAN are bridges, switches and routers.

Bridges connect two LAN segments, acting as a packet filter between two segments. Only packets with an address outside the segment are allowed to pass through the bridge. A bridge is generally a simple and cheap device, but it is limited in terms of support provided towards network security (7).

Switches are more complex hubs, which can handle communication using different port speeds. They are also quicker and more efficient in transferring packets than bridges, while allowing for the connection of a higher number of network nodes (7).

Routers are the most expensive hubs, allowing for management of performance, security and traffic within and between network segments. Due to their efficiency, routers are often used to achieve WAN connectivity (7). Internet routers are an example of the routing concept implemented for a particular protocol (TCP/IP).

Depending mainly on the distance span by the network and the number of hosts, one, all or any combination of bridging, switching and routing may be appropriate for a particular case.

As mentioned before, the network technologies most commonly used in implementing LANs are Ethernet and Token-Ring Networks. A description of both technologies follows.

1.4 Ethernet Networks

Ethernet networks were originally developed by Xerox Corporation at its Palo Alto Research Center. In the late 1970s, in a joint effort with Digital Equipment Corporation (DEC) and Intel Corporation, Xerox promoted Ethernet as an open standard for computing (15).

Two of the early implementations of the Ethernet (10Base5 and 10Base2) used a tapped bus network topology. 10BaseT Ethernet, which is currently the most widely used Ethernet topology, uses a star network topology at its basic structure. All workstations connected to one hub have equal rights in transmitting information. When a station transmits a frame (the form in which information travels across Ethernet networks), the hub receives it and then it retransmits it to all stations connected. The station that is indicated as the destination address processes the information, while the other stations ignore it. If two stations try to transmit at the same time, an event called collision occurs. In this case, the stations try to resend the information after a random time interval, until a transmission is successfully made. If the traffic over the network is intense, the process can result in delays or even transmission being interrupted (17).

If the hub is connected to other hubs (usually in a bus topology), the Ethernet becomes a star-bus topology. 10BaseT Ethernets support transfer of information with speeds of up to 10Mbps (10 Megabytes per second) (15, 17).

In a move to support larger bandwidths, Fast Ethernet networks have been developed, supporting speeds of up to 100Mbps. These new standards are 100BaseTX, 100BaseT4 and 100BaseFx, and are intended to facilitate communication across networks with high traffic. With the exception of 100BaseFx, which is supported by optic cable medium, the other Ethernet types are supported by copper (either coaxial or twisted) cable (15, 17).

1.5 Token-Ring Networks

Token-Ring networks were introduced by IBM in the mid 1980s. They were originally implemented over copper wires and can currently achieve speeds of 4 or 16Mbps. Fiber Distributed Data Interface (FDDI) networks, a more recent development of token-ring networks, are supported by fiber optics and support communication with up to 100Mbps (17).

In a token-ring network, a group of workstations are connected to a common device called a Multistation Access Unit (MAU) in a star topology. MAUs are then connected in a ring topology, resulting in a star-ring configuration for token-ring networks. In a token-passing access method, a token (represented by a unique bit patterns) travels around the network. When a station has to transmit data, it has to seize a free token. The token is transformed to indicate that it is in use and data is added to it. During the time a token is in use, other stations remain idle on the ring. The token travels around the ring until it reaches the destination. The station receiving the data frees the token, which is then taken by the next station that needs to start a transmission. Since a station can only transmit when it has a free token, token passing eliminates the requirement for collision detection. Unlike Ethernet networks in which access to transmission is non-predictable, token-ring networks have a deterministic access method, which allows traffic to increase without significant reduction in performance. Although the cost of token-ring technologies is higher than those used in Ethernet networks, the predictability and consistency of token-ring networks makes them the preferred LAN for organizations with high utilization or a potential requirement for adding applications that can increase network utilization (10).

2. Wide Area Networks

As specified before, an intranet can also be implemented to cover large geographic distances, in the case of organizations with offices in different cities or countries. In such a case, the intranet is supported by a Wide Area Network (WAN). Distance is the key factor that determines the implementation of a WAN, asking for a different technique to establish and maintain reliable communications. Subsequently, more powerful and different equipment is required to support distant communication between computers (7). The main differences between LANs and WANs are summarized below (after held 1994).

Data transmission and error rates. Since LAN cabling is primarily within a building or extends over a small geographic area, it is less susceptible to outside influence, either natural or induced by human activities. This reflects in different error rates: a WAN is susceptible for a transmission error rate of 1 in every 10^6 to 10^7 bits, while the rate error for a LAN is 1 in every 10^7 to 10^8 bits.

Ownership. The implementation of a WAN requires the leasing of transmission facilities from one or more communications carriers. By contrast, the organization implementing a LAN typically owns all the components of the network, including the cables that form the transmission path between devices.

Regulation. Most of the regulations affecting WANs refer to the services communication carriers provide to their customers and the rates they can charge for those services. Regulations affecting LANs are primarily in the area of building codes, regulating the type of wiring that can be installed in a building and whether or not the wiring must run in a conduit.

Data routing and topology. In a LAN, the path along which data is routed defines the network topology (star, bus, ring, loop, tree). A WAN resembles rather to a mesh, with data travelling alternative routes to reach the final destination.

Type of information carried. Most WANs currently support simultaneous transmission of voice, data and video information. By contrast, LANs are typically limited to carrying data.

While WANs are typically implemented to provide connections between distant fixed points, access to a company's intranet often requires employees to log into the network from various locations. In such a case, the typical equipment required by WAN implementation (routers, leased lines) can not be used for cost considerations. Dial-up connections are used in this case to provide for remote connectivity. A user from a remote workstation dials into the company's network through a modem. At the other end, the server that is configured to provide remote access is connected to the phone lines through a modem pool, each modem serving one potential user at any given moment (7). The dial-up procedure is widely used to provide Internet access by dialing local phone numbers, but for employees traveling to remote locations the cost associated with long-distance or international phone rates may be a limiting factor.

Virtual Private Networks (VPNs) have emerged to overcome this disadvantage. A VPN (Virtual Private Network) is a private connection between two machines that sends private data traffic over a shared or public network, the Internet. This emerging technology lets organizations extend their network service over the Internet to branch offices and remote users creating a private WAN (Wide Area Network) via the Internet. The appeal of a VPN is the Internet and its global presence. Communication links can be done quickly, cheaply, and safely across the world (1).

3. Network Computers

An intranet is based on the client-server networking model. Under this model, **client** refers to the program or process that submits a request to a server, and server refers to the program that receives the request from the client, processes it, and returns the results to the client (16). In a similar fashion, computers on an intranet can be categorized into clients and servers. A server computer provides a specific service (printing services, e-mail transmission and receiving, file transfer, etc.) to more client computers. In an intranet, the complexity of tasks that a server has to perform and the number of clients connected to it dictate the computing power that the server has to deploy.

The increase in computing power coupled with a significant decrease in prices makes today's high-end PCs suitable as servers for small to medium intranets. Larger organizations can choose to use enterprise-grade server platforms (such as IBM's AS/400), which have increased processing power compared to PCs. Furthermore, enterprise servers can play multiple server roles at the same time (Web server, e-mail, e-commerce), resulting in ease of administration and increased reliability (3). Mainframe platforms can also be used in implementing an intranet. But, although many mainframe systems today support intranet applications, it is arguable if the use of mainframe computing power over the intranet is cost-effective (2).

Client computers (also referred to as desktop computers) are used by individuals to share information over the intranet. Because TCP/IP allows for communication between computers running on different hardware or software architectures, the choice between the combinations of hardware (PC or Mac) and operating systems (Unix, Windows, Linux, OS/2, etc.) is only limited by the resources available to implement and maintain them and the skills of the employees that are using the desktops.

The subject of intranet hardware is much vaster than it has been described so far in this paper. However, the purpose of the study is to offer a global picture of the functionality of an intranet. Therefore, more detailed discussion of intranet hardware is beyond the scope of the paper. The literature cited could be useful as a starting point for further investigation that may be needed.

INTRANET SOFTWARE

1. Internet Applications

The same applications that were developed for the Internet can be implemented over an intranet. Before presenting intranet software that can facilitate this services, an overview of the Internet tools most likely to be implemented on an intranet is appropriate.

1.1 The World Wide Web

The World Wide Web (WWW or simply the Web) is an Internet tool that allows access to various types of information (text, images, video, databases, etc.) through a highly interactive interface. WWW technology is based on HTML (Hyper Text Mark-up Language) and programming languages (i.e. JavaScript, Java, VisualBasic Script, CGI – Common Gateway Interface). The information is stored on computers acting as servers; users access this information and transfer it to their computers (client computers) using special programs called browsers. The WWW is the fastest growing part of the Internet; this is mainly due to its ease of use and high interactivity.

The protocol used to transfer information between Web clients and servers is HyperText Transfer Protocol (HTTP). An HTTP transaction takes place in four steps (16). **Step 1**: the client establishes a connection to the Web server using the TCP/IP protocol. **Step 2**: after connecting to the server, the client submits a request for a specific information. **Step 3**: the server processes the request and returns either the requested information or a response that it cannot process the request. **Step 4**: the server or the client closes the TCP/IP connection. As a result of the procedure described above, HTTP is a *stateless* protocol, which means that for each transaction of information a separate connection has to be made. This is necessary because in a typical Web browsing session a client can requests information from more than one servers; also, by keeping the TCP connection open just enough for the information to be transmitted to the client computer, the server can process more requests for information in a given amount of time (16).

Web documents are stored in HTML (HyperText Markup Language) format. HTML is a language used to define documents using generic commands, called *tags*, to indicate document elements and corresponding formatting. HTML documents therefore do not contain formatted elements, but *information* about the formatted elements. The distinction has dramatic effects towards cross-platform compatibility: browsers identify the formatting of various elements in an HTML document, and then present the formatted elements to the user. Therefore, the same document can look slightly different on different computers, depending on the screen resolution, browser and other individual settings. This is not to say that the document presents different information, but rather it presents the same information in a slightly different way. For instance, employees could be accessing the same technical brochure on the company's intranet, and each of them would be able to see it in the format best suited for the computer from which they are accessing it (5, 16). HTML documents can also contain *hyperlinks*, which contain information about a destination Web address. By clicking on a hyperlink, users direct their browsers to new information, in a highly interactive fashion.

The cross-platform compatibility of the Web (as a result of the TCP/IP protocol, HTTP and HTML), as well as its interactivity, have made the Web the preferred medium to disseminate information over the Internet, to the point where the two have almost become synonymous in general terms usage. Intranets have also become increasingly Web-centric. With the exception of Telnet, each of the other services presented (e-mail, FTP, Usenet and Gopher) can be configured on an intranet to run through a client-server Web interface, while the same cannot be said about the other applications. This ability to use the same tool for a wide range of tasks, coupled with the ease of use of a Web browser, can be instrumental in implementing intranets in the organization, resulting in reduced employees resistance to change and quicker training (12).

1.2 Electronic Mail

Electronic mail (e-mail) is the most popular tool of the Internet when the exchange of information (as opposed to dissemination of information) is concerned. It allows the interchange of written, non-interactive messages between any two individuals with an e-mail address. In addition, information stored in file formats other than text (i.e. databases, images and sounds) can be shared through e-mail. E-mail offers some advantages that are derived from its simplicity: it is widely used, requires relatively little resources to implement and maintain and provides a cheap and quick (though not instantaneous) means of communication.

The way electronic mail is transmitted over networks is similar to the way regular mail travels through the postal service. Each message has a destination address and a sender's address. An e-mail address identifies a user and a domain name describing the network to which the user is connected (i.e. user@company.com). Through the client e-mail application, the sender contacts its mail server and loads the message onto the server. The mail server reads the destination address of the message and, using a DNS server, translates it into an IP address. The message is then sent to the server thus identified, which stores it until the addressee contacts the server and downloads the message (16).

The process is facilitated by three open protocols: SMTP, POP3 and IMAP. SMTP (Simple Mail Transfer Protocol) is used to send e-mail from the client e-mail software to the server and between e-mail servers. POP3 (Post Office Protocol version 3) and IMAP (Internet Mail Access Protocol) are used to download messages from e-mail servers. IMAP includes a few improvements over POP3. For instance, a client using POP3 has to download all the messages from the server before being capable of identifying which messages are relevant and which are not. Once downloaded, the messages are deleted from the server. By contrast, an IMAP user can contact the server and receive first only summary information about each message: sender, subject, date. Based on this information, the user can choose which messages are to be read and which are to be deleted. Once read, the messages are not deleted from the server. This allows users to read the same messages from different locations, conferring the system increased flexibility (16).

Note: As specified before, Web clients can also be used to send and receive e-mail. However, the mail server must run special software that allows the user to communicate using only by means of the HTTP protocol. Thus, sending, receiving and managing of messages is made using Web pages.

1.3 File Transfer Protocol

File Transfer Protocol (FTP) allows the flow of files bi-directionally between servers and client computers. It gives users access to various resources, especially software. It is also used when downloading files from the WWW, but in a perfectly transparent manner (i.e. the user is not required to input any information and does not interfere with the process after starting it). Compared to the WWW, FTP uses slower navigating procedures (users need to initiate a login procedure for each FTP server) and provides a less interactive interface. However, FTP allows transfer of files **to** and **from** the client computer (the Web allows one-way transfer, from server to client) and still retains the cross-platform compatibility that allows users to share files using different platforms (16).

An FTP server listens to requests for information using the TCP/IP protocol. When a request for information is made, the server prompts the client for a user name and password. If the combination is valid, the server makes files available according to the user's rights. Most FTP servers allow *anonymous* access in order to provide nonregistered users access to public files.

Once users are connected to an FTP servers, they can navigate through the directory structure of the server, view, download or upload files according to their access rights.

Note: Most recent versions of Web browsers have the additional capability to serve as FTP clients. Also, the downloading of files from Web pages is made using the FTP protocol, but in a transparent manner to users (the login procedure - if required - is performed by the Web server in order to provide access to the Web pages from which pages are downloaded).

1.4 News

News is an Internet application that allows users with a specific interest to participate in discussions and exchange of information related to a certain subject. Whereas e-mail is badly suited to share information among multiple users, news was created to facilitate distribution of information to large groups.

A news server maintains a list of newsgroups, each of which containing messages posted by users with interest in a specific subject. After contacting the news server using the TCP/IP, a news client uses the NNTP (Network News Transport Protocol) to solicit information about newsgroups available on the server. Newsgroups that have appeared since the last connection are thus identified. Users can choose to subscribe to one or more newsgroups, from which a list of new messages is then downloaded. A user can choose to read and reply to any of the existing messages, creating a message *thread* (16).

An organization can choose to implement its own news system to restrict its use to the intranet. In this case, a separate news server must be maintained if access to news available on the Internet is also desired. The conversational aspect of news makes it an ideal service for employees to discuss issues or technical problems or to brainstorm on important projects. Also, because the messages can be configured to reside on the server for a limited time, and messages are downloaded only by users to which they are relevant, news provides better usage of network bandwidth and file storage space (16).

Note: Most Web browsers also provide support for news. Therefore, implementing news on an intranet can be a straightforward task, as users must only be informed where to direct their browsers. An additional administrative effort is required if a dedicated server is implemented for intranet news.

1.5 Gopher

Gopher is a user interface initiated in 1991 by students at the University of Minnesota. It was the first widely available, easy-to-use client application for finding information on Internet servers. It is menu-driven and has the main advantage of being easy to use. It is also fast and offers access to a wide area of information. Although the WWW has become significantly more popular, Gopher is still in use. The main disadvantage of Gopher is its less attractive user interface when compared to the WWW (6).

A typical Gopher session resembles a Web session, in that the server does not maintain a permanent connection after the information is sent to the client. If the resource requested by the client is a text file, the content of the file is displayed to the user's screen. If the resource is a binary file (i.e. graphic files, audio files, databases, etc.), the user is prompted to indicate a location where the file is stored. Therefore, Gopher can be used to download files from Gopher servers, providing a more intuitive interface that FTP transfer.

Even if the Web surpasses Gopher in its ease of use, a company may choose to implement Gopher on its intranet for three reasons (16):

- a large number of text documents have to be made available immediately to users. By putting the files on a Gopher server, the text files can be converted into HTML at a later time;

- an easier way to download file than FTP is desired;
- the organization already has existing Gopher server(s) that need to be moved to a new platform.

Note: Most Web browsers also provide support for Gopher. The navigation procedure using a Web browser is not changed for an user when switching from a Web to a Gopher server. However, the information on Gopher servers appears less attractive than on Web pages. Furthermore, Gopher lacks the interactivity the Web provides through hyperlinks.

1.6 Telnet

Telnet is one of the oldest applications of the Internet. It gives access to various databases, chat servers and selected e-mail servers. Telnet is entirely text based and originated from the Unix operating system, although it is available today on servers running other operating systems such as Windows NT (6). Telnet requires users to know specific access commands and procedures; users must continuously input this information using the keyboard, therefore being less "user-friendly". Its use is thus limited to a comparatively small number of users, usually with a broad knowledge of computers. During a Telnet session, the connection client-server is permanently maintained until the user chooses to connect to another server. Telnet, is also an upper-layer protocol, supported by TCP/IP. Because of its ability to provide connection to older systems (mainframe computers), Telnet may become part of a company's intranet (16).

2. Intranet Implementation of Internet Software

An organization chooses to implement an intranet when quick, efficient and cost-effective methods of sharing information and allowing people to communicate are required across the whole organization. Therefore, when implementing Internet applications on an intranet, some modifications have to be made that allow the intranet to fulfill its objectives.

Although all internet application presented above can be implemented using a separate client-server application for each of them, the procedure is not cost and time-efficient. Instead, a stronger focus on Web applications is required, resulting in Web-centric intranets (12). As a result, people in the organization have the ability to use a Web browser as a client application for all intranet services: Web, e-mail, FTP, news. Additionally, the Web browser can interact with other applications to allow employees to access company data that is not necessarily in HTML format, such as order/invoices databases or human resources records.

An overview of software that supports Web-centric intranets is presented below. The purpose of this exercise is not to describe in detail the installation, configuration or use of the products. Rather, the purpose is to describe how a company can benefit from using such software on an intranet.

2.1 Server Software

Servers are an essential part of an intranet. While client computers may be disconnected from the network or shut down during the day, servers must run continuously in order to provide services for users that could connect at any time. As specified before, servers may be based on different hardware platforms (mainframe, AS/400, Mac, PC) and operating systems (Windows NT, Unix, OS/2). For each of these combinations, network administrators have a choice regarding the software that is run on the server to provide Web, FTP, e-mail or news services. Each service may have a dedicated server, or one server may provide more than one (or even all) of the above services.

Since an intranet focuses on the Web, the choice of the Web server software is likely to have the greatest impact on the IT department. Numerous Web server programs run today on the Internet, the most popular being presented in Figure 4 (survey of **6,598,697** Web sites, Netcraft 1999). Although accurate data regarding the market share of Web servers running on intranets is difficult to gather, it is likely that the major players are the same. Because Web servers present their information to users in the standard HTML format, users do not distinguish between different Web server software packages. Thus, the choice of a particular product is influenced more by the pricing structure of the product (some are free, such as Appache Software Foundation Web servers, other require the purchase of a license) and by the familiarity of the IT department with a particular product.

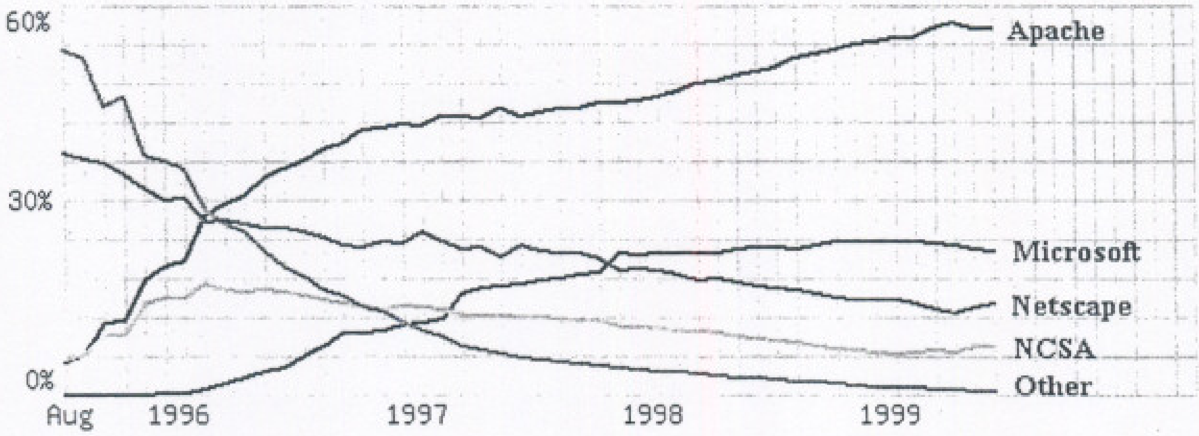

Figure 4. Market share of the most popular Web servers on the Internet (Netcraft 1999)

Of a particular interest for intranets are back-office applications, which dynamically generate HTML pages based on user input. Back-office applications run on the Web server and are accessed through the users' browsers. This allows users to interact with different applications (databases, spreadsheets, etc.) using the familiar Web browser interface. Back-office applications and the Web server communicate through special scripts. One of the most commonly used techniques to create such scripts is using CGI (Common Gateway Interface). A CGI script receives the request from the Web server (for instance, a form filled in on the user's browser) and runs the corresponding back-office application using the variables inputted by the user. After receiving the result, the CGI script converts the response into a standard HTML page, which is then presented by the Web server to the user (5, 12, 16). Using CGI scripts, and other similar techniques, Web servers can be enabled to provide users not only standard Internet services (e-mail, FTP, news) but also customized services (database entry and querying, document editing, etc.).

2.2 Client Software

The Web browser is the application used to interact with the Web server. It can provide employees access to all intranet services therefore requires employees a minimum amount of training. Additionally, two of the widest used browsers (Microsoft's Internet Explorer and Netscape's Navigator) are free products. This means that any number of employees can use the product without the need to purchase a license, reducing the cost of implementing an intranet. The above benefits are a direct result of the HTML being an open standard, allowing any software company to produce HTML-capable applications. Furthermore, the intranet administrator is not constrained by compatibility problems between client software and browser software, meaning that any combination of Web server and Web browser can be used across the organization (5, 12, 16).

In addition to Web browsers, users may use separate client programs for FTP, e-mail or news applications. A separate client program may provide increased functionality for a given application, or some users may use it simply for preference reasons. However, if such an approach is taken across the company's intranet, the direct result is an increase in expenses related to software installation and user support (12).

2.3 Security

A broad definition of network security implies the ability to keep users (inside or outside the organization) from accessing or tampering with network resources (processing power, peripheral devices, proprietary information) that they are not entitled to (8). There are currently seven security levels defined by the United States Department of Defense, ranging form Level D1 (entirely untrusted networks) to Level A (the highest, involving a stringent design, control and verification processes). Implementing the desired security level on a network involves hardware and software configuration, monitoring and auditing procedures, enforcing Acceptable User Policies (AUPs) and corresponding training of personnel (14).

If an intranet is not connected to the Internet, enforcing a security policy is facilitated by the lack of possible points of breaking into the network. However, the benefits of having an intranet are shadowed if access to the outside world is limited or severed. Therefore, specific strategies are typically implemented that allow users inside an intranet to also access the Internet, while protecting network resources from outside access. It is important to note that security measures can be very complex and are also implemented in order to regulate access to resources from **inside** the network. Because an intranet is based on the same technologies as the Internet, the processes and devices used to ensure network security are also the border line between the intranet and the Internet (14).

The network components (both hardware and software) that allow a secure connection between the internal network (intranet) and the outside network (the Internet) comprise a firewall. The amount of hardware resources and administrative efforts that are put into implementing a firewall is dependent on the level of security desired and the budget constraints. However, the principle remains the same, as presented in Figure 5. If the firewall is also used to control how authenticated users from outside the intranet access certain intranet resources, the intranet becomes an extranet, allowing the same communication benefits that the intranet offers to extend to other organizations such as customers or suppliers.

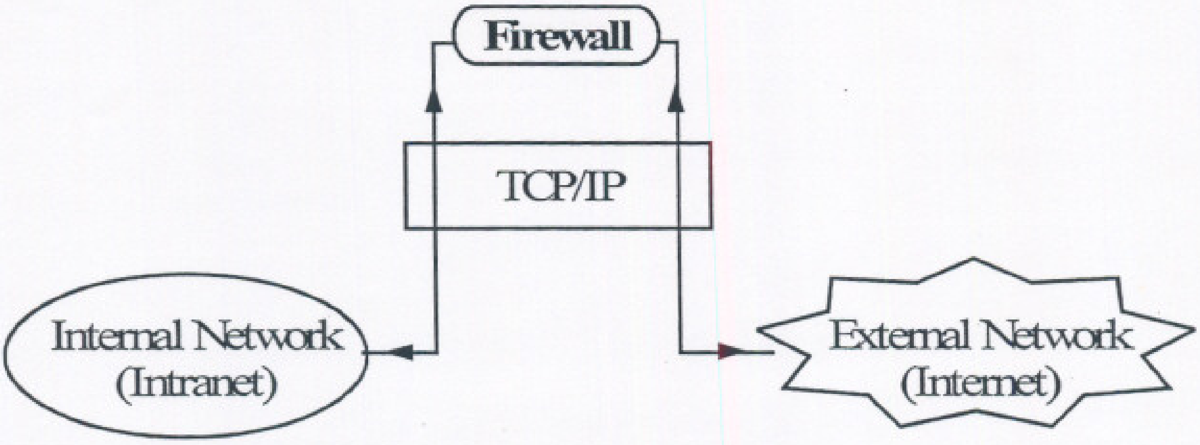

Figure 5. Firewall operation (after Siyan 1995)

REFERENCES

1. Anonymous. 1999. "What is a VPN?", Virtual Private Networks Conference and Exhibition, <http://www.vpncon.com/whatarevpns.htm>, (16 June 1999).

2. Archer, Susan. 1997. "Planning, Installing, and Configuring an Intranet", *Intranet Resource Kit*, Osborne/McGraw-Hill, Berkeley, CA.

3. Biggs, M. 1999. "OS/400 boosts key Web services", InfoWorld, 21(19): 71-86.

4. Black, William C. 1998. Marketing on the Internet. Class Notes. <http://courseweb.bus.lsu.edu/fall98/marketing/4414-1/files/lec1/sld020.htm> (4 July 1999).

5. Casselbery, Rick. 1996. *Running a Perfect Intranet.* Que Corp., Indianapolis, IN.

6. Cedeno, Nancy. 1995. *The Internet Tool Kit.* Sybex, San Francisco, CA.

7. Charles, Gerald T. 1998. *LAN Blueprints.* McGraw-Hill, New York, NY.

8. Cheswick, William R. and steven M. Bellovin. 1994. *Firewalls and Internet Security: repelling the wily hacker*. Addison-Wesley Publishing Co. Reading, MA.

9. Dyson, Peter, Pat Coleman and Len Gilbert. 1997. *The ABCs of Intranets. Plan and Build an Effective Intranet*. Sybex, San Francisco, CA.

10. Held, Gilbert. 1994. *Token-Ring Networks: Characteristics, Operation, Construction and Management*. John Wiley & Sons Ltd. Chichester, West Sussex, UK.

11. Minoli, Daniel. 1997. *Internet & Intranet Engineering. Technologies, Protocols, and Applications*. McGraw-Hill, New York, NY.

12. Moore, Susan and John A. Luoma. 1997. "What Makes Up a Web-Centric Intranet?", *Intranet Resource Kit*, Osborne/McGraw-Hill, Berkeley, CA.

13. Netcraft. 1999. "The Netcraft Web Server Survey", <http://www.netcraft.com/survey/> (8 July 1999).

14. Siyan, Karanjit and Chris Hare. 1995. *Internet Firewalls and Network Security*. New Riders Publishing, Indianapolis, IN.

15. Spurgeon, Charles. 1995. Quick Reference Guide to Ethernet. < http://www.ots.utexas.edu/ethernet/ethernet.html>, (4 July 1999).

16. Stern, Morgan and Tom Rasmussen. 1997. *Building Intranets on NT, NetWare and Solaris: An Administrator's Guide*. Sybex, San Francisco, CA.

17. Trulove, James. 1997. *LAN Wiring. An illustrated guide to network cabling.* McGraw-Hill, New York, NY.