2. INTERNET TECHNOLOGIES OVERVIEW

2.1 CHAPTER INTRODUCTION

The Internet is a collection of computers that communicate using a standard set of protocols. Since there are now millions of computers involved in the Internet, it has grown to be a major means of communication and allows for users to interact with little regard to distance or location. Associated with the Internet is a set of technologies ranging from network protocols to browsers that have been developed to support Internet operations. This Chapter gives a description of the basis of these Internet technologies and how these can be used by corporations to improve their operations.

The structure of the Internet is first described with an overview of network standards and the ISO seven layer network model. TCP/IP is then described together with the use of routers to route the packets and the use of Internet addressing. This leads to a description of the World Wide Web (WWW) and the use of the Hypertext Transfer Protocol (HTTP). The implementation of Internet technologies is then addressed with a description of the growth in the size and usage of the Internet and the future use of Internet technologies. Corporate use of Internet technologies is then described with a description of Intranets and Extranets and the phases that corporations tend to pass through in implementing Internet technologies. An example corporation, Deere & Co., is then profiled on their usage of Internet technologies and how these technologies are beginning to reach into every facet of company operations. A brief overview of the financial justification of one of the Deere & Co. applications of Internet technologies is then given.

2.2 INTERNET NETWORK STRUCTURE

2.2.1 NETWORK STANDARDS

Standards are a necessary part of most technological developments, and have been developed since the early days of the industrial revolution. The use of interchangeable parts by Eli Whitney is an example of the early use of standards, although these standards were necessarily ad-hoc in nature. As the process of industrialization has gathered pace so too has the formulation of standards, ranging from standards of measurements (the metric standard is an example) to that of computer networks. Complete and adequate standards allow for interaction between individuals, groups and corporations since each party can base their operations on the same standards and avoid the needless confusion that will otherwise necessarily result. There would appear to be two main mechanisms whereby standards are formulated:

1. **Standards Organizations.** The first mechanism is where a body, usually international in nature, develops a standard based on a consideration of the multiple factors that are of concern. An example is the Initial Graphics Exchange Specification (IGES) which was adopted as a standard in 1981 to allow for the exchange of Computer Aided Design drawings. This method of developing standards tends to be extremely slow with frequent delays caused by the deliberations of the, usually, many bodies involved. The main advantage of this method of developing standards is that a wide variety of considerations can be brought to bear in the standard and methods are usually developed for maintaining the standard.
2. ***De Facto* Standards.** The second mechanism of developing a standard is what may be thought of as a direct result of widespread use. If, for example, a computer file format is in widespread use then this can become a *de facto* standard. An example is the DXF file format for Computer Aided Design files that was used initially by the AutoCAD computer package. As this package had the main market share of CAD packages on personal computers, then the DXF file format became a *de facto* standard. The main advantage of developing standards in this manner is that of speed since such standards can be very quick to emerge. The main disadvantages are that these standards can change very quickly, they can be proprietary in nature, and there may be no international body that maintains the standard so that it can fragment.

In terms of computer networks then separate standards have been developed by each of these two methods. The standards organization manner of developing standards has resulted in the ISO model while the (approximately) *de facto* manner has resulted in the development of TCP/IP.

2.2.2 ISO MODEL

The International Standards Organization (ISO), based in Geneva Switzerland, is composed of groups for various countries that set standards working towards the establishment of world-wide standards for communication and data exchange. One notable accomplishment has been the development of a reference model that contains specifications for a network architecture for connecting dissimilar computers, with a main goal being that of producing an open and non-proprietary method of data communication. This reference model, called the Open Systems Interconnect Reference Model (OSI RM), was developed in 1981 (?) and revised in 1984.

The OSI RM uses 7 layers, each independent of each other, to allow computers to exchange data. To transfer a message from user A to user B, the data has to pass through the 7 layers on user's A machine, before being transmitted through the selected medium. At the receiving computer of user B, the data must then pass through the 7 layers again, this time in reverse sequence, before being received by user B. For data to be transferred, it must pass through all 7 layers on both computers.

Each layer follows a relatively strict specification and this allows the differing layers to be produced and implemented by different concerns. Each layer can then interface with its neighboring layers even though they may have been developed by different groups.

One way of viewing the activities of the layers is that as the original message passes down the layers towards the medium on the computer of user A, additional information, for both formatting or addressing information, is added to the beginning or end of the message. The additional information is added to ensure correct communication. At the other end (i.e. at user B) this information is gradually stripped off at the data passes through the 7 layers - this time in reverse order.

The layers are arranged in order as follows:

Layer 7, Application Layer. This layer defines network applications such as error recovery, flow control and network access. Note that user applications are not part of the layers.

Layer 6, Presentation Layer. This layer determines the format used to exchange data in such aspects as data translation, encryption and protocol conversion. The data from user A is translated to a common syntax that can be understood by user B. In this way it specifies how applications can connect to the network and to user B.

Layer 5, Session Layer. This layer controls the communication session between computers. It is responsible for establishing and removing communication sessions between computers. Additional address translations and security are also performed. This layer therefore instigates a data transfer session between user A and user B so that an extended data transfer can take place.

Layer 4, Transport Layer. This layer is responsible for ensuring that data is delivered free of error and provides some flow control. This layer ensures that data is transferred as part of the session instigated by the Session Layer.

Layer 3, Network Layer. This layer handles the delivery of data by determining the route for the information to follow. The data is divided into packets with addressing information attached. It also translates address from names into numbers. Intermediate addresses are also attached.

Layer 2, Data Link Layer. This layer defines the network control mechanism and prepares the packets for transmission.

Layer 1, Physical Layer. This layer is concerned with the transmission of binary data between stations and defines the connections. The connection definition includes such aspects as mechanical, electrical,, topology and bandwidth aspects.

At the receiving end the process is reversed so that the binary data that is received is translated back into the original message for User B.

Note that each layer communicates only with it's immediate neighbors. For example the Transport Layer only communicates with the Session Layer and with the Network Layer. Each network architecture can be defined by a set of protocols for each of the layers. This allows for a degree of simplification and modularity in design.

In spite of an enormous amount of worked and effort having been expended on the ISO RM, very little of it is in use compared to TCP/IP (described in the following section). Perhaps one reason is that the OSI RM is extremely complex and it takes a long time to implement all the functions. However a more likely reason is that TCP/IP is in widespread use and has preempted much of the work on implementing the OSI RM.

2.2.3 TCP/IP

The Internet grew out of the Cold War in the 1960s as a response to the issue of making sure that computer networks could survive a nuclear weapons attack. The problem was that a nuclear war could destroy much of the military communications and computer networks and that military control would then be lost. An approach was therefore needed whereby the networks could operate even when substantive portions had been destroyed. A number of possible network structures were proposed, with most using analogue approaches with relatively sophisticated mechanisms for making sure that network connections were maintained. Such approaches were difficult to implement effectively since all possible scenarios of damage to the network had to be preprogrammed into the network control algorithms. This structure was unwieldy and it was difficult to make changes to the network.

Also proposed, but not implemented at that time, was the structure that was to form the basis of the Internet. This approach was based on a simple and elegant digital model of a very decentralized network and is described in more detail below. Such a network is digital in nature and was therefore dependent on readily available computing power. Such power was becoming available in the 1970s and it was then that the University of California at Berkeley received a contract from the United States Department of Defense to develop a computer network that would:

1. Operate on a wide variety of computer hardware with differing communications media
2. Reconfigure itself if portions of the network failed

The earlier proposed network of data packets being directed by routers was implemented under this contract in a network structure formulated as a series of protocols which are described under the general heading of Transmission Control Protocol and Internet Protocol (TCP/IP). TCP/IP is described in more detail later. What helped the adoption of TCP/IP was that Berkeley was also developing a version of UNIX that was available to other academic institutions and TCP/IP was included with the UNIX software tape. Each was available free to academic institutions and the use of TCP/IP gradually grew. Indeed, the widespread adoption of TCP/IP has made it a *de facto* standard.

The TCP/IP model, and hence the Internet, is based on two structures, one for the data being transmitted and the other for the routing computers that would make up the core of the network:

- **Data** is broken down into smaller *packets* for transmittal through the network. Each packet includes the address of the destination computer as well as other information such as the transmitting computer. The packets are reassembled into the data file at the destination computer.
- The **network** is essentially composed of a number of routing computers (or *routers*) that route the packets towards their destination computer by passing them to the next router that is available in the general direction of the destination computer.

The overall approach is therefore that a file of data that is to be transmitted to a destination computer is broken down into packets, each of which contains the address of the destination computer. Each packet is then sent individually through the network, passing from router to router. Each router examines the destination address and passes the packet onto available routers that are in the general direction of the destination computer. At the destination computer, the packets are reassembled into the data file.

An analogy is if you wanted to send a long message using the postal service and you then broke down the long message into postcards. Each postcard (packet) would contain a portion of the message (data) as well as such information as the destination address. The Post Office would then treat each postcard separately, passing them in the direction of their destination. Each time the postcards were examined and sent to the next intermediate stage would be comparable to the operation of the routers. When all postcards arrived at their destination they could be reassembled back into the original message.

For both the Internet and for the postcards, it should be noted that each packet (or postcard) may follow a completely different route and each may arrive at different times. For the digital Internet network there are some steps we can take to improve performance that are difficult to take with the postcard analogy. For example, we are concerned with packets getting lost or delayed, so we can arrange for routers to send duplicate packets on different routes, with the first to arrive at the destination computer being used. Since these are digital signals, duplication can be exact and can be readily done. Such an approach, if done moderately, need not swamp the network and can result in an improvement in network performance.

For the Internet, computers are connected to a router, and each router is connected to several other routers, with each of these other routers being connected in turn to several other routers. The overall structure can be likened to that of a web of routers, linking a extremely large number of computers. The routers continually exchange status information, such as data on transmission delays, with the other routers. In this manner routers can each build up a picture of gaps and delays in the network and can therefore route packets accordingly.

The Internet network model of packets of data being passed from router to router until the destination computer is reached, has some features which have contributed to it's widespread use:

- The model is essentially one of *non-hierarchical control* in character with each router being at the same level of control.
- The model is highly *decentralized* with each router operating quasi-independently. The routers can be programmed to operate on packets in certain ways and they will continue to operate in the programmed manner independently of other routers if necessary.
- The model is also *self-managing* to some extent. As local bottlenecks or gaps in the network occur, then routers are programmed to send packets around problem areas. Obviously, the larger the problem area the greater will be the degradation in overall system performance, so there are some limits on this capability. For transient bottlenecks however this can be an effective approach. For more chronic and permanent bottlenecks then the structure allows for the identification of bottleneck routers or other elements and their replacement with higher capacity elements in a straightforward manner (see *scalable* below).
- The model is also *scaleable* in that we can continue to add (or subtract) routers and computers to the network without changing its essentially characteristics.
- The Internet model is an *open* standard with the specifications being openly and freely available. This has allowed a large number of universities and companies to make Internet capable systems and Internet technologies available without the necessity to undertake licensing or to make royalty payments.

..

2.2.4 EXERCISE 2-1

This is an exercise designed to show how the data routing system inherent in the Internet operates.

You will need one or two packs (two packs are preferred) of ordinary playing cards.

Gather a group of several people together standing in a fairly spread out group (perhaps 1 meter apart). Each will now take on the role of being a router in the Internet.

Choose 1 corner of the group as the destination of the Clubs, and other (different) corners as destinations for the other three suits (diamonds, hearts and spades).

The task of the routers (people) is that when they receive a packet (a card) they should examine it’s address (it’s suit) and pass it to another router (another person) in the general direction of it’s destination. The packets should be passed to a router that is relatively free of other packets. CLEARLY EXPLAIN THIS TO THE GROUP.

Shuffle the packets (cards) so that the suits are completed mixed together, and give perhaps 10 packets (cards) to each of several routers (people) selected randomly throughout the group. These routers have just received the packets from a computer host.

At a given signal instruct the routers (people) to begin their task of routing.

Note especially the following which *should* happen:

- While there appears to be general confusion with packets going in all directions, the packets will usually reach their destination.
- There are many different routes that packets can pass from source to destination
- Local log jams are avoided since the routers pass the packets to routers that are relatively unloaded.
- If a log jam does happen then it can be easily seen. In the Internet a router that becomes a log jam can be replaced by one with a higher capacity.

Repeat the above but this time simulate the breakdown of a router (person) during the above process. Arrange for one router to pass on all the packets they have as usual but to suddenly stop accepting packets from other routers. What should happen is that the other routers pass packets around this now defunct router and that packets still reach their destination, even if it does take a little longer. This shows the self managing aspects of the Internet operation

2.2.5 TCP/IP STRUCTURE

TCP/IP consists of a whole series of protocols, applications and services that together support the powerful capabilities of Internet technologies. Whereas the OSI RM has seven layers the TCP/IP can be thought of as consisting of five layers:

The **application layer** containing such protocols and applications as Simple Mail Transfer Protocol (SMPT), File Transfer Protocol (FTP), Hypertext Transfer Protocol (HTTP) and Telnet.
The **transport layer** contains such protocols as Transmission Control Protocol (TCP) and User Datagram Protocol (UDP).
The **Internet layer** contains such protocols as Internet Protocol (IP), Internet Control Message Protocol (ICMP), Address Resolution Protocol (ARP), and Reverse Address Resolution Protocol (RARP)
The **data link layer** and
The **physical layer** handle the hardware connections. A wide variety of hardware network connections are possible ranging from token ring to Ethernet and from twisted pair cables to fiber optic cables.

As with the OSI RM, each message that is being transmitted must be passed down through the layers to the hardware, while the reverse happens on the receiving machine.
Another way of looking at the TCP/IP set is to categorize it into upper layers (the application layer) mid-layers (the transport and Internet layer) and into lower layers (the data link and physical layers). The upper layers handle the applications, while the lower layers handle the hardware connections. The mid-layers form the core of TCP/IP.
Note that not all protocols, applications and services are used on all sessions, rather each can be used under particular scenarios. For example TCP/IP is often combined with Ethernet for a specific hardware implementation. Ethernet is a protocol for the data link and physical layers that uses carrier-sense multiple-access with collision detection (CSMA-CD). It is simple to install and is available for a wide range of computer hardware. An arrangement such as TCP/IP with Ethernet (with twisted pair cables) therefore covers the mid and lower layers. Alternatively TCP/IP may connect to a token ring network where transmission is only allowed if the node has a token passed to it from another node on the network. In this case the data link and physical layers of the TCP/IP protocol suite are concerned with the interface to the token ring network. It should be noted that this modular approach allows for a wide variety of configurations to be possible. Ironically, the modular approach also allows for the use of Internet technologies that address the upper layers (such as WWW browsers) to use non-TCP/IP networks (Novell networks as an example).

2.2.6 TCP/IP PROTOCOLS

There are several TCP/IP protocols that allow for certain activities to be carried out over the Internet. These protocols tend to be concerned with the mid and upper levels of the TCP/IP suite of protocols.
The Application Layer

File Transfer Protocol (FTP) is a protocol that regulates file transfers between local and remote systems. The two systems connected use a FTP server and FTP client software and FTP supports bi-directional transfer of both binary and ASCII files.
Telnet is a protocol where terminal emulation on one machine is supported to allow for remote login to another machine.
Simple Mail Transfer Protocol (SMPT) is a protocol that determines how e-mail is transmitted on the Internet.
Post Office Protocol (POP or POP3) allows for mail to be handled at the local level and is defined by RFC 1721. A POP3 mailbox stores mail received by SMPT until it is read by the user and also passes mail to the SMPT server for transmission.
Network News Transfer Protocol (NNTP) this is the protocol used by Usenet to deliver newsgroups.
Hypertext Transfer Protocol (HTTP) – see below.

The Transport Layer

Transmission Control Protocol (TCP) allows for packets to be delivered and carries out error checking and sequence numbering. TCP also sends out instructions to the transmitting computer to re-send lost or corrupted data.

User Datagram Protocol (UDP) carries out a similar function to TCP but without the error checking or sequence numbering. Hence faulty or missing packets are not instructed by UDP to be resent.

The Internet Layer

Internet Protocol (IP) handles the addressing of packets.

Internet Control Message Protocol (ICMP) reports problems and handles some network functions. For example ICMP can be used for the ping command to test the network transmission times.

Address Resolution Protocol (ARP) resolves some network addresses from the TCP/IP address to the network interface card address.

Reverse Address Resolution Protocol (RARP) operates by resolving from the network interface card address to the TCP/IP address.

TCP/IP is therefore a suite of protocols and applications, with differing protocols and applications being used for different requirements. Note that it is however necessary to have at least one protocol/application from each layer for each session.

The modular design of TCP/IP helps in being able to develop new protocols and applications. As a consequence new additions are continually being made to the TCP/IP protocol suite. The TCP/IP protocols are overseen by the Internet Activities Board. This Board works closely with the Internet Society which has three main groups that are concerned with various aspects of the Internet. These are the Internet Engineering Task Force (IETF), which is concerned with the day-to-day running of the Internet, the Internet Engineering Steering Group (IESG), which is concerned with setting strategic goals for the Internet, and the Internet Research Task Force, which is concerned with research in the core technologies of the Internet.

2.2.7 INTERNET ADDRESSING

The Internet Protocol (IP) uses numbers to identify host computers and uses these address numbers to route data between them. The IP addresses are 32 bit (or 4 byte) binary values, for example:

10000000.11111111. 00010111.10111100

These are usually expressed in decimal with a period between the bytes for convenience. The above IP address would therefore be expressed in decimal as follows:

128.255.23.188

Each site connected to the Internet has it's own IP address and messages can be addressed using this number. The routers on the Internet then pass the message through to its address.

The maximum value of each byte is 11111111, which is 255 in decimal. Theoretically, at least, there are therefore 255^4 (which equals 4,228,250,625) possible Internet addresses. This would seem to be more than enough but, because of the way in which blocks of Internet addresses have been allocated, some sub-sections are running out of numbers and some changes will have to be made in the future.

The above numbering scheme was thought to be somewhat difficult to use and early in the development of the Internet a parallel naming scheme was begun. This uses descriptive words for the site address, so that *microsoft.com* for example could be used instead of it's IP address of *207.68.137.53*. The system that operates this is called the Domain Name System (DNS). Lookup tables are incorporated into the Internet to convert from the more descriptive form, the DNS name, to the IP number address. A message can be addressed by the user to, for example, microsoft.com, and as it enters the Internet, this address is converted to it's numeric form which then used for the remainder of it's transmission through the Internet. Additional parameters can be added to the address so that the message can be properly handled at its destination. For example, www.microsoft.com would indicate a message that would be handled by the WWW server at *207.68.137.53*.

In the DNS, domain types are allocated to particular categories. For example, in U.S. then the categories are allocated as follows:

.com Commercial (e.g. intel.com)

.edu Education (e.g. uiowa.edu)
.org Organization (e.g. ims.org)
.gov Government (e.g. nsf.gov)
.mil Military (e.g. navy.mil)

The countries are also identified. For example:

.jp Japan
.kr Korea
.uk United Kingdom
.de Germany
.nl Netherlands

The DNS name is usually cascaded. For example,

www.eng.cam.ac.uk

refers to the WWW site in the Engineering Department (eng) at Cambridge University (cam) which is an academic institution (ac) in the United Kingdom (uk).

The lookup table referred to above is, in fact, a relatively large database which is replicated at points in the Internet. For the U.S. there is, at present, a central operation called InterNIC, that registers computer hosts with their corresponding DNS name and IP address. The updated databases are distributed at frequent intervals to the points on the Internet where the DNS name to IP address conversions are carried out.

2.3 THE WORLD WIDE WEB

While the Internet provided powerful capabilities in such utilities as telnet and FTP, it was not particularly easy to use. This began to change in 1993 when researcher at CERN in Switzerland developed a means of sharing data using *hypertext*, where codes in the document being examined allowed users to jump to another document merely by clicking on a *hyperlink*. FTP and telnet capability were added so that they could also be invoked merely by clicking on a hyperlink. This type of program became known as a browser and while the CERN browser was limited to text documents, a team at University of Illinois at Urbana-Champaign (specifically the National Center for Supercomputer Applications – NCSA) developed a more powerful browser called Mosaic which allowed for the inclusion of graphics. Mosaic was freely available and led to a huge increase in the use of the Internet and WWW. Some of those involved in the development of Mosaic helped form Netscape Corporation, which has developed commercial versions of both browsers and servers.

2.3.1 WWW CLIENTS AND SERVERS

The WWW, in it's early form, is a very large collection of clients and servers that support the HyperText Transfer Protocol (HTTP) on the Internet. This is an open standard and is implemented on a wide variety of platforms. The operation of this early WWW can be thought of as being divided into two portions: clients/browsers on one side and servers on the other.

The clients help users form a request, send it to a server, and present the users with the results from the server. This is most frequently done by the user clicking on a hyperlink containing hidden codes that allow the client/browser to formulate a request for a document from a server. The request is in the HTTP format and usually consists of a GET command (see below).

A server receives and validates the request, retrieves data, and delivers them to the requesting client. The server therefore usually serves as a data storage site and then becomes responsible for checking the request for security permissions, retrieving the document from its disk drives and sending the document to the requesting client.

The browser/client-server model describes communication between service consumers (clients) and service providers (servers). This model allows for processes at different sites and on different computer systems to exchange messages interactively. Note that the interaction is intermittent: data is exchanged in bursts of data packets and as the last one is received all communication is halted between the sites until one of the sites instigates another communication session.

Since the data can be plain documents, sounds, or images, proper viewers need to be launched on the client side. WWW clients use WWW browser software, such as NSCA Mosaic, Microsoft Explorer or Netscape Navigator, to access and view such data. The usually manner in which information is displayed is by using Hypertext Markup Language (HTML) which consists of ASCII text that indicates the text or graphics to be

displayed by the browser as well as commands, hidden from normal view, that dictates such aspects as formatting, hyperlinks or higher level information.. The complete range of data types are however possible. For example, VRML, introduced in 1995, is a language for describing multi-participant interactive simulation (Bell et al., 1995). VRML is capable of creating virtual worlds networked via the global Internet and hyperlinked with the WWW. The aspects of virtual worlds, including display, interaction, and Internet working, can be specified using VRML.
VRML viewers are companion applications to standard WWW browsers for navigating and visualizing. The objects contained in the viewers have hyperlinks to other worlds, Hyper-Text Markup Language (HTML) documents, or other valid Multipurpose Internet Mail Extensions (MIME) types (Grand, M., 1995). When a user selects an object having a hyperlink, the appropriate MIME viewer is automatically launched. For example, when the user selects an anchor to a VRML within a correctly configured WWW browser, the corresponding VRML viewer is launched. The user at the client is then able to manipulate a 3-D graphical representation (Fig 2-1.).

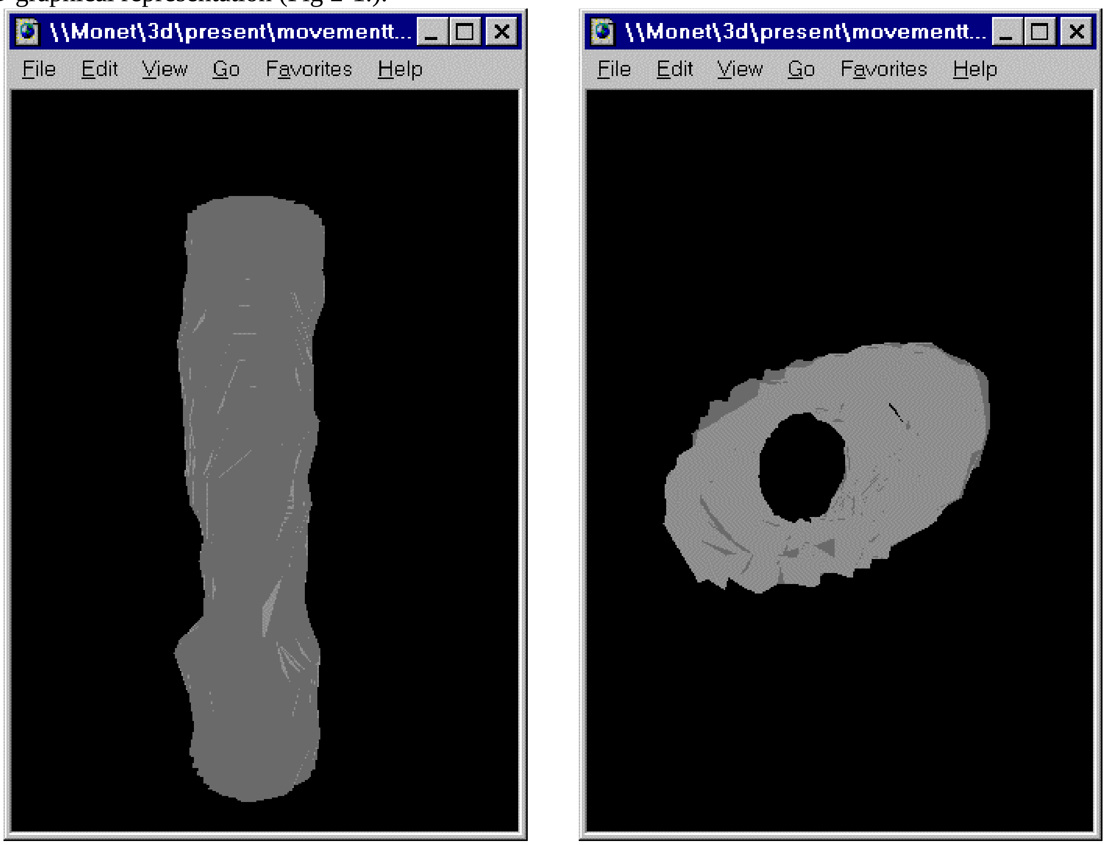

Figure 2-1: Examples of VRML manipulation

2.3.2 WWW DEVELOPMENTS

The interaction between WWW servers and clients can be classified as follows:

1. The Web server sends a static file to the client as a result of a HyperText Transmission Protocol (HTTP) request from the client. This static document can be in any format but the formats that are recognized readily by browsers include Hyper Text Mark Up Language (HTML), Virtual Reality Modeling Language (VRML), and image files that are in the standard formats. Other formats may invoke software on the client.
2. The WWW server can process data in response to input from client browser. Such processes can include, for example, extracting information from the corporate databases in response to the client browser requests.

3. A program can be downloaded from the WWW server to a client which can then carry out the programmed actions on the client. Programs written in Java are one example. Alternatively the program may reside on the client computer but is invoked by a command from the server. This later mode is frequent with ActiveX components.

2.3.3 HTTP 1.0 AND 1.1

HyperText Transfer Protocol is the main protocol that underlies the WWW, operating with TCP/IP to allow hyperlinked text and images to be requested by clients from servers.

HTTP operates on a request- deliver basis. When a hyperlink is clicked, a TCP/IP connection is opened with the server concerned and data is requested. The TCP/IP connection is closed when a response has arrived (or after a timeout period has elapsed). The main HTTP commands are GET and POST. GET is a request for data from a server, while POST is used to send, for example, the response from a HTML form to the server.

Under HTTP 1.0, requests are made for each element of a WWW page so that multiple graphics on a WWW page, for example, will each generate their own individual request. Each will therefore need their own, separate connection. This operation can be easily seen on the status bar of a WWW browser if a page with multiple graphics files is opened. This may not be a problem with simple WWW pages and with good Internet capacity but as WWW pages grow in complexity and as the WWW traffic is growing exponentially (Chapter 3), the efficiency of this opening and closing TCP/IP connections for each request has recently been the subject of much scrutiny.

Three main problems with this HTTP 1.0 procedure are evident. The first is that establishing a connection is time consuming and may not be successful. This lowers the overall efficiency of the HTTP procedures. The second is that each HTTP request has to wait for the previous request to receive a response before it can begin. The multiple waiting periods and time spent establishing and severing multiple connections can become onerous for even moderately complex WWW pages. The third is that WWW traffic is increased since each request and response is sent as a separate message.

HTTP 1.1 aims to address these difficulties by having a more efficient interaction with TCP/IP that permits persistent connections that remain open for multiple requests. Under HTTP 1.1 a WWW page with, for example, nine graphics files would require just a single connection, compared with nine connections with HTTP 1.0. In HTTP 1.1, requests are buffered and then sent together over the (single) connection, responses similarly are made together. This procedure is known as pipelining.

There is another advantage of HTTP 1.1 that is associated with the packet size sent. As indicated above, data is sent over the Internet in packets. The size of the packets sent can range from 1 byte to several kilobytes, with 576 bytes usually used as the smallest packet to avoid packet fragmentation. Using persistent connections allows for larger packets to be sent with an overall increase in network efficiency.

One other aspect should also increase efficiency and that is to do with the Internet naming system, the DNS: instead of typing the IP address directly (e.g. 128.255.23.188), most will type the address as www.iil.ecn.uiowa.edu/internetlab and the Internet translation system translates the written address into the IP numbers. This however takes time and is another drain on the Internet capacity. Persistent connections in HTTP 1.1 will reduce this need for multiple translations by caching of the domain name.

HTTP 1.1 will require that browsers and servers are updated but does promise an overall increase in efficiency in WWW operations. Initial tests by the W3 Consortium (http://www.w3.org) show that WWW page download times will be reduced by about 30% and that this will be combined with a decrease in the net WWW traffic generated.

2.3.4 COMMON GATEWAY INTERFACE (CGI)

CGI is a standard for interfacing external applications with WWW servers (Common Gateway Interface, 1995). CGI is the interface between the server of Web site's HTTP (HyperText Transport Protocol) and the other server's host computer. CGI is not really a language or a protocol in the strictest senses of those terms but rather is a set of commonly-named variables and agreed-upon conventions for passing information back and forth between the client (the user's WWW browser) and the server (the computer that sends Web pages to the client). CGI gives programmers a way for HTML WWW pages to call external programs and get back the results. A WWW daemon in general can retrieve static files such as HTML documents, which typically change only infrequently. On the other hand, a CGI program can provide a way to handle dynamic information in real-time. For example, consider a UNIX database which users query on the Internet. This requires the dynamic features, such that 1) a number of users may query the database with each requesting different data, 2) each of the users using the same data may require different results based on his/her own context, 3) the users may want to perform interactive operations with the information servers on the WWW, and 4) the database itself is continuously updated. A CGI program can support these dynamic features.

Once a CGI program is created, a WWW daemon can execute it. The daemon transmits the queries provided by users to the CGI program. This program first retrieves the requested data from the database. It then converts the retrieved raw data into a proper MIME type. Finally, the daemon receives the resulting information and displays it on a client site.

Information can be passed from the Web browser to the CGI program in a variety of ways, and the program can return the results with embedded HTML tags, as plain text, or as image. The Web browser interprets the returned results just like it does any other document. This feature of CGI is helpful in retrieving database information from WWW. The procedure to retrieve the database with CGI can be described as the following (see Figure 2-2):

1. Step 1. The user calls a gateway program using CGI, usually by clicking on a hyperlink or pressing a button in the Web browser.
2. Step 2. The Web browser collects the information entered by the user.
3. Step 3. The browser contacts the HTTP server where the CGI program resides, asking the server to find the CGI program and pass it the information.
4. Step 4. The CGI program is executed.
5. Step 5. The CGI program processes the passed-in information into the format which the database recognizes.
6. Step 6. The CGI program passes the database query to the database interface.
7. Step 7. The database interface sends the database query to the database.
8. Step 8. The database performs the query and passes back any results to the CGI program through the database interface.
9. Step 9. The CGI program formats the results and sends them back through the CGI to the server.
10. Step 10.The Web browser displays the results.

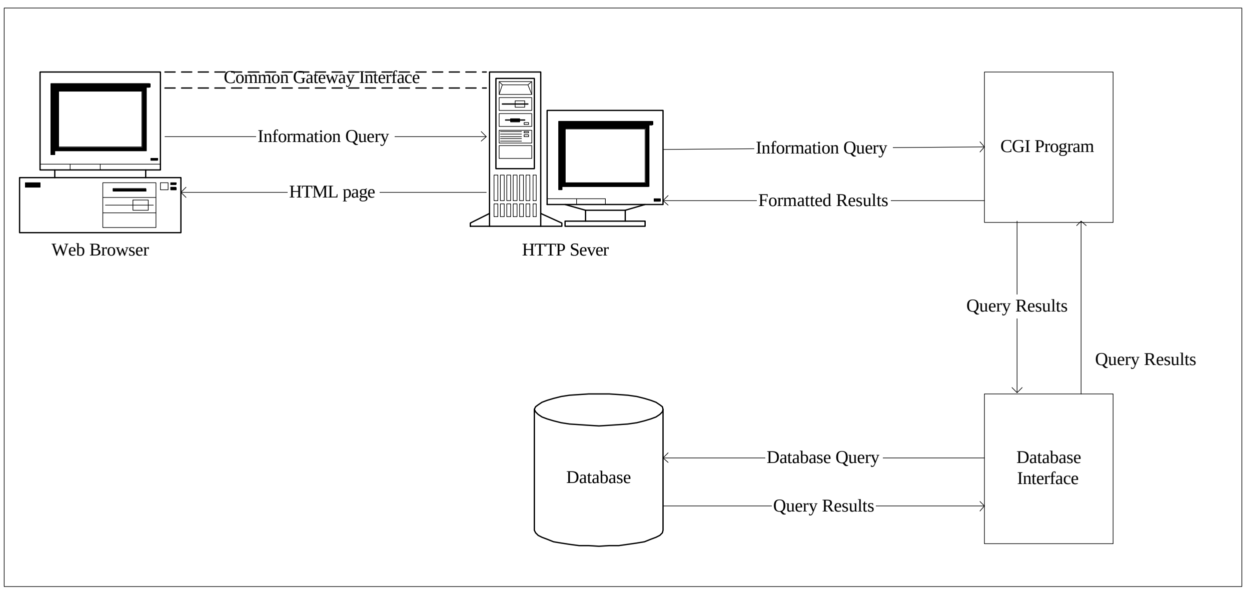

Figure 2-2: Retrieving Database Information Using CGI

2.3.5 USE OF APIS

Each time a new CGI program is started a new process is begun and this is notoriously inefficient. For a server with multiple CGIs executing simultaneously, the result can be very slow response times. Recently, therefore, much attention has begun to be focussed on more efficient methods of executing programs on the server. Much of the developments has been on application program interfaces (API) which, in their pure form, are a set of functions in the operating system that programmers can use. Server software that includes an API therefore permit a more direct use of the operating system. The result can be that programs that use the API are much more efficient that those that use CGI. Much effort is underway in developing both the API and programs that use the API to take advantage of this increased efficiency.

However APIs tend to be proprietary and therefore are particular to server/operating system combinations. An example is the Internet Services Application Programming Interface (ISAPI) which is proprietary to Microsoft. ISAPI allows for programs to be dynamic link libraries (DLLs) for use on the Windows NT server operating system and only one copy of the DLL needs be loaded no matter how many programs use it.

Results would indicate that servers that use APIs have substantially shorter response times than those that use CGIs and we can therefore expect that CGIs will tend to fall into relative disuse.

2.4 GROWTH OF THE INTERNET

The Internet has experienced a very rapid growth in popularity in terms of the number of sites that are connected to the Internet, the number of users and in the amount of data that is sent over the Internet.

2.4.1 NUMBER OF INTERNET SITES

The number of sites that are connected to the Internet is surprisingly difficult to pin down. Part of the reason for this is that the number is changing with rapid, exponential growth being the norm. The second is that the Internet is now very far reaching so that it is difficult to estimate the number of computers connected, for example, behind corporate firewalls or in the military. The global nature of the Internet also posses difficulties in estimating the number of sites since management of the Internet is decentralized to each country.

Perhaps what is of more relevance is to look at the *growth* of the Internet, to discern trends and to begin to forecast the Internet as it will be in the future (modeling growth is discussed in Chapter 3). The growth in the number of domains is shown in Fig. 2-3 (source: Network Wizards). This shows the number of registered sites for the com, org and edu domains, for U.S. It does not include for example overseas domains, nor the gov or mil domains, nor does it include internal Intranets. Even with these exclusions we can see that there has been very rapid growth from around 100,000 domains in July 1995 to around 800,000 domains in November 1996. This represents an 8-fold growth in 16 months. Of course such an exponential growth rate is difficult to maintain so that we should expect a slowing of growth in the future as the number of companies and other organizations connected to the Internet reaches saturation.

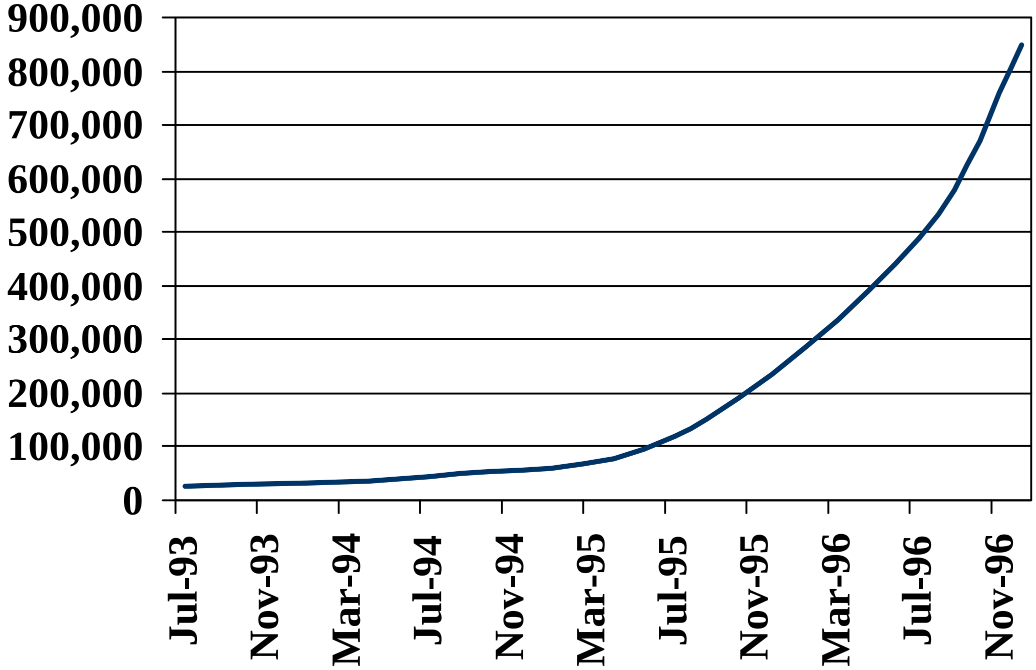

Figure 2-3: growth in the number of domains

Much of the growth has been fuelled by the WWW. Some data on reasons for setting up a corporate WWW site is shown in Fig. 2-4 (source: Computerworld) with the reasons given being cost savings (35%), customer service (32%), revenue generation (18%) and marketing (13%). When more detailed figures are examined then it can be seen that there is a correlation between size of company and their having a WWW site (Fig. 2-5, source: Gilder Technology Report: Feb. 1997). With larger companies represented by the Fortune 500 group then over 80% have a corporate web site. This figure drops with decreasing size of company with just under 70% of Fortune 1000 companies having a web site, just under 40% of companies with greater than 50 employees and with just over 20% of all companies having a web site. It should be remembered that these figures are changing rapidly since it is likely that all larger companies will soon have a WWW site and that WWW sites will be the norm for most companies of all sizes.

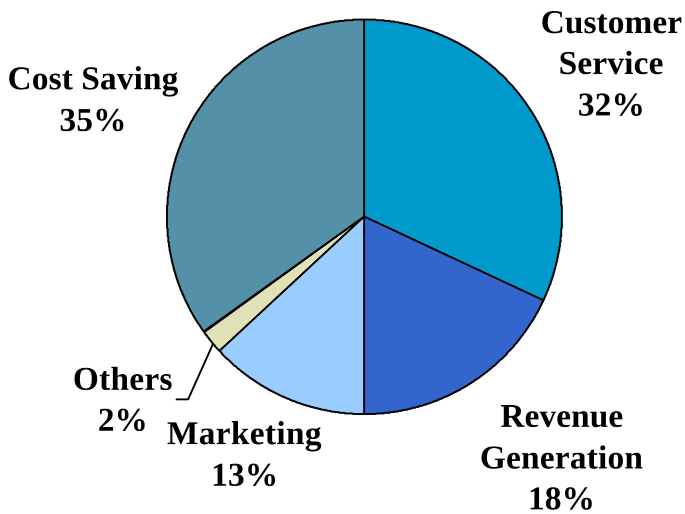

Figure 2-4: reasons for setting up a corporate WWW site

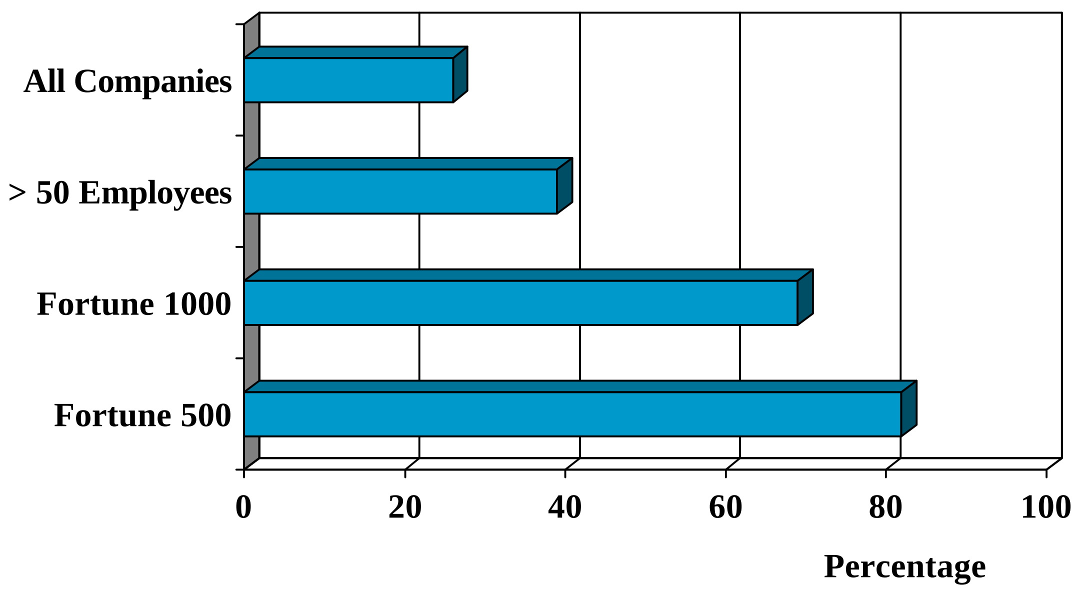

Figure 2-5: Correlation between size of company and their having a WWW site

2.4.2 INTERNET TRAFFIC

As with estimating the number of sites connected to the Internet, it is difficult to estimate the amount of traffic using the Internet. The reasons are similar: the numbers are changing rapidly, much of the traffic is hidden behind firewalls and the overseas traffic is difficult to estimate because operations are decentralized to each country. However the data traffic in the core U.S. Internet is shown in Fig. 2-6 (Source: Gilder Technology Report, Feb. 1997) where it can be seen that the core traffic has jumped from 50 terabytes per month in August 1995 to 450 terabytes per month in December 1996. This represents a 9-fold increase in 16 months and can be thought of as due to an increase usage by those connected, an increase in the number connected, and an increase in the data transmitted as bulky graphics, sound and video data increasingly take the place of undemanding text data.

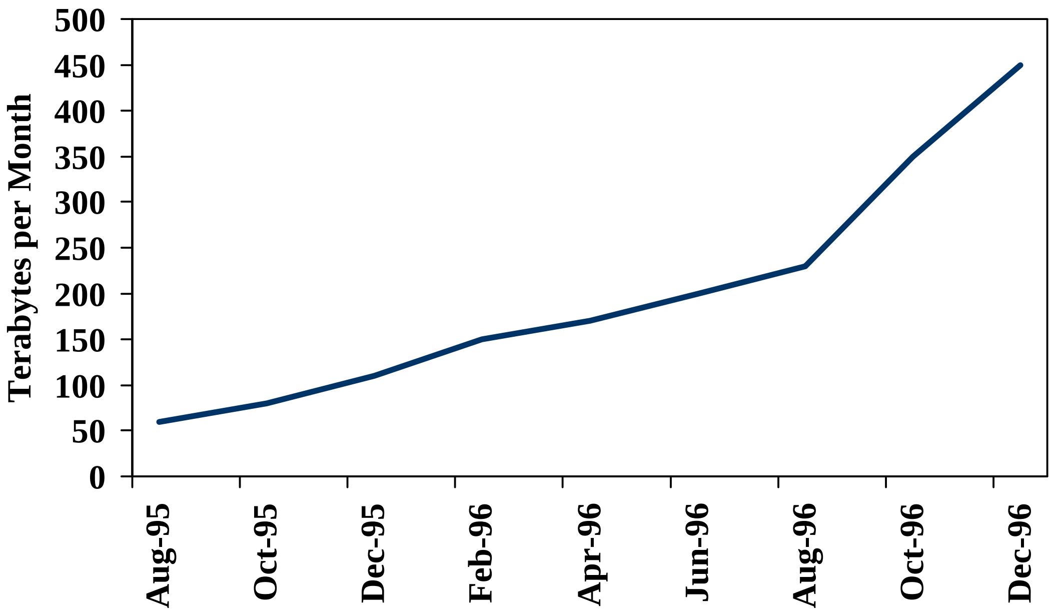

Figure 2-6: Data Traffic in Core Internet

Whereas with the number of Internet domains we can see a natural limit as saturation takes place, it is not clear that such a limit applies to the amount of Internet traffic. The requirements of video in particular promise to mean a long term substantial growth in Internet traffic. One aspect that will limit the growth of Internet traffic is that of congestion on the Internet itself as routers become overloaded. There have been two schools of thought on the issue of Internet overload.

- The first suggests that the Internet capacity is a finite resource but that users are not charged for each use. There is therefore every incentive for users to overload the capacity of the Internet causing substantial delays. An analogy can be used of common grazing land where users can allow their animals to graze free of charge. What happens is that each user has an incentive to increase the number of animals they put on the common grazing land. The result will be over grazing and a deterioration of the common grazing land. Similarly, there will be the tendency for users to increase their own use of the Internet with the inevitable result of the Internet being bogged down.
- The second school of thought more optimistically suggests that the Internet can continue to grow and response time can remain reasonable. This would suggest that the analogy with common grazing land is limited in that the common land is fixed in sized, whereas the Internet can continue to grow.

What evidence there is would tend to support the second school of thought: the fact that the Internet is based on a very decentralized model allows for bottlenecks to be relatively easily identified and rectified and one would therefore expect that because of it's elegant decentralized and modular structure, the Internet will be resilient in the face of substantially increased traffic. Routers on the Internet continually collate the return time for test packets and this is called the ping delay. Figure 2-7 (Source: Gilder Technology Report, Feb. 1997) shows the average ping delay in the core U.S. Internet. This shows a reduction from 40 milliseconds in October 1995 to under 20 milliseconds in October 1995, in spite of a considerable increase in data traffic (Figure 2-6). The implication of this is that the capacity of the Internet is rising *faster* than the demand and this would tend to refute suggestions that the Internet is starting to deteriorate.

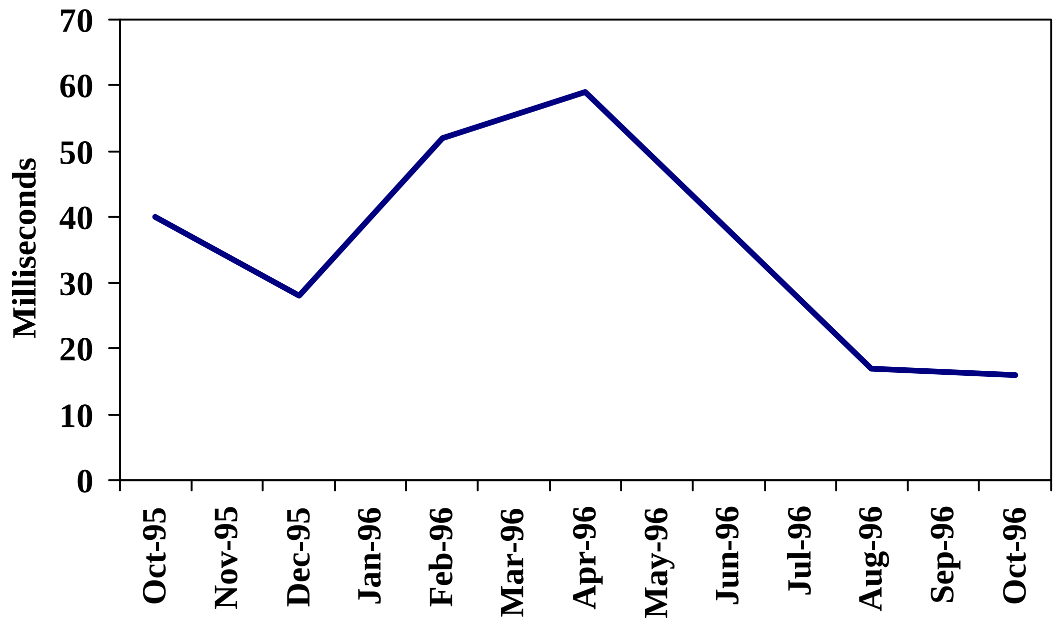

Figure 2-7: Average Internet Ping Delay

Whereas Moore's Law (Chapter 3) suggests a doubling of computer processing power every two years, others have suggested that communications capacity is increasing by a factor of 3 each year, the evidence suggests that the capacity of the Internet is increasing at a much faster rate. Combining the data from Figure 2-6 and Figure 2-7 would suggest that the capacity of the Internet is increasing at a rate approaching 8 times per year. If the rate of 8 times per year remains in effect, then in five years the capacity of the Internet will be a factor of 8^5 (= 32,768) times the present capacity. Such rapid rates of growth may moderate as usage saturates and the capacity may never reach this figure, but there appears to be no evidence of any slowdown in growth rate as yet. Since the use of Internet standards in Intranets is not included in the above analysis, then it may be that actual growth may be larger than the figures above.

In some ways the Internet may be self limiting since if capacity limitations are reached and delays increase, then users will be frustrated and start to moderate their usage. The converse can also be valid: if the Internet increases in capacity faster than usage, then response times will decrease and usage will increase as users are happier with the performance of the Internet.

2.4.3 FUTURE USE OF THE INTERNET

The above has suggested that the Internet may well undergo a massive increase in capacity. Such an increase in capacity will be driven by a (rough) matching in usage and this begs the question: what will the Internet be used for, to account for such a large increase in usage (as measure by the amount of data)? The increase in usage will be composed of a number of overlapping elements:

- An increase in the number of users - at present a minority of the population in USA, and the world, use the Internet regularly but the number of users will grow as the Internet becomes more ubiquitous.
- An increase in the number of computers connected to the Internet - this includes regular computers but also telephones, TVs, radios and other devices that will increasingly move to a digital operation.
- An increase in audio and video data which is much more data intensive than conventional text data.

One vision of an Internet household of the future would be where there would be an Internet connection (with the capacity of a fiber optic cable) to the house, with the household containing multiple machines that are Internet compatible including TVs, radios, cameras, phones, videophones, and computers. Such a vision would mean, for example, that conventional TV delivered via antenna or coax cable would be replaced by an Internet based system that would permit users to receive TV from any Internet connected source so that they can receive TV programs either live or on demand. Consumers would therefore be able to have much more flexibility in their TV viewing.

Another vision for a commercial company is of much of the business being conducted via Internet connections both within the company but also with suppliers and customers. All company activities would be carried out using Internet technologies (with much of it removed from the public Internet by being on corporate Intranets) including requests for bids, proposals, purchasing, customer ordering, financial accounting, scheduling and production planning and control. Communications within corporation and with suppliers/customers will be substantially improved and one effect will be the shrinking of distance. It may no longer be substantial disadvantage to be some distance from the customers and so one effect may be that customers seek out suppliers that are better performers even if they are some distance away. From the suppliers viewpoint, one effect may be that the better suppliers increase their market share as they can now supply to a much larger constituency.

2.5 USING INTERNET TECHNOLOGIES

2.5.1 INTERNET COLLABORATION

One area that is developing quickly is a category of computer software known as “groupware”. Groupware can be fundamentally defined as computer-based systems that provide support for groups engaged in a common task within a shared environment. Thus, groupware could assist individuals even if team members are geographically dispersed (Grantt, 1993; Marca and Bock, 1992).

Groupware can be divided into two categories, namely proprietary and open. Proprietary groupware uses, as the name suggests, proprietary formats while the open type of groupware uses open formats. Examples of groupware include Lotus Notes (proprietary) and the Internet/World Wide Web (open). Note that Lotus Notes does include the capability to interface with the Internet.

We can categorize communication approaches into three main categories:

Unidirectional - where information flows, primarily, in one direction. Conventional WWW pages based on text and graphics are one example. One Rockwell facility, for example, has their design for manufacturing guidelines posted as WWW pages allowing for Concurrent Engineering team members to access this information when needed. Such unidirectional approaches do not readily allow much in the way of collaboration but can be useful in providing background material.

Asynchronous two-way - where information is passed asynchronously. E-mail is an example of asynchronous two-way communication. Another example is the use of news groups, where users can post documents to a central location allowing other users to read the documents, comment on them, or post their own documents. Such asynchronous two-way communication can allow for improved communication, especially where team members are geographically separated. However, the degree that such asynchronous two-way communication can aid collaboration is limited.

On-line Collaboration - where team members can interact synchronously to allow team based activities to take place over a computer network.

It should be noted that the base Internet technology of network communication allows for the transmission of packets of data. Such packets can be of <u>any</u> form of data, including text, graphics, voice, and video. On-line collaborative systems are being developed that allow for rich collaboration including the sharing of text, graphics, voice, and video information (Fig. 2-8).

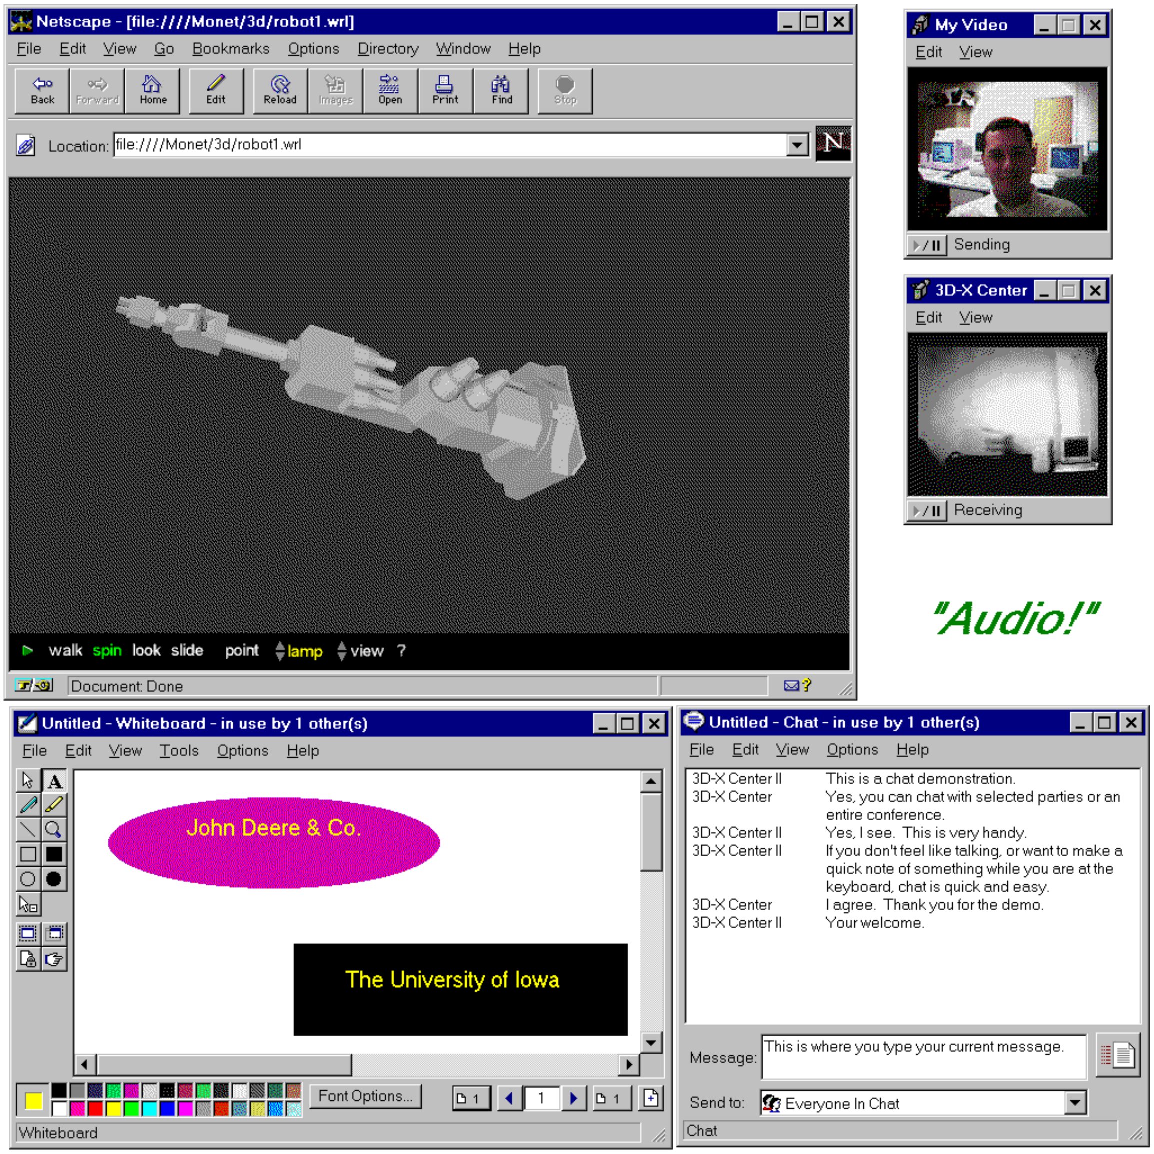

Figure 2-8: Example screen of Internet technology for collaboration

2.5.2 The Internet, Intranets and Extranets

Much of the above descriptions have concentrated on describing the Internet, which is the publicly accessible global network. The tremendous growth in all aspects of the Internet has meant that a number of technologies have been developed that can be very attractive for implementing in *private* networks. These Internet technologies include routers, network cards, TCP/IP protocols, browser software, server software, and conferencing software. Because of the high volume of use of these technologies they tend to be of high quality and low cost. They are therefore very attractive to corporation for implementing private networks (called Intranets) that use Internet technologies but which are kept private from the public Internet, although they are usually connected to the Internet. Privacy can be maintained by several different methods including the use of firewall routers that prevent outsiders from gaining access to the Intranet.

The implementation of such Intranets has shown remarkable recent growth in recent years often with special provisions in the firewall to allow suppliers and customers access under controlled conditions from the public Internet. An Intranet that includes access by outsiders such as suppliers and customers is called an Extranet.

2.5.3 IMPLEMENTATION PHASES FOR INTERNET TECHNOLOGIES

Corporate use of Internet Technologies can be viewed as going through several phases:
Phase I – The Brochure Stage - Establishing a WWW Presence. The corporation establishes a WWW site on the public Internet and uses this site as an electronic brochure to aim to emulate their printed publicity material. This is unidirectional communication. The direct impact of this stage is relatively small – while the cost of the site is low, the direct financial return is negligible, but can be directly justified as changing, or maintaining, the corporate image of a corporation that uses high-technology approaches. However the main benefits for Phase I are indirect: The corporation begins to understand the Internet technology and begins to see how the technology can be used in the future.
Phase II - The Handbook Stage – Disseminating Information Within the Corporation. The corporation establishes an internal WWW site (an Intranet) and uses this to disseminate relatively static information such as design guidelines, company policies/procedures, and press releases. This is unidirectional communication. During this stage e-mail often begins to be used within the corporation (if it is not previously in place) and as the number of users connected to the network increase so does the use of e–mail and is usually the first use of asynchronous two-way communication. The direct benefits from the Handbook Stage are usually enough to justify the investment (see, for example, the John Deere example in Section 2.6.4). However, as is the case with the Brochure Phase, many of the benefits are indirect: Much of the infrastructure required for networked computers is implemented and the employees begin to understand how Internet technologies can be used for more ambitious activities. A foundation is therefore being put in place for more substantive use of Internet technologies within the corporation.
Phase III – The Business Form Stage – The corporation begins to use Internet technologies as a means of improving two way communication. To begin with, such two way communication is limited and mainly consists of on-line HTML forms that can be filled out and submitted by the user. Such forms can be used, for example, for vacation requests, for fringe benefit changes to retirement or medical plans, or for ordering items. As experience in using Internet technologies grows, then the forms are used as means for on-line database searches, where the servers are interfaced to the databases and the users can obtain information from the databases according to their entries on the on-line forms. The Business Form stage is a type of asynchronous two-way communication. The benefits, although sometimes difficult to quantify, are usually sufficient to justify the investment but again many of the benefits are more indirect and more long-term. Perhaps the main benefit is *educational*: users gradually develop an understanding of Internet technologies and confidence grows in the use of these technologies.
Phase IV – The Coalesce Stage – Collaboration using Internet Technologies. The corporation uses Internet technologies to enable a close working relationship between users within the corporation and also with users in suppliers and customers. Such collaboration means the use of such Internet technologies as video conferencing, audio, information sharing, 3D visualization, discussion sites, chat, and whiteboard. The result of this is that the activities of the corporation can begin to merge and coalesce into more streamlined activities.

2.6 EXAMPLE CORPORATE USE - DEERE AND CO.

2.6.1 DEERE AND CO. BACKGROUND

Deere & Co. is a corporation that has specialized in agricultural equipment. It was founded as a one man blacksmith shop in 1837 and now employs 34,000 people in 160 countries with revenues of around $11 billion per year. Although the primary focus of the company is agricultural equipment, Deere & Co. are also engaged in insurance and health care activities. Deere & Co. is a useful example company since it is a mainstream corporation engaged in a wide range of activities.
The global nature of Deere & Co. activities means that communications and coordination are of central importance. Other business aspects of Deere & Co. also add to the importance of communications and coordination. For example the 1996 business plan called for the largest introduction of new models in the history of Deere & Co. The successful introduction of such new models involves much in the way of communication and collaboration.
Deere & Co. have undertaken an ambitious program for the use of Internet technologies for business communication, coordination and collaboration. They have, by and large, followed the Phases I- III outlined in 2.5.3, and are moving into the early stages of Phase IV.

Phase I – The Brochure Stage - Establishing a WWW Presence was completed early in the program for the use of Internet technologies. The corporate WWW site (http://www.deere.com) includes links to the agricultural equipment, ground care, parts, credit and health care parts of the corporation. Information is contained about, for example, the Deere & Co. product range and so the WWW site operates in part as an on-line brochure.
Phases II and III is in progress at Deere & Co. with some activities already in use while others are in development. Below is described the use of Internet technologies in differing areas of Deere & Co.

2.6.2 Deere & Co. - Use of Internet Technologies for Sales Support

Deere & Co. has several Internet technologies applications for dealers, with several in use and others scheduled for use in the near future. These applications allow for a number of activities to carried out using Internet technologies, including:

- An application to allow dealers to order service parts via a WWW browser. This removes the need for faxes to be sent and is especially helpful for overseas dealers who now only need local Internet access, and who now do not need to pay international phone charges.
- Applications that support the management of dealerships. These include the provision of on-line statistics for dealers on sales histories, current inventories and equipment orders. This helps the dealers in their ordering decisions.
- An application that shows manuals on-line for servicing equipment. This is being extended to include parts catalogs and parts drawings.
- An application that gives market share status for dealers. The use of Internet technologies has reduced the time taken to receive this report from eight hours to three minutes.
- An application is being written to process loans over the Internet. Users will then contact the Deere & Co. credit bureau on-line to obtain loan approval, while another application will show the status of the loan including the amount owed, and the payment schedule.
- An application that will show dealers their warranty claim information.
- An application that will allow major customers to receive information on the inventory dealer levels and allow them to order on-line

These applications allow for much improved communications with dealers and customers. The overall result should be better customer service and hence increased sales.

2.6.3 Deere & Co. - Use of Internet Technologies for Internal Applications

Deere & Co. have been using Internet technologies on the corporate Intranet for many internal operations. These applications include:

- The on-line benefits book which allows Deere & Co. employees to peruse up-to-date information concerning employee benefits.
- The development of an on-line application to allow employees to change options in their 401k retirement plans
- An on-line integrity review process to certify application which will be deployed on the external WWW site. The developer will be able to fill out forms on-line to supply the integrity review team with the information needed to certify the application. This should substantially reduce the time taken to certify an application.
- A photo lab application (described in more detail below in Section 2.6.4) that allows others to download picture of Deere & Co. equipment, parts etc. This greatly reduces the time to process requests, helps ensure that the pictures used are up-to-date, and reduces the overall costs of the operation.
- The development of an on-line magazine called JDOnline that provides communications to all Deere & Co. employees about company activities.

2.6.4 DEERE & CO. – EXAMPLE FINANCIAL ANALYSIS FOR INTERNET TECHNOLOGIES

We can examine in this section the financial impact of the Deere & Co. photo lab application (described briefly in the previous section).

The old photo lab process involved the following:

- A designer who needed a picture for a publication would have to look through books of contact sheets to identify a desirable picture. Note that this required that the person actually have the books of contact prints to examine – this can be very difficult for a distributed or global operation.
- The contact sheet would then be scanned and this scanned image was then placed in the draft document as a place holder awaiting the delivery of the high resolution image.
- The photo lab staff would then print a copy of the photograph from the negative and send this to the designer.

The use of Internet technologies has allowed this inefficient process to be streamlined. The new process simply involves the designer using a WWW browser to search through and retrieve the image at the required resolution. This means that there is no longer the necessity for the designer to have physical access to the contact sheets so that the designer can be anywhere there is an Internet connection. Three security levels have been implemented to ensure that more sensitive information, such as pictures of upcoming products, are kept secure.

This Internet-based process has a number of benefits:

1. The designer has a more efficient access to images.
2. The use of Internet technologies means that there is cross-platform access to the images.
3. There is increased feasibility for new product releases since changes in the designs can occur frequently.
4. The images are available immediately for inclusion in the publication.

The overall economic benefits per year have been as follows:

- There has been a $85,000 reduction in employee hours.
- There have been savings of $228,000 in pre-press costs.
- Photo-lab fees have been reduced by $175,000.
- There has been a reduction of $5,000 in film costs.

A more detailed spreadsheet is shown below (from IDC) **to be converted to conventional table form **.

Deere & Company*

Annual Savings	Base	Year 1	Year 2	Year 3
Personnel Savings	$0	$85,000	$80,000	$80,000
Pre-Press Savings	$0	$228,000	$228,000	$228,000
Lab Fees	$0	$175,000	$175,000	$175,000
B+W Film	$0	$5,000	$5,000	$5,000
Total Savings Per Period	$0	$493,000	$488,000	$488,000

Depreciation Schedule**	Initial	Year 1	Year 2	Year 3***
Software	$373,000	$74,600	$74,600	$74,600
Hardware	$63,000	$0	$0	$0
Total Per Period	$436,000	$74,600	$74,600	$74,600

Expensed Costs	Initial	Year 1	Year 2	Year 3
Maintenance	$0	$101,000	$101,000	$101,000
Personnel	$66,000	$0	$0	$0
Photographers Manuals	$0	$2,000	$2,000	$2,000
Training	$6,000	$6,000	$5,000	$5,000
Total Per Period	$72,000	$109,000	$108,000	$108,000

Basic Financial Assumptions	
All Federal and State Taxes	50%
Discount Rate	15%
Depreciation - Straight Line (Years)	5

Net Cash Flows	Initial	Year 1	Year 2	Year 3
Total Benefits		$493,000	$488,000	$488,000
Less: Total Costs	$508,000	$109,000	$108,000	$108,000
Less: Depreciation		$74,600	$74,600	$74,600
Net Profit Before Tax	$508,000	$309,400	$305,400	$305,400
Net Profit After Tax	$254,000	$154,700	$152,700	$152,700
Add: Depreciation		$74,600	$74,600	$74,600
Net Cash Flow After Taxes	($254,000)	$229,300	$227,300	$227,300

Financial Analysis	Results	Year 1	Year 2	Year 3
Annual ROI		79%	146%	205%
3-Year ROI	**205%**			
Payback (Years)	**1.32**			
3-Year IRR	72%			
3-Year NPV	$266,716			

* All calculations are based on International Data Corporations analysis of Deere & Company's application using information gathered during on-site interviews and other data collection methods. These numbers have not been generated by Deere & Company.

** Hardware and software costs totalling more than $100,000 are depreciated over 5 years on a straight line basis. All other costs are treated as expenses in the initial year.

*** Any software upgrade is treated as a depreciable asset if greater than $100,000, otherwise it is included as an expensed cost for the year.

2.7 ACTIVITY

(suggested homwork)

IMPLEMENTING A WEB SERVER

You suddenly realize that it's the middle of June and your thoughts start to wander to a gentle winding down of activities before the late Summer break. You joined the AWOI Corporation just about 18 months ago and you have risen (thanks to your excellent education) very quickly to your present level. Abruptly your self congratulatory mood is rudely interrupted by a phone call from the President of AWOI Corporation.

"We need a web site to advertise our material handling products and we need it up and running as soon as possible,", the President says. "I know that since you are a UoI I.E. graduate, you have the capability to do this. We need it completed by June 18 and you are the only person who can pull this off."

"Er…Yes, of course", you slowly reply and thereby commit yourself to the project.

After the phone call ends, you sit in stunned silence and slowly a plan begins to develop in your mind ………

2.7.1 TASK

The aim of this task is to learn about WWW Server side programming, building a WWW site, and network security issues.

It is preferable to work in sections with a maximum of three in each section. Each section will need to obtain server software [for example, download the Netscape Fast track Server (http://www.netscape.com)], install it and then to configure your server system.

Each group will then have to set up a WWW site for AWOI. This will involve:

1. create your group directory under c:\isd\ with the name *groupi*
2. create the home page under c:\isd*groupi*\ and name index.htm
3. specify your home page URL as http://IP address/*groupi*/index.htm
4. Write other(simple) pages to describe the material handing equipment that AWOI sells (robots, conveyors and fork lift trucks).

Then you need to carry out the following:

1. Specify access control to allow one page to be read only by a user with a password (user name= ogrady, password= ogrady1)
2. Specify access control to allow another page to be read only by a user

(hostname: w-ie04, IP address: 128.255.23.176) .

3. In your group home page, allows a link to a CGI program (test.bat):
    - you need to specify your cgi directory and URL as http://IP address/groupi/cgi/test.bat
    - create the simple cgi test program:

```
@echo off
@echo content-type: text/plain
echo.
echo This is a cgi test
```

Reference:

1. FastTrack installation and reference guide (http://home.netscape.com/comprod/server_central/support/fasttrack_man/index.htm)

…..SOMETIME LATER……………..

June 16 1997 2:00 p.m.

Playa Tamarindo, Costa Rica

The warm waters of the Pacific engulf you as you swim lazily away from shore. You again feel a sense of smugness as you congratulate yourself on having completed the project early and then flying down to Costa

Rica. The President was REALLY pleased with the work you had done and she suggested that you take a few days off to make up for all of your hard work. Of course you forgot to mention that thanks to your excellent education, the project hardly took any time at all. Still the handsome pay rise, promotion and stock options that the President announced were very welcome.

The massive Pacific breakers look inviting and you start thinking that perhaps it's now time to learn how to surf ………..

2.8 REFERENCES

Common Gateway Interface, 1995, http://hoohoo.ncsa.uiuc.edu/cgi/intro.html.

Grand, M., 1995, *MIME overview and technical summary*

Bell, G., Parisi, A., and Pesce, M., 1995, The Virtual Reality Modeling Language Version 1.0 Specification.

2.9 QUESTIONS