COMPUTER NETWORKS LECTURES

DR. PROF. P. G. GYARMATI

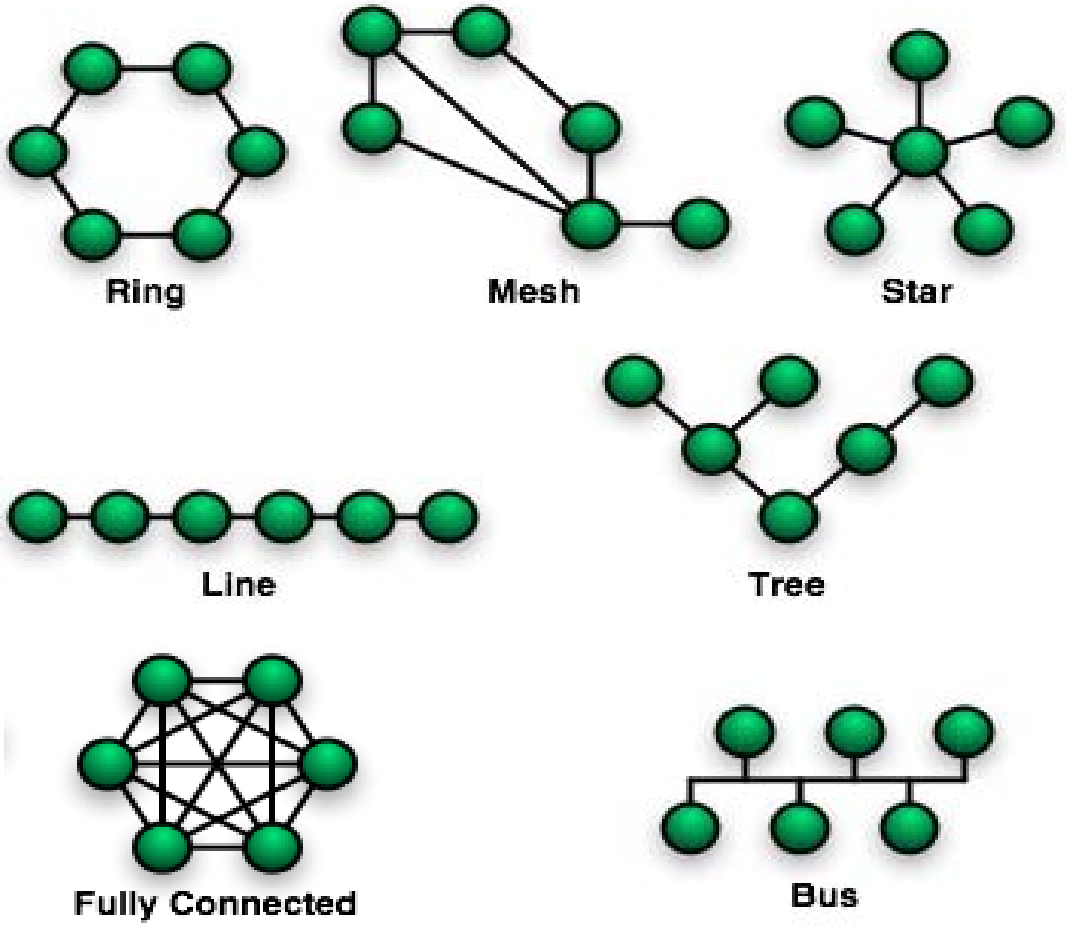

STANFORD UNIVERSITY
EE-CSL
1999-2000

COMPUTER NETWORKS
LECTURES

DR. PROF. P. G. GYARMATI
STANFORD UNIVERSITY
EE-CSL

1999.-2000.

Table of contents

Anatomy of a Network

1.1 The basic idea

Have you ever installed an Ethernet card and wondered how it all works. You install the hardware and then the driver and suddenly you are in touch with everybody on your network. How is it possible to send messages from your PC to a UNIX machine? What does the software need to do to work with the hardware? Put another way, where is the border between what the software does and what the hardware does? There must be rules laid down somewhere that define the requirements of the hardware and software to allow one computer to communicate with another computer of perhaps a different type. There are.

In order to allow computers on a network to communicate, the information has to get from the user's application to the network interface card and across the network into the user's application at the destination. This gets more complex when there may be computers of completely different architectures on the network, and even more complex when there are networks of different types connected together. To allow this, the process of passing information down from the user's application to the network interface card was broken down into a number of different steps called layers.
Various functions are built into each layer and each layer has a well-defined interface with clearly defined entry and exit points. This means that any of the layers can be modified without changing the behavior of the layer above or below it.
The layers are arranged in a stack as shown in Figure 11. Information originates at the top and is passed down through the stack until it reaches the network interface. It is then passed through the network and the process is reversed at the destination. The information is passed up through the stack until it reaches the application running on the destination machine.
The intention is that the lower layers should be transparent; the application layers of the client and server should appear to have a virtual connection.

1.2 The TCP/IP stack

The TCP/IP stack is detailed in Figure 1.1. and shows examples of applications that could be running at the top layer, such as a mailer application, a news reader application, a Telnet application, and a web application.

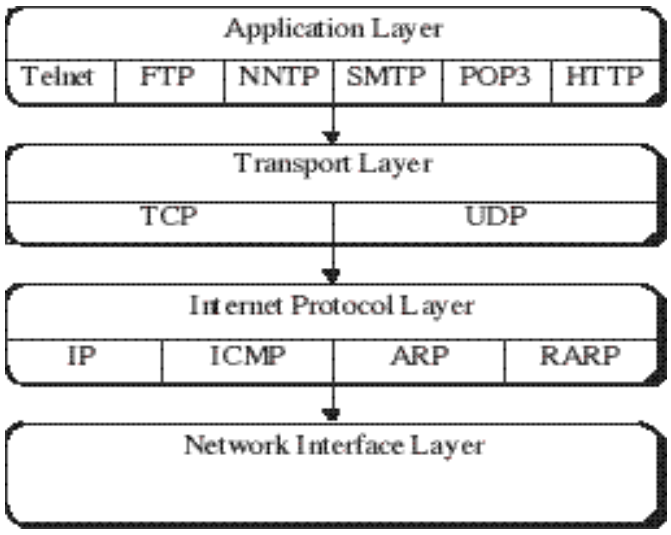

The TCP/IP stack

1.3 How It Works

When two machines are connected across the network, it will appear that the connection is between corresponding layers, as shown in Figure 1.2. An application running on one machine will employ the services of the lower layers of the stack to communicate with the application layer on the remote machine. At the top layer, the application will employ a protocol to arrange the sequence of data transfer; it will send data, wait for a response, then send new data depending on the response it received. The responses will come from the corresponding Application Layer at the other end. Similarly, the transport layer will set up the transport service.

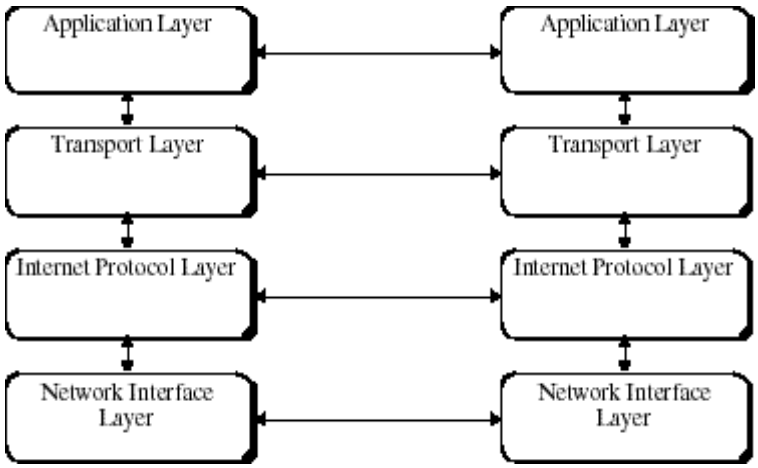

When two machines are connected

At each step along the way each layer will appear to be talking to its peer layer on the remote end. So what purpose do the various layers serve? The functions of the various layers are summarized below.

1.3.1 Application Layer

At the top is the user application. The application will be the front end; it could be, for example, a mail client, a web browser, or a newsreader. It could also be a proprietary client application transferring and retrieving information from a server over a network. The application layer will operate in accordance with a set of specifications called a protocol. Data will be passed to the corresponding layer at the remote end using the services of the Transport Layer in much the same way as two people talking on the phone. They converse as if they were physically talking to each other, when in fact there are some intermediaries in the way.
The protocol used by an application will dictate how the application establishes a connection with its peer and the sequence of data transfers. For example, on successful connection, a protocol implementation may offer a welcome message and ask for a user name. Subsequently it might ask for a password.

1.3.2 Transport Layer

The Transport Layer defines the type of transport service. The possibilities are either a connection-oriented service or a connectionless datagram-oriented service. With the first type, a connection is opened and maintained for the duration of the session, and then closed. The channel is bidirectional; data can be transmitted and received at the same time. Incoming packets are error-checked and resequenced as necessary. This type of service is referred to as a "reliable" connection. The other alternative is a datagram-oriented channel whereby each packet is dispatched on its way without establishing a dedicated channel. As a result, each packet may take a different route to the recipient so that some packets may arrive out of sequence or, indeed, not at all. This type of connection is sometimes called unreliable. So why would you ever use an "unreliable" transport service? If you intend sending binaries of any kind, then you should certainly think about using a reliable transport service; if, on the other hand, you are sending audio or video signals it may not be that important. Suppose you are trying to transmit video signals over a network; whether or not a few individual pixels arrive corrupted is relatively unimportant; what is more important is the speed with which the information reaches the remote end since any incorrect pixels will be updated moments later anyway. The overhead of error correction and resequencing is excessive and completely unnecessary in this case; speed is far more important.

1.3.3 Internet Protocol Layer

Next in line comes the Internet Protocol (IP) Layer. This layer provides, among other things, routing information and a best effort delivery service. Data is passed from the Transport Layer to the IP Layer in the form of IP Datagrams. The destination IP address is resolved into a physical address called the MAC (Media Access Control) address and the packet is dispatched. There is no error checking or resequencing at this layer; packets arriving at the remote end may take different routes and arrive out of sequence as a result. It is also possible for packets to be lost.

1.4 So What Does Winsock Offer?

INTRODUCTION

Network Protocols

TCP/IP

To provide the best network interoperability over WANs and network routers, include TCP/IP in your ideal configuration. The 32-bit TCP/IP stack included with Windows NT Workstation incorporates both SLIP and PPP. Microsoft TCP/IP, in combination with Windows NT Workstation, provides a scalable solution for enterprise networks that include a mix of system types.

When TCP/IP is used as the enterprise networking protocol, an IP addressing scheme is needed for your organization. If your servers run Windows NT Server, you can use Dynamic Host Control Protocol (DHCP) and Windows Internet Naming Service (WINS) for easy TCP/IP address management. With DHCP, administrators can centrally define global and subnet TCP/IP parameters for interconnected networks. The DHCP service dynamically configures the IP address and subnet mask of each workstation. For name resolution on TCP/IP internetworks, use WINS Servers or LMHosts files.

RFCs can be obtained by means of FTP from nis.nsf.net, nisc.jvnc.net, venera.isi.edu, wuarchive.wustl.edu, src.doc.ic.ac.uk, ftp.concert.net, ds.internic.net, or nic.ddn.mil.

Internetworking

Using TCP/IP for Scalability

TCP/IP delivers a scalable internetworking technology widely supported by hardware and software vendors. These solutions are:

- Wide area networks (WAN), TCP/IP printing, FTP, Telnet, DHCP, WINS, and DNS client software, Windows Sockets, and extended LMHOSTS file.
- DHCP Server, WINS Server, and DNS Server software.
- Microsoft Windows 95, with enhancements to support wide area networks (WAN), DHCP, WINS, and DNS client software, extended LMHOSTS file, and Windows Sockets.
- Microsoft TCP/IP-32 for Windows for Workgroups, with Windows Sockets support, can be used to provide access for Windows for Workgroups computers to Windows NT, LAN Manager, and other TCP/IP systems. Microsoft TCP/IP-32 includes DHCP, WINS and DNS client software.
- Microsoft LAN Manager — including both client and server support for Windows Sockets — and MS-DOS–based connectivity utilities. The Microsoft Network Client 2.0 software on the Windows NT Server compact disc includes new Microsoft TCP/IP support with DHCP and WINS clients.

As shown in Figure 30.2, the current version of TCP/IP for Windows NT also supports IP routing in systems with multiple network adapters attached to separate physical networks (*multihomed systems*).

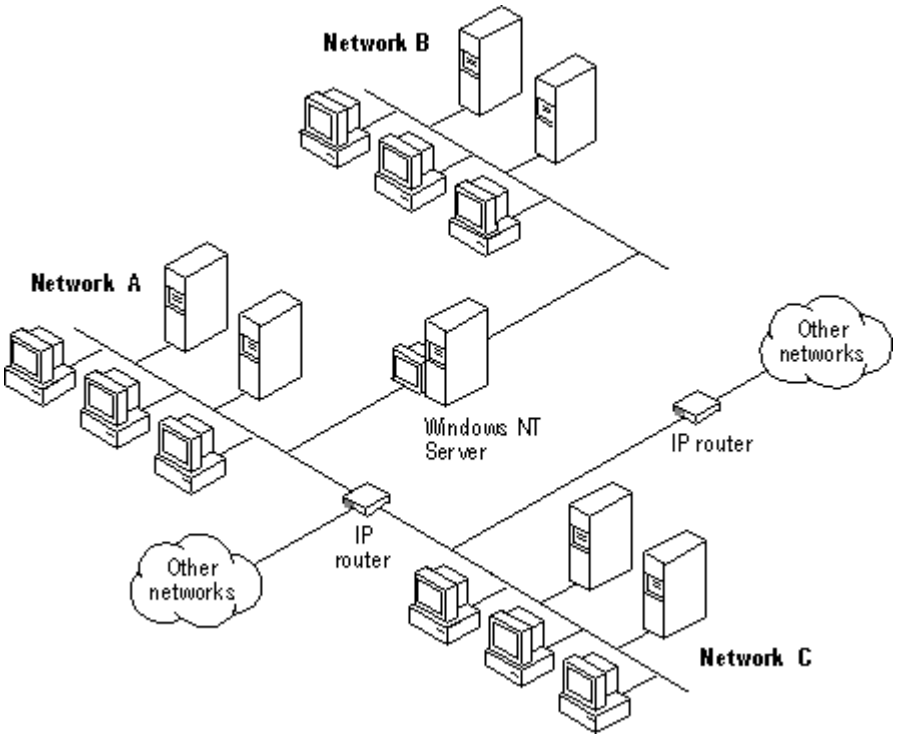

Figure 30.2 TCP/IP for Windows NT Supports IP Routing for Multihomed Systems

Benefits of Intelligent Network

As businesses increase their use of various forms of networking-the Internet, as well as corporate intranets and extranets-Intelligent Networks become more and more important. This enables two important benefits to be realized: differentiated services and more efficient sharing of networked resources between users and applications.

There are three main problems faced in building Intelligent Networks:

- Managing the increasing configuration complexity of devices
- Ensuring that consistent policies are applied to all elements of the system
- Enabling the needs of applications to be related to the services the network provides

This is shown conceptually in Figure 1.5.

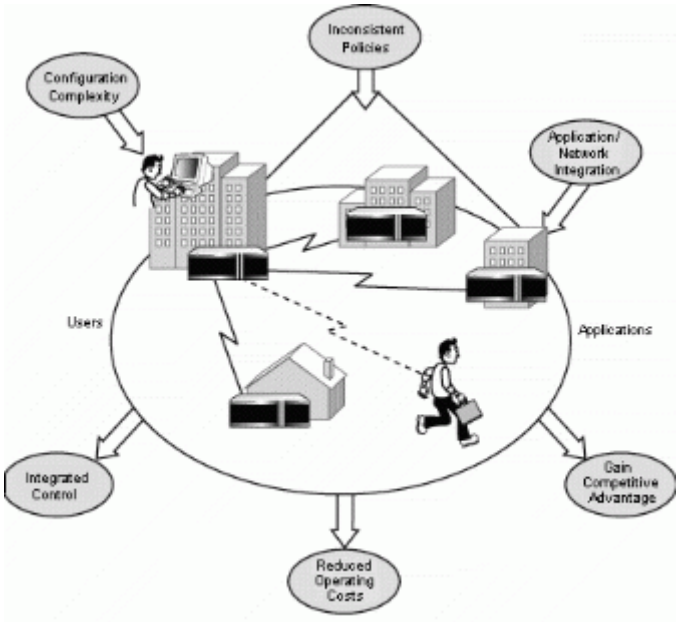

Figure 1.5: Problems the Intelligent Network solves, and derived benefits.

See full-sized image.

Key to all these is the capability to manage and maintain a comprehensive Information Model that enables each of these components to be related to each other. The Information Model presents a single mechanism to represent how network elements provide services, and how policies can be used to control them. It also links to existing models of users, applications, printers, and other resources, and associates their use with network services.

This results in a new philosophy that can be used to build Intelligent Networks: a service-oriented philosophy. Enterprises and service providers alike can use this.

Once connected, the client can send and receive packets over the Internet. The network access server uses the TCP/IP protocol for all traffic to the Internet.
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 virtual private networking (VPN) connection to a PPTP server on the private enterprise LAN, this is referred to as a *tunnel*. This is shown in the following figure:

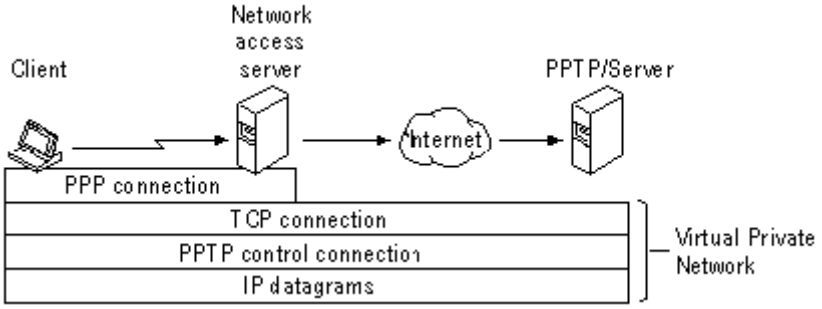

Figure 1: - The PPTP Tunnel

Tunneling is the process of sending packets to a computer on a private network by routing them over some other network, such as the Internet. The other network routers cannot access the computer that is on the private network. However, tunneling enables the routing network to transmit the packet to an intermediary computer, such as a PPTP server, that is connected to the both the routing network and the private network. Both the PPTP client and the PPTP server use tunneling to securely route packets to a computer on the private network by using routers that only know the address of the private network intermediary server.
When the PPTP server receives the packet from the routing network, it sends it across the private network to the destination computer. The PPTP server does this by processing the PPTP packet to obtain the private network computer name or address information in the encapsulated PPP packet. 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.
The following figure 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.

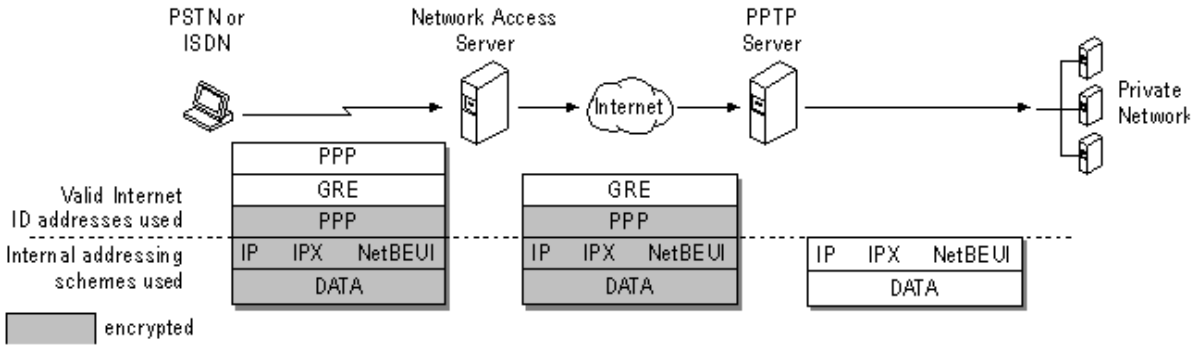

Figure 2: - Connecting a Dial-Up Networking PPTP Client to the Private Network

Note: An IPX internetwork path is recorded in a similar fashion to the MAC-sublayer routing information in a Token Ring source routing Explorer frame. However, unlike Token Ring source routing, the IPX internetwork path is not used in the subsequent communication. The IPX internetwork path is only used to prevent the broadcast packet from being forwarded on the same IPX network more than once.

Tunneling

Tunneling, also known as *encapsulation*, is a method of using an internetwork infrastructure of one protocol to transfer a payload, the frames (or packets) of another protocol (see Figure 8). Instead of sending the frame as it is produced by the originating host, the frame is encapsulated with an additional header. The additional header provides routing information so the encapsulated payload can traverse an intermediate internetwork (also known as a *transit internetwork*). The encapsulated packets are then routed between tunnel endpoints over the transit internetwork. Once the encapsulated payload packets reach their destination on the transit internetwork, the frame is de-encapsulated and forwarded to its final destination.
The entire process of encapsulation, transmission, and de-encapsulation of packets is tunneling. The logical path through which the encapsulated packets travel through the transit internetwork is called a *tunnel*.

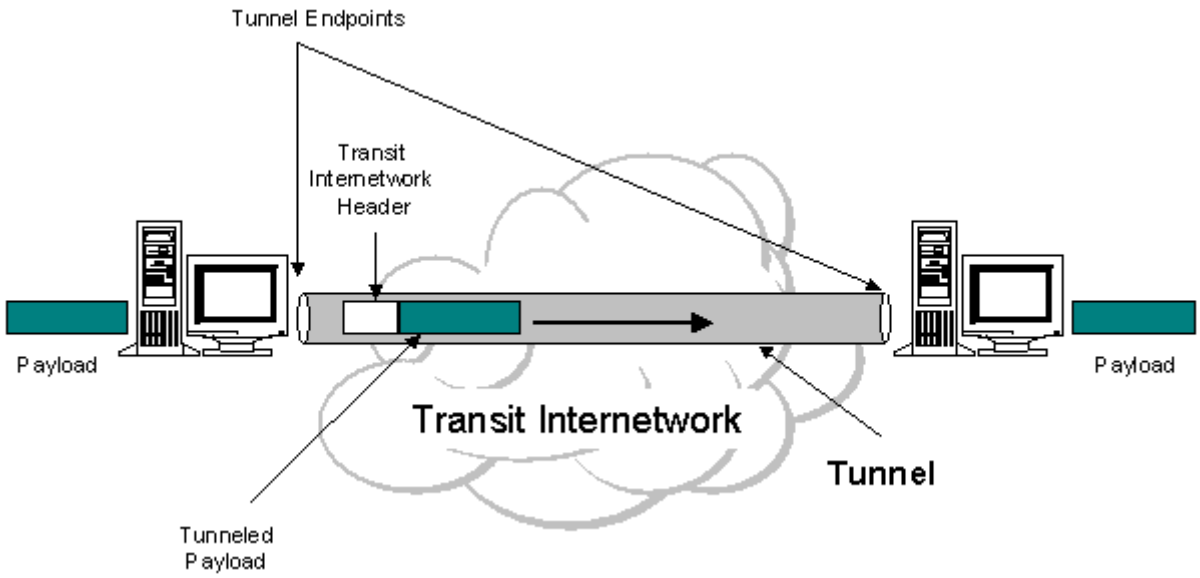

Figure 8: Tunneling.

The transit internetwork can be any internetwork. The Internet is a good example as the most widely known public internetwork. There are also many examples of tunnels that are carried over corporate internetworks.
Some common types of tunneling:

- **SNA Tunneling over IP Internetworks.** To send System Network Architecture (SNA) traffic across a corporate IP internetwork, the SNA frame is encapsulated with a User Datagram Protocol (UDP) and IP header. This is known as Data Link Switching (DLSw) and is described in RFC 1434.

- **IP Tunneling for Novell NetWare.** IPX packets are sent to a NetWare server or IPX router that wraps the IPX packet with a UDP and IP header and sends them across an IP internetwork.

User Manager for Domains A Windows NT Server tool used to manage security for a domain or an individual computer. User Manager for Domains administers user accounts, groups, and security policies.

user name A unique name identifying a user account to Windows NT. An account's user name cannot be identical to any other group name or user name of its own domain or workgroup. *See also* user account.

user password The password stored in each user's account. Each user generally has a unique user password and must type that password when logging on or accessing a server. *See also* password; volume password.

User privilege One of three privilege levels you can assign to a Windows NT user account. Every user account has one of the three privilege levels (Administrator, Guest, and User). Accounts with User privilege are regular users of the network; most accounts on your network will probably have User privilege. *See also* user account; Administrator privilege; Guest privilege.

user profile Configuration information that can be retained on a user-by-user basis, and is saved in user profiles. This information includes all the per-user settings of the Windows NT environment, such as the desktop arrangement, personal program groups and the program items in those groups, screen colors, screen savers, network connections, printer connections, mouse settings, window size and position, and more. When a user logs on, the user's profile is loaded and the user's Windows NT environment is configured according to that profile. *See also* personal groups; program item.

user rights Define a user's access to a computer or domain and the actions that a user can perform on the computer or domain. User rights permit actions such as logging onto a computer or network, adding or deleting users in a workstation or domain, and so forth.

user rights policy Manages the assignment of rights to groups and user accounts. *See also* user account; user rights.

users In the Macintosh environment, a special group that contains all users who have user permissions on the server. When a Macintosh user assigns permissions to everyone, those permissions are given to the groups users and guests. *See also* everyone category; guests.
Top of page

5.22 V

value entry The string of data that appears in the right pane of a registry window and which defines the value of the currently selected key. A value entry has three parts: name, datatype, and the value itself. *See also* key; subkey.

Van Jacobsen header compression A TCP/IP network layer compression technique, VJ compression reduces the size of IP and TCP headers. *See also* IP; TCP; TCP/IP.

variables In programming, a variable is a named storage location capable of containing a certain type of data that can be modified during program execution. System environment variables are