

Teresa Piliouras et al. "Intranets"

The CRC Handbook of Modern Telecommunications

Ed. Patricia Morreale and Kornel Terplan

Boca Raton, CRC Press LLC. 2001

2

Intranets

Teresa Piliouras

TCR, Inc.

Andrew Resnick

Citicorp

John Braun

Endre Sara

Goldman, Sachs & Co.

Karen M. Freundlich

TCR, Inc.

Dermot Murray

Iona College

Mihir Parikh

Polytechnic University

Introduction

2.1 Internet and Intranet Management Concepts .

Management Overview of the Internet and Intranets • Intranet Planning and Management • Technical Overview • Intranet Components • Intranet Implementation • Intranet Deployment • Intranet Security Issues • Summary

2.2 Internet Security

Physical Security • Modems • Data Security • Passwords • Workstation Security • TCP/IP Security

2.3 Virtual Private Networking Solutions.

Layer 2 Protocols • Layer 3 Tunneling Protocols • Frame Relay • Layer 2 or Layer 3 Comparison

2.4 Effective Website Design.

Goals Defined • Production and Maintenance Efforts • A System for Measuring Effectiveness • Intuitive Layout

2.5 Web-enabled Data Warehousing

Introduction • Data Warehousing Overview • Web-enabled Data Warehousing • Vendors • Future Trends • Conclusion

2.6 E-commerce Technologies: A Strategic Overview

E-commerce Technologies • Strategic Challenges • Emerging Trends for the Future

2.7 Internet Protocols

Addressing for Internet • Communication Protocols in Internet • Information Transfer in Internet • Types of Internet Access • Internet E-mail • Telnet in Internet • File Transfer in Internet • News and Usenet • Mailing Lists in Internet • Information Search in Internet • Netscape and Microsoft

Introduction

Teresa Piliouras

The Internet started as a technological revolution, designed to protect national interests by ensuring redundancy and resiliency in governmental networks, particularly in time of war. It has spawned world-wide cultural revolution, fostering universal communication exchange with limitless geographic, time, and subject matter boundaries. The extent and ease of the Internet's adoption has had profound implications on all — including personal, business, and governmental — aspects of life. There is no place on earth that cannot be reached by the Internet.

In this chapter, we review the basic technological underpinnings of the Internet and discuss why it is so flexible. As we explore the evolution of the Internet, which continues at an ever-increasing pace, we also examine corresponding effects on communication paradigms, particularly in a business context.

Once caution was the order of the day. Now, businesses small and large alike are racing to join the Internet bandwagon and to have a “Web” presence.

In a dynamic environment that changes faster than words can be put into print, we can only hope to scratch the surface on the Internet’s development and major trends in this Chapter.

2.1 Internet and Intranet Management Concepts

Teresa Piliouras

2.1.1 Management Overview of the Internet and Intranets

An Intranet is a company specific, private network based on Internet technology, and as such, it is a form of local area network (LAN). However, one of the major distinctions between traditional LANs and Intranets is the reliance of the latter on TCP/IP, packet switching, and Internet technologies. In the case of the Internet, the technology is deployed over a public network, while in the case of Intranets, the technology is deployed within a private network.

According to George Eckel, author of *Intranet Working*, one of the important benefits of Intranets is that they provide a cost-effective vehicle for communication, since the expense of reaching one person or one million people is essentially the same. Intranets are becoming the corporate world’s equivalent of a town hall where people can meet, chat, and exchange information.

The emergence of Intranets promises to change the way companies communicate with their employees and how they conduct their business. For example, after years of using satellite feeds to disseminate information to its 208 network affiliates, CBS News now uses an Intranet to provide affiliates with point-and-click access to information on upcoming news stories. Access to this information is provided through the CBS Newspath World Wide Web home page.

2.1.1.1 Benefits of Intranets

Intranets offer many potential benefits, including:

- Reduced operating costs
- Improved employee productivity
- Streamlined processing flows
- Improved internal and external communication
- New and improved customer service
- Cross-platform capability

We will discuss some of the ways these benefits can be achieved.

2.1.1.1.1 The Paper-less Office

Many companies find that Intranets simplify corporate-wide communications, and reduce printed material costs by eliminating the need for many paper-based processes. For example, some organizations offer complete manuals on their corporate Web site in electronic form, instead of distributing the information in printed form. Companies can benefit immediately from an Intranet by replacing their printed materials, little by little, with electronic versions. Electronic media is cheaper to produce, update, and distribute than printed material. Often times, printed material is out of date by the time it is distributed. Electronic documents, however, can be easily modified and updated as the need arises.

2.1.1.1.2 Improved Customer Service

For many organizations, having the right information at the right time can make a significant difference in their ability to close a sale or meet a deadline. In today’s competitive business environment, companies are also under constant pressure to improve productivity while reducing costs. To achieve these productivity

gains, companies must constantly improve their relationships with employees, customers, vendors, and suppliers. Intranets provide an important avenue for making these advancements.

Using an Intranet, vendors, employees, and customers can access information as it is needed, alleviating delays associated with mailing or distributing printed materials. For example, Intranets have been used to:

- Distribute software updates to customers, reducing the need to send printed materials and floppy diskettes or CD-ROMs
- Handle customer orders on-line
- Process and respond to customer inquiries and questions about products and services
- Collect customer and survey data

Using an Intranet, all these activities can be completed electronically in a matter of minutes.

2.1.1.1.3 Improved Help Desks

Intranets have been used to augment help desk services. For example, when someone in the organization learns about a new technology or how to perform a new task (for example, running virus software), he/she can put information and instructions for others on a personal Web page. Others within the organization, including help desk staff, can then access this information as needed. In an organization empowered by an Intranet, all employees can leave the imprints of their expertise.

2.1.1.1.4 Improved Corporate Culture

Intranets help to cultivate a corporate culture that encourages the free flow of information. Intranets place information directly into the hands of employees, promoting a more democratic company structure. The danger of “information democracy” is that once it is in place and taken for granted, management can not easily revert to older, more controlled forms of communication without seriously damaging employee morale and cooperation. Every individual in an Intranet environment is empowered to access and distribute information, both good and bad, on a scale heretofore unknown in the corporate realm.

Intranets dissolve barriers to communication created by departmental walls, geographical location, and decentralized organizations. Placing information directly in the hands of those who need it allows organizations to decentralize and flatten decision making and organizational processes, while maintaining control over the information exchange. Individuals and groups can distribute ideas freely, without having to observe traditional channels of information (i.e., an individual, a printed document, etc.) that are far less effective in reaching geographically dispersed individuals.

2.1.1.1.5 Cross-platform Compatibility

Since the early 1980s, organizations with private networks have struggled with connecting and disseminating information between different types of computers — such as PCs, Macintoshes, and Unix-based machines. To help manage potential barriers to electronic communication posed by hardware and software incompatibilities, many companies have instituted strict standards limiting corporate users to specific hardware and software platforms. Even today, if a company uses PCs, Macs, and Unix-based machines, sharing a simple text document can be a challenge.

Intranets provide a means to overcome many of these software and hardware incompatibilities, since Internet technologies (such as TCP/IP) are platform independent. Thus, companies using Intranets no longer need to settle on one operating system, since users working with Macintosh, PC, or Unix-based computers can freely share and distribute information. In the sections that follow, we will explain why this is so.

2.1.2 Intranet Planning and Management

To implement an Intranet, a company needs a dedicated Web server, communications links to the Intranet, and browser software. *Unfortunately, Intranets do not come prepackaged and fully assembled.* They require careful planning and construction, if they are to be effective in meeting the needs of the organization. In the sections that follow, we discuss recommendations for planning and implementing an Intranet.

2.1.2.1 Gaining Support

The first step toward a successful Intranet implementation is to obtain company-wide support for the project, including endorsement from upper management. A quality presentation should be made to both management and staff to explain the benefits of the Intranet project. Some of these are tangible and easy to measure, while others are intangible and difficult to measure. To gain widespread support for the Intranet project, decision makers must be shown what an Intranet is and how it will benefit the organization. There are many resources (including complete presentations) available on the World Wide Web to help promote the Intranet in a corporate environment.

2.1.2.2 Planning the Intranet Strategy

After selling upper management on the idea of an Intranet, the next step is to define the goals, purpose, and objectives for the Intranet. This is an essential part of the Intranet project planning.

The Intranet project plan should include an overview of the organizational structure and its technical capabilities. The current communication model used to control information flows within the organization should be examined with respect to its strengths and weaknesses in supporting workflow processes, document management, training needs, and other key business requirements. It is important to understand and document existing systems within the organization before implementing the Intranet.

The Intranet plan should clearly define the business objectives to be achieved. The objectives should reflect the needs of the Intranet's potential users. Conducting interviews with employees and managers can help identify these needs. For example, the human resource department may wish to use the Intranet to display job opportunities available within the organization. If this need is to be satisfied, the Intranet should be designed to display job information and job application forms on a Web server, so applicants can apply for positions electronically. The human resource department might also wish to offer employees the ability to change their 401K information by using the Intranet. Each identified goal shapes and defines the functionality that the Intranet must support. An employee survey is also an excellent way to collect ideas on how to employ the Intranet within the organization.

In summary, the following questions are helpful in defining the requirements of the Intranet project:

- Will Intranet users need to access existing (legacy) databases?
- What type of training and support will Intranet users require?
- Who will manage, create, and update the content made available through the Intranet?
- Will individual departments create their own Web pages autonomously?
- Will there be a central authority that manages changes to the content offered on the Intranet?
- Do users need remote access to the Intranet?
- Will the Intranet need to restrict access to certain users and content?
- Will a Webmaster or a team of technicians/managers be assigned to coordinate and manage the maintenance of the Intranet?
- Will the Intranet be managed internally or will it be outsourced?

2.1.2.3 Selecting the Implementation Team

After the Intranet project plan has been developed and approved, the implementation team should be assembled. If the organization does not have an infrastructure in place that is capable of implementing the Intranet, additional staff and resources will need to be hired or the project will need to be outsourced to a qualified vendor.

It is important that the Intranet team has the requisite skills to successfully execute the project plan. A number of skill assessment checklists are provided below to help evaluate the resources available within an organization and their abilities to successfully support the Intranet implementation.

Technical Support Skills Checklist

The Intranet project will require staff with the technical skills needed to solve network problems, understand network design, troubleshoot hardware and software compatibility problems, and implement

client-server solutions (such as integrating network databases). Thus, the following skills are required to support an Intranet:

- Knowledge of network hardware and software
- Understanding of TCP/IP and related protocols
- Experience implementing network security
- Awareness of client-server operations
- Practice with custom programming
- Abilities of database management

Content Development and Design Checklist

A typical organization has many sources of information: human resource manuals, corporate statements, telephone directories, departmental information, work instructions, procedures, employee records, and much more. To simplify the collection of information that will be made available through the Intranet, it is advisable to involve people familiar with the original documentation and also those who can author content for Intranet Web pages. If possible, the original authors of the printed material should work closely with the Intranet content developers to ensure that nothing is lost in translation.

The following technical skills are needed to organize and present information (content) in browser-readable format:

- Experience in graphic design and content presentation
- Basic understanding of copyright law
- Knowledge of document conversion techniques (to convert spreadsheet data, for example, into a text document for HTML editing)
- Experience in page layout and design
- Experience with Web browsers and HTML document creation
- Knowledge of image-conversion techniques and related software
- Knowledge of programming languages and programming skills
- CGI programming and server interaction

Management Support Skills Checklist

As previously discussed, the company's management should be involved in the planning and implementation of the Intranet. Ideally, management should have a good understanding of the Intranet benefits, and the expected costs and time frames needed for the project completion. Managers with skills relating to quality-control techniques, process-management approaches, and effective communication are highly desirable. Thus, the following management skills are recommended:

- Understanding of the organization's document flow
- Experience with the re-engineering process
- Knowledge of quality-control techniques
- Knowledge of the company's informal flow of information
- Experience with training and project coordination

2.1.2.4 Funding Growth

The initial cost of setting up a simple Intranet is often quite low and may not require top management's approval. However, when complex document management systems are needed to integrate database access, automate workflow systems, implement interactive training, and other advanced features, the Intranet should be funded with the approval of top management. To gain approval for the project, upper management must be convinced that the Intranet is an integral part of the company's total information-technology deployment strategy. This involves quantifying the tangible benefits of the Intranet to the organization. Management also needs to understand how the Intranet will change the way people work and communicate.

2.1.2.5 Total Quality Management (TQM)

Effective deployment of an Intranet often involves re-engineering current process flows within the organization. Employees are usually most receptive to changes that make their jobs easier. To avoid perceptions that the Intranet is an intimidating intrusion of yet another technology, it is advisable to involve staff as early on as possible in the deployment planning. This will facilitate the transition to the Intranet, and encourage employee participation in the Intranet's success.

After migrating the company's work processes to the Intranet, it is up to managers and employees to adhere to the procedures that have been put in place to improve productivity and teamwork. Management should not assume that because employees have a new tool — the Intranet — that this alone is sufficient to ensure that the desired attitudes and service levels will be attained. Instead, managers should view the Intranet as one aspect of their quest for total quality management (TQM).

TQM involves creating systems and workflows that promote superior products and services. TQM also involves instilling a respect for quality throughout the organization. TQM and the successful deployment of Intranets represent a large-scale organizational commitment, which upper management must support.

2.1.2.6 Training Employees

If employees are expected to contribute content to the Intranet, they will need to be given tools and training so they can author HTML and XML documents. In general, it is a good idea to encourage employees to contribute to the content on display through the Intranet. To do otherwise means that the organization may have to depend on only a few people to create HTML and XML documents.

After initial training, users should be surveyed to determine if the tools they have been provided satisfy their needs. Many users find that creating HTML/XML documents is difficult. If so, then they may also need special training. In corporations, this training is often provided by one person in each department who has been given responsibility for training the rest of the department.

In summary, the following actions are recommended to help develop an effective program for training employees to author high-quality HTML/XML documents:

- Conduct a survey to assess user training needs and wants
- Train users how to develop HTML/XML content
- Provide users with HTML/XML authoring tools that complement what they already know (for example, the Internet Assistant for Microsoft Word is a good choice for users already familiar with Microsoft Word)
- Review the design and flow of material that will be “published” on the Intranet
- Give feedback to HTML/XML authors on ways to improve the site appearance and ease of use

2.1.2.7 Organizational Challenges

In addition to technological challenges, companies may also face the following organizational challenges after the initial release of an Intranet:

- Marketing the Intranet within the organization so that all employees will support its growth and continued use
- Obtaining additional funding on an ongoing basis to implement new capabilities
- Encouraging an information-sharing culture within the company so that all employees will contribute toward building a learning organization
- Merging a paper-based culture with the new culture of electronic documentation
- Ensuring that the content on the Intranet is updated on a regular basis
- Preventing one person or group from controlling (monopolizing) the content on the Intranet
- Instructing employees to author HTML/XML content so they can contribute material to the Intranet
- Informing employees on Intranet etiquette, thereby facilitating courteous on-line discussion forums and other forms of user interaction on the Intranet

- Using the Intranet as an integral part of working with customers and vendors
- Measuring the Intranet's overall effectiveness and contribution to the organization

As is the case when introducing any new information technology to an enterprise, Intranet deployment requires careful planning, effective implementation, and employee training. In the short term, most of the organizational focus is usually on the technical aspects of the Intranet deployment. But as time goes on, organizational issues relating to how the Intranet is used within the organization must be managed. When an organization actively examines and works toward resolving these issues, they are better able to achieve a culture of teamwork and collaboration.

2.1.2.8 Management Summary

The following list summarizes key points surrounding the use of Intranets:

- An Intranet is a company-based version of the Internet. Intranets provide an inexpensive solution for information sharing and user communication.
- An Intranet provides an easy way for users to communicate and share common documents, even if they are using different machines, such as IBM compatible and Macintosh personal computers.
- Some organizations have expanded their Intranet to allow customers to access internal databases and documents.
- Many companies can establish a functional Intranet using in-house personnel with a minimal amount of new equipment.

Internet technology adheres to open standards that are well documented. This, in turn, encourages the development of cost-effective and easy-to-implement Intranet solutions. As the popularity of Intranets has increased, so has the demand for new tools and Web-based solutions. This demand has fueled competition among software manufacturers which, in turn, has resulted in better and less expensive Intranet products.

In summary, Intranets can be used to improve productivity, simplify workflows, and gain a competitive advantage over those who have yet to learn how to capitalize on the benefits of Intranets.

2.1.3 Technical Overview

2.1.3.1 Internet Basics

2.1.3.1.1 Packet Switching

Packet switching was introduced in the late 1960s. In a packet-switched network, programs break data into pieces, called packets, which are transmitted between computers. Each packet contains the sender's address, the destination address, and a portion of the data to be transmitted. For example, when an e-mail message is sent over a packet-switched network, the e-mail is first split into packets. Each packet intermingles with other packets sent by other computers on the network. Network switches examine the destination address contained in each packet, and route the packets to the appropriate recipient. Upon reaching their destination, the packets are collected and reassembled to reconstitute the e-mail message.

2.1.3.1.2 TCP/IP

The U.S. Advanced Research Projects Agency (ARPA) was a major driving force in the development and adoption of packet-switched networking. The earliest packet-switched network was called the ARPAnet. The ARPAnet was the progenitor to today's Internet. By the early 1980s, ARPA needed a better protocol for handling the packets produced and sent by various network types. The original ARPAnet was based on the Network Control Protocol (NCP). In January 1983, NCP was replaced by the Transport Control Protocol/Internet Protocol (TCP/IP). TCP/IP specifies the rules for the exchange of information within the Internet or an Intranet, allowing packets from many different types of networks to be sent over the same network.

2.1.3.1.3 Connecting to the Internet

One way to connect to the Internet is to install a link from the company network to the closest computer already connected to the Internet. When this method is chosen, the company must pay to install and

maintain the communications link (which might consist of a copper wire, a satellite connection, or a fiber optic cable) to the Internet. This method was very popular with early adopters of the Internet, which included universities, large companies, and government agencies. However, the costs to install and maintain the communications link to the Internet can be prohibitive for smaller companies.

Fortunately, specialized companies — called Internet Service Providers (ISPs) — are available to provide a low-cost solution for accessing the Internet. ISPs pay for an (expensive) connection to the Internet, which they make accessible to others through the installation of high-performance servers, data lines, and modems. Acting as middlemen, the ISPs rent time to other users who want to access the Internet.

Two important decisions must be made when deciding what type of Internet connection is the most appropriate. The first decision is the company budget allocated for Internet connectivity, and the second is the Internet connection speed needed to support the business requirements. Both decisions are inter-related. ISPs offer a variety of options for connecting to the Internet, ranging from a simple dial-up account over phone wires to high-speed leased lines from the company to the ISP. Dial-up accounts are typically available for a low, flat monthly fee, and are generally much cheaper than a leased line connection. However, the leased line connection is usually much faster than the dial-up connection.

When a dial-up account is used, a modem and a phone line are used to call and log into the ISP server (or computer), which, in turn, acts as the doorway to the Internet. The transmission speed of the connection is limited by the speed of the modems employed by the user and the ISP. A modem is unnecessary when a leased line connection is available to the ISP. Leased lines are offered in many different configurations with a variety of options. The most common link types are ISDN (which support transmission speeds from 56 Kbps to 128 Kbps), T1 (transmitting at speeds up to 1.54 Mbps), and T3 (transmitting at speeds up to 45 Mbps).

If a company only needs to make an occasional connection to the Internet — for example, less than 20 to 50 hours per month for all users — a dial-up account should be sufficient. However, if a company needs faster data transfer speeds or has several users who must access the Internet for substantial periods of time over the course of a month, a leased line connection should be considered.

The fastest growing segment of Internet users are those who connect to the Internet through an ISP via an ordinary telephone connection. There are two major protocols for connecting to the Internet in this way: Serial Line Internet Protocol (SLIP) and Point-to-Point Protocol (PPP). SLIP is the older protocol and is available in many communications packages. The faster PPP is newer and therefore it is not as widely supported.

Principles of queuing analysis can be applied to the problem of sizing the links needed to support the Internet access, whether or not that access is to an ISP or to a direct Internet connection. The reader is referred to Chapter 2, *Network Design: Management and Technical Perspectives* by Mann-Rubinson and Terplan, for specific techniques on how to estimate the throughput and performance characteristics associated with using different size link capacities. This analysis can be used to determine whether or not a dial-up or leased line connection is sufficient to support the bandwidth requirements with tolerable transmission delays.

2.1.3.1.4 Basic Terminology

In this section, we define commonly used Internet and Intranet terminology.

The World Wide Web

The World Wide Web — or Web — is a collection of seamlessly interlinked documents that reside on Internet servers. The Web is so named because it links documents to form a web of information across computers worldwide. The documents available through the Web can support text, pictures, sounds, and animation. The Web makes it very easy for users to locate and access information contained within multiple documents and computers. “Surfing” is the term used to describe accessing (through a Web browser) a chain of documents through a series of links on the Web.

Web Browsers

To access and fully utilize all the features of the Web, special software — called a Web browser — is necessary. Its main function is to allow the user to traverse and view documents on the Web. Browser

software is widely available for free, either through a download from the Internet or from ISPs. Commercial on-line services — such as America Online and Prodigy — also supply browsers as part of their subscription products. The two most commonly used browsers are Netscape Navigator and Microsoft Internet Explorer. Some of the common tasks which both support include:

- Viewing documents created on a variety of platforms
- Creating and revising content
- Participating in threaded discussions and news groups
- Watching and interacting with multimedia presentations
- Interfacing with existing legacy data (non-HTML based data) and applications
- Gaining seamless access to the Internet

It should be noted that the same Web browser software used for accessing the Internet is also used for accessing documents within an Intranet.

Uniform Resource Locator (URL)

The Web consists of millions of documents that are distinguished by a unique name called a URL (Uniform Resource Locator), or more simply, a Web address. The URL is used by Web browsers use to access Internet information. Examples of URLs include:

`http://www.netscape.com`
`ftp://ftp.microsoft.com`

A URL consists of three main parts:

1. A service identifier (such as `http`)
2. A domain name (such as `www.ups.com`)
3. A path name (such as `www.ups.com/tracking`)

The first part of the URL, the service identifier, tells the browser software which protocol to use to access the file requested. The service identifier can take one of the following forms:

- `http://` — This service identifier indicates that the connection will use the hypertext transport protocol (HTTP). HTTP defines the rules that software programs must follow to exchange information across the Web. This is the most common type of connection. Thus, when Web addresses start with the letters “http” it indicates that the documents are retrieved according to the conventions of the HTTP protocol (hypertext transport protocol).
- `ftp://` — This service identifier indicates that the connection will use the file transfer protocol (FTP). This service identifier is typically used to download and copy files from one computer to another.
- `gopher://` — This service identifier indicates that the connection will utilize a gopher server to provide a graphical list of accessible files.
- `telnet://` — This service identifier indicates that a telnet session will be used to run programs from a remote computer.

The second part of the URL, the domain name, specifies which computer is to be accessed when running server software. An example of a domain name is: `www.tcrinc.com`.

The final part of the URL, the path name, specifies the directory path to the specific file to be accessed. If the path name is missing from the URL, the server assumes that the default page (typically, the homepage) should be accessed. Large, multi-page Web sites can have fairly long path names. For example, these URLs request specific pages within a given Web site:

- `http://www.apple.com/documents/productsupport.html`
- `http://www.bmwusa.com/ultimate/5series/5series.html`
- `ftp://ftp.ncsa.uiuc.edu/Mac/Mosaic`
- `http://www.microsoft.com/Misc/WhatsNew.htm`

Home Pages

Companies, individuals, and governments that publish information on the Internet usually organize that information into “pages,” much like the pages of a book or a sales brochure. The first page that people see in a sales brochure is the cover page, which may contain an index and summary of the brochure contents. Similarly, a home page is the first page that users see when they access a particular Web site. The home page is to the Web site what the cover page is to a sales brochure. Both must be appealing, concise, informative, and well organized to succeed in maintaining the reader’s interest. The home page is usually used to convey basic information about the company and what it is offering in the way of products and/or services.

Many companies publish the Internet address (or URL) of their home page on business cards, television, magazines, and radio. To access a Web site, a user has merely to type the URL into the appropriate area on the Web browser screen.

Client Programs and Browsers

Across the Internet, information (i.e., programs and data) is stored on the hard disks of thousands of computers called servers. These are so named because, upon request, they serve (or provide) users with information. A server is a remote computer that may be configured to run several different types of server programs (such as Web server, mail server, and ftp server programs).

A client program is used to initiate a session with a server. Client programs are so named because they ask the server for service. In the case of the Web, the client program is the Web browser. All client–server interactions take the same form. To start, the client connects to the server and asks the server for information. The server, in turn, examines the request and then provides (serves) the client with the requested information. The client and server may perform many request–response interactions in a typical session.

Software programs — such as a browser — use HTTP commands to request services from an HTTP server. An HTTP transaction consists of four parts: a connection, a request, a response, and a close.

Where Web Documents Reside

When users publish Web pages, they actually store the pages as files that are accessible through a file server. Typically, Web pages reside on the same computer on which the server program is running, but this is not necessarily true. For security reasons, it may be necessary to limit accessibility to various files on the Web server. Obviously, it might be disastrous if internal documents and data were made available to competitors. To prevent this type of security risk, a WebMaster (or Systems Administrator) can configure the Web server so it only allows specific clients to access confidential information, based on a need-to-know basis. The WebMaster can control access to the server by requiring users to log-in with a username and password that has predetermined access privileges.

HTML — The Language of the World Wide Web

The European Particle Physics Laboratory at CERN, in Geneva, Switzerland, developed Hypertext Markup Language (HTML) in the late 1980s and early 1990s. HTML is the language of the World Wide Web. Every site on the Web uses HTML to display information.

Each Web document contains a set of HTML instructions that tell the browser program (e.g., Netscape Navigator or Microsoft Internet Explorer) how to display the Web page. When you connect to a Web page using a browser, the Web server sends the HTML document to your browser across the Internet. Any computer running a browser program can read and display HTML, regardless of whether that computer is a personal computer running Windows, a Unix-based system, or a Mac.

If word processor formatted files — such as Microsoft Word — were used to create Web pages, only users with access to Microsoft Word would be able to view the Web page. HTML was designed to overcome this potential source of incompatibility. All users can access Web pages from their browser since all Web pages conform to HTML standards. A HTML Web page is a plain text file (i.e., an ASCII text file) that can be created and read by any text editor. There are many software programs available to convert document files to HTML equivalents. In addition, many standard presentation and word processing

packages offer built-in routines to convert a standard document into a Web-ready HTML file. This type of conversion might be helpful, for example, if you wanted to convert a Microsoft PowerPoint presentation into a set of HTML files for display on the Web.

After HTML files are transferred to a Web site, anyone with a browser can view them. HTML provides the browser with two types of information:

1. “Mark-up” information that controls the text display characteristics and specifies Web links to other documents.
2. “Content” information consisting of the text, graphics, and sounds that the browser displays.

Hypertext and Hyperlinks

Documents on the Web can be interconnected by specifying links (called hyperlinks) that allow the user to jump from one document to another. The HTML code, which drives all Web pages, supports hypertext. Hypertext, in turn, supports the creation of multimedia documents (containing pictures, text, animation, sound, and links) on the Web.

Hyperlinks (or simply, links) are visually displayed on the Web pages as pictures or underlined text. When a user clicks on a hyperlink displayed on their browser screen, the browser responds by searching for and then loading the document specified by the hyperlink. The document specified in the hyperlink may reside on the same computer as the Web page on display or it may reside on a different computer on the other side of the world. Much of the Web’s success has been attributed to the simplicity of the hyperlink point-and-click user interface.

There are four basic layouts for linking Web pages with hyperlinks: linear, hierarchical, Web, and combination. Which layout is the most appropriate depends on the type of information that is being presented and the intended audience.

FTP — The File Transfer Protocol

The FTP (file transfer protocol) is a standard protocol for transferring and copying files from one computer to another. Depending on the configuration of the FTP server program, you may or may not need an account on the remote machine to access system files. In many cases, you can access a remote computer with FTP by logging on with a username of “anonymous,” and by entering your e-mail address as the password. This type of connection is referred to as “anonymous FTP session.”

After logging in to the remote FTP server, it is possible to list a directory of the files that are available for viewing and/or copying. The systems administrator determines which files can be accessed on the remote server, and who has access privileges. When system security is a major concern, the system administrator may require a specific username and password (as opposed to allowing an anonymous log-on procedure) to gain access to system files.

FTP is very useful in accessing the millions of files available on the World Wide Web. Most browsers have built-in FTP capabilities to facilitate downloading files stored at FTP sites. To access a FTP site using your browser, you type in the FTP site address, much like entering a Web address. For example, to access the Microsoft FTP site, the address “ftp://ftp.microsoft.com” would be entered into the browser address window.

Java

Java is a new programming language released by Sun Microsystems that closely resembles C++. Java is designed for creating animated Web sites. Java can be used to create small application programs, called applets, which browsers download and execute. For example, a company might develop a Java applet for their Web site to spin the company’s logo, to play music or audio clips, or to provide other forms of animation to improve the appeal and effectiveness of the Web page.

Network Computers

Although personal computers (PCs) are becoming more and more common, the number of households with a PC is still only about one third the number of households with a television. The main reason is that PCs are still too expensive for the masses. The network computer is a scaled-down, cheaper version (under \$500) of the PC. A network computer is designed to operate exclusively with the Internet and Java applets.

2.1.4 Intranet Components

This section will provide an overview of the components necessary to create an Intranet. The final selection of the Intranet components depends upon on the company's size, level of expertise, user needs, and future Intranet expansion plans. In addition, we also examine some of the costs associated with the various Intranet components.

An Intranet requires the same basic components found on the Internet, including:

1. A computer network for resource sharing.
2. A network operating system that supports the TCP/IP protocol.
3. A server computer that can run Internet server software.
4. Server software that supports hypertext transport protocol (HTTP) requests from browsers (clients).
5. Desktop client computers equipped with network software capable of sending and receiving TCP/IP packet data.
6. Browser software installed on each client computer.

It should be noted that if a company does *not* want to use an internal server, an ISP can be used to support the Intranet. It is very common for organizations to use an ISP, especially when there is little information content or interest in maintaining a corporate-operated Intranet server. ISPs are also used when the organizational facilities can not support the housing of an Intranet server.

In addition to the software and hardware components listed above, HTML/XML documents must be prepared to provide information displays on the Intranet. The creation and conversion of documents to HTML/XML format is very easy using commercial software packages, such as Microsoft's FrontPage. Third-party sources are also available to provide this service at a reasonable cost.

2.1.4.1 Network Requirements

The first requirement for an Intranet is a computer network. For the purpose of this discussion, we assume that a basic computer network is in place. We now focus on the hardware and software modifications needed to support an Intranet.

Most computer networks are local-area networks (LANs). LANs are based on a client-server computing model that uses a central, dedicated computer — called the server — to fulfill client requests. The client-server computing model divides the network communication into two sides: a client side and a server side. By definition, the client requests information or services from the server. The server, in turn, responds to the client's requests. In many cases, each side of a client-server connection can perform both client and server functions.

Network servers are commonly used to send and receive e-mail, and to allow printers and files to be shared by multiple users. In addition, network servers normally have a storage area for server programs and to backup file copies. Server applications provide specific services. For example, a corporate-wide e-mail system typically uses a server process that is accessible from any computer within the company's network.

A server application (or server process) usually initializes itself and then goes to sleep, spending much of its time simply waiting for a request from a client application. Typically, a client process will transmit a request (across the network) for a connection to the server, and then it will request some type of service through the connection. The server can be located at either a local or remote site.

Every computer network has a physical topology by which it is connected. The most common topologies used to connect computers are the star, token ring, and bus topologies.

A network-interface card (NIC) is needed to physically connect a computer to the network. The network-interface card resides in the computer and provides a connector to plug into the network. Depending on the network, twisted-pair wiring, fiber optic, or coaxial cable may be used to physically connect the network components. The network-interface card must be compatible with the underlying network technology employed in the network (e.g., Ethernet or Token Ring).

2.1.4.2 Network Operating Systems

The Internet supports connectivity between various hardware platforms running various operating systems. In theory, there is no reason why an organization must stay with one type of machine or operating system when implementing an Intranet. However, in practice, many organizations use only one network operating system to simplify the task of managing the network.

The primary choices for network operating systems are: UNIX, Windows NT, and Novell NetWare. We now discuss each of these operating systems and important considerations surrounding their use.

UNIX

Many larger companies use UNIX-based machines as their primary business application server platform. UNIX is a proven operating system that is well suited for the Internet's open system model. Unfortunately, learning how to use UNIX is not easy. Also, using a UNIX-based machine limits the choices available for developing interactive Intranets and other software applications. Many programmers, for example, prefer to develop applications using Windows-based machines and programming languages (such as Microsoft's Visual Basic or Borland's Delphi).

Windows NT

Many companies choose Windows NT over UNIX because NT is easy to install, maintain, and administer. Windows NT, like UNIX and OS/2, provides a high-performance, multi-tasking workstation operating system. It also supports advanced server functions (including HTTP, FTP, and Gopher) and communications with clients running under MS-DOS, Windows 3.1, Windows 95, Windows for Workgroups, Windows NT Workstation, UNIX, or Macintosh operating systems. The latest version of Windows NT Server includes a free Internet Information Server (IIS) and a free Web browser (Internet Explorer). Microsoft designed the IIS so that it can be installed and up and running on a Windows NT workstation in less than 10 minutes. The Windows NT Server comes with a built-in remote access services feature that supports remote access to the Intranet through a dial-up phone connection.

Novell NetWare (IPX/SPX)

The NetWare operating system provides network-wide file and printer sharing for Ethernet or token ring networks. It runs on all major computer platforms, including UNIX, DOS, Macintosh, and Windows. However, behind the scenes, NetWare sends and receives data packets based on the Internetwork Packet Exchange/Sequenced Packet Exchange (IPX/SPX) protocol. Like TCP/IP, the IPX/SPX protocol defines a set of rules for coordinating network communication between network components.

Many companies use Novell network products within LANs to operate file and print servers. Therefore, it is important to understand how Novell's NetWare products can be used in an Intranet implementation strategy. If a company has an existing NetWare network, it might choose not to provide TCP/IP software to its network clients. Instead, the local-area network features provided by the NetWare software can be used.

A true Intranet uses Internet technology. This implies that an Internet Protocol (IP) address is assigned to each network computer (actually to each network-interface card), and that the TCP/IP protocol is used in the network. However, it is possible to run an Intranet on top of a NetWare LAN using various software products that translate IPX to IP. Many of these software packages provide IPX to IP translation while leaving the existing LAN infrastructure unchanged.

For example, the Novell product, IntranetWare, does not require assignment of an IP address to each client on a network. Instead, an IP address is only assigned to the NetWare Web server. The software performs IPX to IP translation on the client side, translating the TCP/IP protocols used by a Web browser to IPX protocols. After the protocol translation on the client side, the messages travel across the network until they reach the NetWare Web Server. At this point, the IntranetWare Server running on the NetWare Web server translates the IPX messages back into TCP/IP so they can be sent on to other servers on the network.

Another way to create an Intranet using a NetWare-based LAN is to use the Internet gateway product Inetix from Micro Computer Systems, Inc. Inetix runs on NetWare, Windows NT, or UNIX servers. Inetix does not require that the NetWare client machines support TCP/IP, only that a single IP address be assigned to the Intranet Web server. The Inetix client software allows a Web browser to execute on the client side, even though the client machine does not use TCP/IP-based software.

In a Mac-based network, NetWare for Macs, or AppleTalk¹ can be used to support the Intranet.

2.1.4.3 Server Hardware

Server machines run the network operating system and control how network computers share server resources. Large businesses with thousands of users typically use high-speed Unix-based machines for their servers. Small and medium-sized companies normally use less expensive Intel-based machines. The load (i.e., the number of users and the amount of network traffic) on the Intranet server machine will influence the selection of a specific processor type.

There is considerable debate in the industry as to which machine makes a better Intranet server: a Unix workstation, an Intel-based machine, or a PowerPC-based system. In general, the server choice depends on the plans for the Intranet and the level of familiarity the company has with each of these platforms. The server hardware selection also depends on the network operating system in use.

If a Unix-based server is chosen, a company will pay more for an equivalent amount of computing power provided by an Intel-based machine. Unix machines still carry a price premium over PCs, because they are made from custom parts, while Intel-based machines are made from commodity components available from many hardware vendors and suppliers. For example, a high-end Pentium machine with the capacity to serve over 1000 client machines can be bought for about one tenth of the cost of a comparable Unix server. Macintosh-based systems are more expensive than comparable Intel-based machines, but they are still much less expensive than Unix-based machines.

The decision to use a Unix-based machine vs. an Intel-based machine as the Intranet server is also influenced by maintenance costs. Maintaining a Unix-based machine requires more resources than maintaining an Intel-based machine. Hardware upgrades for Intel-based machines are also cheaper than hardware upgrades for Unix workstations. A Macintosh server will cost more to upgrade than an Intel-based machine. However, these costs are still lower than a comparable upgrade on a Unix-based machine.

The debate over Unix- and Intel-based machines focuses primarily on their performance in supporting business application servers. For example, companies that use large accounting and financial software packages often use Unix servers. On the other hand, companies that do not want to pay the price premium for Unix machines and/or are not familiar with Unix machines often select Intel-based machines as their business application servers.

A Pentium-class machine supports a vast array of software applications and server software. Many industry experts believe that Pentium-class machines will take a significant amount of the market share away from Unix workstations. This, in turn, means that more and more Intranet-based applications will use Pentium-class machines in the future.

2.1.4.4 Web Server Software

A working Intranet requires server software that can handle requests from browsers. In addition, server software is needed to retrieve files and to run application programs (which might, for instance, be used to search a database or to process a form containing user-supplied information).

For the most part, selecting a Web server for an Intranet is similar to selecting a Web server for an Internet site. However, Internet servers must generally handle larger numbers of requests and must deal with more difficult security issues. The performance of the Web server has a major impact on the overall

¹AppleTalk — AppleTalk is Apple's proprietary network operating system for supporting a LAN. Unlike AppleTalk, TCP/IP supports both LANs and WANs.

performance of the Intranet. Fortunately, it is fairly easy to migrate from a small Web server to a larger, high-performance Web server as the system usage increases over time.

Web servers are available for both Windows NT and Unix from Microsoft, Netscape, and a number of other companies. The Web server selection is limited by the server operating system. In the discussion that follows, we describe Web servers for Unix, Windows NT, Windows 95, and Macintosh operating systems.

It is expected that in the next two to five years, stand-alone Web servers will be replaced by servers that are an integral part of the operating system. In addition, these Web servers will handle many tasks that require custom programming today, such as seamless connection to databases, video and audio processing, and document management.

UNIX Web Servers

One of the best and oldest Web servers for Unix-based machines is the National Center for Supercomputing Application's (NCSA) HTTP Web Server. Much of the Internet's growth is primarily due to the popularity of this server, which is free. NCSA is committed to the continued development of its Web server, which provides both common gateway scripting (CGI) capabilities and server side includes (SSI) software. SSI software is used by Web servers to display and/or capture dynamic (changing) information on an HTML/XML page. For example, SSI can be used to display a counter showing the number of visitors to a Web site. The NCSA Web server also allows the creation of virtual servers on the same machine. The virtual servers can have their own unique universal resource identifier (URL). This is useful, for example, for assigning a different IP address to different departments using the same machine.

Netscape Communications Corporation offers the most popular commercial Unix-based Web servers: Enterprise Server and FastTrack Server. Enterprise Server software is designed for building large Intranets. FastTrack is easier to install than Enterprise Server but it offers less functionality. FastTrack is well suited for companies that plan to build small- to medium-sized Intranets. These Web servers are also available for Windows NT.

Windows NT Web Servers

As discussed in the previous section, both Netscape's Enterprise Server and FastTrack Server are available for Windows NT. These servers should be considered when both Windows NT and Unix-based machines are used as servers.

If cross-platform Web server software is not needed, Microsoft's IIS should be considered. This server is used to drive Microsoft's Internet site, and it works well for a large organization. At the time of this writing, Microsoft offers this server free of charge. It also comes bundled with Microsoft's Windows NT Server software. IIS is easy to install and allows new users to be added to the Intranet with minimal effort. IIS comes with an FTP server.

WebSite Professional, from O'Reilly & Associates, Inc., is another popular Windows NT Web server. WebSite Professional is very easy to install. This server has built-in search capabilities, a Web site management tool, and the popular HTML editing tool, HotDog.

Macintosh Web Server

Not many Web servers are available for Macintosh computers. If your organization does not have a Macintosh network already in place, then it is advisable to avoid installing a Mac-based Intranet server.

The largest market share of the Macintosh Web server market belongs to the WebStar server from Quarterdeck Corporation. This server is a mature product and is a good choice for a Mac-based Web server. WebStar is very easy to install and maintain.

NetWare Web Server

The NetWare Web Server from Novell is an excellent choice for companies that have a NetWare network already in place. When using a NetWare Web server, IPX to IP translation software must be installed on each client machine. When this is done, an IP address does not need to be assigned to each client machine.

2.1.4.5 Desktop Clients Running TCP/IP

TCP/IP must be installed on each client machine running on the Internet.² To use an Internet-based application (such as a Web browser) on a Windows-based machine, a TCP/IP stack must be present. Windows 95, Windows NT, and IBM's OS/2 Warp operating systems include the TCP/IP protocol suite. Most Unix-based systems use TCP/IP as their main network communication protocol.

If a company has Windows 3.1 clients, they should consider upgrading the clients to Windows 95 or Windows NT. Many of the advanced Internet and Intranet applications are only available for Unix, Windows 95, and Windows NT operating systems. If the Windows 3.1 clients can not be upgraded, then TCP/IP software must be installed for each client. One of the more popular TCP/IP software applications for Windows 3.1 is Trumpet Winsock. Trumpet Winsock can be downloaded from the Internet for free.

2.1.4.6 Web Browsers

The last component needed to make a functional Intranet is a Web browser. There are two major choices for a browser: Microsoft's Internet Explorer or Netscape's Navigator.

It is expected that Netscape will retain its dominance in the Unix market, since Microsoft has chosen not to support a Unix platform. Therefore, if the Intranet must support Unix, Macintosh, and Windows clients, and the company requires standardization on a single browser, Netscape's Navigator product is the only viable option.

Browsers, as we know them today, will probably not exist in a few years. For example, it is expected that eventually Microsoft will integrate the browser's functionality into its business application software (such as Word, Excel, etc.) and operating system.

2.1.4.7 Intranet Component Summary

In this section, the basic components of an Intranet were examined. We recapitulate below some of the key concepts covered in this section:

- Intranets are based on a client-server network computing model. By definition, the client side of a network requests information or services and the server side responds to a client's requests.
- The physical components of an Intranet include network interface cards, cables, and computers.
- Suites of protocols, such as TCP/IP and IPX/SPX, manage data communication for various network technologies, network operating systems, and client operating systems.
- IPX to IP translation programs provide NetWare users with the ability to build an Intranet without running the TCP/IP suite of protocols on their network.
- Windows-based Intranets are easier and less expensive to deploy than Unix-based Intranets.
- Netscape and Microsoft provide both low- and high-end server software products designed to meet the needs of large, medium, and small organizations.
- Netscape's Navigator and Microsoft's Internet Explorer provide advanced browser features for Intranet applications.

2.1.5 Intranet Implementation

2.1.5.1 Information Organization

After the physical components of the Intranet are in place, the next step is to design the information content of the Intranet and/or Internet Web pages. This task involves identifying the major categories and topics of information that will be made available on the Intranet. Information can be organized by department, function, project, content, or any other useful categorization scheme. It is advisable to use

²If the network does not support TCP/IP, a gateway application that translates TCP/IP for the network operating system protocol must be used.

cross-functional design teams to help define the appropriate informational categories that should be included on the corporate Web site. The following types of information are commonly found on corporate Intranet homepages:

- What's new
- Corporate information (history and contacts)
- Help desk and technical support
- Software and tools library
- Business resources
- Sales and marketing information
- Product information
- Human resources related information (benefits information, etc.)
- Internal job postings
- Customer feedback
- Telephone and e-mail directory
- Quality and system maintenance records
- Plant and equipment records
- Finance and accounting information
- Keyword search/index capability

2.1.5.2 Content Structure

After the main topics of information to be displayed on the corporate Web page(s) have been identified, the flow and manner of presentation on the Intranet must be developed. Four primary flow models are used to structure the flow of presentation at an Intranet Web site: linear, hierarchical, non-linear (or Web), and combination information structures.

A linear information structure is similar in layout to a book in that information is linked sequentially, page by page. When a linear layout is used, the Web pages are organized in a "slide show" format. This layout is good for presenting pages that should be read in a specific sequence or order. Since linear layouts are very structured, they limit the reader's ability to explore and browse the Web page contents in a non-sequential (or non-linear) manner.

When a hierarchical layout is used to structure the information, all the Web pages branch off from the home page or main index. This layout is used when the material in the Web pages does not need to be read in any particular order. A hierarchical information structure creates linear paths that only allow up and down movements within the document structure.

A non-linear, or Web, structure links information based on related content. It has no apparent structure. Non-linear structures allow the reader to wander through information spontaneously by providing links that allow forward, backward, up and down, diagonal, and side-to-side movement within a document. A non-linear structure can be confusing, and readers may get lost within the content, so this structure should be chosen with care. The World Wide Web uses a non-linear structure. The advantage of a non-linear structure is that it encourages the reader to browse freely.

The combination Web page layout, as the name implies, combines elements of the linear, Web, and hierarchical layouts. Regardless of the type of flow sequence employed, each Web page typically has links that allow the user to move back and forth between pages and back to the home page.

Over the lifetime of the Intranet, it is likely that the layout and organization of information on the corporate Web pages will change many times. It is often helpful to use flow charting tools to help manage and document the updated information flows. Visio for Windows by Visio Corp. and ABC Flowcharter by Micrografx are two excellent tools for developing flowcharts. In addition, some of the Web authoring tools offer flowcharting and organizational tools to help design and update the information structure on the Web pages.

2.1.5.3 Interface Design

After defining the Intranet's structure, the next step is to define the functionality and user interface. The Intranet design should be consistent with the organization's corporate image. For example, items such as corporate images, logos, trademarks, icons, and related design themes add a familiar look and feel to the content. Where possible, they should be included in the Web page design. It is also advisable to work with the marketing department when designing the web page layouts to ensure that a consistent theme is maintained in all the company communications that will be viewed by the outside world.

A technique called storyboarding is frequently used to design the web page layout. Storyboards are used by film producers, story writers, and comic strip artists to organize the content and sequence of their work. A storyboard depicts the content, images, and links between pages of the Intranet in the form of a rough outline.

Software — such as Microsoft PowerPoint or a similar presentation program — can be used to develop a storyboard and sample Web pages. It is a good idea to test the interface design to ensure that the icons, buttons, and navigational tools are logical and intuitive. An Intranet without intuitive navigational tools is like a road without signs. Just as it would be difficult for drivers to find their way from one city to another without the aid of signs, street names, and directional information, Intranet users will find it difficult to retrieve information without easy to follow categories, buttons, and links. It is often helpful to employ graphic designers and marketing-communications staff to create effective graphics and images for the web site.

2.1.5.4 Free Clip Art and Images

Many icons and navigational signs are available as clip art that comes with word processing, page layout, and presentation software programs. In addition, many Web sites offer images and clip art, which can be downloaded for free. However, the licensing agreements for downloading free images may have restrictions and requirements that should be observed.

2.1.5.5 Intranet Functionality

The required functionality of an Intranet dictates many of the design and user interface features. One of the goals in designing the Intranet should be to improve existing systems and infrastructures. After examining the current information structure, it may become clear which aspects of the structure work well and which ones need improvement.

Workflow processes, document management, and work collaboration are areas that the organization should strive to improve through the use of an Intranet. A workflow analysis should consider ways in which the Intranet can automate various organizational tasks and processes. For example, if a company has a geographically dispersed project team, the Intranet might be used to post and update project information as various tasks are completed. Other team members could then visit the Intranet page at any time to check the project status.

The following checklist is helpful when developing a list of the functions that need to be supported by the Intranet:

Functionality Checklist

- The user interface must be intuitive and tested
- The Intranet's design should support continuous updates
- The Intranet may need to be integrated with database management systems to allow users to access information (such as customer and product data)
- The Intranet should support existing (legacy) applications, as needed
- The Intranet should have built-in directories, such as corporate telephone numbers and e-mail addresses
- The Intranet should incorporate groupware applications
- Support (or future expansion) for on-line conferencing should be considered
- The Intranet should provide division-specific and corporate-wide bulletin boards for electronic postings

- The Intranet should be designed with a document sharing and management process in mind
- The Intranet should foster teamwork and collaboration by enhancing channels of information distribution
- Search engines — which simplify a user's ability to locate and access information — should be made available
- The Intranet should support e-mail
- The Intranet should support (for future expansion) multimedia applications that use text, images, audio, and video
- Automated real-time Web-page generation should be encouraged
- The Intranet should be designed so it can interface, at least potentially, with factory equipment, other manufacturing devices, or other critical legacy systems
- The Intranet should support the automation of organization workflows

2.1.5.6 Content Management

Many organizations struggle with the tasks of information creation, management, and dissemination. They are time consuming and difficult to control. The Intranet alone cannot solve information management problems unless specific Intranet solutions are implemented that directly address the need for document management. The following list identifies content-management tasks that should be considered in the Intranet plan:

- Users must have the ability to easily add or update content on a regular basis
- Users must have the ability to protect their content from changes by other users
- A content-approval process should be defined and in place. This process should encompass ways to manage and control document revisions, especially changes to shared documents

As policies and procedures relating to content management are formulated, it is important to designate responsibilities to specific individuals to ensure that they are put into place and followed. An Intranet style guide should be developed that provides page layout, design elements, and HTML/XML code guidelines. The style guide will help the organization to maintain a consistent look and feel throughout the Intranet's Web pages. The style guide should contain information on where to obtain standard icons, buttons, and graphics, as well as guidelines on page dimensions and how to link to other pages. As part of the style guide, it is helpful to create Web page templates. These templates consist of HTML/XML files, and are used to provide a starting point for anyone interested in developing Web pages or content for the Intranet. Although it is very easy to create a working Web page and to publish it for mass viewing, the real challenge is in producing a well-conceived Web page.

2.1.5.7 Training and Support

After the Intranet is up and running, efforts should be focused on how to maintain the information content and on employee training. Part of the document-management strategy should encompass the selection of content stakeholders. Content stakeholders are individuals in different departments or work groups who are responsible for the creation and maintenance of specific content. Stakeholders can be department managers, team leaders, or content authors and publishers.

Some organizations create a position called a Webmaster. This position is responsible for maintaining and supporting the content published on the Intranet. A good Webmaster should have the following skills:

- Basic Internet skills, including an understanding of e-mail, FTP, and Telnet
- A thorough understanding of HTML/XML document creation
- Experience with CGI programming
- Programming experience with languages such as Perl, C/C++, and Java
- Experience with content creation and the conversion of text and images
- Knowledge of client-server processing
- Experience with server setup and maintenance

- Knowledge of your organization's structure and inner workings
- Organizational and training skills

It is possible that the organization may choose to decentralize the maintenance of the information content. In this case, individuals from various departments might be selected to maintain the content relating to their respective department. These individuals should be trained to handle a variety of maintenance issues. A decentralized approach depends on having more than one individual with the necessary skills available to maintain the web pages. A decentralized support structure gives authors and content owners direct control and responsibility for publishing and maintaining information. This can help prevent bottlenecks in making information available in a timely fashion.

Training for stakeholders, Webmasters, and Intranet users is an important part of an Intranet strategy. Intranet customers and content stakeholders should be trained to understand the Intranet and how it will improve the organization and the way the company does business. They should also be given training on how to create, utilize, and maintain content on the Web page(s). Companies that invest in the education and training of their employees will have a better chance of creating and maintaining a successful Intranet.

2.1.6 Intranet Deployment

Since Intranets are easy to set up, many companies do not realize what the true resource requirements are to maintain the Intranet with up-to-date information. The goal of this section is to provide a realistic perspective on how organizations are most likely to achieve long-lasting benefits from the Intranet.

Some companies invest much more than their competitors in information technology, such as an Intranet, but still fail to effectively compete in the marketplace. Computers alone do not, and cannot, create successful companies. A good start is to empower all employees to contribute to the Intranet. As is true for any collaborative effort, every member is responsible for the overall success of the team.

2.1.6.1 Technological Considerations

The major technological challenges facing the organization after the initial implementation of an Intranet include:

- Converting existing paper documents into electronic documents that employees can access electronically via the Intranet.
- Connecting existing databases to the Intranet so they are accessible by a wide range of computing platforms (such as Windows- and Mac-based systems).
- Coordinating the use of multiple servers used across departmental lines.
- Continuously enhancing the Intranet's features and capabilities to keep employees motivated to use the Intranet.
- Installing security features within the Intranet to prevent unauthorized access to confidential or sensitive information.

Intranet technology, and information technology in general, is changing so fast that keeping up with the latest software and hardware solutions requires a substantial ongoing organizational commitment.

Conversion of Paper Documents Into Electronic Form

The first issue facing companies after the initial Intranet release is how to convert large numbers of existing paper documents into electronic format ready for distribution on an Intranet. There are many tools, such as HTML Transit, that can be used to convert documents from most electronic formats to HTML or XML format. Microsoft's Internet Assistant for Microsoft Word can also be used to easily convert existing Word documents into HTML or XML documents. After paper documents have been converted to HTML or XML and placed on the Intranet, the next challenge is to keep the documents up to date.

TABLE 2.1.1 Intranet Document Tracking Information

Data For Tracking Intranet Documents

Name of document
Document description
Page owner
Type of document (i.e., official, unofficial, personal)
Confidentiality status (i.e., confidential, non-confidential, etc.)
Original publish date
Date document last modified
Frequency of update (i.e., daily, weekly, monthly, etc.)

Obsolete information can frustrate Intranet users and may encourage them to revert to old ways of information gathering (i.e., calling people, walking to various offices, and writing memos). One way to minimize this problem is to create a database containing the document title, date of last change, and frequency of update in a database. Other useful information that can be used to track the status and nature of documents on the Intranet is shown in [Table 2.1.1](#). A program can then be written to search the Intranet for documents that have not been updated recently. The program can then issue e-mail to the document owner to request an update.

Interface to Legacy Database(s)

Connecting databases to the Intranet is not an easy task, and may require additional staff or reassignment of current programming staff. Legacy database vendors are currently working on various Intranet solutions to facilitate the implementation of this requirement.

Companies may need to connect the Intranet to legacy databases in order to access:

- Financial reports (regarding project costs, product costs, the overall financial health of the enterprise, etc.)
- Document-management systems
- Human resources information (e.g., so employees can review details on health care and benefits)

Use of Multiple Servers

As the Intranet becomes more complex, multiple servers will be needed. This is especially true for companies that have a large number of divisions and business units using the Intranet. For example, a product-development group may need to provide team members the ability to search project-specific databases, submit forms to various databases, and to use a private on-line discussion group. The Webmaster may find it impossible to support these service needs in a timely manner. When this happens, companies frequently relegate the task of server maintenance to each respective department.

Over the next few years, installing and using a Web server will become as easy as installing and using word processor software. Web servers will probably become part of the Windows NT server operating system. When each department is responsible for maintaining their own Web server, it is particularly important to choose server software that is easy to install and maintain. A Pentium-class machine running Windows NT server software and Microsoft's IIS is a good choice for small departments. Another way to provide departments with their own domain name and disk space is to use a virtual domain name. Companies use virtual servers to reduce hardware costs. In the case of the Web, an HTTP-based server runs on a server computer. For example, a company may need two types of Web servers, one that allows easy access and one that requires usernames and passwords. In the past, the company would have purchased two different computers to run the Web server software. Today, however, the company can run both servers on the same system — as virtual servers.

Standardizing Hardware and Software

To avoid supporting multiple hardware and software components, it is important to standardize the server software, hardware, HTML and XML editing tools, and browser software. This will help to minimize the potential for unexpected network errors and incompatibilities.

2.1.6.2 Maintaining the Information Content on the Intranet

One of the major challenges organizations must face is how to transition from paper-based systems to computer-based systems, while keeping information up to date.

Automating HTML/XML Authoring

After establishing a policy for the distribution of Intranet documents, it is advisable to develop a set of guidelines that clearly specifies who is responsible for keeping them current. Inaccurate information greatly reduces the effectiveness of the Intranet. If employees lose confidence in the accuracy of the on-line information, they will revert to calling people to find information. Unfortunately, many people tend to ignore the need to update information, irrespective of its form (i.e., electronic or print).

In some cases, the Intranet will contain information that employees must update daily, weekly, or monthly. Spreadsheets can be used to capture highly time-sensitive data. Macros can be written (typically, by a staff programmer) that automatically convert the spreadsheet data into HTML or XML format.

Managing Document Links

In a traditional document-management system, documents often reference one another. In most cases, authors list the applicable references at the top of each new document. Intranets, unfortunately, create a situation where organizations cannot easily control the accuracy of links in documents.

HTML or XML document developers can use links freely and, in many cases, without checking the accuracy of those links. Even if employees test the initial accuracy of their document links, it is difficult to maintain and check the accuracy of those links after the document is released. If you have ever encountered a “broken” link when surfing the Web, you know that it can be frustrating. People depend on links within a Web document to find information. Today, however, there are a few mechanisms available to assure the accuracy of document links. Employees must understand that other people may link to their pages, and that they should not freely move the location of their documents. Employees must view their needs in the total organizational context.

2.1.6.3 Centralized vs. Distributed Control

The implementation of an Intranet is a major change for any organization. Although change is not easy, people are more inclined to modify their behavior when leaders have a clear sense of direction, involve employees in developing that direction, and are able to demonstrate how the Intranet will positively affect the employees’ well being. Managers should work with their employees to show that Intranets can free them from the routine aspects of their job. This, in turn, will allow employees to spend more time learning and developing new ideas for the corporation.

Some of the benefits that can be obtained using a distributed model of Intranet control are:

- Employees can tap into the knowledge of everyone in the organization, making everyone a part of a solution.
- The power of any one Webmaster to dictate the Intranet’s form and function is limited.
- It empowers departments to create their own information databases and to work with outside customers and vendors.

2.1.7 Intranet Security Issues

By their very nature, Intranets encourage a free flow of information. This means that it is also very easy for information to flow directly from the Intranet to the desktops of those who might seek to gain access to information they should not have. To guard against this situation, adequate security measures should be in place when the Intranet is deployed. In the discussion that follows, we review various security techniques to protect an Intranet from unauthorized external and internal use.

2.1.7.1 Firewalls

The Internet was designed to be resistant to network attacks in the form of equipment breakdowns, broken cabling, and power outages. Unfortunately, the Internet today needs additional technology to prevent attacks against user privacy and company security. Luckily, a variety of hardware and software

solutions exist to help protect an Intranet. The term *firewall* is a basic component of network security. A firewall is a collection of hardware and software that interconnects two or more networks and, at the same time, provides a central location for managing security. It is essentially a computer specifically fortified to withstand various network attacks. Network designers place firewalls on a network as a first line of network defense. It becomes a “choke point” for all communications that lead in and out of an Intranet. By centralizing access through one computer (known as a *firewall-bastion host*), it is easier to manage the network security and to configure appropriate software on one machine. The bastion host is also sometimes referred to as a *server*.

The firewall is a system that controls access between two networks. Normally, installing a firewall between an Intranet and the Internet is a way to prevent the rest of the world from accessing a private Intranet. Many companies provide their employees with access to the Internet long before they give them access to an Intranet. Thus, by the time the Intranet is deployed, the company has typically already installed a connection through a firewall. Besides protecting an Intranet from Internet users, the company may also need to protect or isolate various departments within the Intranet from one another, particularly when sensitive information is being accessed via the Intranet. A firewall can protect the organization from both internal and external security threats.

Most firewalls support some level of encryption, which means data can be sent from the Intranet, through the firewall, encrypted, and sent to the Internet. Likewise, encrypted data can come in from the Internet, and the firewall can decrypt the data before it reaches the Intranet. By using encryption, geographically dispersed Intranets can be connected through the Internet without worrying about someone intercepting and reading the data. Also, a company’s mobile employees can use encryption when they dial into your system (perhaps via the Internet) to access private Intranet files.

In addition to firewalls, a router can be used to filter out data packets based on specific selection criteria. Thus, the router can allow certain packets into the network while rejecting others.

One way to prevent outsiders from gaining access to an Intranet is to physically isolate it from the Internet. The simplest way to isolate an Intranet is to not physically connect it to the Internet. Another method is to connect two sets of cables, one for the Intranet and the other for the Internet.

Even without a connection to the Internet, an organization is susceptible to unauthorized access. To reduce the opportunity for intrusions, a policy should be implemented that requires frequent password changes and keeping that information confidential. For example, disgruntled employees, including those who have been recently laid off, can be a serious security threat. Such employees might want to leak anything from source code to company strategies to the outside. In addition, casual business conversations, overheard in a restaurant or other public place, may lead to a compromise in security. Unfortunately, a firewall cannot solve all these specific security risks.

It should be noted that a firewall can not keep viruses out of a network. Viruses are a growing and very serious security threat. Prevention of viruses from entering an Intranet from the Internet by users who upload files is necessary. To protect the network, everyone should run anti-virus software on a regular basis.

The need for a firewall implies a connection to the outside world. By assessing the types of communications expected to cross between an Intranet and the Internet, one can formulate a specific firewall design. Some of the questions that should be asked when designing a firewall strategy include:

- Will Internet-based users be allowed to upload or download files to or from the company server?
- Are there particular users (such as competitors) that should be denied all access?
- Will the company publish a Web page?
- Will the site provide telnet support to Internet users?
- Should the company’s Intranet users have unrestricted Web access?
- Are statistics needed on who is trying to access the system through the firewall?
- Will a dedicated staff be implemented to monitor firewall security?
- What is the worst-case scenario if an attacker were to break into the Intranet? What can be done to limit the scope and impact of this type of scenario?
- Do users need to connect to geographically dispersed Intranets?

There are three main types of firewalls: network level, application level, and circuit level. Each type of firewall provides a somewhat different method of protecting the Intranet. Firewall selection should be based on the organization's security needs.

Network, Application, and Circuit-level Firewalls

Network-level Firewall

A network-level firewall is typically a router or special computer that examines packet addresses, and then decides whether to pass the packet through or to block it from entering the Intranet. The packets contain the sender and recipient IP address, and other packet information. The network-level router recognizes and performs specific actions for various predefined requests. Normally, the router (firewall) will examine the following information when deciding whether to allow a packet on the network:

- Source address from which the data is coming
- Destination address to which the data is going
- Session protocol such as TCP, UDP, or ICMP
- Source and destination application port for the desired service
- Whether the packet is the start of a connection request

If properly installed and configured, a network-level firewall will be fast and transparent to users.

Application-level Firewall

An application-level firewall is normally a host computer running software known as a proxy server. A proxy server is an application that controls the traffic between two networks. When using an application-level firewall, the Intranet and the Internet are not physically connected. Thus, the traffic that flows on one network never mixes with the traffic of the other because the two network cables are not connected. The proxy server transfers copies of packets from one network to the other. This type of firewall effectively masks the origin of the initiating connection and protects the Intranet from Internet users.

Because proxy servers understand network protocols, they can be configured to control the services performed on the network. For example, a proxy server might allow ftp file downloads, while disallowing ftp file uploads. When implementing an application-level proxy server, users must use client programs that support proxy operations.

Application-level firewalls also provide the ability to audit the type and amount of traffic to and from a particular site. Because application-level firewalls make a distinct physical separation between an Intranet and the Internet, they are a good choice for networks with high-security requirements. However, due to the software needed to analyze the packets and to make decisions about access control, application-level firewalls tend to reduce the network performance.

Circuit-level Firewalls

A circuit-level firewall is similar to an application-level firewall in that it, too, is a proxy server. The difference is that a circuit-level firewall does not require special proxy-client applications. As discussed in the previous section, application-level firewalls require special proxy software for each service, such as ftp, telnet, and HTTP.

In contrast, a circuit-level firewall creates a circuit between a client and server without needing to know anything about the service required. The advantage of a circuit-level firewall is that it provides service for a wide variety of protocols, whereas an application-level firewall requires an application-level proxy for each and every service. For example, if a circuit-level firewall is used for HTTP, ftp, or telnet, the applications do not need to be changed. You simply run existing software. Another benefit of circuit-level firewalls is that they work with only a single proxy server, making it easier to manage, log, and control a single server than multiple servers.

2.1.7.1.1 Firewall Architectures

Combining the use of both a router and a proxy server into the firewall can maximize the Intranet's security. The three most popular firewall architectures are the dual-homed host firewall, the screened host firewall, and the screened subnet firewall. The screened-host and screened-subnet firewalls use a combination of routers and proxy servers.

Dual-homed Host Firewalls

A dual-homed host firewall is a simple, yet very secure configuration in which one host computer is dedicated as the dividing line between the Intranet and the Internet. The host computer uses two separate network cards to connect to each network. When using a dual-home host firewall, the computer routing capabilities should be disabled, so the two networks do not accidentally become connected. One of the drawbacks of this configuration is that it is easy to inadvertently enable internal routing.

Dual-homed host firewalls use either an application-level or a circuit-level proxy. Proxy software controls the packet flow from one network to another. Because the host computer is dual-homed (i.e., it is connected to both networks), the host firewall can examine packets on both networks. It then uses proxy software to control the traffic between the networks.

Screened-host Firewalls

Many network designers consider screened-host firewalls more secure than a dual-homed host firewall. This approach involves adding a router and placing the host computer away from the Internet. This is a very effective and easy-to-maintain firewall. A router connects the Internet to your Intranet and, at the same time, filters packets allowed on the network. The router can be configured so that it sees only one host computer on the Intranet network. Users on the network who want to connect to the Internet must do so through this host computer. Thus, internal users appear to have direct access to the Internet, but the host computer restricts access by external users.

Screened-subnet Firewalls

A screened-subnet firewall architecture further isolates the Intranet from the Internet by incorporating an intermediate perimeter network. In a screened-subnet firewall, a host computer is placed on a perimeter network which users can access through two separate routers. One router controls Intranet traffic and the second controls the Internet traffic. A screened-subnet firewall provides a formidable defense against attack. The firewall isolates the host computer on a separate network, thereby reducing the impact of an attack to the host computer. This minimizes the scope and chance of a network attack.

2.1.7.2 CGI Scripting

Web Sites that provide two-way communications use CGI (common gateway scripting). For example, if you fill in a form and click your mouse on the form's Submit button, your browser requests the server computer to run a special program, typically a CGI script, to process the form's content. The CGI script runs on the server computer, which processes the form. The server then returns the output to the browser for display.

From a security perspective, the danger of CGI scripts is that they give users the power to make a server perform a task. Normally, the CGI process works well, providing an easy way for users to access information. Unfortunately, it is also possible to use CGI scripts in ways they were never intended. In some cases, attackers can shut down a server by sending potentially damaging data through the use of CGI scripts. From a security perspective, it is important to make sure that users cannot use CGI scripts to execute potentially damaging commands on a server.

2.1.7.3 Encryption

Encryption prevents others from reading your documents by “jumbling” the contents of your file in such a way that it becomes unintelligible to anyone who views it. You must have a special key to decrypt the file so its contents can be read. A key is a special number, much like the combination of a padlock, which the encryption hardware or software uses to encrypt and decrypt files. Just as padlock numbers have a certain number of digits, so do encryption keys. When people talk about 40-bit or 128-bit keys, they are simply referring to the number of binary digits in the encryption key. The more bits in the key, the more secure the encryption and less likely an attacker can guess your key and unlock the file. However, attackers have already found ways to crack 40-bit keys.

Several forms of encryption can be used to secure the network, including: link encryption, document encryption, secure-sockets layer (SSL), and secure HTTP (S-HTTP). The following sections describe these encryption methods in more detail.

2.1.7.3.1 Public-key Encryption

Public-key encryption uses two separate keys: a public key and a private key. A user gives his/her public key to other users so anyone may send them encrypted files. The user activates his/her private key to decrypt the files (which were encrypted with a public key).

A public key only allows people to encrypt files, not to decrypt them. The private user key (designed to work in conjunction with a particular public key) is the only key that can decrypt the file. Therefore, the only person that can decrypt a message is the person holding the private key.

2.1.7.3.2 Digital Signatures

A digital signature is used to validate the identity of the file sender. A digital signature prevents clever programmers from forging e-mail messages. For example, a programmer who is familiar with e-mail protocols can build and send an e-mail using anyone's e-mail address, such as BillGates@microsoft.com.

When using public-key encryption, a sender encrypts a document using a public key, and the recipient decodes the document using a private key. With a digital signature, the reverse occurs. The sender uses a private key to encrypt a signature, and the recipient decodes the signature using a public key. Because the sender is the only person who can encrypt his or her signature, only the sender can authenticate messages. To obtain a personal digital signature, you must register a private key with a certificate authority (CA), which can attest that you are on record as the only person with that key.

2.1.7.3.3 Link Encryption

Link encryption is used to encrypt transmissions between two distant sites. It requires that both sites agree on the encryption keys that will be used. It is commonly used by parties that need to communicate with each other frequently. Link encryption requires a dedicated line and special encryption software. It is an expensive way to encrypt data. As an alternative to this, many routers have convenient built-in encryption options. The most common protocols used for link encryption are PAP (Password Authentication) and CHAP (Challenge Handshake Authentication Protocol). Authentication occurs at the data link layer and is transparent to end-users.

2.1.7.3.4 Document Encryption

Document encryption is a process by which a sender encrypts documents that the recipient(s) must later decrypt. Document encryption places the burden of security directly on those involved in the communication. The major weakness of document encryption is that it adds a step to the process by which a sender and receiver exchange and receive documents. Because of this extra step, many users prefer to save time by skipping the encryption. The primary advantage of document encryption is that anyone with an e-mail account can use document encryption. Many document encryption systems are available free or for little cost on the Internet.

2.1.7.3.5 Pretty Good Privacy (PGP)

Pretty good privacy (PGP) is a free (for personal use) e-mail security program developed in 1991 to support public-key encryption, digital signatures, and data compression. PGP is based on a 128-bit key. Before sending an e-mail message, PGP is used to encrypt the document. The recipient also uses PGP to decrypt the document. PGP also offers a document compression option. Besides making a document smaller, compression enhances the file security because compressed files are more difficult to decode without the appropriate key. According to the PGP documentation, it would take 300 billion years for someone to use brute force methods to decode a PGP-encrypted, compressed message.

2.1.7.3.6 Secure Socket Layer (SSL)

The Secure Socket Layer (SSL) was developed by Netscape Communications to encrypt TCP/IP communications between two host computers. SSL can be used to encrypt any TCP/IP protocol, such as HTTP, telnet, and ftp. SSL works at the system level. Therefore, any user can take advantage of SSL because the SSL software automatically encrypts messages before they are put onto the network. At the recipient's end, SSL software automatically converts messages into a readable document.

SSL is based on public-key encryption and works in two steps. First, the two computers wishing to communicate must obtain a special session key (the key is valid only for the duration of the current

communication session). One computer encrypts the session key and transmits the key to the other computer. Second, after both sides know the session key, the transmitting computer uses the session key to encrypt messages. After the document transfer is complete, the recipient uses the same session key to decrypt the document.

2.1.7.3.7 Secure HTTP (S-HTTP)

Secure HTTP is a protocol developed by the CommerceNet coalition. It operates at the level of the HTTP protocol. S-HTTP is less widely supported than Netscape's Secure Socket Layer. Because S-HTTP works only with HTTP, it does not address security concerns for other popular protocols, such as ftp and telnet.

S-HTTP works similarly to SSL in that it requires both the sender and receiver to negotiate and use a secure key. Both SSL and S-HTTP require special server and browser software to perform the encryption.

2.1.7.4 Intranet Security Threats

This section examines additional network threats that should be considered when implementing Intranet security policies.

2.1.7.4.1 Source-routed Traffic

As discussed earlier, packet address information is contained in the packet header. When source routing is used, an explicit routing path for the communication can be chosen. For example, a sender could map a route that sends packets from one specific computer to another, through a specific set of network nodes. The road map information contained in the packet header is called *source routing*, and it is used mainly to debug network problems. It is also used in some specialized applications. Unfortunately, clever programmers can also use source routing to gain (unauthorized) access into a network. If a source-routed packet is modified so that it appears to be from a computer within your network, a router will obediently perform the packet routing instructions, permitting the packet to enter the network, *unless special precautions are taken*. One way to combat such attacks is simply to direct your firewall to block all source-routed packets. Most commercial routers provide an option to ignore source-routed packets.

2.1.7.4.2 Protecting Against ICMP Redirects (Spoofing)

ICMP stands for Internet Control Message Protocol. ICMP defines the rules routers use to exchange routing information. After a router sends a packet to another router, it waits to verify that the packet actually arrived at the specified router. Occasionally, a router may become overloaded or may malfunction. In such cases, the sending router might receive an ICMP-redirect message that indicates which new path the sending router should use for transmission.

It is fairly easy for knowledgeable "hackers" to forge ICMP-redirect messages to reroute communication traffic to some other destination. The term *spoofing* is used to describe the process of tricking a router into rerouting messages in this way. To prevent this type of unauthorized access, it may be necessary to implement a firewall that will screen ICMP traffic.

2.1.8 Summary

Intranets are being used to improve the overall productivity of the organization. Important Intranet concepts covered in this chapter are summarized below:

- TCP/IP was created because of the need for reliable networks that could span the globe. Because of its reliability and ease of implementation, TCP/IP has become the standard language (or protocol) of the Internet. TCP/IP defines how programs exchange information over the Internet.
- An Intranet is based on Internet technology. It consists of two types of computers: a client and a server. A client asks for and uses information that the server stores and manages.
- Telnet, ftp, and gopher are widely used network programs that help users connect to specific computers and to transfer and exchange files.
- The World Wide Web (or Web) is a collection of interlinked documents that users can access and view as "pages" using a special software program called a browser. The two most popular browser programs are Netscape Navigator and Microsoft Internet Explorer.

- HTML (hypertext markup language) and XML (extended markup language) are used to describe the layout and contents of pages on the Web.
- Java is a new computer programming language that allows users to execute special programs (called applets) while accessing and viewing a Web page.
- A network computer is a low-cost, specialized computer designed to work in conjunction with the Internet and Java application programs.

To be effective, the Intranet must deliver quality information content. To ensure this, management must be in a proactive role in assigning staff who will keep the corporate information reservoirs on the Intranet current and relevant. The following is a checklist of some of the ways to encourage the development of a high-quality Intranet:

- Give users access to document management systems and various corporate databases.
- Distribute the responsibility of maintaining the Intranet to increase the number of staff involved in developing and enhancing Intranet content.
- Create a corporate culture based on information sharing.
- Place employee training at the center of an Intranet deployment strategy.
- Design and implement appropriate security measures as soon as possible.
- Use firewalls to control access to the network.
- Use anti-virus software.
- Implement a security plan that controls the access that employees and outsiders have to the network.
- Design and implement CGI scripts with security in mind.
- Encourage users to encrypt files before sending confidential data across the Internet/Intranet.

2.2 Internet Security

John Braun

The Internet offers the ability for anyone, from an individual computer owner to a major multinational corporation, to make information and computing resources available to the world. Alas, since the Internet was initially designed with ease of communications, rather than security, in mind, care must be taken to ensure that sensitive information is not made available to the world as well. This section will discuss several aspects of protection information that one makes available via a network or the Internet.

2.2.1 Physical Security

Although not as trendy or exciting as some of the exotic attacks that can be made against a network on the protocol level, physical security is nonetheless important. You can have the best security software in the world installed on your network, but it may do you little good if an attacker can waltz up to a key piece of network or computing equipment and disable it!

Key pieces of network hardware, such as routers, firewalls, and servers, should be stored in a secure room with some sort of access control such as a traditional or electronic lock, card reader, or other means which can limit access to authorized individuals. Access to especially sensitive devices should be further restricted by placing them in a locked cabinet. Also, access to buildings that contain these rooms should also be controlled with security guards or other access control so that visitors can't wander about.

Exposed network cables, especially those connected to routers, hubs, and other devices which carry traffic for several other users, should also be physically protected. If an attacker has access to these cables, they could physically cut them and insert equipment that could monitor and even generate network traffic. For maximum protection, network cables should be placed inside of pressurized pipes with sensors placed on various locations along the piping. Network monitoring tools should also be used so that a break in a cable can be identified quickly.

2.2.2 Modems

Modems present two security threats. First, modems offer a channel for data to leave your premises, circumventing security and auditing measures that may be in place for the rest of the network. A review of services that are accessed by modem should be made, and, if possible, this access should be rerouted over a secure internal network. Second, modems offer a potential method for unauthorized individuals to access your network from the outside. Since there may be a need for users who are on the road or working from home to access a network remotely, additional security measures need to be taken for these connections. One measure is a dial-back modem, which will call back the user at a predetermined number before allowing access to the network. Another is a system where each user is provided with an electronic card that displays a random number every few minutes. A similar device that performs the same calculation to produce this number is located on the network one wishes to access. Without the card, and the ability to produce this number, the remote user is denied access.

2.2.3 Data Security

There are many aspects to data security. One is to prevent files from being viewed by someone other than the owner. On a multi-user system, this concern can be addressed by proper system administration. Users should not be allowed access to directories or files that do not belong to them. On a single-user system, the ability to share files over the network should be closely monitored, so that users don't inadvertently allow access of their hard drive contents to anyone who cares to look. Another aspect of data security is to protect the contents of file or network data via encryption. This can be especially important for data travelling over a network, since the data can be broadcast to other terminals or network devices which are not the intended recipients of the data. By using encryption, data that falls into the wrong hands will be unusable unless an encryption key or password is also known. Although security at the TCP/IP level is still a ways off, several third-party products provide the ability to protect and encrypt data.

2.2.4 Passwords

Since passwords usually comprise an initial layer of defense against an attack, they should be chosen and implemented with care. A written policy and/or enforcement by the operating system can help. Passwords should not be dictionary words, should be as long as possible, contain a series of letters, numbers, and other characters, and be changed on a regular basis. When deciding on a policy, care should be taken to balance security needs vs. ease of use. If a policy is too tedious to follow, users may end up writing their passwords down somewhere near their terminal, eliminating any sort of benefit the password policy could have offered.

2.2.5 Workstation Security

Unattended workstations could be a great danger to the entire system, and a security system could be completely wasted if an unauthorized person could access someone else's logged-in workstation. For that reason, users need to be aware of this danger and be properly trained how to secure an unattended workstation, either by logging off or by using a screen saver or screen lock which activates after a short amount of inactivity.

2.2.6 TCP/IP Security

Since the current version of TCP/IP was designed to provide a robust, standard method of moving data on a network, rather than security, one should be aware of several attacks which could compromise network security or availability.

2.2.6.1 IP Spoofing

Spoofing is the act of altering the contents of a TCP or IP packet header in order to trick the remote system into thinking the packet is valid. One trick is to change the source IP address of a packet to an

address that is valid on a network behind a firewall or router. Older equipment that would have otherwise blocked the packet will allow it to go through since it appears to be coming from a friendly network. There are attacks where a connection can be hijacked or terminated by combining IP address spoofing with the spoofing of the SEQ and ACK fields in a TCP header. The SEQ and ACK fields help synchronize traffic between two hosts. If these fields are modified by attackers, the attackers can take over connections, while legitimate hosts lose the connection since their packets now appear to be out of order. Additional fields could be activated so that the connection is terminated prematurely.

2.2.6.2 Denial of Service (DoS)

Many DoS attacks take advantage of nuances in the method used to establish a TCP/IP connection. Since connections may take a while to establish, portions of the TCP/IP establishment process include timeouts so that slow equipment or busy networks will not cause a connection attempt to fail. However, a program which intentionally completes only a portion of this negotiation will result in a host waiting for a connection to complete, when it never will. While the host is waiting for the connection attempt to time out, system resources are being used. If enough of these bogus attempts are made, the host will run out of resources, and future connection attempts will be refused.

Another major type of attack involves the sending of single packets, whose contents have been modified to some unexpected or invalid data. This can result in the remote system crashing with a nasty Blue Screen of Death, system bomb, or core dump. One attack sends an ICMP (a special type of IP packet, with no TCP) echo request, also known as a ping, whose data payload is very large. Since a ping packet normally has no data associated with it, some implementations that don't expect this data will grind to a halt when receiving this type of packet. Another attack interferes with the data offset field in the TCP header, so that the remote host is tricked into trying to read packet data where none exists, also causing a crash.

2.3 Virtual Private Networking Solutions

Endre Sara

Today's large corporate networks are geographically distributed and clients or employees need access to corporate information from different locations. The cost of a long-distance dialup session is very high, and it is also not efficient to deploy point-to-point connection between each possible location.

Virtual private network (VPN) is a concept of securely transferring sensitive corporate information between various geographically dispersed sites over a public network, such as the Internet. The market numbers for these services tell a success story and a win/win situation for both providers and users. But, while the market numbers are good, numerous concerns remain. They are:

- There are not enough integration services to help users deploy VPNs
- The products are not yet interoperable
- Security standards are not yet unique
- There are different protocol standards
- Many users are not yet fully comfortable using Internet technologies in the mainline business

Typical users of VPNs are driving the implementation of VPN services. The most important benefits and points to consider are summarized in [Table 2.3.1](#). Weights are not included in this consideration.

This document describes and compares the available standards and products for VPN solutions.

The VPN is a network that uses a private address space which operates over another network infrastructure. It means that the VPN will use the same physical cabling, switches, bridges, and routers, but it uses a different address space. This is accomplished by encapsulating the VPN traffic (which doesn't have to be IP) into secure protocols. The emerging standards concentrate around Layer 2 and Layer 3 protocols.

TABLE 2.3.1 VPN Benefits and Concerns

VPN Applications	Benefits	Points to Consider
Dial access for remote users	Outsource modems reduce dial-in costs Eliminate access lines	Client software Are appropriate tunneling protocols supported in client software? Encryption performance issues
Connecting branch offices	Reduces number of dedicated lines Lets IT managers consolidate central-site WAN equipment	Does VPN access-control system integrate with existing user access privileges?
Extranet	Gives trading partners and customers access to intranet Makes collaborating with contractors and consultants much easier	Does system scale well? Are there tools to handle the administrative burden of adding new users?
New business	Can create just-in-time networks for short-term projects Can give worldwide sites access much sooner than waiting for leased lines	Interoperability of different VPN equipment Management of mixed equipment environment is not easy

2.3.1 Layer 2 Protocols

Layer 2 protocols enable the transfer of data from a remote client to the private network of an enterprise by creating a virtual private network most often across a TCP/IP-based data network. Layer 2 protocols support on-demand, multi-protocol virtual private networking over public networks, such as the Internet. Internet access is provided by an internet service provider (ISP), who wishes to offer services other than traditional registered IP address-based service to dial-up users of the network.

This architecture is transparent to the end systems. In case of connecting two distant local area networks (LANs) through a VPN, the users will notice no difference while their traffic is being encapsulated in IP packets and transmitted to the remote VPN access server, which puts them back to the remote LAN. If a remote user wants to connect to the private network of the enterprise through a VPN connection, his/her computer has to support the implemented VPN protocol to be able to encapsulate the traffic. Although this encapsulation provides some security against intercepting the actual data, additional encryption should be implemented to provide secure communication.

2.3.1.1 Point-to-Point Tunneling Protocol (PPTP)

The PPTP networking technology is defined as an extension to the remote access Point-to-Point Protocol (RFC1171). PPTP is a network protocol that encapsulates PPP packets into IP datagrams for transmission over the Internet or other public TCP/IP-based networks.

After the client has made the initial PPP connection to the ISP, a second dial-up networking call is made over the existing PPP connection. Data sent using this second connection is in the form of IP datagrams that contain PPP packets, referred to as *encapsulated PPP packets*.

The second call creates the VPN connection to a PPTP server on the private enterprise LAN, referred to as a *tunnel*. This is shown in [Figure 2.3.1](#).

The secure communication using the PPTP protocol typically involves three processes, each of which requires successful completion of the previous process:

PPP Connection and Communication — The PPTP client uses PPP to connect to an ISP by using a standard phone line or ISDN line. This connection uses the PPP protocol to establish the connection and encrypt data packets.

PPTP Control Connection — Using the connection to the Internet established by the PPP protocol, the PPTP protocol creates a control connection from the PPTP client to the PPTP server over the Internet. This connection uses TCP to establish the connection and is called a PPTP tunnel.

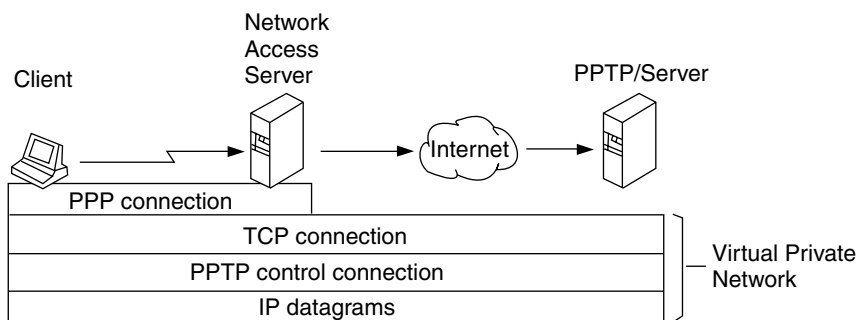


FIGURE 2.3.1 The PPTP tunnel.

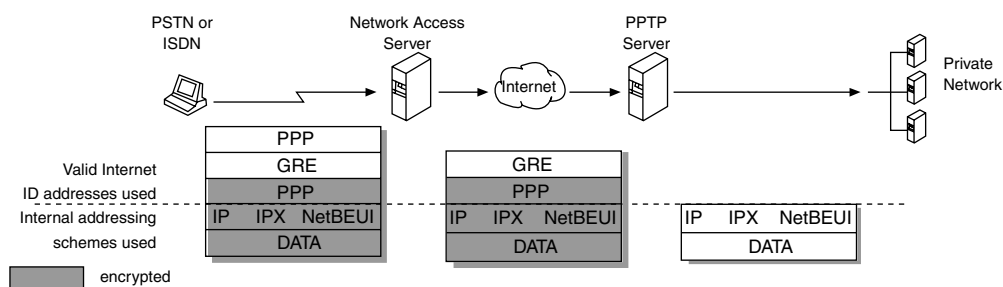


FIGURE 2.3.2 Connecting a dial-up networking PPTP client to the private network.

PPTP Data Tunneling — Finally, the PPTP protocol creates IP datagrams containing encrypted PPP packets, which are sent through the PPTP tunnel to the PPTP server. The PPTP server disassembles the IP datagrams and decrypts the PPP packets, and then routes the decrypted packets to the private network.

Note that the encapsulated PPP packet can contain multi-protocol data such as TCP/IP, IPX, or NetBEUI protocols. Because the PPTP server is configured to communicate across the private network by using private network protocols, it is able to read multi-protocol packets.

Figure 2.3.2 illustrates the multi-protocol support built into PPTP. A packet sent from the PPTP client to the PPTP server passes through the PPTP tunnel to a destination computer on the private network.

PPTP encapsulates the encrypted and compressed PPP packets into IP datagrams for transmission over the Internet. These IP datagrams are routed over the Internet until they reach the PPTP server that is connected to the Internet and the private network. The PPTP server disassembles the IP datagram into a PPP packet and then decrypts the PPP packet using the network protocol of the private network. As mentioned earlier, the network protocols on the private network that are supported by PPTP are IPX, NetBEUI, or TCP/IP.

An ISP's network access server may require initial dial-in authentication. If this authentication is required, it is strictly to log on to the ISP network; it is not related to the PPTP server authentication.

There are different options to provide data encryption between the PPTP client and the server. Microsoft uses the RAS “shared secret” encryption process. It is referred as a shared secret, because both ends of the connection share the encryption key. Under Microsoft's implementation of RAS, the shared secret is the user password. Other encryption methods base the encryption on some key available in public; this method is known as *public key encryption*. Microsoft's PPTP uses the PPP encryption and PPP compression schemes called Microsoft Point-to-Point Encryption (MPPE). The Compression Control Protocol (CCP) used by PPP is used to negotiate encryption. The encryption key is derived from the hashed password stored at both the client and the server. The RSA RC4 standard is used to create the 40-bit session key based on the client password. This key is used to encrypt all data that is passed over the Internet, keeping the connection private and secure.

PPTP is aimed primarily at Internet-based remote access. The main advantages are multi-protocol support and simplicity, because it functions on the Layer 2 level. This can be a preferred solution in a multi-protocol environment, but data security is a concern. The proposed standard does not provide a data encryption solution, although Microsoft has a vendor-specific solution as discussed earlier. The other difference is, compared to Layer 3 solutions, PPTP only provides a single point-to-point connection. It means that there can be no simultaneous Internet access while using a VPN connection. With multi-point tunneling, such as a Layer 3 solution discussed later, a user could have an Internet session at the same time as several VPN connections. This is also an inherent consequence from the PPTP architecture being a client–server model-based solution, while Layer 3 solutions are based on a more general host-to-host model.

2.3.1.2 Layer 2 Forwarding (L2F)

L2F achieves private network access through a public system by building a secure “tunnel” across the public infrastructure that connects directly to a user’s home gateway. Multiple corporate networks can use a single local telephone number terminated on a service provider’s dialup switch or access server. The access server establishes identity, sets up a private tunnel to the user’s home gateway router, and tunnels clients to that gateway. The gateway is responsible for authentication of the remote user, thereby ensuring client control of access security and addressing.

A key component of the virtual dialup service is tunneling, a vehicle for encapsulating packets inside a protocol that is understood at the entry and exit points of a given network. These entry and exit points are defined as a tunnel interfaces. The tunnel interface itself is similar to a hardware interface, but is configured in software.

Figure 2.3.3 shows the format in which a packet would traverse the network within a tunnel.

Tunneling involves the following three types of protocols:

- The passenger protocol is the protocol being encapsulated; in a dialup scenario, this protocol could be PPP, SLIP, or text dialog
- The encapsulating protocol is used to create, maintain, and tear down the tunnel. Cisco supports several encapsulating protocols including the L2F protocol, which is used for virtual dialup services
- The carrier protocol is used to carry the encapsulated protocol; IP will be the first carrier protocol used by the L2F protocol, because of IP’s robust routing capabilities, ubiquitous support across different medias, and deployment within the Internet

No dependency exists between the L2F protocol and IP. In subsequent releases of the L2F functionality, Frame Relay, X.25 VCs, and asynchronous transfer mode (ATM) switched virtual circuits (SVCs) could be used as a direct Layer 2 carrier protocol for the tunnel.

Cisco’s L2F implementation provides several management features. End system transparency ensures that neither the remote end system nor its corporate hosts should require any special software to use this service. Authentication is provided by dialup PPP, Challenge Handshake Authentication Protocol (CHAP), or Password Authentication Protocol (PAP), including Terminal Access Controller Access Control System Plus (TACACS+) and Remote Authentication Dial-In User Service (RADIUS) solutions, as well as support for smart cards and one-time passwords; the authentication will be manageable by the user independent of the ISP. Addressing will be as manageable as dedicated dialup solutions; the address will be assigned by the remote user’s respective corporation, and not the ISP. Authorization will be

IP/UDP	L2F	PPP (Data)
Carrier Protocol	Encapsulator Protocol	Passenger Protocol

FIGURE 2.3.3 Tunneling packet format.




FIGURE 2.3.4 Virtual dialup topology.

managed by the corporation's remote users, as it would be in a direct dialup solution. Accounting will be performed both by the ISP (for billing purposes) and by the user (for charge back and auditing).

Figure 2.3.4 illustrates the topology of a virtual private connection using Cisco's L2F.

In a traditional dialup scenario, the ISP using the NAS in conjunction with a security server follows an authentication process by challenging the remote user for both the username and password. If the remote user passes this phase, the authorization phase can begin.

For the virtual dialup service, the ISP pursues authentication to the extent required to discover the users' apparent identity (and by implication, their desired corporate gateway). No password interaction is performed at this point. As soon as the corporate gateway is determined, a connection is initiated with the authentication information gathered by the ISP. The corporate gateway completes the authentication by either accepting or rejecting the connection. (For example, the connection is rejected in a PAP request in which the username or password are found to be incorrect.) Once the connection is accepted, the corporate gateway can pursue another phase of authentication at the PPP layer. These additional authentication activities are outside the scope of the specification, but might include proprietary PPP extensions, or textual challenges carried within a TCP/IP Telnet session.

For each L2F tunnel established, L2F tunnel security generates a unique random key to resist spoofing attacks. Within the L2F tunnel, each multiplexed session maintains a sequence number to prevent the duplication of packets.

Cisco provides the flexibility of allowing users to implement compression at the client end. In addition, encryption on the tunnel can be done using IPSEC.

There are similar advantages of L2F as of PPTP, because the two solutions are similar, and being merged to a common standard called L2TP. L2F also has a nice advantage of connecting multi-protocol networks, because of its Layer 2 functionality. But again there is no support in the proposed standard for VPN data encryption. Cisco refers to the IPsec standard as a possible encryption method for IP traffic carried with L2F. For non-IP traffic L2F lacks the solution for security. In comparison with Layer 3 solutions, L2F provides only a single point-to-point connection, which makes parallel Internet access impossible while being connected to the VPN.

2.3.1.3 L2TP

The IETF draft for PPTP titled as "Point-to-Point Tunneling Protocol," draft-ietf-pptp-00.txt was submitted to the Internet Engineering Task Force (IETF) in June 1996 by the companies of the PPTP Forum, which includes Microsoft Corporation, Ascend Communications, 3Com/Primary Access, ECI Telematics and US Robotics. Cisco's proposal for L2F was submitted to the IETF for approval as a proposed standard. Northern Telecom, Inc., and Shiva Corporation have announced their support for L2F. At the June 1996 IETF meeting in Montreal, the IETF PPP Extensions working group agreed to combine Cisco's proposal with PPTP proposed by Microsoft Corporation. The emerging proposed standard, Layer 2 Tunneling Protocol (L2TP) is currently drafted by Cisco Systems, Microsoft, Ascend, 3Com, and US Robotics. The latest L2TP draft (09) was submitted in January 1998.

A typical connection scenario would start with the remote user initiating a PPP connection to an ISP via either the PSTN or ISDN. The ISP's access point (LAC) accepts the connection and the PPP link is established. The ISP may now undertake a partial authentication of the end system/user. Only the username field would be interpreted to determine whether the user requires a Virtual dial-up service. It

is expected — but not required — that usernames will be structured (e.g., username@company.com). Alternatively, the ISP may maintain a database mapping users to services. In the case of Virtual dial-up, the mapping will name a specific endpoint, the L2TP Network Server (LNS). If a virtual dial-up service is not required, standard access to the Internet may be provided.

If no tunnel connection currently exists to the desired LNS, one is initiated. L2TP is designed to be largely insulated from the details of the media over which the tunnel is established; L2TP requires only that the tunnel media provide packet-oriented point-to-point connectivity. Obvious examples of such media are UDP, Frame Relay PVCs, or X.25 VCs. Once the tunnel exists, an unused slot within the tunnel, a “Call ID,” is allocated, and a connect indication is sent to notify the LNS of this new dial-up session. The LNS either accepts the connection, or rejects it. The initial connect notification may include the authentication information required to allow the LNS to authenticate the user and decide to accept or decline the connection. In the case of CHAP, the set-up packet includes the challenge, username, and raw response. For PAP or text dialog, it includes username and clear text password. The LNS may choose to use this information to complete its authentication, avoiding an additional cycle of authentication.

If the LNS accepts the connection, it creates a “virtual interface” for PPP in a manner analogous to what it would use for a direct-dialed connection. With this virtual interface in place, link layer frames may now pass over this tunnel in both directions. Frames from the remote user are received at the POP, stripped of CRC, link framing, and transparency bytes, encapsulated in L2TP, and forwarded over the appropriate tunnel. The LNS accepts these frames, strips L2TP, and processes them as normal incoming frames for the appropriate interface and protocol. The virtual interface behaves very much like a hardware interface, with the exception that the hardware in this case is physically located at the ISP POP. The other direction behaves analogously, with the LNS encapsulating the packet in L2TP, and the LAC stripping L2TP before transmitting it out via the physical interface to the remote user.

At this point, the connectivity is a point-to-point PPP session whose endpoints are the remote user’s networking application on one end and the termination of this connectivity into the LNS’s PPP support on the other. Because the remote user has become simply another dial-up client of the LNS, client connectivity can now be managed using traditional mechanisms with respect to further authorization, protocol access, and packet filtering.

Accounting can be performed at both the L2TP Access Concentrator (LAC) as well as the LNS. This document illustrates some accounting techniques that are possible using L2TP, but the policies surrounding such accounting are outside the scope of this specification.

For the virtual dial-up service, the ISP pursues authentication only to the extent required to discover the user’s apparent identity (and by implication, their desired LNS). This may involve no more than detecting DNS information when a call arrives, or may involve full LCP negotiation and initiation of PPP authentication. As soon as the apparent identity is determined, a call request to the LNS is initiated with any authentication information gathered by the ISP. The LNS completes the authentication by either accepting the call, or rejecting it. The LNS may need to protect against attempts by third parties to establish tunnels to the LNS. Tunnel establishment can include authentication to protect against such attacks.

L2TP, like other Layer 2 protocols, does not provide any further security for data encryption, but rather refers to Layer 3 encryption techniques for IP traffic, such as IPSec.

Although L2TP seems to be the result of different Layer 2 Tunneling initiatives, it still does not have the widest acceptance in the industry. Its predecessor PPTP is popular, because of the large number of Windows NT users, but IPSec, a general initiative to add security to the IP protocol, has the strongest support by manufacturers and suppliers. In a multi-protocol environment there will still be a need for Layer 2 Tunneling with the addition of encryption for IP traffic (using IPSec). But in an IP-only environment Layer 3 solutions are more effective. Unless augmented with IPSec these Layer 2 solutions cannot support extranets, because extranets require keys and key management.

In comparison with other Layer 2 protocols, L2TP has better features, such as an addition to PPTP and L2F, and the support for ATM or SONET as an underlying transmission medium, which is only planned for the other two protocols.

2.3.2 Layer 3 Tunneling Protocols

In the previously discussed architecture the PPP connection begins at the remote client and terminates at the corporate network's L2TP server, going through the ISP access point and the corporate gateway.

Layer 3 tunneling proposes a different scenario initiating the PPP connection from the remote client, but terminating it at the ISP. This requires the re-encapsulation of the PPP frame and transmitting the Layer 3 information only to the corporate access router. This access router does not need to support any additional standard, but it acts only as a simple router. The difference from traditional dial-up services is that the ISP's IP Gateway will provide the IP address to the client. It can roam with this address as long as the IP Gateway does the reframing of the IP PPP packet and sends it to the corporate gateway as a standard IP packet.

The advantage of the latter scheme is that there is no need to support L2TP at either the remote client end or at the corporate gateway end. The remote client only needs a standard IP stack, and the corporate gateway acts as a simple IP router. In this case if the remote node is a router connecting a sub-network to the corporate network, the packets can be routed just like any other traffic through the ISP network. In either case an additional functionality is needed to provide security on the IP Layer (network layer). Although L2TP can encrypt the PPP packets on Layer 2, it does not claim to be secure against denial of service attacks or man-in-the-middle attacks (someone modifying the PPP frames in the tunnel).

2.3.2.1 IPSec

IPSec is a protocol suite defined by the IETF working group on IP security to secure communication at the Layer 3 (network layer) between communicating peers. The goal of the IPSec protocol suite is to provide secure tunneled transport of IP data only. Essentially, it takes private IP packets, performs data security functions such as encryption, authentication, and integrity, then wraps these secured packets in other IP packets for transport across the Net. Key management functions also will be a part of the IPSec protocol suite. The IETF has issued five requests for comments — RFC 1825 through 1829. An interesting note is that if IPv6 succeeds in replacing IPv4, IPSec will be the automatic Internet VPN standard since it is integrated into the IPv6 specifications.

Like the Layer 2 VPN protocols, IPSec works as a LAN-to-LAN and dialup-to-LAN solution. It is designed to support multiple encryption protocols, a feature that allows users to choose a desired amount of data privacy. Obviously, IPSec will only be of value to companies that want to tunnel IP exclusively since it doesn't support other data protocols.

There are several different scenarios where IPSec can be used. In case two hosts A and B want to communicate with each other through a firewall, the host A can tunnel packets to the firewall, the firewall can decrypt/authenticate the packets, and send them to B based on its rules. In a different setup there can be a secure tunnel between host A and host B, where the firewall is authorized to act as a key management proxy, and has the capability to decrypt the packets and apply its packet filtering policy. A third setup is a combination of the previous two, where the inner payload is secured from host A to B, and the outer payload is secured and tunneled through the firewall. The advantage of this scheme is that the firewall is able to authenticate packets and decide whether to allow the packet without applying its filtering rules. This is typical of what happens today, where an employee gets into the network via dialup PPP.

There is a different scheme, where packets have to be secured while travelling the Internet. In this case IPSec will secure packets between two or more border routers of a topologically distributed organization. In this case since security associations are set up between the border routers, any traffic should go through these routers. All packets between the two routers must contain valid IPSec, otherwise they will be dropped.

A growing number of VPN, security, and major network companies either support or plan to support IPSec. It is also strongly supported by a user group consisting of manufacturers and suppliers. Although it deals with IP-only traffic, it is the most often recommended or chosen solution to ensure privacy in VPN communication. It can be implemented as a single Layer 3 solution, but it can be implemented over a Layer 2 solution to provide data encryption for IP traffic. It is capable of maintaining multiple

tunnels, including simultaneous VPN and public access connection inherently from its general host-to-host model. It can also support extranets with its built-in key management functionality, which is missing in other Layer 2 solutions.

2.3.2.2 Mobile IP

Mobile IP is intended to enable nodes to move from one IP subnet to another. It is just as suitable for mobility across homogeneous media as it is for mobility across heterogeneous media. That is, Mobile IP facilitates node movement from one Ethernet segment to another as well as it accommodates node movement from an Ethernet segment to a wireless LAN, as long as the mobile node's IP address remains the same after such a movement.

Mobile IP introduces the following new functional entities:

Mobile node — A host or router that changes its point of attachment from one network or subnetwork to another. A mobile node may change its location without changing its IP address; it may continue to communicate with other Internet nodes at any location using its (constant) IP address, assuming link-layer connectivity to a point of attachment is available.

Home agent — A router on a mobile node's home network which tunnels datagrams for delivery to the mobile node when it is away from home, and maintains current location information for the mobile node.

Foreign agent — A router on a mobile node's visited network which provides routing services to the mobile node while registered. The foreign agent detunnels and delivers datagrams to the mobile node that were tunneled by the mobile node's home agent. For datagrams sent by a mobile node, the foreign agent may serve as a default router for registered mobile nodes.

A mobile node is given a long-term IP address on a home network. This home address is administered in the same way as a permanent IP address is provided to a stationary host. When away from its home network, a care-of address is associated with the mobile node and reflects the mobile node's current point of attachment. In this case Mobile IP uses protocol tunneling to hide a mobile node's home address from intervening routers between its home network and its current location. The tunnel terminates at the mobile node's care-of address. The care-of address must be an address to which datagrams can be delivered via conventional IP routing. At the care-of address, the original datagram is removed from the tunnel and delivered to the mobile node.

The Mobile IP standard contains a very strong authentication between the mobile node and the home agent to authenticate themselves. The default algorithm is keyed MD5, with a key size of 128 bits. This will result in the mobile node's traffic being tunneled to its care-of address. But the standard does not provide privacy protection; it rather refers to other IP encryption standards, such as IPSec.

The Mobile IP standard as with other Layer 3 standards has the advantage of scalability, security, and reliability, but they are more complex to develop and, inherent from their Layer 3 functionality, they only support a specific protocol, which is IP in this case.

2.3.3 Frame Relay

Traditionally, VPNs were provided in the form of broadband packet switched services, such as Frame Relay, X.25, or ATM. Now, with growth of the Internet as a viable service infrastructure, it is possible to run VPNs over an alternative protocol. As it can be seen from the previously discussed standards, these latter solutions tend to be more popular. The reason is fairly simple: while traditional VPNs are very useful for fixed LAN-to-LAN connectivity, they do not easily accommodate individual users whose only access to the outside world is in the form of their PC, a modem, and the public switched telephone network. VPNs that run over IP are easily accessed by these users.

With any choice of the above-mentioned protocols, virtual circuit connections can be defined between remote locations, and the LAN traffic can be bridged over these circuits. This usually provides an emulated Layer 2 network for the users, which can be used to transmit multiprotocol traffic. Although proper authentication and encryption also has to be taken care of in these solutions, these networks are not as

sensitive to security threats, as they are usually a private ATM, Frame Relay backbone of the provider, but in any case less public than the Internet in the previous solutions.

The advantages of this solution are the built-in quality of service (QoS) guarantees that are part of the virtual circuit definitions. Where the bandwidth availability could be argued in the past for an Internet-based VPN solution compared to these solutions, nowadays the ISPs can also provide high-speed Internet connections, especially when they utilize only their backbone network to provide VPN services. The disadvantage of the Frame Relay, ATM, or X.25-based VPN services is that these means of access should be available at each location where the VPN service has to be used. It usually means special equipment, wiring, and additional management needs. The user mobility is not solved with these solutions.

2.3.4 Layer 2 or Layer 3 Comparison

The goal of Layer 2 tunneling protocols is to transport Layer 3 protocols such as AppleTalk, IP, and IPX across the Internet. To achieve this, the architects of PPTP and L2F leveraged the existing Layer 2 PPP standard, which is designed to transport different Layer 3 protocols across serial links. In these schemes, Layer 3 packets are encased in PPP frames, which are then encased in IP packets for transport across the Internet.

From a security standpoint, Layer 2 tunneling protocols are insufficient to be secure VPN solutions on their own. None of these protocols provide the data encryption, authentication, or integrity functions that are critical to maintaining VPN privacy. The L2TP specification disclaims any data security functions and refers IP data security to IPSec, but no serious security provisions or references are made for the other Layer 2 protocols. In addition, none of these protocols provide a mechanism for key management, which limits their scalability.

PPTP and L2F are vendor-specific, proprietary protocols, so interoperability is limited to products from supporting vendors. In contrast, L2TP is a multivendor effort, so interoperability is not as much of a problem. It is important to note that when utilizing tunneling protocols besides IP, users will have to rely on vendor-specific data security features. On the upside, PPTP, L2F, and L2TP can transport multiple protocols. They also function both in LAN-to-LAN and dial-up-to-LAN tunneling modes, allowing them to cover the applications most desired for VPN.

In case of Layer 3 Tunneling there is no need for a globally unique address space, which is a requirement with L2TP for remote client address assignments. This global address space doesn't have to be registered, since it is seen from the corporate network but not visible from the public network (Internet).

Another difference is that the tunneling causes an additional overhead as opposed to Layer 3 Tunneling, which only sends regular IP packets after the ISP's IP Gateway to the corporate network. (The tunneling takes place only between the remote client and the IP Gateway.) This might make Layer 3 solutions more scalable, but with the loss of the additional features, such as the freedom of network protocols that can be used over the PPP layer.

References

Bay Networks: <http://www.baynetworks.com/Solutions/vpn/>

Cisco: <http://www.cisco.com/warp/public/779/servpro/solutions/vpn/>

Microsoft: <http://www.microsoft.com/ntserver/nts/commserv/exec/feature/VPNFeatures.asp>

Shiva: <http://www.shiva.com/remote/vpn.html>

3Com: <http://www.3com.com/enterprise/vpn/>

2.4 Effective Website Design

Karen M. Freundlich

The goal of an effective website should be to achieve the desired results using the best available technology for the job. Many webmasters fall prey to the seduction of fitting the goals of their site to the sophisticated

tools now freely available. A measurable objective should be narrowly defined before the site is even designed. As in any good novel, the plot must be carefully laid out before the writing and editing take place. This segment will review four steps to consider when creating a website to provide a business solution.

In today's world of "what you see is what you get" website creation tools it is easy to construct sites similarly to what we did as children, creating anything out of the available blocks or Tinkertoys in the bag. Instead, really effective websites, especially ones focused on financially rewarding e-commerce, must avoid the temptation to "*build*." Measurable goals must be defined, the production and maintenance efforts determined, a system for measuring effectiveness put in place, and the intuitive layout tested and determined before the "building" tools are contracted. This methodical process will more readily guarantee the site's success.

2.4.1 Goals Defined

The goals of the business behind the site must be strictly defined. The website should be regarded as a tool to run the business; a store, not just a storefront. E-commerce software is often accompanied by e-commerce consulting for just this reason. It is simply not enough to have a clean, attractive, well-filled store; the products must be suited to the market and carefully distributed according to supply and demand. It is the same in the virtual world of the internet. There is a destructive misconception that everything on the internet occurs in shortened time, for example, that a "web year" is only four months. Like any computer, the web is run by humans, and just like the chain is as strong as the weakest link, the speed of the web is only as fast as the humans that run it. Admittedly, the web can reduce communication times among the parties, but the thinking and planning process to design the systems takes the same time on the part of the human. Trying to rush this is guaranteed to cause errors, especially in business decisions. The web cannot make informed, intelligent business decisions more quickly. You must engage the best minds to provide the best solutions. Only when the goals are clearly defined to achieve the business solutions should one proceed to the smaller details of implementing them.

2.4.2 Production and Maintenance Efforts

Once the goals are defined, the next step is to determine what levels of production and maintenance are needed to meet the goals. Since, presumably, the website is blazing a trail, it is prudent to use the least complex solutions to meet your goals, then provide a term for evaluating the results and to better define the needs, and finally upgrade according to your growth. It is useful to review and critique the approaches used by similar or parallel industries. The web makes competitor research very accessible. The site should be produced with maintenance objectives easily implemented. For instance, an administrator page can be included which allows password-protected users to perform a variety of maintenance tasks on-line. The inclusion of the maintenance goal in the site design will assure that a comprehensive and complete job is achieved. Quality control should be an important objective, best planned for in advance.

2.4.3 A System for Measuring Effectiveness

In order to determine if the site is meeting its goals, there must be inherent systems in place to measure the results. It takes much less effort to build your site with systems to track measurable objectives than it is to try to figure out how and what to measure once the site and systems are designed. The types of measures will vary based on the goals, but one way to approach the problem is to design the reports that will communicate the results. Once the report's answers to questions are put in writing, it is much easier to determine what pieces of information will be needed. The pieces will usually reside in two places. First, they will consist of direct responses from the users, such as orders and return mail. This type of information is designed during the production of the site and should be carefully constructed to provide measurement. Second, there will be data stored on the server's log report of all activity on the site. Know

in advance what these data are because they can provide important demographic and navigational information that can be strategically used to provide measurement of the site's objectives. If the site's hyperlinks are carefully designed, much information about the user's thought process while navigating the site can be gleaned from the log report's data regarding the order in which the pages were requested from the server. A much more efficient evaluation of the site's effectiveness can be achieved when measurable objectives are included as part of the site design.

2.4.4 Intuitive Layout

Finally, no web browser likes to be “lost in the funhouse.” Chances are, there was enough work expended just to find your site via a search engine, word of mouth, or reputation. Once there, the user wants to find solutions with speed and comfort. If they came to browse, the information must be presented and found quickly. Everyone knows how easy it is to get lost, frustrated, and eventually one leaves without much thought of returning. Intuitive, concise, and neatly laid out site design will make the user feel comfortable there. Remember, unlike traditional stores, it is much easier to “walk out” of a website. And even though good help is hard to find in those traditional stores, web users don't expect that they will even need help. The design should be uncluttered, the directions and navigation should be intuitive, and there should always be a link that will allow the user to return to the home page. Remember that the original goal behind the hyperlink was to allow non-linear access to information. All of your website's offerings should not be shown on one page. A carefully designed linking system will allow the user to access information in an intuitive way, without having to muddle through the clutter. Overuse of graphics and animation should be avoided. Remember to use the best available tools to meet the objective without burdening the user or the computer resources with extraneous efforts.

2.5 Web-enabled Data Warehousing

Dermot Murray

2.5.1 Introduction

Since the early 1980s, business analysts have identified the inherent value in analyzing the huge amounts of data generated in production online transaction processing (OLTP) systems. Hidden deep inside such data repositories lies key information that can make a product or service more marketable, that can make a customer more profitable, and that can make processes more efficient. This process of analyzing production data in order to unearth that critical business intelligence is known as decision support systems (DSS). The problem has always been getting at that information in such a way that it adds value to the business as a whole. William H. Inmon promoted the concept of data warehousing³ in the 1980s as a means of separating operational systems from DSS in order to extract the data necessary for business intelligence without impacting mission critical processing systems. Since then, there has been a huge growth in the market for data warehousing and decision support tools, with all of the major database vendors such as Oracle, IBM, and Sybase developing products to satisfy the demand.

This paper examines the next logical step in the evolution of data warehousing technology; i.e., Web-enabling the applications that provide access to this key business data. With the growth in the use of intranets, extranets, and the Internet in general, such Web-enabled data warehousing products have revolutionized the way business analysts generate the reports and charts they need in order to analyze the trends and patterns in the operational data. We will explore the concepts behind Web-enabled data warehousing and look at the technology that makes it all happen.

³*The Essential Client/Server Survival Guide* — Orfali, Harkey, and Edwards, (John Wiley, 1996).

2.5.2 Data Warehousing Overview

2.5.2.1 Concepts

The concept behind data warehousing first emerged in the client/server platform environment of the 1980s. Although Inmon first started writing about data warehousing in 1981, it wasn't until the early 1990s that industry giants such as IBM, Oracle, and Sybase started taking the technology seriously. Today, the worldwide market for data warehousing products is estimated at over \$30 Billion.⁴ What drives the popularity of data warehousing is the need for executives and business analysts to gain competitive advantage from analyzing trends and patterns hidden within the mountains of data that companies generate in their day-to-day transactions. Data warehousing is the process of extracting data from various OLTP systems into a centralized format that can be analyzed using DSS and executive information systems (EIS) tools, which provide more business-specific and powerful queries for higher level managers and executives. Collectively, these tools are known as online analytical processing (OLAP) or multidimensional analysis (MDA). The data warehouse itself can either be a single or distributed specialized database management system (DBMS) that contains replicated data from different sources within the organization. This data is usually extracted from internal data production sources such as OLTP and enterprise resource planning (ERP) systems but increasingly from external sources such as Dow Jones, Reuters, and even the Internet itself. The data is typically cleansed and transformed to a format that can be analyzed by DSS tools. Information about the content and format of the data is stored as "metadata" in information directories that can be accessed by both business analysts and database administrators (DBA) alike. [Figure 2.5.1](#) demonstrates the various steps involved in data warehousing.

2.5.2.2 Advantages

The main advantage of a data warehouse is that it provides executives with up-to-date data that can be queried for reporting and analysis in order to assist them in making strategic decisions about the operations of their organizations. The SQL queries generated by OLAP tools are typically very sophisticated and can contain multiple criteria such as sales by region, state, and customer. The results of these queries are displayed as reports or charts that can then be presented in a concise form to executives who typically don't want all of the detail that is contained in the original production data. Due to their multiple-criteria nature, these queries are usually long-lived and may even result in lost or stray processes that could jeopardize the data integrity of a real-time production environment such as OLTP or ERP. Therefore, data warehouses are usually separate DBMSs that store replicated and transformed copies of the production data and are not required to provide the same fast response times that are necessary for mission-critical OLTP systems. However, the new genre of operational data stores⁵ provides faster response to DSS/EIS queries by using near real-time transactional data for the most current query and analysis. In addition, DSS/EIS tools provide drill down capabilities that allow executives to get more detailed information on a particular trend or subject matter by generating more detailed queries on that subject. Data warehouses also support data mining tools that provide analysts with the ability to discover unexpected patterns in data by using fuzzy logic searches.

2.5.2.3 Data Marts

A very popular form of data warehousing is the data mart, which is typically a subset of the data warehouse that contains data of specific interest to a particular department, such as marketing, sales, or human resources. These data marts are less expensive to build and maintain than enterprise-wide data warehouses and offer immediate business value to the departments that they service. However, the very independence that data marts provide tends to act as an obstacle to true enterprise integration in information reporting and analysis. While the debate between centralized enterprise-wide data warehousing and decentralized data marts has consumed industry advocates for the last ten years, a compromise concept of dependent

⁴*Database Solutions White Paper* — Palo Alto Management Group, Inc., July 1998.

⁵*Data Warehousing Management/ Productivity Tools* — Datamation White Paper, 1997.

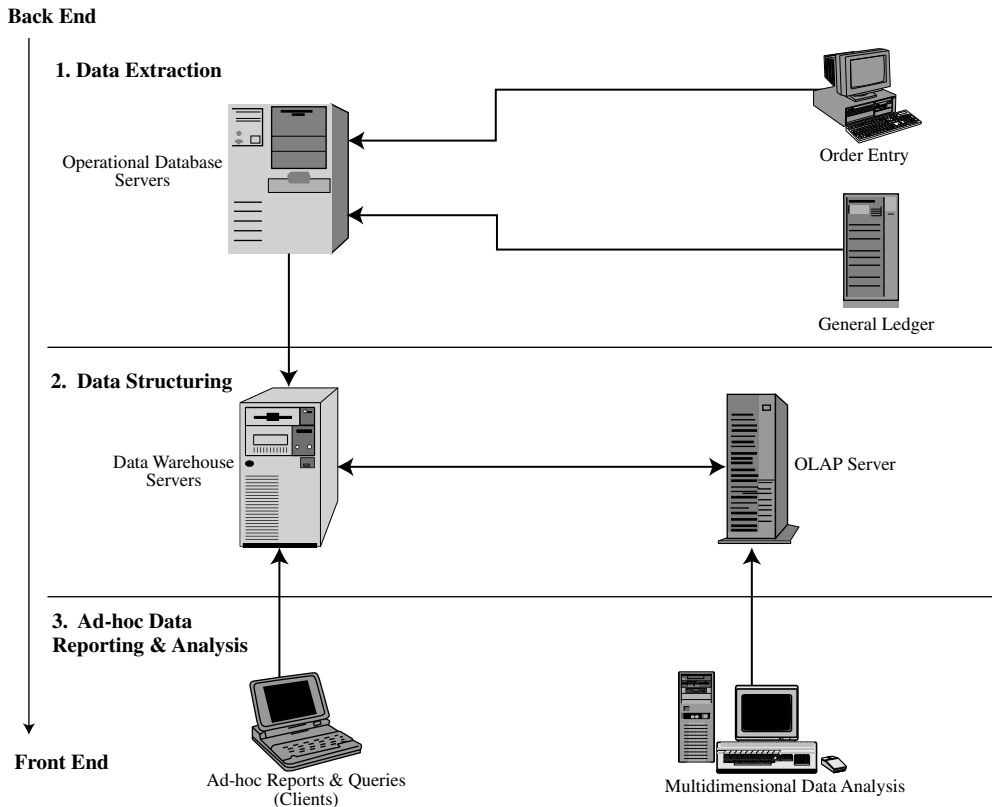


FIGURE 2.5.1 Data warehouse/OLAP environment.

data marts has potentially bridged the gap. This hybrid solution provides for the extraction of data from the centralized data warehouse to the departmental data marts so that each department is working off the same enterprise data, albeit formatted to their individual needs.⁶

2.5.2.4 Future Growth

In spite of the technical and organizational challenges that data warehouse implementations pose, the ultimate benefits that the technology offers has led to a huge growth in its adoption. For instance, 90% of the Global 2000 companies had either implemented or planned to implement data warehouses by June 1998.⁷ The future of data warehousing also seems pretty secure with market estimates expecting to grow exponentially to over \$100 billion by 2002 (Figure 2.5.2).⁸ One of the biggest factors in the expected growth of data warehouse implementation is the ease of access to timely reports and queries that Web-based data warehouse tools provide.

2.5.3 Web-enabled Data Warehousing

2.5.3.1 Benefits

Perhaps the most perplexing aspects of implementing a data warehouse strategy have been the escalating costs of installing and maintaining the GUI-based clients that are used to generate the queries. According

⁶The Middle Ground — *CIO Magazine*, January 1999.

⁷Data Warehousing Is Worth The Investment — *InternetWeek*, June 1998.

⁸Database Solutions White Paper — Palo Alto Management Group, Inc., July 1998.




FIGURE 2.5.2 Worldwide data warehousing solutions market (Source: Palo Alto Management Group).

to a Meta Group, Inc. study in 1997, the average cost of implementing a data warehouse project was \$1.9 million,⁹ which accounted for up to 19% of the total IT budgets of those companies surveyed. One solution to minimizing the costs associated with data warehouse implementations is Web-enabled data warehousing, a concept that has grown in popularity since it first appeared in the mid 1990s. These tools have given IT departments a more cost-effective option of allowing access to the enterprise-wide data warehouses for a larger group of users. By allowing users to generate queries and reports through their Web browsers, IT departments can roll out data warehouse installations in a much shorter period of time by simply allowing the users to have the appropriate level of access to the DSS application server. With the advent of virtual private networking (VPN) and subsequent extranet technology, various levels of data warehouse access can even be provided to external users such as suppliers and customers. This is particularly critical in today's supply chain environment where companies have to work closely with their strategic partners, i.e., suppliers and customers, and therefore have to provide a certain level of access to their enterprise data warehouses for reporting and analysis. This Web-enabled approach has had the effect of reducing the costs associated with hardware and software by using the "thin" client approach, as opposed to the "fat" client requirements of the previous client/server model. For instance, the Aberdeen Group was able to cut costs from \$1000 per seat for client/server DSS tools to \$50 for Web-based DSS access, primarily by consolidating most of the software and hardware costs at the server.¹⁰ This server-centric model also has the effect of reducing software licensing costs by switching from a per-seat basis to a more cost-effective server-based licensing model. The move away from dependence on fat desktop clients to the thin browser-based client model may also allow the new breed of Internet appliances, such as mobile phones and network computers, to be able to access information from data warehouses over the Internet. For instance, the advent of "push" technology could allow a sales executive to be paged if sales figures for a certain region fall below a predefined threshold.

2.5.3.2 Making it Happen

Web-enabled data warehousing involves a combination of HTML, XML, HTTP, and mobile component-based technology, typically Java or ActiveX. In this model, a user can generate a SQL query using a HTML/XML form that embeds the control information for the search criteria into a HTML/XML query and sends it to the remote Web server that in turn passes the request to a Web Gateway server. The Gateway server converts the HTML/XML query into an OLAP-specific request and passes it to the OLAP

⁹The Middle Ground — *CIO Magazine*, January 1999.

¹⁰Warehousing And The 'Net — Marriage Made In Heaven — *Insurance & Technology*, July 1997.




FIGURE 2.5.3 Web-based client/server architecture.

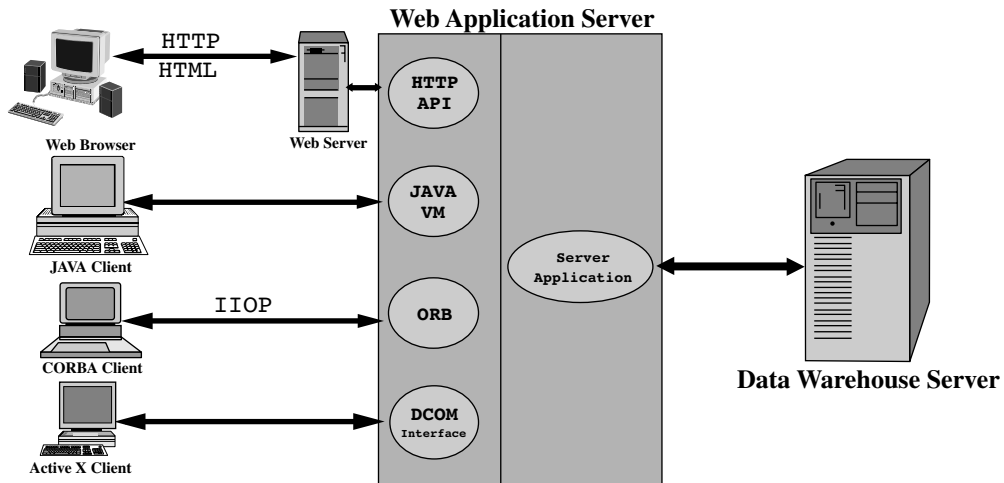


FIGURE 2.5.4 Object-based architecture (Source: Adapted from *DB2 Magazine*).

server, which executes the query directly against the data warehouse. Typically the Web gateway and OLAP Server are bundled into the vendor-supplied DSS Server application package. When that application retrieves the resulting data, it sends the data embedded in an HTML/XML file back to the Web server which forwards it back to the Web client.¹¹ Some DSS servers also store predefined DSS objects such as reports and queries that can be executed by the user to generate more standardized reports. This type of interaction is similar to three-tier client/server architecture and in fact most industry analysts would argue that Web-enabling traditional client/server applications are just the next step in the evolution of client/server — the intergalactic model (see [Figure 2.5.3](#)).

The interaction between the DSS server application and data warehouse is provided using distributed-object-computing protocols, such as CORBA, IIOP, or DCOM, which provide the client/server state management unlike HTTP, which is a stateless protocol.¹² In some cases, the DSS server downloads client-based applications, such as Java applets or ActiveX controls, that actually run within the Web browser's memory space and interact directly with the application server. The purpose of this is to provide an improved graphical user interface to the end-user in situations where plain old HTML is not sufficient for multi-dimensional data visualization, such as graphs and charts. These applets also provide more robust object-based communications between client and application server by using CORBA Internet Inter-ORB Protocol ("IIOP") and Java remote method invocation ("RMI"), with HTTP only being used for downloading the applets to the client. [Figure 2.5.4](#) demonstrates the different levels of interaction

¹¹*A Methodology For Client/Server And Web Application Development* — Roger Fournier (Yourdon Press, 1998).

¹²*Building Web Information Systems* — *Byte Magazine*, July 1998.

between the client, application server, and data warehouse server using both the Web and object-based protocols such as CORBA and DCOM.

The advent of Web-enabled data warehousing has spawned the growth of information delivery, whereby specific business information is “pushed” to information consumers at predefined intervals. In effect, consumers subscribe to information by searching the metadata stored in the information directories using Java/ActiveX agents or Web-based search engines in order to locate the information that is of interest to them.¹³ The information delivery itself can either be schedule-driven, meaning that the appropriate decision support objects are executed at user-defined intervals and the resulting information delivered to a Web server, or event-driven when particular business events occur. These user-defined events, also known as triggers, could also indicate that new data has arrived from the source DBMS. “Push” technology is also being incorporated into the channel method of delivering information, whereby a business analyst can subscribe to a channel, which in turn delivers the most up-to-date information on a particular subject.¹⁴

2.5.3.3 Obstacles and Limitations

Although Web-enabled data warehousing has revolutionized the distribution of enterprise-wide business information, it has not been the panacea for all data warehouse access problems. There are limitations to what “power users” can do with Web-enabled DSS tools and in those cases, the traditional desktop approach is still being used. Such functionality as drill-down analysis and multidimensional queries are still not as effective with the Web-based tools and so users who need this functionality still require the more mature client/server tools to perform these analyses. As Web-enabled tools become more sophisticated, this problem will be rectified, but in the meantime, IS managers must deal with the reality that both Web-enabled and desktop data warehouse access must coexist for the foreseeable future.

Ironically, the ActiveX and Java components that are designed to make user interaction with the data warehouse environment possible can cause incompatibility problems if the end-user’s Web browser cannot support them. This might be the case if the browser is an older version or if the particular downloadable component only works with either Netscape or Internet Explorer. Of course, the solution to this problem is to standardize Web browsers and plug-in components, a daunting task for organizations with potentially thousands of users, but nevertheless manageable.

Security has always been a problem with enabling external users to access corporate information. As mentioned earlier, the growth in supply-chain management has dictated that companies work in tandem with their suppliers and customers in order to develop their products and services. This trend has forced companies to open up their corporate information systems, including data warehouses, to these external users. Of course, this has led to fears that such mission-critical and confidential information could be compromised in transit between users across a public network such as the Internet. Another potential security hole is the remote access required by mobile users, such as sales agents, who need to pull up reports using dial-up connections over the Internet. Improvements in corporate firewall and VPN technology, coupled with better encryption algorithms and Public Key Infrastructure, have eased these fears. However, organizations implementing Web-enabled data warehousing must constantly be vigilant not only with potential hackers from the outside but also with internal users who should only have the level of access to reports and data that they require. One potential solution to this problem that has been promoted since Web-enabled data warehousing began is to distribute the data among data marts, thereby limiting the users’ access to only the data stored in those data marts.¹⁵

Like every new technology, Web-enabled data warehousing has its downside as illustrated above. However, with vendors addressing these issues all the time, the usage of this technology has had a phenomenal growth rate and while traditional client/server warehouse access is not going away any time soon, Web-enabled DSS tools are making strong inroads into the marketplace.

¹³Building Web Information Systems — *Byte Magazine*, July 1998.

¹⁴Warehouses Webicum: Evolution Of A Species — *DB2 Magazine*, April 1998.

¹⁵Just Browsing Thanks — *CIO Magazine*, October 1996.

2.5.4 Vendors

This burgeoning sector of the industry has produced a number of leading products from both established and startup vendors. While most of the major database vendors have been very proactive in providing Web-based access to their data warehouse products, the market for packaged applications has come mostly from startup vendors like MicroStrategy. As there are too many vendors to mention in this market segment, I will discuss a sample of the products being offered by these different categories of vendors.

2.5.4.1 Major Database Vendors

Oracle Corp., arguably the biggest name in the relational database market, has its presence thanks to the Oracle Express suite, which includes the Express Server and Express Web Agent modules. This product was actually acquired from Information Resources in 1995 and accounted for 21% of the OLAP market in 1998.¹⁶ The Oracle Express Server is the actual back-end OLAP engine that performs the end-user queries. This is the component that provides the interaction between the Express OLAP Server and the Web server. This module takes advantage of the “Network Computing Architecture” model, which is Oracle’s blueprint for a three-tier thin client environment, and the Express Stored Procedure Language cartridge to provide the communication between the Web server and the Express Server. The Express Agent Developer’s Toolkit, which comes with the Agent, supports the development of both HTML pages and components such as Java and ActiveX to produce customizable reports and analyses for the end user.

For their part, IBM provides Web access to the DB2 OLAP Server through the use of its Net.Data development platform. This product supports standard SQL statements as well as C++ and Java enablers that allow developers to write macros that automate SQL queries. On the server side, it supports FastCGI, which is a high-performance Web server interface that provides better performance for Net.Data applications. In addition, IBM integrated its OLAP Server with Hyperion’s Essbase Web Gateway, which is a Web application server component similar to Oracle’s Web Agent.

As for the other major database vendors, it appears that they have been slow to fully integrate their OLAP products with Web access, including Microsoft. For their part Informix has tried to buy its way into the Web access market by acquiring Red Brick Systems, who in turn had teamed up with Web application specialists Caribou Lake¹⁷ to deliver Web-based OLAP products. Therefore, this is still a market segment where the niche DSS vendors have been able to take the lead, at least for now.

2.5.4.2 Startup DSS Vendors

MicroStrategy Inc., the market leader in Web-based OLAP products, develops a suite of DSS applications that interacts with DBMS platforms such as Oracle, Informix, and DB2 among others. This suite includes the DSS Web Server and DSS Broadcaster products that use the Internet and the World Wide Web as the communications medium. DSS Web Server, at version 5.5 as of this writing, is a Web-based interface that works with the company’s DSS Server OLAP engine to allow users to generate reports and analyses over the Web. Because the company uses ActiveX and Java components in DSS Web Server 5.5, they have been able to provide features such as drill down analysis and report pivoting, that were only available in their Windows-based DSS Agent product. Its DSS Broadcaster product is an application that allows for customized information delivery to be pushed to Web-based clients such as browsers and Internet appliances. The product won the “IT Manager’s Choice” Product of the Year in the data warehouses and data marts category organized by *Datamation* magazine.¹⁸ In addition, the company announced 1999 first quarter revenue of \$35 million, an increase of 80% over the corresponding quarter in 1998, gaining 55 new clients, including First USA Bank, France Telecom, and Kmart.¹⁹

InfoSpace, Inc. is the developer of SpaceOLAP, a fully Java-compliant solution that provides client/server-like interfaces for reporting and analysis. It includes a Java Application Server that resides on the Web server

¹⁶End-User Query and Reporting Tools — *ComputerWire PLC.*, March 1998.

¹⁷PR Newswire Association, Inc, December 1998.

¹⁸*Datamation* — February 1999.

¹⁹MicroStrategy’s Revenue, Profit Soars — *InformationWeek*, April 1999.

and supports data extracts from Oracle Express, DB2 OLAP, and Hyperion Essbase OLAP servers. Administration is provided through the SpaceOLAP Administrator module, while the SpaceOLAP Designer lets users and developers create customized reports and graphs based on the results of their queries. Reports, or “Presentations” as they are referred to by the company, can be displayed in HTML or Java applet supporting format, supporting pivot and drill down capabilities as well as multidimensional analyses.²⁰

Information Builders services the Web-enabled data warehousing market through its WebFOCUS product line. WebFOCUS provides reporting and publishing via the Web from both legacy databases and data in ERP applications. The company takes advantage of Java applets that provides for customized reporting and a Java Developers Workbench for designing customized reports from a Web browser. Among the other startup companies that have made a name for themselves in the Web-enabled data warehousing market are Cognos Corporation, the makers of Impromptu and PowerPlay, and Information Advantage, Inc., developers of the DecisionSuite OLAP product line.

2.5.4.3 Established DSS Vendors

The more established data warehousing companies such as Prism Solutions and Brio Technology have also developed Web-based interfaces into their OLAP products. The Brio Enterprise Server suite provides two modules that allow the user to extract information from data warehouses via the Web — the OnDemand Server and the Broadcast Server, competitor products to MicroStrategy’s DSS Web Server and DSS Broadcaster, respectively. Prism Solutions added Web-enabled functionality to its Warehouse Directory product in 1997 when it introduced the Web Access module, which allows users to view and query the Directory from their Web browsers. Having been acquired by Ardent Software, Inc. in April 1999, this module has since been integrated as a standard feature of the Warehouse Directory package and indeed, Ardent’s Vice President Peter Fiore sees Extensible Markup Language (XML) as the future conduit for sharing access to corporate data warehouses.²¹

The following table summarizes the list of vendors and products mentioned above.

Vendor	Product(s)	Technologies Supported
Oracle Corp.	Express Server (Web Agent)	HTML, Java, ActiveX
IBM	Net.Data, DB2 OLAP (Hyperion Essbase)	SQL, C++, Java, CGI
MicroStrategy	DSS Server OLAP, DSS Web Server, DSS Broadcaster	HTML, Java, ActiveX
InfoSpace	SpaceOLAP	Java, HTML
Information Builders	WebFOCUS	Java
Brio Technology	Enterprise Server (OnDemand Server, Broadcast Server)	CGI, Java Runtime Environment (JRE)
Prism Solutions	Warehouse Directory (Web Access module)	Java, XML

2.5.5 Future Trends

The trend to add more functionality to Web-based DSS tools will continue to grow as more companies realize the benefits of implementing Web-based access to their corporate data warehouses. However, don’t expect to see the demise of pure client/server-based DSS tools anytime soon as analysts predict that a coexistence of both forms of access will prevail for a time.²² As mentioned earlier, the main reason is that there are still features and functionality available with mature client/server-based tools and not with Web access tools. DSS vendors are working hard to add that functionality to their Web-based offerings so that in the future, even power users will be able to get the reports and analyses they need via the Web. Overall,

²⁰SpaceOLAP — *DBMS*, August 1997.
²¹Ardent’s Peter Fiore Is Passionate About Data Warehousing — *InfoWorld*, May 1999.
²²Self Storage; Lower Data Warehouse-Management Costs With Web Access — *Communications News*, August 1998.

the emphasis will shift away from desktop access and more toward the centralized “intergalactic” Web-based model as data warehouse implementations become larger and more users require access.

2.5.5.1 Web Farming

One area of interaction between data warehouses and the Web that is set to grow steadily in the next few years is the concept of Web Farming. This is defined as “systematic business intelligence by farming the information resources of the Web, so as to enhance the contents of a data warehousing system.”²³ Essentially, it is the process of using the Web not as a means of distributing data warehouse information to the end user but as a means of obtaining the raw data that goes into the warehouse itself. Such external data sources include commercial databases, e.g., the IBM database of patents, and public databases such as The Electronic Data Gathering, Analysis, and Retrieval (EDGAR) operated by the U.S. Security and Exchange Commission to track publicly traded companies.²⁴

Similar to extracting production data from OLTP systems into data warehouses, Web data must be refined to be suitable for use by the warehouse applications. This process, known as acquisition can vary depending on the sources of the Web content and is crucial in ensuring the usability and reliability of the data content as it becomes part of the decision support structure. While this sector is still in its infancy, it won’t be long before the major vendors turn their expertise to developing products that will make this form of data collection more reliable and efficient.

2.5.6 Conclusion

The value of using the World Wide Web to disseminate decision support information is evident from the growth in the use of Web-based DSS and OLAP tools. By providing access to corporate data warehouses over the Web, companies are empowering a larger user base with the necessary tools to analyze the business data. This in turn has opened up the decision-making process to more people within the organization. The lower costs of rolling out and maintaining Web-based data warehouse access has made this medium very attractive to companies that wish to gain the most value from their data warehouses. In addition, the trend to open up information systems to the customers and suppliers who are part of the supply chain has made easier access to data warehouse information more critical to these external users.

While this technology has grown in leaps and bounds and looks set to take over as the predominant form of data warehouse access, IS departments will still have to grapple with coexistence between Web-based and client/server-based access. But with the improving feature set of the various products being offered by the many vendors in this arena, Web-enabled data warehouse access will become the *de facto* access medium for this powerful business information.

References

- Carrickhoff, Rich — SpaceOLAP (*DBMS*, August 1997).
- ComputerWire PLC — End-User Query and Reporting Tools (March 1998).
- Datamation White Paper — Data Warehousing Management/Productivity Tools. (Datamation, 1997).
- Davis, Beth — MicroStrategy’s Revenue, Profit Soars (*InformationWeek*, April 1999).
- Fournier, Roger — *A Methodology for Client/Server and Web Application Development* (Yourdon Press, 1998).
- Hackathorn, Richard — *Web Farming For The Data Warehouse* (The Morgan Kaufmann Series in Data Management, 1998).
- Hackathorn, Richard — Routing the Web for Your Data Warehouse (*DBMS*, August 1998).
- Koch, Christopher — The Middle Ground (*CIO Magazine*, January 1999).
- Orfali, Robert. Harkey, Dan. Edwards, Jeri — *The Essential Client/Server Survival Guide* (John Wiley & Sons, Inc., 1996).

²³*Web Farming For The Data Warehouse* — Richard D. Hackathorn (Morgan Kaufmann Publishers, 1998).

²⁴Routing The Web For Your Data Warehouse — *DBMS*, August 1998.

Palo Alto Management Group, Inc., — Database Solutions White Paper (July 1998).
 PR Newswire Association, Inc. — Red Brick and Caribou Lake Software Team Up (December 1998).
 Reinauer, Rob — Self Storage; Lower Data Warehouse-Management Costs with Web Access (*Communications News*, August 1998).
 Row, Heath — Just Browsing Thanks (*CIO Magazine*, October 1996).
 Schroeck, Mike — Data Warehousing is Worth the Investment (*InternetWeek*, June 1998).
 Schwartz, Susana — Warehousing And The ‘Net — Marriage Made In Heaven (*Insurance & Technology*, July 1997).
 Scott, Jim — Warehousing Over the Web (Association for Computing Machinery, *Communications of the ACM*, September 1998).
 Vizard, Michael — Ardent’s Peter Fiore Is Passionate About Data Warehousing (*InfoWorld*, May 1999).
 White, Colin — Building Web Information Systems (*Byte Magazine*, July 1998).
 White, Colin — Warehouses Webicum: Evolution of a Species (*DB2 Magazine*, April 1998).

2.6 E-commerce Technologies: A Strategic Overview

Mihir Parikh

The changes sweeping through electronic communications will transform the world’s economies, politics, and societies — but they will first transform companies.

—Frances Cairncross, in *The Death of Distance*, 1997, p. 119.

E-commerce is a result of these changes. In simple terms, e-commerce is defined as a business conducted over the Internet with the use of computer and communications technologies. There are three major types of e-commerce: business to business (B2B) e-commerce, business to consumer (B2C) e-commerce, and electronic markets. *B2B e-commerce* deals with one business providing products and services to another business typically as a part of the supply chain or as an enabler of business processes. Examples of B2B e-commerce include an auto-part manufacturer supplying an auto company, or a bank providing credit card payments and other financial services to a retailer. For many years, electronic data interchange (EDI) handled B2B transactions in many companies. Now, Web-based open system applications are replacing proprietary EDI systems. *B2C e-commerce* deals with a business providing products and services to consumers at the end of the supply chain. Examples include a book retailer selling books to a reader (Amazon.com) and a broker providing financial trade executions to an individual investor (E*Trade). *Electronic markets* provide marketplaces on the Internet, as opposed to marketplaces in the physical world, where buyers and sellers can meet and exchange products and services. Electronic markets are of two types: consumer markets and business hubs. Ebay, priceline.com, and accompany.com are typical examples of electronic markets for consumers. Chemdex, Metalsite.com, and Ultraprise.com are typical examples of electronic markets for businesses.

Recently, International Data Corporation (IDC) estimated that yearly worldwide e-commerce would increase to more than \$1 trillion by year 2003.²⁵ IDC estimated that by then non-U.S. countries would account for half of worldwide e-commerce. Are the current businesses ready for it? Well, not really. A recent survey conducted by the Cutter Consortium found that 65% of companies did not have an overall e-commerce strategy and nearly 25% lacked even a basic business and implementation plan for e-commerce.²⁶ A major reason for this is that a few companies have aligned their business strategies with information technology (IT) strategies. A large percentage of companies still do not involve IT in high-level strategic planning. IT in these companies is relegated as a support function. The companies have failed to recognize the changing role of IT from business support to business enabler. The strategic implications of IT and the use of IT as a strategic weapon are not well understood.

²⁵International Data Corporation Report, NET062899PR.htm, June 28, 1999.

²⁶*InternetWeek*, September 6, 1999. Page 29.

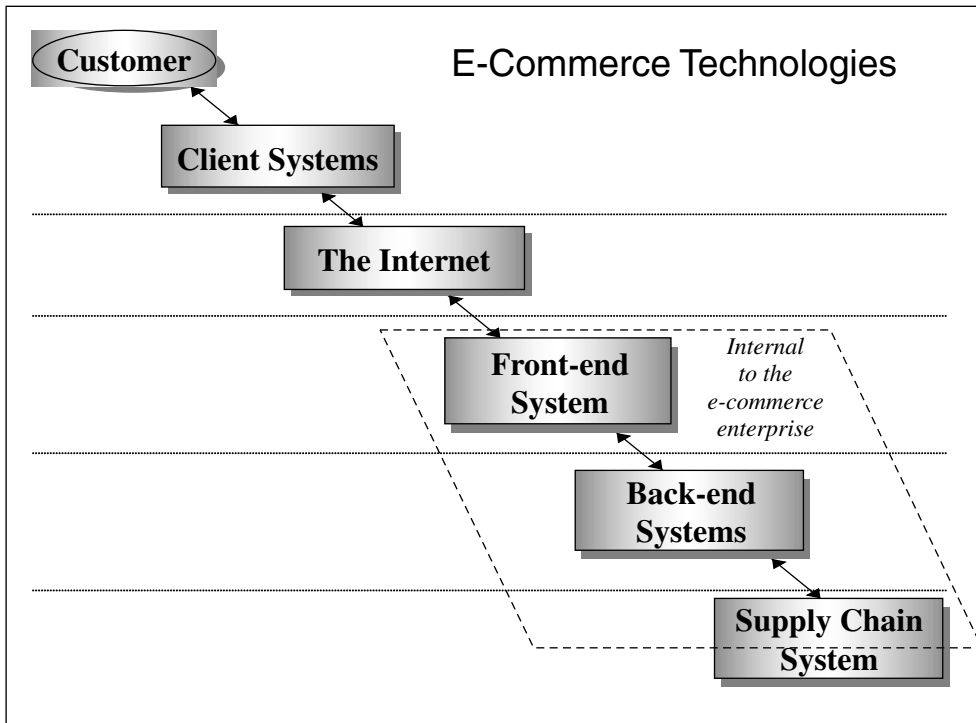


FIGURE 2.6.1 E-commerce technologies

Information technologies are key enablers of e-commerce. A combination of many information technologies ranging from a Web browser to logistics management systems makes e-commerce possible. Identifying and implementing the right technologies to execute business processes is critical to succeed in the e-commerce space.

2.6.1 E-commerce Technologies

Figure 2.6.1 shows different categories of e-commerce technologies. All of these technologies are required and play significant roles in managing and conducting business over the Internet.

2.6.1.1 Client Systems

Client systems reside on customers' computers. Customers utilize these technologies to participate in e-commerce activities.

Web browser and cookies: The most common of all client systems is a Web browser. Web browsers provide an interface through which a user can view information on the Internet. In the last five years, Web browsers have evolved from a small software using simple hypertext markup language (HTML) to a very sophisticated software that uses new technologies such as Java, ActiveX, VRML, XML, and different types of plug-ins and Web applications in addition to a much advanced version of HTML. Most web browsers utilize cookies whereby they transmit basic information about a user to the Web server on the other end for identification purposes. Such identification can be used to personalize services for the user. Multiple levels of cookies are used. Low-level cookies simply provide user name and password, while high-level cookies may include information about credit card, mailing address, previous purchase patterns, and browsing habits.

Communications software and hardware: Communications software and hardware help customers' computers to connect to the Internet via one of the multiple modes including modem, cable modem, satellite links, and local area networks.

Plug-ins: Plug-ins are independent software applications utilized to show special data files within a Web browser. Plug-ins enable a Web browser to show a multimedia presentation or to play a streamed audio piece. Shockwave Flash, RealPlayer, and Adobe Acrobat Reader are plug-ins.

Software agents: Agents are other types of independent software that can assist users in carrying out some specific activities. Such activities include filtering information, searching the Web to find the right information, and comparison shopping.

Biometric identification: Biometric identification is a technology that uses a measurable physical characteristic to recognize the identity, or verify the claimed identity, of an enrollee.²⁷ It utilizes physical characteristics such as fingerprints, facial design, iris patterns, retina patterns, hand geometry, signature verification, and voice recognition instead of keys, passwords, and plastic cards. Several advantages, including reduction in fraud and never losing the identification, have prompted several large companies such as IBM and Motorola to invest heavily in biometric identification. Biometric identification is one of the fastest-growing areas on client systems and security.

Push technologies: These publish and subscribe technologies enable delivery of possibly useful information without the recipient asking for it. Push technologies utilize extensive user profiles containing preferences of each user to match and deliver information over the Internet. Push evolved as an alternative to the Web (pull technology) from PointCast's personalized broadcasting technology in 1996. In the last three years, it has gone through the full length of the hype curve. Recently, it has returned with reasonable expectations and clear understanding of what the technologies can and cannot do. Several companies, including Marimba and BackWeb, provide push technology solutions.

2.6.1.2 The Internet

Many advanced telecommunications technologies have been utilized to create and operate the Internet. Most prominent are optical fiber, routers, digital switches, Synchronous Optical Network (SONET) rings, asynchronous transfer mode (ATM), frame relay, Transmission Control Protocol (TCP), and Internet Protocol (IP). These hardware and software technologies provide the backbone lines and the transmission rules to carry out information exchange over the Internet. As the Internet is growing, new technologies are emerging to increase speed and volume of data transmission. The continuing convergence of textual data with audio and video data over the Internet is prompting new technologies to support quality of service (QoS) which recognizes the differences between data types and assigns appropriate priority for transmission.

2.6.1.3 Front-end Systems

Front-end systems are the ones with which the customers interact. They provide a face to an e-commerce business.

Web pages: Web pages are the data files containing HTML-coded information. In the early days of e-commerce, Web pages were static and generated by human Web programmers. Now, in most e-commerce sites, Web pages are dynamic and are generated by Web page management systems. These systems work with other front-end systems and back-end systems to develop HTML-coded content that the customers receive.

Traffic management: Due to a special sales event or some top news, sometimes an unexpected number of visitors go to a Web site all at once. Such an unexpected load puts pressure on the Web server and dramatically reduces its speed and increases download time. If the load extends for a longer period of time, it crashes the system. To avoid such a mishap, traffic management tools are used. These reduce massive congestion by spreading traffic load on multiple servers and increase overall network efficiency.

Search engines: Several types of search engines are utilized by e-commerce companies. Some search engines provide capabilities to search a specific product based on description, features, or other information. Some search engines provide capabilities to search and locate Web content based on key words or phrases. Search engines can be used to search a specific Web site, a regional part of the Web, or the

²⁷Digital Imaging: Connecticut Biometric Imaging Project. <http://www.dss.state.ct.us/digital.htm>

whole Web. Often directory engines are used along with search engines to automatically categorize Web content for future searches.

Site servers: Site servers are the most comprehensive tools for e-commerce. They provide support ranging from creating an electronic storefront to controlling and managing it. Site servers include support for site building, standard code sharing, code library development and management, dynamic Web page generation, product promotion, product cataloging, order taking and processing, securing transactions, and managing payment systems. Site servers help provide product information, dynamic pricing information, marketing and promotion, shopping cart services, tax calculations, shipping and handling calculations, and automated post-sales follow-up. In addition, they capture market demographic information and coordinate with back-end systems. Several off-the-shelf products are available to support small business shopping services including IBM's Net.Commerce and Microsoft's Site Server. However, a major e-commerce site requires custom application development to support the above-discussed activities.

Shopping engines: These types of software enable customers to find and compare different products or services on features and prices. Some shopping engines also include product reviews from independent agencies such as *Consumers Digest*. Some shopping engines also provide reviews from and discussion with current users of the products to help others make more informed purchase decisions. It is a very useful tool in e-retailing. Amazon.com's Junglee division and Inktomi are the two leading companies that provide shopping engines.

Customer relationship management (CRM): Through a marketing agreement, E-commerce merchants pay between \$0.90 to \$2.67 per visitor referred by a portal.²⁸ As the search cost for every new customer is staggering, increasing customer satisfaction and loyalty is crucial to maintain the current customer base and for e-commerce success. CRM systems provide the critical customer support and services. They provide a database of frequently asked questions (FAQ), searchable knowledge base, multiple way (e-mail, Internet telephony, video conferencing, etc.) to assist shoppers in real time, follow-up support, order tracking, return processing, after-sales services, and warranty processing. They also help maintain profiles of buyers that contain their shopping behaviors and preferences. The currently available CRM-related technologies include enterprise portal, mobile computing, net telephony, desktop video conferencing, speech recognition, call center systems, and data warehousing.

Personalization: As the number of e-commerce businesses increases, a key differentiator would be how well an e-commerce business customizes its storefront for each customer. Knowing the customer and his or her preferences will improve customer service and retention. This personalization or mass customization requires creating user profiles from purchase patterns and browsing patterns, and applying business rules and inference on the information collected in the user profiles to create new knowledge about the users. Key technologies used for personalization are databases, dynamic Web pages, business rules, inference engines, cookies, and push technologies.²⁹ In some cases, data warehousing and mining technologies can also be used for personalization.

Security: More than half of B2C e-commerce transactions are paid with credit cards. Protecting transmission of credit card and other private, confidential information over the Internet and securing the information on merchants' Web servers are two of the most pressing issues in e-commerce. Encryption of data through secure sockets layer (SSL) and private and public keys are the most commonly used technologies to secure Internet transmission of confidential data.

2.6.1.4 Back-end Systems

While the front-end systems manage the interface with customers, the back-end systems carry out operations and manage e-commerce organizations.

Enterprise Resource Planning (ERP): ERP systems are commercial software packages utilized to integrate information flowing through different functions of an organization, such as financial and accounting, human resources, sales and customer service, supply chain, etc. ERP systems coordinated

²⁸International Data Corporation Report, May 4, 1999. NET050499PR.htm

²⁹The Web gets personal. *Byte*, special section on E-business Technology, June 1998.

with front-end systems can capture orders, provide order confirmation, accept payment, check credit cards for approval, process coupons and other promotions, handle billing and invoicing, control inventory and procurement, integrate with payment systems, and coordinate with fulfillment systems for order execution. Before ERP systems, these processes were handled by various, independent information systems indigenous to different functional divisions in the organization. Information systems from one functional division were often not compatible with the information systems in other functional divisions. This brought inefficiencies in business processes and overall higher costs. ERP systems promise seamless integration and easy information flow among different functional divisions. However, successful implementation of an ERP system has been one of the major critical issues in the utilizing ERP systems.

Databases/Data Warehouses/Data Mining: Databases and data warehouses are at the center of running a business in this information economy, especially an e-commerce business. They provide repositories for information collected through business processes. This information is the lifeblood of organizations and its optimum use is very important. Several emerging database technologies such as multidimensional databases provide holistic perspective and better understanding of the information. When used in conjunction with data mining technologies, e-commerce businesses can find and exploit hidden relationships and buying behaviors to increase market share and sales.

2.6.1.5 Supply Chain Systems

These systems work with business alliance partners and other enablers. They provide smooth transfer of information with partners to carry out outsourced business processes. Some of these systems are external to e-commerce organizations and implemented by the partners.

Supply chain management: Most emerging e-commerce companies are not vertically integrated. They depend on many partners and intermediaries on both sides of the supply chain (a value addition sequence through which raw material flows to become a finished product). Supply chain management systems help businesses coordinate their processes with suppliers, manufacturers, raw material providers, shippers, distributors, and associated retailers. In e-commerce businesses, greater efficiencies are achieved by moving information rather than actual products along the supply chain. Actual products are generally delivered directly to the consumer by the product manufacturers without any supply chain intermediaries handling or storing the products. Supply chain management systems help control product life cycle, forecast demand, arrange advanced scheduling, plan manufacturing and distribution, and enable order promising and processing. Several companies provide supply chain software. The leaders are i2 Technologies, Manugistics, and Numetrix. Several ERP companies are also moving into this area by extending ERP capabilities.

Payment systems: A majority of payments, almost 89%, over the Internet is conducted through credit cards and checks. Payment systems help e-commerce businesses coordinate with banks and credit card companies to approve credit card purchases and clear checks. Some e-commerce companies also utilize these systems to work with soft cash providers such as CyberCash and electronic money.

Fulfillment/Logistics management: These systems coordinate with logistics partners such as FedEx, UPS, and independent warehouse operators. These systems work with back-end and front-end systems to help determine shipping and handling charges, delivery terms, delivery schedule, order tracking, freight management, custom and excise duty clearance, and other fulfillment issues.

2.6.2 Strategic Challenges

In the first chapter of the e-commerce storybook, the technology largely drove business models. Now the business models are driving technology.

—Peter G. Keen, *Computerworld*, September 13, 1999

2.6.2.1 High Cost of Small Errors

In e-commerce, there is little room for errors. In most cases, you do not get a second chance. Integrating right strategies with right technologies and continuously improving competitive position are important for survival and success. Any error in technology or strategy implementation can lead to a serious loss in market position and the ability to carry out the business in future. Once, eBay's stock price dipped more than 50%

largely due to recurring, unexpected shutdowns of its Web site. Most of these shutdowns were not more than a few hours long. However, persistence and a strong commitment by the company's top management to improve technological infrastructure lead to a rebound in the stock price. E-commerce businesses pay a very high cost for small errors. As an e-commerce business expands, the probability of making such errors increases. While most managers and leaders assume that being there first is the key to success, a study has found that many pioneers fail and most current leaders are not pioneers.³⁰ The study found that five factors (vision, persistence, commitment, innovation, and asset leverage) are critical to success. These factors often lead to making fewer errors and quickly correcting the errors when they are made.

2.6.2.2 Building Relationships

Technology is a double-edged knife. While it provides unprecedented advantages to you, it also provides the same benefits to your competitors and future competitors. It enables your existing competitors to quickly react to your moves. It also reduces the barriers to entry for new competitors. Assuming that if you build it, they will come and stay is one of the fallacies of e-commerce. To pre-empt this, you have to build a strong customer base and retain it. Building a strong customer base requires building relationships and providing useful services and content to your customers. Often it is measured in terms of "stickiness" of the Web site, that is, how long a visitor stays on the Web site. Media Metrix reports eBay (125.5 average minutes per user), E*Trade (66.5), Microsoft sites (66.0), and Yahoo sites (64.6) are the top four stickiest Web sites.³¹ It is not a coincidence that the companies with the stickiest sites are the most successful in e-commerce.

2.6.2.3 Speed

The speed of doing business has increased tremendously with the Internet. Speed is required in growth, in decision making, in adapting to the changing conditions, and in supporting and servicing customers. An e-commerce company has to continuously innovate and improve its business processes and Web-based storefront. It is always vulnerable to the quick imitation of its processes and its innovative shopping features by its competitors. Often not only the Web site designs but also business models are copied overnight by competitors. This also requires building flexibility in the front-end and back-end systems to adapt to the continuously changing conditions and stay ahead of competitors. Therefore, industry experts recommend an open and adaptive architecture for enterprise information systems.³²

Operationally, when a customer visits an electronic storefront, performance of the Web site becomes an important issue. First of all, the customer wants quick downloading of the Web pages and immediate response to any search queries. If the Web site is slow, the customer will very likely move on to another competing store. This invariably happens when an unexpected number of customers come to the Web site at the same time. Several technologies can help improve speed of the Web site. Use of traffic management tools can balance the load on a Web server and increase the speed of interaction. Scalability of the hardware and software provide quick integration of additional resources. The use of better search engines and shopping engines can also increase the speed of searching the requested product from millions of product profiles stored in the databases. Good shopping engines can also help identify related, complementary products for cross selling and up selling.

2.6.2.4 Security

With the growth of e-commerce, more and more business processes and databases are put on the Internet. This is required to improve customer service and increase organizational efficiency. However, this also makes the processes and databases vulnerable to malicious forces including business spies, computer hackers, and disgruntled former employees. Having complete control over who gets to see what and who

³⁰First to Market, First to Fail? Real Causes of Enduring Market Leadership. By Gerard Tellis and Peter Golder. *Sloan Management Review*, 32(2), 1996, 65-75.

³¹Snapshot: Sticky Sites. *Computerworld*, September 13, 1999. Page 42.

³²Designing a growing back end. *InfoWorld*, August 23, 1999. Pages 34-35.

gets to change what is extremely important. Security management software, virus protection software, and intrusion detection software can help increase security of the Web site.

2.6.2.5 Technology Evolution

E-commerce technologies are in a constant state of flux. E-commerce businesses have to continuously evaluate emerging technologies and adapt them quickly to stay competitive. As the technologies are constantly evolving, few standards exist. The ability to choose a right technology with a strong future becomes a critical skill in managing e-commerce.

2.6.3 Emerging Trends for the Future

What is my ROI (Return on Investment) on e-commerce? Are you crazy? This is Columbus in the New World. What was his ROI?

—Andy Grove, Intel Chairman

New technologies are emerging every day. Some of these technologies may become a killer app for e-commerce. While it is almost impossible to predict them, using standard measures (such as ROI) to evaluate may fail, too. However, several underlying trends may help to identify and evaluate the right technologies.

The drop in the cost of computing continues. Storage, processing, and distribution cost of information is decreasing to a level where the cost is less than the value of information. This has led to the development of new enterprises that provide free products and services in return for information and loyalty. On the hardware side, we have FreePC, PeoplePC, eMachine+Compuserve alliance, etc. On the applications side, we have HotMail, when.com, Yahoo, etc. On the Internet access side, we have NetZero, Freeserve, etc. On the Web site hosting side, we have GeoCity, Tripod, etc. More and more of these types of e-commerce businesses will continue to rise, making it difficult for the current businesses to compete.

While the power of computer processors is going up, their sizes are decreasing. In addition, computer processors are now used in many devices and products (such as autos, refrigerators, washing machines, dishwashers, etc.). As computing technologies continue to expand to household products and appliances, communications technologies will soon follow. These products and appliances, when connected to the Internet, will create new e-commerce opportunities. For example, a refrigerator in the future may be able to automatically buy groceries for you. The jar of milk is put in the refrigerator in one location where there is a sensor that notices how much milk is left. As soon as milk reaches the reorder level (determined based on your consumption pattern), the refrigerator will automatically connect with a grocery store on the Web (maybe NetGrocer, PeaPod, or WebVan) and place an order for milk.

New software applications are emerging every day. Each one makes e-commerce business processes more efficient and effective. This will enable more and more people to go online for their shopping, entertainment, business and home management needs. New e-commerce models will rise to support these changes in customer behavior.

2.7 Internet Protocols

John Braun

The purpose of this section is to describe concepts that are essential to understanding how Internet protocols work. Basic addressing at the device or client level will be described, followed by protocols that are used to exchange information among these devices. Although the term device may seem vague at first, the types of devices that communicate via the Internet have matured from simple text-based information, to the use of audio, video, animation, and other forms of communication. Once the basics of communication are laid out, some of the specific protocols and applications that utilize these basics will be discussed. Security aspects of these protocols and applications will be covered. Useful search tools that can help locate information on the Internet will be covered, and a discussion of some of the major industry players will conclude this section.

testing purposes, or as a sanity check to determine if one’s TCP/IP implementation is working properly. The typical value used for this purpose is 127.0.0.1.

References

Postel, J., “Internet Protocol,” RFC 791, USC/Information Sciences Institute, September 1981.
Reynolds, J. and J. Postel, “Assigned Numbers,” RFC 943, USC/Information Sciences Institute, April 1985.
*** Latest is RFC 990

2.7.1.1 DNS

The domain name system (DNS) is a global network of computers that can translate a numerical IP address to a human-readable name, and also translate a name to the corresponding IP address. This makes navigation of a TCP/IP network much easier. For example, www.crcpress.com is much easier to remember than 199.29.24.3, the IP address which corresponds to this address.

Before DNS, each computer on a network would have to maintain a large file (typically called hosts) with all known IP address and name pairs. This is obviously impossible to do now with the sheer number of hosts on the Internet, but can still be useful for small, private TCP/IP networks which are not directly connected to the Internet.

When configuring a client to access DNS services, multiple DNS servers should be specified, if available. Modern TCP/IP implementations are smart enough to try another DNS server if the initial one is unavailable.

A DNS client will submit a DNS request to a server, and receive one of three types of reply. The client will be told that the lookup was successful and be given the name, that the server couldn’t perform the lookup but knows another server that may, or that the lookup failed.

References

RFC 1035 — Domain names — implementation and specification. P.V. Mockapetris. Nov-01-1987

2.7.2 Communication Protocols in Internet

There are many protocols used on the Internet. Most are classified at either the lower layers (Layer 3, Network and Layer 4, Transport) of the OSI model, or at the application level (Layer 9).

IP

RFC 791

The IP in TCP/IP refers to the Internet Protocol used at the network layer. The basic purpose of this protocol is to try to deliver packets. It does not offer such services as acknowledgment, retransmission, error correction, flow control or guarantee of order of delivery. This is the job of higher-level protocols. The advantage of IP is that it offers a common framework for devices to communicate.

An IP header looks like this:

Version	IHL	Type of Service		Total Lengh
Identification		Flags		Fragment Offset
Time to Live	Protocol		Header Checksums	
Source Address				
Destination Address				
TCP Header, Then your data *****				

UDP

RFC 768

User Datagram Protocol (UDP) is a connectionless protocol that provides a means to send data with a low protocol overhead. Only the source port, destination port, length, and checksum are added to the raw data. However, it does not guarantee delivery or protection from duplicates. For applications where performance is critical, a small loss of data is not critical, and link is known to be reliable, UDP may be used. Streaming audio or video are examples of applications where high performance is more important than a possible, but potentially correctable, loss of data.

TCP

RFC 793

Transmission Control Protocol (TCP) offers a more robust, connection-oriented method for reliably sending data. Flow control and multiplexing are supported. Unlike UDP, TCP supports the concept of a continuous stream of data between two hosts. Unlike IP, TCP will make multiple attempts to deliver data if the initial attempt fails. If the integrity of data is critical, TCP should be used.

A TCP header looks like this:

Source Port		Destination Port	
Sequence Number			
Acknowledgment Number			
Data Offset	Reserved		Windows
Checksum		Urgent Point	
Your Data *** next 100 octets *****			

2.7.3 Information Transfer in Internet

A socket is a virtual communications channel that is established between two hosts. It can use either UDP or TCP for the transport protocol. Each socket has a unique descriptor, and multiple sockets can reside on the same port. This allows multiple clients to take advantage of a service on a single port without the server getting confused. In order to prevent congestion on a single port, many application-level protocols have a control connection on a known port, and then negotiate another port for subsequent data transfer.

2.7.4 Types of Internet Access

There are two basic types of Internet access. One is a full-time direct connection; the other is via a telephone line with a modem. The disadvantage of a direct connection is mostly cost in both dollars and network maintenance. The advantage is speed, where T1 (1.5 Mb/sec) or T3 (45 Mb/sec) rates are common measures. The disadvantage of a modem connection is reliability (line noise) availability (busy signals) and speed. Certain schemes such as v.90 can achieve up to 56k bps download speeds. The advantage of a modem connection is the cost of both equipment and service, with 56k bps modems available for under \$300 U.S., and unlimited service for around \$20 U.S. a month.

Integrated services digital network (ISDN) can provide 128 kb/sec transfer rates. It never seemed to catch on due to the difficulty in configuring the equipment and inconsistent pricing plans across the country. Many providers charge a flat rate, with some adding charges for each unit of time or unit of data sent, whereas a local telephone line typically allows unlimited usage.

Asymmetric digital subscriber line (ADSL) is a relative newcomer that takes advantage of existing twisted-pair wiring, and can reach speeds of 6 Mb/sec for downloading, and 640 kb/sec for uploading. It is being deployed in major cities, but it remains to be seen how widespread the service will become.

Cable modems are slowly becoming available, and offer up to 10 Mb/sec transfer rates. Typically, the upstream connection consists of high-speed fiber, with the final connection to the cable modem being made with coax.

There are hybrid solutions, such as a satellite (with speeds of 400 kb/sec and higher) that uses a phone line for uploading data, and a satellite for data download. This can be a good solution for scenarios such as surfing the web, where the amount of data sent to request information is much less than bandwidth-intensive data types, such as graphics and sound, which comprise the data.

There are two major protocols used for establishing a TCP/IP connection over phone lines. Point to Point Protocol (PPP) is the more modern method, and PPP software is included with nearly every major operating system. Serial Link Internet Protocol (SLIP) is an older standard that is being phased out in favor of PPP.

2.7.5 Internet E-mail

SMTP — STD10

POP — STD53

IMAP4 — RFC2060

There are a few standards for sending and receiving Internet e-mail. The most popular protocol for receiving e-mail is Post Office Protocol (POP) which defaults to TCP port 110. A newer protocol, Internet Message Access Protocol (IMAP4), used for receiving e-mail, resides on TCP port 143.

The most common protocol used for sending e-mail is Simple Mail Transfer Protocol (SMTP) which defaults to TCP port 25. Note that POP can also be used for sending e-mail, but this feature is an optional extension of the POP protocol, and is not supported by many e-mail clients and servers.

SMTP

The SMTP service allows the sending of e-mail by providing relevant information, in a specific order, to a SMTP server. The conversation between the client (sender) and the server consists of human-readable text commands that are assigned four-letter codes, and three-digit code numbers as responses. A typical exchange may look like this:

```
C: HELO megacorp.com                (sender identification)
S: 250 smtp.conhugeco.com           (server identification)
C: MAIL FROM:<johnbraun@megacorp.com> (who is this from?)
S: 250 OK
C: RCPT TO:<dilbert@anotherdomain.net> (who is the recipient?)
S: 250 OK
C: DATA                            (ready for data?)
S: 354 Go ahead...                  (sure)
C: Hey Dilbert get to work!
C: .                                (end of input)
S: 250 OK
C: QUIT                             (sender all done)
S: 221 Bye...                       (server says seeya)
```

This system works well enough, with servers cooperating when mail needs to be forwarded to a domain outside of their own. Messages can be delivered in minutes, if all the servers in a message's path respond in a timely manner.

Although standard SMTP does not have any provision for confirmation of receipt (there are proposed standards, but they are not widely implemented yet) problems encountered along the path of the message are usually reported to the sender. If a server finds that another server is unavailable, it will usually try

to send the message several times before giving up. If a server crashes while processing mail, the message may get lost forever. For critical documents, it may be wise to ask the receiver to confirm receipt.

One problem with current SMTP implementations is that they don't confirm the validity of the sender's address. This has led to massive abuse by junk e-mail senders, who really don't want you to respond via e-mail, anyway. An attempt to reply will result in getting a message saying that the (bogus) return address doesn't exist, or that it has already been shut down.

There are proposed methods of authenticating an entity wishing to use an SMTP server, in hopes of reducing spam and other evils. There are also solutions, mainly using public key encryption and digital signatures, to confirm the identity of the sender. These features are not available at the protocol level yet.

POP

The POP protocol is used to retrieve e-mail for a specific user. Session establishment can be done with a simple username and password scheme, or can optionally use more sophisticated means like APOP. APOP never sends the user's password over the connection, instead applying the MD5 algorithm to the password and other time-sensitive data. Not all POP servers support APOP, and will return an error message if this method of login is attempted but not supported.

After a message is retrieved, it is usually deleted. Some POP servers allow the user to keep their e-mail on the server after it has been retrieved, but this is not guaranteed. The IMAP4 protocol is better suited for keeping e-mail on a server.

Some POP servers offer the ability to send e-mail as well as receive it. The advantage of using POP for sending e-mail is that it requires users to identify themselves, thereby reducing the generation of spam and other unauthorized e-mail. It also eliminates the need to maintain a separate SMTP server. The disadvantage of this method is that not all e-mail clients and servers support sending via POP.

IMAP4

IMAP4 is a newer protocol for retrieving e-mail. It offers a richer set of features than POP, including keeping e-mail on a remote server, searching e-mail before retrieval, and a greater number of authentication schemes.

The option of keeping e-mail on a server helps reduce resource requirements on the client, but increases the need for more disk space, as well as regular backups to ensure the historical data is preserved. This scheme can benefit devices such as portable computers, network computers, and portable digital assistants, which may have limited storage capacity. It can also offer convenience to mobile users, since their mail can be stored on the server, rather than being spread among multiple clients.

2.7.6 Telnet in Internet

Telnet — STD8

Telnet is a service that resides on TCP port 23 and offers terminal services to remote users. The client and server can negotiate features of the connection, which can range from a simple ASCII exchange, to enhanced services such as cursor control and styled text.

A telnet client can be used to interact with other TCP services that exchange text or binary data, provided that the telnet client allows one to specify the port one wants to connect to. For example, a telnet client can connect to an SMTP server on port 25 and send e-mail. This can be handy for debugging or investigative purposes.

2.7.7 File Transfer in Internet

FTP

Request for Comments: 959

RFC 990 — Port #

File Transfer Protocol (FTP) provides a reliable, standard way to transfer both text and binary files between hosts. There are two types of connections when using FTP, a control connection and a data

connection. A control connection is used for the client and server to exchange commands and status information. The default port for a control connection is 21. The default client data port is 21, and the default server data port is 20.

In most cases, a data port with a value outside the range of commonly known services is selected for security reasons. The client can request a specific data port via the PORT command, or can ask the server to select one with the PASV command.

Although FTP makes every attempt to deliver a file, circumstances beyond the user's control (modem disconnect, network failure) may cause the transfer to be interrupted. Most, but not all, FTP implementations support a restart mode where file transfer can begin at a point other than the beginning of the file, allowing data transfer to begin at the point of interruption and complete.

2.7.8 News and Usenet

NNTP — RFC977

The most common way of distributing news on the Internet is by using the Network News Transfer Protocol (NNTP). The system, originally created to exchange messages in a university environment, has evolved to a worldwide messaging system with thousands of distinct topics of interest, called newsgroups.

Each newsgroup name consists of several elements, with the general classification at the beginning of the name, and the specific subject matter at the end. For example, the group comp.sys.mac.hardware.video is a computer-related group about video cards for Mac systems.

The most popular hierarchies are sometimes referred to as the “big seven” and include comp (computing), misc (miscellaneous), news (newsgroup-related), rec (recreation), sci (science), soc (social), and talk (discussion). There are many other hierarchies for alternative and localized content.

Clients to allow one to read and post news are available as both stand-alone applications, or integrated with other Internet clients such as a web browser. Sophisticated clients can allow one to organize messages about a similar topic into threads, or filter messages based on certain criteria. Although most messages pertain to the topic of the newsgroup, there are those who “spam” the newsgroups with ads or other material unrelated to the group.

Care should be taken when configuring your news client where it asks for your e-mail address. Your address is normally added to any messages that you post, with the intent of making it easier for others to make a personal reply. Alas, there are those who scan the newsgroups for e-mail addresses, which are then used for the purpose of sending junk e-mail. A good strategy is to add some text to your address that a human would know to delete before making a personal reply, such as john@ihatespam.mega-corp.com.

A dedicated news host is required before clients can post and retrieve messages. Typically, a news host will server a large group of clients, and will exchange articles with another upstream host. The connection between hosts is referred to as a feed, and is meant for propagating articles and control messages, and not meant for direct client connections. If all hosts along a certain path agree to exchange the same group, a message posted in a particular group on one host will eventually propagate to all other news hosts. Hosts can also choose to restrict the groups they carry, which can help conserve disk space and bandwidth.

Care must be taken in deciding which groups a host should carry, and how long the articles should remain on the server before being deleted or expired. Inaccurate estimates can result in a dreaded disk full error, which causes many hosts to reject articles until someone clears some disk space.

2.7.9 Mailing Lists in Internet

A mailing list is a method of allowing Internet users to communicate about specific topics via e-mail. To join a mailing list, one sends an e-mail to a special address, and indicates a desire to join the list, as well as what e-mail address should receive further information from the list. The subscriber will then receive information on how to submit messages to the list, and how to perform administrative functions,

such as being removed from the list. If the list is unmoderated, any message submitted to the list will be sent to all other subscribers. If the list is moderated, an administrator will decide which submissions should be distributed to other members. Messages may be sent one-by-one, or grouped and sent in a digest. Choosing to receive messages in digest form is a good idea for busy lists, lest your mailbox gets filled with hundreds of messages.

<http://www.lsoft.com/listserv-hist.stm>

<http://www.wrt.tcu.edu/wrt/wri/files/barnes.htm>

listserv distribute protocol information — [rfc1429.txt](http://www.ietf.org/rfc/rfc1429.txt)

2.7.10 Information Search in Internet

2.7.10.1 Spiders

There are several tools, referred to as “spiders” which basically “crawl” through a web site, acquiring some or all of the information on each page. They then allow users to search through this information, in hopes of finding a resource relating to the topic of interest.

Many spiders have evolved into portals, which not only allow you to find information drawn from the Internet, but also provide links to other commonly used services, from maps to news headlines to package tracking.

Some of the more popular spiders are:

Alta Vista <http://www.altavista.digital.com>

Excite <http://www.excite.com>

Infoseek <http://www.infoseek.com>

Lycos <http://www.lycos.com>

WebCrawler <http://www.webcrawler.com>

2.7.10.2 Category Indexes

Another method of finding what you are looking for is to use a tool where sites are scanned by actual humans, and then placed into one or more categories. This method of locating information is a good complement to using a spider, since a spider may return too much information to be useful. The most popular of these services is Yahoo at <http://www.yahoo.com>.

2.7.11 Netscape and Microsoft

Netscape (NASDAQ: NSCP) and Microsoft (NASDAQ: MSFT) are two of the major players in the commercial Internet client and server markets. Netscape, established in 1994, was the first company to offer a commercial Internet browser, followed by server products. Microsoft was late to the party, but now also offers a full suite of Internet client and server products. Whereas Microsoft server products run on Windows, and their client programs span Windows, Mac, and UNIX, Netscape offers server products on both the Windows and UNIX platforms, and client programs for Windows, Mac, and UNIX.

Netscape has evolved from being a browser-only company, to one that now depends on server software and access to their NetCenter site for revenue. This is no doubt due to Microsoft’s giving away their Internet Explorer browser. Netscape eventually made the source code to their browser available, in an attempt to gain back some browser market share by opening up their product.

Although the tight integration of Microsoft’s browser and server products with the Windows operating system can offer increased functionality, it also tends to lock one into Windows. This can be a problem when attempting to use Microsoft products with more standards-based solutions like those from Netscape. Proprietary technologies like ActiveX controls and VBScript don’t always work well with Microsoft products.

In the fast-moving world of the Internet, it is hard to predict which, if either, company will win. The strength of Netscape is that they provide solutions for a wider variety of systems, especially in the UNIX

realm, and that they don't lock you into a single environment. The strength of Microsoft is that their products are on almost all desktop systems. For a solution where both client and server products are guaranteed to come from Microsoft, the increased functionality of being closely linked to Windows can be worth it.

To put things in perspective, the most popular web server application as of this writing is the freeware Apache web server, with almost 50% market share.

Intranet References

The following web sites provide comprehensive sources of information and links to other sites for reference material relating to the Internet:

<http://www.cisco.com>
<http://www.intel.com>
<http://www.microsoft.com>
<http://www.sun.com>
<http://www.wcom.com>

Glossary of Intranet Acronyms and Terms

ARPAnet	Earliest packet switched network; the progenitor to today's Internet.
ASCII	Plain text file, containing only regular keyboard characters.
BCP	Bridging Control Protocol, used to configure bridge protocol parameters on both ends of a point-to-point link.
C/C++	High-level programming language allowing program control of system hardware, and designed for code portability on all types of computers. C++ is a programming language that extends the Object Oriented capabilities of C.
CA	A Certificate Authority: a third-part which can attest that you are on record as the only person with the key associated with a personal digital signature.
CCITT	International Telegraph and Telephone Consultative Committee.
CHAP	Challenge Handshake Authentication Protocol: a commonly used protocol for link encryption. Authentication occurs at the data link layer and is transparent to end users.
CGI	Common Gateway Scripting: used to support two-way browser communication.
Digital Signature	Used to validate the identity of the file sender's e-mail.
DNS	Domain Name Server protocol, part of TCP/IP.
ECP	Encryption Control Protocol, part of the PPP suite.
Ethernet	LAN protocol as specified in the IEEE 802.3 standard.
Firewall	Collection of hardware and software that interconnects two or more networks and, at the same time, provides a central location for managing security.
FTP	The File Transfer Protocol: a standard protocol for transferring and copying files from one computer to another.
Kbps	1,024 bits per second.
Home Page	The first page that users see when they access a particular Web site.
HTML	HyperText Markup Language: used to describe the layout and contents of pages on the Web. The Hypertext Markup Language (HTML) is the language of the World Wide Web.
ICMP	Internet Control Message Protocol: defines the rules routers use to exchange routing information.

ICMPv6	Internet Control Message Protocol Version 6.
IGMP	The Internet Group Management Protocol is used by IP hosts to report host group clusters to neighboring multicast routers.
IGRP	Interior Gateway Routing Protocol, part of TCP/IP.
IPCP	IP Control Protocol, used to configure IP parameters on both ends of the PPP link.
IPX	Internetwork Packet Exchange: Novell's implementation of the Xerox Internet Datagram Protocol (IDP), used to define a set of rules for coordinating network communication between network components.
ISDN	Integrated Services Digital Network: provides digital communications circuit that allows transmission of voice, data, video, and graphics at very high speeds (from 56 Kbps to 128 Kbps), over standard communication lines.
ISP	Internet Service Providers act as middlemen, renting time to other users who want to access the Internet.
JAVA	A computer programming language that allows users to execute special programs (called applets) while accessing and viewing a Web page. Java is designed for creating animated Web sites.
LAN	Local Area Network.
LCP	Link Control Protocol: configures and tests the data link connection, and is part of the PPPsuite.
MARS	Multicast Address Resolution Server.
Mbps	1,000,000 bytes per second.
MPEG	Motion Picture Experts Group defining standards for handling video and audio compression.
OS/2	A high-performance, multi-tasking workstation operating system.
OSPF	Open Shortest Path First, link-state routing protocol used for IP routing.
PAP	Password Authentication Protocol: used for transparent session authentication occurring at the data link layer.
PGP	Pretty Good Privacy is a free (for personal use) e-mail security program developed in 1991 to support public-key encryption, digital signatures, and data compression. PGP is based on a 128-bit key.
PPP	Point-to-Point Protocol is one of the major protocols used to connect to the Internet. PPP is newer and faster than SLIP.
PKE	Public Key Encryption allows a sender to encrypt a document using a public key, which the recipient decodes using a private key.
POP3	Post Office Protocol version 3, allowing dynamic workstation access to mail drops on a server host.
PPTP	Point-to-Point Tunneling Protocol.
RIP	Routing Information Protocol: maintains network exchange and topology information.
RFC	Request for Comments: discussion notes, recommendations, and specifications for the Internet.
Router	This hardware device can be used to filter out data packets-based specific selection criteria. Thus, the router can allow certain packets into the network while rejecting others.
S-HTTP	Secure HTTP: a protocol developed by the CommerceNet coalition that operates at the level of the HTTP protocol.

SLIP	Serial Line Internet Protocol: one of the major protocols used to connect to the Internet. It predates PPP.
SMTP	Simple Mail Transfer Protocol: specifies the format and delivery handling of electronic messages.
SPX	The Sequenced Packet Exchange protocol defines a set of rules for coordinating network communication between network components.
SSI	SSI software is used by Web servers to display and/or capture dynamic (changing) information on an HTML page.
SSL	The Secure Socket Layer (SSL) was developed by Netscape Communications to encrypt TCP/IP communications between two host computers.
Subnet	Partition on an IP network based on class, involving use and movement of a subnet mask.
Surfing	Term used to describe accessing (through a Web browser) a chain of documents through a series of links on the Web.
T1	A leased line which can support transmission speed to 1.54 Mbps.
TCP/IP	The Transmission Control Protocol/Internet Protocol: specifies the rules for the exchange of information within the Internet or an Intranet, allowing packets from many different types of networks to be sent over the same network.
TCM	Total quality management involves creating systems and workflows that promote superior products and services.
Telnet	This service allows connection to a remote Internet host so that programs can be executed from a remote computer.
UDP	User Datagram Protocol provides message service for TCP/IP.
UNIX	UNIX is an operating system well suited to the Internet's open system model.
URL	Millions of documents that are distinguished by a unique name called a URL (Uniform Resource Locator), or more simply, a Web address. The URL is used by Web browsers use to access Internet information
UUCP	Unix to Unix Copy: allows files to be copied from the Unix system to another.
XML	Extended Markup Language: designed to provide a self-descriptive, platform-independent mechanism for exchanging management information between applications.
Web Browser	Allows you to traverse and view documents on the World Wide Web.
Windows NT	Provides a high-performance, multi-tasking workstation operating system.
WWW	The World Wide Web — or Web — is a collection of seamlessly interlinked documents that reside on Internet servers. The Web is so named because it links documents to form a web of information across computers worldwide.