

COMPUTER NETWORKS

C-138

Self Learning Material



Directorate of Distance Education

**SWAMI VIVEKANAND SUBHARTI UNIVERSITY
MEERUT-250005
UTTAR PRADESH**

SIM Module Developed by :

Ramesh Bangia has been writing computer books for the one decade. He has written books on various topics of computers, mainly, related to the software. He has to his credit books pertaining to school levels, both middle and secondary level; books for Department of Electronics' O Level papers; books for various polytechnics; books for general software usage. Prior to this he edited a software magazine called Software Today.

Reviewed by the Study Material Assessment Committee Comprising:

1. **Dr. V.B. Sahai, Vice Chancellor**
2. **Dr. G.S. Bhatnagar, Pro-Vice Chancellor**
3. **Dr. Mohan Gupta**
4. **Mr. Sumit Bhatnagar**
5. **Mr. Rohit Kumar**

Copyright © Laxmi Publications Pvt Ltd

No part of this publication which is material protected by this copyright notice may be reproduced or transmitted or utilized or stored in any form or by any means now known or hereinafter invented, electronic, digital or mechanical, including photocopying, scanning, recording or by any information storage or retrieval system, without prior permission from the publisher.

Information contained in this book has been published by Laxmi Publications Pvt Ltd and has been obtained by its authors from sources believed to be reliable and are correct to the best of their knowledge. However, the publisher and its author shall in no event be liable for any errors, omissions or damages arising out of use of this information and specially disclaim and implied warranties or merchantability or fitness for any particular use.

Published by : Laxmi Publications Pvt Ltd., 113, Golden House, Daryaganj, New Delhi-110 002.
Tel: 43532500, E-mail: info@laxmipublications.com

CONTENTS

Units	Page No.
1. Introduction to Computer Networks	1-54
2. Introduction to IPV6	55-76
3. Mobility in Network	77-102
4. TCP Extensions	103-130
5. Network Security	131-162

SYLLABUS

COMPUTER NETWORKS

C-138

Unit 1:

Introduction: Overview of computer network, seven-layer architecture, TCP/IP suite of protocol, etc. Mac protocols for high-speed LANs, MANs & WIRELESS LANs. (For example, FDDI, DQDB, HIPPI, Gigabit Ethernet, Wireless Ethernet etc.)
Fast access technologies. (For example, ADSL, Cable Modem, etc.)

Unit 2:

IPv6: why IPv6, basic protocol, extension & option, support for QoS, security, etc, neighbor discovery, autoconfiguration, routing. Change to other protocols. Application programming interface for Ipv6. 6bone.

Unit 3:

Mobility in network. Mobile. Security related issues. IP Multicasting. Multicasting routing protocols, address assignments, session discovery, etc.

Unit 4:

TCP extensions for high-speed networks, transaction-oriented application, other new option in TCP.

Unit 5:

Network security at various layers. Secure-HTTP, SSL, ESP, Authentication header, Key distribution protocols. Digital signatures, digital certificates.

UNIT 1

NOTES

INTRODUCTION TO COMPUTER NETWORKS

STRUCTURE

- 1.1 Overview of Computer Networks
- 1.2 Seven Layer Architecture
- 1.3 TCP/IP suite of Protocol
- 1.4 Mac protocols for high speed
- 1.5 LANs
- 1.6 MANs
- 1.7 Wireless LANs
- 1.8 FDDI
- 1.9 DQDB
- 1.10 HIPPI
- 1.11 Gigabit Ethernet
- 1.12 Wireless Ethernet
- 1.13 Fast access technologies
- 1.14 ADSL
- 1.15 Cable Modem
 - Summary
 - Self Assessment Questions
 - Further Readings

Learning Objectives

After going through this unit, students will be able to:

- know about the seven layers of computer architecture
- know about the various networks like LAN, MAN and WAN.
- learn about Cable Modem.
- learn about the fast access technologies
- learn about Gigabit Ethernet

1.1 OVERVIEW OF COMPUTER NETWORKS

When more than two computers are connected to each other and sharing information, resources and remote systems then this is called Networking.

NOTES

Technical definitions:

"A network of data processing nodes that are interconnected for the purpose of data communication".

"An interconnection of three or more communicating entities".

1.1.1 Classification of Computer Networks

1.1.1.1 Network Layer

In this computer networks follow the industry standards of OSI reference model and TCP/IP model. Where as OSI is of seven layers and TCP/IP is defined in five layers.

1.1.1.2 Scale

It can be classified as:

- Local area network (LAN)
- Campus area network (CAN)
- Metropolitan area network (MAN)
- Wide area network (WAN).
- Personal area network (PAN)

1.1.1.3 Connection Method

The connection methods available are:

- Ethernet
- Power line communication
- Wireless LAN
- HomePNA

1.1.1.4 Functional Relationship

This exists between the network elements:

- Peer-to-peer
- Client-server

1.1.1.5 Network Topology

It is a logical layouts of the network. Topologies are:

- Star network
- Ring network
- Bus network

- Tree network
- Star-bus network
- Mesh network

1.1.1.6 Services

It provides following services:

- Wireless community network
- Server
- Storage area networks
- Process control networks
- Value-added network

1.1.1.7 Protocol

On network protocols are used as communication language. Several types of protocols are available:

- TCP/IP
- Network IPX/SPX

1.1.2 Types of Networks


Following is the list of the most common types of computer networks in order of scale.

1.1.2.1 Personal Area Network (PAN)

PAN is used for communication among the personal devices (intrapersonal communication), or for connecting to a higher level network and Internet. The reach of a PAN is typically a few meters.

A personal area network (PAN) is used for communication among computer devices. For example,

- Telephones



NOTES

- Personal digital assistants

Personal area networks may be wired with computer buses such as USB and FireWire. A wireless personal area network (WPAN) can also be made possible with network technologies such as IrDA and Bluetooth.

NOTES

1.1.2.2 Local Area Network (LAN)

A network covering a small geographic area, like a home, office, or building. Current LANs are most likely to be based on Ethernet technology. The defining characteristics of LANs, in contrast to WANs (wide area networks), include their much higher data transfer rates, smaller geographic range, and lack of a need for leased telecommunication lines.

1.1.2.3 Campus Area Network (CAN)

A network that connects two or more LANs but that is limited to a specific (possibly private) geographical area such as a college campus, industrial complex, or a military base. A CAN, may be considered a type of MAN (metropolitan area network), but is generally limited to an area that is smaller than a typical MAN.

1.1.2.4 Metropolitan Area Network (MAN)

A network that connects two or more Local Area Networks or CANs together but does not extend beyond the boundaries of the immediate town, city, or metropolitan area. Multiple routers, switches & hubs are connected to create a MAN

1.1.2.5 Wide Area Network (WAN)

A WAN is a data communications network that covers a relatively broad geographic area and that often uses transmission facilities provided by common carriers, such as telephone companies. WAN technologies generally function at the lower three layers of the OSI reference model: the physical layer, the data link layer, and the network layer.

1.1.2.6 Internetwork

Two or more networks or network segments connected using devices that operate at layer 3 (the 'network' layer) of the OSI Basic Reference Model, such as a router. Any interconnection among or between public, private, commercial, industrial, or governmental networks may also be defined as an internetwork.

1.1.2.7 Internet

A specific internetwork, consisting of a worldwide interconnection of governmental, academic, public, and private networks based upon the Advanced Research Projects Agency Network (ARPANET) developed by ARPA of the U.S. Department of Defense – also home to the World Wide Web (WWW) and referred to as the 'Internet' with a capital 'I' to distinguish it from other generic internetworks.

1.1.2.8 Extranet

A network or internetwork that is limited in scope to a single organization or entity but which also has limited connections to the networks of one or more other usually, but not necessarily, trusted organizations or entities (e.g., a company's customers may be provided access to some part of its intranet thusly creating an extranet while at the same time the customers may not be considered 'trusted' from a security standpoint). Technically, an extranet may also be categorized as a CAN, MAN, WAN,

or other type of network, although, by definition, an extranet cannot consist of a single LAN, because an extranet must have at least one connection with an outside network.

Intranets and extranets may or may not have connections to the Internet.

If connected to the Internet, the intranet or extranet is normally protected from being accessed from the Internet without proper authorization. The Internet itself is not considered to be a part of the intranet or extranet, although the Internet may serve as a portal for access to portions of an extranet.

1.1.3 Basic Hardware Components

All networks are made up of basic hardware building blocks to interconnect network nodes, such as Network Interface Cards (NICs), Bridges, Hubs, Switches, and Routers. In addition, some method of connecting these building blocks is required, usually in the form of galvanic cable (most commonly Category 5 cable). Less common are microwave links (as in IEEE 802.11) or optical cable ("optical fiber").

1.1.3.1 Network Interface Cards

A **network card**, **network adapter** or **NIC** (network interface card) is a piece of computer hardware designed to allow computers to communicate over a **computer network**. It provides physical access to a networking medium and provides a low-level addressing system through the use of MAC addresses. It allows users to connect to each other either by using cables or wirelessly.

1.1.3.2 Bridges

A **network bridge** connects multiple network segments at the data link layer (layer 2) of the OSI model. Bridges are similar to repeaters or network hubs, devices that connect network segments at the physical layer, however a bridge works by using bridging where traffic from one network is managed rather than simply rebroadcast to adjacent network segments.

1.1.3.3 Hubs

A **hub** is a piece of hardware which provides the connectivity of a segment of a network by directing traffic through the network. It does this in a rudimentary way, it simply copies the data to all of the Nodes connected to the hub. Hubs are commonly used to connect segments of a LAN. A hub contains multiple ports. When a packet arrives at one port, it is copied to the other ports so that all segments of the LAN can see all packets.

1.1.3.4 Switches

Switches are the device of networking that directs traffic to the correct node by filtering and forwarding packets between Nodes. Switches operate at the data link layer (layer 2) and sometimes the network layer (layer 3) of the OSI Reference Model and therefore support any packet protocol. LANs that use switches to join segments are called switched LANs or, in the case of Ethernet networks, switched Ethernet LANs. In a circuit-switched data network, a switch is used to create a virtual circuit between the pairs of endpoints. This means that it creates a path to the destination node from the source node.

NOTES

1.1.3.5 Routers

Routers are the networking device that forwards data packets along networks by using headers and forwarding tables to determine the best path to forward the packets. Routers also provide interconnectivity between like and unlike devices on the network. This is accomplished by examining the Header of a data packet. They use protocols such as ICMP to communicate with each other and configure the best route between any two hosts. A router is connected to at least two networks, commonly two LANs or WANs or a LAN and its ISP's network.

Routers are usually located at gateways, the places where two or more networks connect. Many household DSL and Cable Modems are also routers.

1.1.4 Building a Simple Computer Network

A simple computer network may be constructed from two computers by adding a network adapter (Network Interface Controller (NIC)) to each computer and then connecting them together with a special cable called a crossover cable. This type of network is useful for transferring information between two computers that are not normally connected to each other by a permanent network connection or for basic home networking applications. Alternatively, a network between two computers can be established without dedicated extra hardware by using a standard connection such as the RS-232 serial port on both computers, connecting them to each other via a special crosslinked *null modem* cable.

Practical networks generally consist of more than two interconnected computers and generally require special devices in addition to the Network Interface Controller that each computer needs to be equipped with. Examples of some of these special devices are hubs, switches and routers.

1.2 SEVEN LAYER ARCHITECTURE

Ans 2
In 1977, the International Organization for Standardization (ISO), began to develop its OSI networking suite. OSI has two major components: an abstract model of networking (the Basic Reference Model, or seven-layer model), and a set of concrete protocols. The standard documents that describe OSI are for sale and not currently available online.

Parts of OSI have influenced Internet protocol development, but none more than the abstract model itself, documented in ISO 7498 and its various addenda. In this model, a networking system is divided into layers. Within each layer, one or more entities implement its functionality. Each entity interacts directly only with the layer immediately beneath it, and provides facilities for use by the layer above it.

In particular, Internet protocols are deliberately not as rigorously architected as the OSI model, but a common version of the TCP/IP model splits it into four layers. The Internet Application Layer includes the OSI Application Layer, Presentation Layer, and most of the Session Layer. Its End-to-End Layer includes the graceful close function of the OSI Session Layer as well as the Transport Layer. Its Internetwork Layer is equivalent to the OSI Network Layer, while its Interface layer includes the OSI Data Link and Physical Layers. These comparisons are based on the original seven-layer protocol model as defined in ISO 7498, rather than refinements in such things as the Internal Organization of the Network Layer document.

Ans 1

Protocols enable an entity in one host to interact with a corresponding entity at the same layer in a remote host. Service definitions abstractly describe the functionality provided to a (N)-layer by an (N-1) layer, where N is one of the seven layers inside the local host.

1.2.1 Layer 7: Application Layer

The application layer is the seventh level of the seven-layer OSI model. It interfaces directly to and performs common application services for the application processes; it also issues requests to the presentation layer. Note carefully that this layer provides services to user-defined application processes, and not to the end user. For example, it defines a file transfer protocol, but the end user must go through an application process to invoke file transfer. The OSI model does not include human interfaces.

The common application services sublayer provides functional elements including the Remote Operations Service Element (comparable to Internet Remote Procedure Call), Association Control, and Transaction Processing (according to the ACID requirements).

Above the common application service sublayer are functions meaningful to user application programs, such as messaging (X.400), directory (X.500), file transfer (FTAM), virtual terminal (VTAM), and batch job manipulation (JTAM).

1.2.2 Layer 6: Presentation Layer

The Presentation layer transforms the data to provide a standard interface for the Application layer. MIME encoding, data encryption and similar manipulation of the presentation are done at this layer to present the data as a service or protocol developer sees fit. Examples of this layer are converting an EBCDIC-coded text file to an ASCII-coded file, or serializing objects and other data structures into and out of XML.

1.2.3 Layer 5: Session Layer

The Session layer controls the dialogues/connections (sessions) between computers. It establishes, manages and terminates the connections between the local and remote application. It provides for either full-duplex or half-duplex operation, and establishes checkpointing, adjournment, termination, and restart procedures. The OSI model made this layer responsible for "graceful close" of sessions, which is a property of TCP, and also for session checkpointing and recovery, which is not usually used in the Internet protocols suite. Session layers are commonly used in application environments that make use of remote procedure calls (RPCs).

iSCSI, which implements the Small Computer Systems Interface (SCSI) encapsulated into TCP/IP packets, is a session layer protocol increasingly used in Storage Area Networks and internally between processors and high-performance storage devices. iSCSI leverages TCP for guaranteed delivery, and carries SCSI command descriptor blocks (CDB) as payload to create a virtual SCSI bus between iSCSI initiators and iSCSI targets.

1.2.4 Layer 4: Transport Layer

The Transport layer provides transparent transfer of data between end users, providing reliable data transfer services to the upper layers. The transport layer controls the reliability of a given link through flow control, segmentation/desegmentation, and

NOTES

error control. Some protocols are state and connection oriented. This means that the transport layer can keep track of the segments and retransmit those that fail.

The best known example of a layer 4 protocol is the Transmission Control Protocol (TCP).

NOTES

The transport layer is the layer that converts messages into TCP segments or User Datagram Protocol (UDP), Stream Control Transmission Protocol (SCTP), etc. packets.

Perhaps an easy way to visualize the Transport Layer is to compare it with a Post Office, which deals with the dispatch and classification of mail and parcels sent. Do remember, however, that a post office manages the outer envelope of mail. Higher layers may have the equivalent of double envelopes, such as cryptographic Presentation services that can be read by the addressee only.

Roughly speaking, tunneling protocols operate at the transport layer, such as carrying non-IP protocols such as IBM's SNA or Novell's IPX over an IP network, or end-to-end encryption with IPsec. While Generic Routing Encapsulation (GRE) might seem to be a network layer protocol, if the encapsulation of the payload takes place only at endpoint, GRE becomes closer to a transport protocol that uses IP headers but contains complete frames or packets to deliver to an endpoint. L2TP carries PPP frames inside transport packets.

1.2.5 Layer 3: Network Layer

The Network layer provides the functional and procedural means of transferring variable length data sequences from a source to a destination via one or more networks while maintaining the quality of service requested by the Transport layer. The Network layer performs network routing functions, and might also perform fragmentation and reassembly, and report delivery errors. Routers operate at this layer—sending data throughout the extended network and making the Internet possible. This is a logical addressing scheme – values are chosen by the network engineer.

The addressing scheme is hierarchical. The best known example of a layer 3 protocol is the Internet Protocol (IP).

Perhaps it's easier to visualize this layer as managing the sequence of human carriers taking a letter from the sender to the local post office, trucks that carry sacks of mail to other post offices or airports, airplanes that carry airmail between major cities, trucks that distribute mail sacks in a city, and carriers that take a letter to its destinations. Think of fragmentation as splitting a large document into smaller envelopes for shipping, or, in the case of the network layer, splitting an application or transport record into packets.

1.2.6 Layer 2: Data Link Layer

The Data Link layer provides the functional and procedural means to transfer data between network entities and to detect and possibly correct errors that may occur in the Physical layer. The best known example of this is Ethernet. This layer manages the interaction of devices with a shared medium. Other examples of data link protocols are HDLC and ADCCP for point-to-point or packet-switched networks and Aloha for local area networks. On IEEE 802 local area networks, and some non-IEEE 802

networks such as FDDI, this layer may be split into a Media Access Control (MAC) layer and the IEEE 802.2 Logical Link Control (LLC) layer. It arranges bits from the physical layer into logical chunks of data, known as frames.

This is the layer at which the bridges and switches operate. Connectivity is provided only among locally attached network nodes forming layer 2 domains for unicast or broadcast forwarding. Other protocols may be imposed on the data frames to create tunnels and logically separated layer 2-forwarding domain.

The data link layer might implement a sliding window flow control and acknowledgment mechanism to provide reliable delivery of frames; that is the case for SDLC and HDLC, and derivatives of HDLC such as LAPB and LAPD. In modern practice, only error detection, not flow control using sliding window, is present in modern data link protocols such as Point-to-Point Protocol (PPP), and, on local area networks, the IEEE 802.2 LLC layer is not used for most protocols on Ethernet, and, on other local area networks, its flow control and acknowledgment mechanisms are rarely used. Sliding window flow control and acknowledgment is used at the transport layers by protocols such as TCP.

1.2.7 Layer 1: Physical Layer

The Physical layer defines all the electrical and physical specifications for devices. In particular, it defines the relationship between a device and a physical medium. This includes the layout of pins, voltages, and cable specifications. Hubs, repeaters, network adapters and Host Bus Adapters (HBAs used in Storage Area Networks) are physical-layer devices.

To understand the function of the physical layer in contrast to the functions of the data link layer, think of the physical layer as concerned primarily with the interaction of a single device with a medium, where the data link layer is concerned more with the interactions of multiple devices (i.e., at least two) with a shared medium. The physical layer will tell one device how to transmit to the medium, and another device how to receive from it, but not, with modern protocols, how to gain access to the medium. Obsolescent physical layer standards such as RS-232 do use physical wires to control access to the medium.

The major functions and services performed by the physical layer are:

- Establishment and termination of a connection to a communications medium.
- Participation in the process whereby the communication resources are effectively shared among multiple users. For example, contention resolution and flow control.
- Modulation, or conversion between the representation of digital data in user equipment and the corresponding signals transmitted over a communications channel. These are signals operating over the physical cabling (such as copper and fiber optic) or over-a radio link.

Parallel SCSI buses operate in this layer, although it must be remembered that the logical SCSI protocol is a transport-layer protocol that runs over this bus. Various physical-layer Ethernet standards are also in this layer; Ethernet incorporates both this layer and the data-link layer. The same applies to other local-area networks, such as Token ring, FDDI, and IEEE 802.11, as well as personal area networks such as Bluetooth and IEEE 802.15.4.

NOTES

1.3 TCP/IP SUITE OF PROTOCOL

The Internet Protocol Suite is the set of communications protocols used for the Internet and other similar networks. It is commonly also known as TCP/IP, named from two of the most important protocols in it: the Transmission Control Protocol (TCP) and the Internet Protocol (IP), which were the first two networking protocols defined in this standard. Modern IP networking represents a synthesis of several developments that began to evolve in the 1960s and 1970s, namely the Internet and local area networks, which emerged during the 1980s, together with the advent of the World Wide Web in the early 1990s.

The Internet Protocol Suite, like many protocol suites, is constructed as a set of layers. Each layer solves a set of problems involving the transmission of data. In particular, the layers define the operational scope of the protocols within.

Often a component of a layer provides a well-defined service to the upper layer protocols and may be using services from the lower layers. Upper layers are logically closer to the user and deal with more abstract data, relying on lower layer protocols to translate data into forms that can eventually be physically transmitted.

The TCP/IP model consists of (RFC 1122). From lowest to highest, these are the Link Layer, the Internet Layer, the Transport Layer, and the Application Layer.

History

The Internet Protocol Suite resulted from research and development conducted by the Defense Advanced Research Projects Agency (DARPA) in the early 1970s. After initiating the pioneering ARPANET in 1969, DARPA started work on a number of other data transmission technologies. In 1972, Robert E. Kahn joined the DARPA Information Processing Technology Office, where he worked on both satellite packet networks and ground-based radio packet networks, and recognized the value of being able to communicate across both. In the spring of 1973, Vinton Cerf, the developer of the existing ARPANET Network Control Program (NCP) protocol, joined Kahn to work on open-architecture interconnection models with the goal of designing the next protocol generation for the ARPANET.

By the summer of 1973, Kahn and Cerf had worked out a fundamental reformulation, where the differences between network protocols were hidden by using a common internetwork protocol, and, instead of the network being responsible for reliability, as in the ARPANET, the hosts became responsible. Cerf credits Hubert Zimmerman and Louis Pouzin, designer of the CYCLADES network, with important influences on this design.

The design of the network included the recognition that it should provide only the functions of efficiently transmitting and routing traffic between end nodes and that all other intelligence should be located at the edge of the network, in the end nodes. Using a simple design, it became possible to connect almost any network to the ARPANET, irrespective of their local characteristics, thereby solving Kahn's initial problem. One popular expression is that TCP/IP, the eventual product of Cerf and Kahn's work, will run over "two tin cans and a string."

A computer called a router (a name changed from gateway to avoid confusion with other types of gateways) is provided with an interface to each network, and forwards

packets back and forth between them. Requirements for routers are defined in (Request for Comments 1812).

The idea was worked out in more detailed form by Cerf's networking research group at Stanford in the 1973-74 period, resulting in the first TCP specification. The early networking work at Xerox PARC, which produced the PARC Universal Packet protocol suite, much of which existed around the same period of time, was also a significant technical influence; people moved between the two.

DARPA then contracted with BBN Technologies, Stanford University, and the University College London to develop operational versions of the protocol on different hardware platforms. Four versions were developed: TCP v1, TCP v2, a split into TCP v3 and IP v3 in the spring of 1978, and then stability with TCP/IP v4 — the standard protocol still in use on the Internet today.

In 1975, a two-network TCP/IP communications test was performed between Stanford and University College London (UCL). In November, 1977, a three-network TCP/IP test was conducted between sites in the US, UK, and Norway. Several other TCP/IP prototypes were developed at multiple research centres between 1978 and 1983. The migration of the ARPANET to TCP/IP was officially completed on January 1, 1983, when the new protocols were permanently activated.

In March 1982, the US Department of Defense declared TCP/IP as the standard for all military computer networking. In 1985, the Internet Architecture Board held a three day workshop on TCP/IP for the computer industry, attended by 250 vendor representatives, promoting the protocol and leading to its increasing commercial use.

The OSI model is a good place to start to learn more about the differences between physical and logical addressing. Think of the physical address as the 48-bit MAC address that manufacturers encode in their network interface cards (NICs). This type of address is unique, referred to as the Ethernet or hardware address, and cannot be changed but can be spoofed. The MAC or Ethernet address is associated with Layer 2 (data Link) of the OSI Model. The logical address is a 32-bit IP address that is not embedded in the network card but it is assigned to it for the purpose of routing between networks. This type of address operates at Layer 3 (network) of the OSI Model.

The Internet Protocol (IP), in combination with Transmission Control Protocol (TCP), forms the TCP/IP suite, which is the de facto protocol (i.e., universal computer language) that connects the network of networks — that is, the Internet.

The OSI Model is a standard developed by the International Standards Organization (OSI) to provide a blueprint for conformity for software development and network communications.

1.3.1 Different classes of IP addressing special IP address

Classful networking is the name given to the first round of changes to the structure of the IP address in IPv4. Classful networking is obsolete on the modern Internet. There is no longer any such thing as a class A/B/C network. The correct modern representation for what would have been referred to as a "Class B" prior to 1993 would be "a set of /16 addresses", under the Classless Inter-Domain Routing (CIDR) system.

NOTES

STUDENT ACTIVITY 1.1

1. What is a computer network?

2. What is Local Area Networks?

3. Describe the various layers.

4. Define the various types of networks.

NOTES

as a multiple access protocol. This makes it possible for several stations connected to the same physical medium to share it. Examples of shared physical media are bus networks, ring networks, hub networks, wireless networks and half-duplex point-to-point links. The multiple access protocol may detect or avoid data packet collisions if a packet mode contention based channel access method is used, or reserve resources to establish a logical channel if a circuit switched or channelization based channel access method is used. The channel access control mechanism relies on a physical layer multiplex scheme.

The most widespread multiple access protocol is the contention based CSMA/CD protocol used in Ethernet networks. This mechanism is only utilized within a network collision domain, for example an Ethernet bus network or a hub network. An Ethernet network may be divided into several collision domains, interconnected by bridges and switches.

A multiple access protocol is not required in a switched full-duplex network, such as today's switched Ethernet networks, but is often available in the equipment for compatibility reasons.

In computing, a **protocol** is a convention or standard that controls or enables the connection, communication, and data transfer between two computing endpoints. In its simplest form, a protocol can be defined as **the rules governing** the syntax, semantics, and synchronization of communication. Protocols may be implemented by hardware, software, or a combination of the two. At the lowest level, a protocol defines the behaviour of a hardware connection.

A protocol is a convention or standard that controls or enables the connection, communication, and data transfer between two computing endpoints.

1.4.1 Typical properties

It is difficult to generalize about protocols because they vary so greatly in purpose and sophistication. Most protocols specify one or more of the following properties:

- Detection of the underlying physical connection (wired or wireless), or the existence of the other endpoint or node
- Handshaking
- Negotiation of various connection characteristics
- How to start and end a message
- How to format a message
- What to do with corrupted or improperly formatted messages (error correction)
- How to detect unexpected loss of the connection, and what to do next
- Termination of the session or connection.

1.4.2 Importance

The widespread use and expansion of communications protocols is both a prerequisite to the Internet, and a major contributor to its power and success. The pair of Internet Protocol (or IP) and Transmission Control Protocol (or TCP) are the most important of these, and the term TCP/IP refers to a collection (or protocol suite) of its most

used protocols. Most of the Internet's communication protocols are described in the RFC documents of the Internet Engineering Task Force (or IETF).

Object-oriented programming has extended the use of the term to include the programming protocols available for connections and communication between objects.

The pair of Internet Protocol (or IP) and Transmission Control Protocol (or TCP) are the most important of these, and the term TCP/IP refers to a collection (or protocol suite) of its most used protocols.

NOTES

Generally, only the simplest protocols are used alone. Most protocols, especially in the context of communications or networking, are layered together into protocol stacks where the various tasks listed above are divided among different protocols in the stack.

Whereas the protocol stack denotes a specific combination of protocols that work together, a reference model is a software architecture that lists each layer and the services each should offer. The classic seven-layer reference model is the OSI model, which is used for conceptualizing protocol stacks and peer entities. This reference model also provides an opportunity to teach more general software engineering concepts like hiding, modularity, and delegation of tasks. This model has endured in spite of the demise of many of its protocols (and protocol stacks) originally sanctioned by the ISO. The *OSI model* is not the only reference model, however.

1.4.3 Common Protocols

- IP (Internet Protocol)
- DHCP (Dynamic Host Configuration Protocol)
- TCP (Transmission Control Protocol)
- HTTP (Hypertext Transfer Protocol)
- FTP (File Transfer Protocol)
- POP3 (Post Office Protocol 3)
- SMTP (Simple Mail Transfer Protocol)
- IMAP (Internet Message Access Protocol)

1.5 LANS

A local area network (LAN) is two or more computers directly linked within a small well-defined area such as a room, building, or group of closely placed buildings. A LAN may be made up of only microcomputers or any combination of microcomputers and large systems.

A LAN usually consists of the following:

1. Two or more computers
2. Peripheral devices such as printers and hard-disk drives
3. Software to control the operation of the computers or other devices connected to the LAN

4. Special cables, usually coaxial or fibre optic, to connect the computers and other devices
5. A plug-in board to handle the data transmissions.

Some computers require that all the computers be of a certain brand, while others allow a variety of brands to be connected. The number of computers in a LAN varies from smaller LANs that typically connect 2 to 25 computers, to large LANs that can connect as many as 10,000 computers.

The length of the cable connecting a computer to a LAN also varies depending on the LAN. Most LANs allow cables of about 1,000 feet, but some allow cables of several miles to be used. The data transfer speeds range from several thousand bits per second to around 10 million bits per second.

The programs that control the LANs also vary in the features they offer.

Some programs allow the user of more than one operating system; others allow only one. On some LANs file access is limited to one user at a time; on others, more than one user can access a file simultaneously.

1.5.1 Hardware for LAN

The following are the major hardware components/devices for establishing LAN:

1. Transmission Channel
2. Network Interface Unit or NIU
3. Servers
4. Workstations

1.5.2 Server and Workstations

One of the major benefits of implementation of LAN is sharing expensive resources such as storage devices, printers, etc. This is achieved through providing servers on the LAN. It is a dedicated computer that controls one or more resources. This contains both hardware and software interface for LAN. Three major categories of services used in LANs are:

1. File Server
2. Printer Server
3. Modem Server

In networking, file server is used to share storage for files. Besides providing storage space for files in a LAN environment, it is used for taking periodical backup, and also to provide gateway to other servers with and between LANs. Similarly printer server is used to handle printing works of all workstation connected in the network. In LAN environment also modem is required to get connected to other network or simply to use a telephone. A modem server is used to share this expensive resource by all connected workstations in a network ring.

1.5.3 Software for LAN

LAN operating system is required to operate on the LAN system. It has basically two aspects:

1. Server Software
2. Workstation Software

LAN operating system facilitates :

1. Sharing of expensive resources, e.g., printer, storage space, etc.
2. Security for data
3. Connection to other network

There are various types of LAN operating system. Some popular LAN operating systems are :

- Novel Netware
- Ethernet
- Corvus
- ArcNet
- LAN Server
- Omni Net
- PC Net
- IBM PC LAN
- Etherlink Plus, etc.

1.5.4 Features of Local Area Networks

- Typically connects computer in a single building or campus.
- Developed in 1970s.
- Medium : optical fibres, coaxial cables, twisted pair, wireless.
- Low latency (except in high traffic periods).
- High speed networks (0.2 to 100-Mb/sec).
- Speeds adequate for most distributed systems
- Problems : Multi media based applications
- Typically buses or rings.
- Ethernet, Token Ring

1.5.5 Wide Area Networks or WAN

A wide area network is two or more computers that are geographically dispersed, linked by communication facilities such as telephone system or microwave relays. This type of network is usually limited to use by large corporations and government agencies because of the high cost involved in building and maintaining them.

A WAN is a network that links separate geographical locations and this network can be a public system or any of the various packet switched services provided by the public telecommunication agencies.

The main difference between a WAN and LAN is under the complete control of the

NOTES

NOTES

owner, whereas the WAN needs the involvement of another authority like telecom department. LANs are successful and capable of handling very high data transfer rates at low cost because of the small area covered. Besides LANs have a lower error rate than WANs.

1.5.5.1 Hardware Requirements for WAN

There are mainly 4 hardware devices which are required to establish linkage between geographically separated computers. These are:

1. Bridges
2. Routers
3. Gateways
4. X.25 Standard Interface

These are explained below.

Bridges: These are used to connect two LANs that use the identical protocols over a wide area. The bridge acts as an address filter which picks up packets from one LAN that are intended for a destination on another LAN and passes these packets on the network. The amount of processing required at the bridge is minimal because all the devices use the same protocols. In case the distance between two LANs is very large then the user is required to employ two identical bridges at either end of the communication link.

Routers: It is a special type of device that can be used to connect networks that may not be similar. Such type of devices provide connectivity between two LANs and two WANs over large geographical distances. These devices operate at the Network Layer of the OSI model. These devices participate in a routing protocol to access the network topology, and on the basis of this information routers compute the best route from a sender to the receiver.

Gateways: These are used to connect two dissimilar LANs. The terms Gateways and Routers are used interchangeably, though there is discriminating difference between the two. A router operates at the network layer whereas a gateway operates on the application layer of the OSI model. A gateway is required to convert data packets from one protocol format to another before forwarding it, as it connects two dissimilar networks.

X.25 Standard Interface: X.25 is a protocol for interfacing to a Public Packet Switched Network. It is not a protocol used for implementing a network. Two systems that support X.25 can't necessarily be connected back-to-back. They can only be connected through a DCE in a Public Packet Switched. International Telegraph and Telephone Consultative Committee (ITTCC) developed X.25 as the standard interface between the Data Terminal Equipment (DTE) and Data Circuit Terminating Equipment (DCTE). This recommendation has been widely accepted as the industry standard for public packet switched networks.

1.5.5.2 Types of Wide Area Networks

There are mainly two types of WAN:

Public Networks: These are those networks which are installed and run by the telecommunication authorities and are made available to any organization or individual who go for their subscription.

Private Networks: The basic technique used in all forms of private WAN is to use private circuits to link the locations to be served by the network. Between these fixed points the owner of the network has complete freedom to use the circuits in any way they want. They can use the circuits to carry large quantities of data or for high speed transmissions. Private wide area networks can be built using whatever standard technology is available.

1.5.5.3 Features of Wide Area Networks

- Developed in 1960s.
- Generally covers large distances (states, countries, continents).
- Medium : communication circuits connected by routers.
- Routers forwards packets from one to another following a route from the sender to the receiver. Store-and-Forward
- Hosts are typically connected (or close to) the routers.
- Typical latencies : 100ms - 500ms.
- Problems with delays if using satellites.
- Typical speed : 20 - 2000 Kbits/s.
- Not (yet) suitable for distributed computing.
- New standards are changing the landscape.

NOTES

1.6 MANS

A metropolitan area network (MAN) is a large computer network that usually spans a city or a large campus. A MAN usually interconnects a number of local area networks (LANs) using a high-capacity backbone technology, such as fiber-optical links, and provides up-link services to wide area networks (or WAN) and the Internet.

The IEEE 802-2001 standard describes a MAN as being:

- A MAN is optimized for a larger geographical area than a LAN, ranging from several blocks of buildings to entire cities. MANs can also depend on communications channels of moderate-to-high data rates. A MAN might be owned and operated by a single organization, but it usually will be used by many individuals and organizations. MANs might also be owned and operated as public utilities. They will often provide means for internetworking of local networks.

Authors Kenneth C. Laudon and Jane P. Laudon of Management Information Systems: Managing the Digital Firm 10th ed. define a metropolitan area network as:

- A Metropolitan Area Network (MAN) is a large computer network that spans a metropolitan area or campus. Its geographic scope falls between a WAN and LAN. MANs provide Internet connectivity for LANs in a metropolitan region, and connect them to wider area networks like the Internet.

It can also be used in cable television.

Implementation

Some technologies used for this purpose are Asynchronous Transfer Mode (ATM), FDDI, and SMDS. These technologies are in the process of being displaced by Ethernet-based connections (e.g., Metro Ethernet) in most areas. MAN links between local area networks have been built without cables using either microwave, radio, or infra-red laser links. Most companies rent or lease circuits from common carriers due to the fact that laying long stretches of cable can be expensive.

DQDB, Distributed Queue Dual Bus, is the metropolitan area network standard for data communication. It is specified in the IEEE 802.6 standard. Using DQDB, networks can be up to 20 miles (30 km) long and operate at speeds of 34 to 155 Mbit/s.

Several notable networks started as MANs, such as the Internet peering points MAE-West, MAE-East, and the Sohonet media network.

It is a larger network that usually spans several buildings in the same city or town. The IUB network is an example of a MAN.

1.6.1 Features of Metropolitan Area Networks

- Generally covers towns and cities (50 kms)
- Developed in 1980s.
- Medium : optical fibres, cables.
- Data rates adequate for distributed computing applications.
- A typical standard is DQDB (Distributed Queue Dual Bus).
- Typical latencies : < 1 msec.
- Message routing is fast.

1.7 WIRELESS LANS

A wireless LAN or WLAN is a wireless local area network, which is the linking of two or more computers without using wires. WLAN utilizes spread-spectrum or OFDM (802.11a) modulation technology based on radio waves to enable communication between devices in a limited area, also known as the basic service set. This gives users the mobility to move around within a broad coverage area and still be connected to the network.

For the home user, wireless has become popular due to ease of installation, and location freedom with the gaining popularity of laptops. Public businesses such as coffee shops or malls have begun to offer wireless access to their customers; some are even provided as a free service. Large wireless network projects are being put up in many major cities. Google is even providing a free service to Mountain View, California and has entered a bid to do the same for San Francisco. New York City has also begun a pilot program to cover all five boroughs of the city with wireless Internet access.

1.7.1 History

In 1970 University of Hawaii, under the leadership of Norman Abramson, developed

the world's first computer communication network using low-cost ham-like radios, named ALOHAnet.

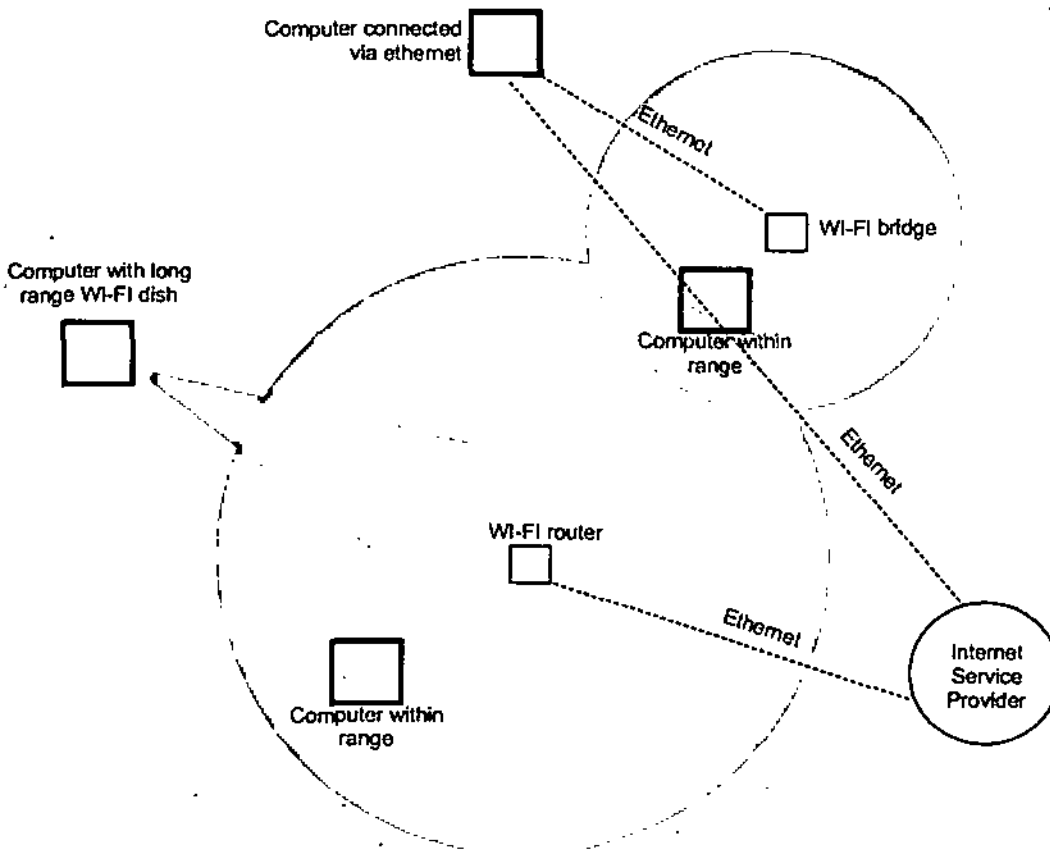
The bi-directional star topology of the system included seven computers deployed over four islands to communicate with the central computer on the Oahu Island without using phone lines.

NOTES

"In 1979, F.R. Gfeller and U. Bapst published a paper in the IEEE Proceedings reporting an experimental wireless local area network using diffused infrared communications. Shortly thereafter, in 1980, P. Ferrert reported on an experimental application of a single code spread spectrum radio for wireless terminal communications in the IEEE National Telecommunications Conference. In 1984, a comparison between Infrared and CDMA spread spectrum communications for wireless



office information networks was published by Kaveh Pahlavan in IEEE Computer Networking Symposium which appeared later in the IEEE Communication Society Magazine. In May 1985, the efforts of Marcus led the FCC to announce experimental ISM bands for commercial application of spread spectrum technology. Later on, M.



NOTES

Kavehrad reported on an experimental wireless PBX system using code division multiple access. These efforts prompted significant industrial activities in the development of a new generation of wireless local area networks and it updated several old discussion in the portable and mobile radio industry.

The first generation of wireless data modems was developed in the early 1980's by amateur communication groups. They added a voice band data communication modem, with data rates below 9600 bit/s, to an existing short distance radio system such as a walkie talkie. The second generation of wireless modems was developed immediately after the FCC announcement in the experimental bands for non-military use of the spread spectrum technology. These modems provided data rates on the order of hundreds of kbit/s. The third generation of wireless modem then aimed at compatibility with the existing LANs with data rates on the order of Mbit/s. Several companies developed the third generation products with data rates above 1 Mbit/s and a couple of products had already been announced by the time of the first IEEE Workshop on Wireless LANs."

The first of the IEEE Workshops on Wireless LAN was held in 1991. At that time early wireless LAN products had just appeared in the market and the IEEE 802.11 committee had just started its activities to develop a standard for wireless LANs. The focus of that first workshop was evaluation of the alternative technologies. By 1996, the technology was relatively mature, a variety of applications had been identified and addressed and technologies that enable these applications were well understood. Chip sets aimed at wireless LAN implementations and applications, a key enabling technology for rapid market growth, were emerging in the market. Wireless LANs were being used in hospitals, stock exchanges, and other in building and campus settings for nomadic access, point-to-point LAN bridges, ad-hoc networking, and even larger applications through internetworking. The IEEE 802.11 standard and variants and alternatives, such as the wireless LAN interoperability forum and the European HIPERLAN specification had made rapid progress, and the unlicensed PCS [Unlicensed Personal Communications Services] and the proposed SUPERNet [later on renamed as UNII] bands also presented new opportunities."

On July 21, 1999, AirPort debuted at the Macworld Expo in New York City with Steve Jobs picking up an iBook supposedly to give the cameraman a better shot as he surfed the Web. Applause quickly built as people realized there were no wires. This was the first time Wireless LAN became publicly available at consumer pricing and easily available for home use. Before the release of the Airport, Wireless LAN was too expensive for consumer use and was used exclusively in large corporate settings.

Originally WLAN hardware was so expensive that it was only used as an alternative to cabled LAN in places where cabling was difficult or impossible.

Early development included industry-specific solutions and proprietary protocols, but at the end of the 1990s these were replaced by standards, primarily the various versions of IEEE 802.11 (Wi-Fi). An alternative ATM-like 5 GHz standardized technology, HIPERLAN, has so far not succeeded in the market, and with the release of the faster 54 Mbit/s 802.11a (5 GHz) and 802.11g (2.4 GHz) standards, almost certainly never will.

In November 2006, the Australian Commonwealth Scientific and Industrial Research Organisation (CSIRO) won a legal battle in the US federal court of Texas against Buffalo Technology which found the US manufacturer had failed to pay royalties on

a US WLAN patent CSIRO had filed in 1996. CSIRO are currently engaged in legal cases with computer companies including Microsoft, Intel, Dell, Hewlett-Packard and Netgear which argue that the patent is invalid and should negate any royalties paid to CSIRO for WLAN-based products.

1.7.2 Benefits

The popularity of wireless LANs is a testament primarily to their convenience, cost efficiency, and ease of integration with other networks and network components. The majority of computers sold to consumers today come pre-equipped with all necessary wireless LAN technology.

The benefits of wireless LANs include:

- **Convenience:** The wireless nature of such networks allows users to access network resources from nearly any convenient location within their primary networking environment (home or office). With the increasing saturation of laptop-style computers, this is particularly relevant.
- **Mobility:** With the emergence of public wireless networks, users can access the internet even outside their normal work environment. Most chain coffee shops, for example, offer their customers a wireless connection to the internet at little or no cost.
- **Productivity:** Users connected to a wireless network can maintain a nearly constant affiliation with their desired network as they move from place to place. For a business, this implies that an employee can potentially be more productive as his or her work can be accomplished from any convenient location.
- **Deployment:** Initial setup of an infrastructure-based wireless network requires little more than a single access point. Wired networks, on the other hand, have the additional cost and complexity of actual physical cables being run to numerous locations (which can even be impossible for hard-to-reach locations within a building).
- **Expandability:** Wireless networks can serve a suddenly-increased number of clients with the existing equipment. In a wired network, additional clients would require additional wiring.
- **Cost:** Wireless networking hardware is at worst a modest increase from wired counterparts. This potentially increased cost is almost always more than outweighed by the savings in cost and labour associated to running physical cables.

1.7.3 Disadvantages

Wireless LAN technology, while replete with the conveniences and advantages described above, has its share of downfalls. For a given networking situation, wireless LANs may not be desirable for a number of reasons. Most of these have to do with the inherent limitations of the technology.

- **Security:** Wireless LAN transceivers are designed to serve computers throughout a structure with uninterrupted service using radio frequencies. Because of space and cost, the antennas typically present on wireless networking cards in the end computers are generally relatively poor. In order

NOTES

NOTES

to properly receive signals using such limited antennas throughout even a modest area, the wireless LAN transceiver utilizes a fairly considerable amount of power. What this means is that not only can the wireless packets be intercepted by a nearby adversary's poorly-equipped computer, but more importantly, a user willing to spend a small amount of money on a good quality antenna can pick up packets at a remarkable distance; perhaps hundreds of times the radius as the typical user. In fact, there are even computer users dedicated to locating and sometimes even hacking into wireless networks, known as wardrivers. On a wired network, any adversary would first have to overcome the physical limitation of tapping into the actual wires, but this is not an issue with wireless packets. To combat this consideration, wireless network users usually choose to utilize various encryption technologies available such as Wi-Fi Protected Access (WPA). Some of the more older encryption methods, such as WEP are known to have weaknesses that a dedicated adversary can compromise.

- **Range:** The typical range of a common 802.11g network with standard equipment is of the order of tens of meters. While sufficient for a typical home, it will be insufficient in a larger structure. To obtain additional range, repeaters or additional access points will have to be purchased. Costs for these items can add up quickly. Other technologies are in the development phase, however, which feature increased range, hoping to render this disadvantage irrelevant.
- **Reliability:** Like any radio frequency transmission, wireless networking signals are subject to a wide variety of interference, as well as complex propagation effects (such as multipath, or especially in this case Rician fading) that are beyond the control of the network administrator. In the case of typical networks, modulation is achieved by complicated forms of phase-shift keying (PSK) or quadrature amplitude modulation (QAM), making interference and propagation effects all the more disturbing. As a result, important network resources such as servers are rarely connected wirelessly.
- **Speed:** The speed on most wireless networks (typically 1-108 Mbit/s) is reasonably slow compared to the slowest common wired networks (100 Mbit/s up to several Gbit/s). There are also performance issues caused by TCP and its built-in congestion avoidance. For most users, however, this observation is irrelevant since the speed bottleneck is not in the wireless routing but rather in the outside network connectivity itself. For example, the maximum ADSL throughput (usually 8 Mbit/s or less) offered by telecommunications companies to general-purpose customers is already far slower than the slowest wireless network to which it is typically connected. That is to say, in most environments, a wireless network running at its slowest speed is still faster than the internet connection serving it in the first place. However, in specialized environments, the throughput of a wired network might be necessary. Newer standards such as 802.11n are addressing this limitation and will support peak throughputs in the range of 100-200 Mbit/s.

1.7.4 Architecture

1.7.4.1 Stations

All components that can connect into a wireless medium in a network are referred to

as stations. All stations are equipped with wireless network interface cards (WNICs). Wireless stations fall into one of two categories: access points and clients.

- Access points

Access points (APs) are base stations for the wireless network. They transmit and receive radio frequencies for wireless enabled devices to communicate with.

- Clients

Wireless clients can be mobile devices such as laptops, personal digital assistants, IP phones, or fixed devices such as desktops and workstations that are equipped with a wireless network interface.

1.7.4.2 Basic service set

The basic service set (BSS) is a set of all stations that can communicate with each other. There are two types of BSS: independent BSS and infrastructure BSS. Every BSS has an identification (ID) called the BSSID, which is the MAC address of the access point servicing the BSS.

- Independent basic service set

An independent BSS is an ad-hoc network that contains no access points, which means they can not connect to any other basic service set.

- Infrastructure basic service set

An infrastructure BSS can communicate with other stations not in the same basic service set by communicating through access points.

1.7.4.3 Extended service set

An extended service set (ESS) is a set of connected BSSes. Access points in an ESS are connected by a distribution system. Each ESS has an ID called the SSID which is a 32-byte (maximum) character string. For example, "linksys" is the default SSID for Linksys routers.

1.7.4.4 Distribution system

A distribution system connects access-points in an extended service set.

1.7.5 Wireless MAN Standard (WiMAX)

The IEEE 802.16 defines the wireless metropolitan area network (MAN) technology which is branded as WiMAX. The 802.16 includes two sets of standards, 802.16-2004 (802.16d) for fixed WiMAX and 802.16-2005(802.16e) for mobile WiMAX. The WiMAX wireless broadband access standard provides the missing link for the "last mile" connection in metropolitan area networks where DSL, Cable and other broadband access methods are not available or too expensive. WiMAX also offers an alternative to satellite Internet services for rural areas and allows mobility of the customer equipment.

IEEE 802.16 standards are concerned with the air interface between a subscriber's transceiver station and a base transceiver station. The fixed WiMax standard IEEE 802.16-2004 (also known as 802.16d) is approved by the IEEE in June 2004, which provides fixed, point-to-multi point broadband wireless access service and its product profile utilizes the OFDM 256-FFT (Fast Fourier Transform) system profile.

NOTES

The fixed WiMAX 802.16-2004 standard supports both time division duplex (TDD) and frequency division duplex (FDD) services - the latter of which delivers full duplex transmission on the same signal if desired. In Dec. 2005, IEEE approved the mobile WiMax standard, the 802.16-2005 (also known as 802.16e). IEEE 802.16e, based on the early WiMax standard 802.16a, adds mobility features to WiMAX in the 2 to 11 GHz licensed bands. 802.16e allows for fixed wireless and mobile Non Line of Sight (NLOS) applications primarily by enhancing the OFDMA (Orthogonal Frequency Division Multiple Access).


1.8 FDDI

Fiber Distributed Data Interface (FDDI) provides a 100 Mbit/s optical standard for data transmission in a local area network that can extend in range up to 200 kilometers (124 miles). Although FDDI logical topology is a ring-based token network, it does not use the IEEE 802.5 token ring protocol as its basis; instead, its protocol is derived from the IEEE 802.4 token bus timed token protocol. In addition to covering large geographical areas, FDDI local area networks can support thousands of users. As a standard underlying medium it uses optical fiber, although it can use copper cable, in which case it may be refer to as CDDI. FDDI offers both a Dual-Attached Station (DAS), counter-rotating token ring topology and a Single-Attached Station (SAS), token bus passing ring topology.

FDDI was considered an attractive campus backbone technology in the early to mid 1990s since existing Ethernet networks only offered 10 Mbit/s transfer speeds and Token Ring networks only offered 4 Mbit/s or 16 Mbit/s speeds. Thus it was the preferred choice of that era for a high-speed backbone, but FDDI has since been effectively obsoleted by fast Ethernet which offered the same 100 Mbit/s speeds, but at a much lower cost and, since 1998, by Gigabit Ethernet due to its speed, and even lower cost, and ubiquity.

FDDI, as a product of American National Standards Institute X3T9.5 (now X3T12), conforms to the Open Systems Interconnection (OSI) model of functional layering of LANs using other protocols. FDDI-II, a version of FDDI, adds the capability to add circuit-switched service to the network so that it can also handle voice and video signals. Work has started to connect FDDI networks to the developing Synchronous Optical Network SONET.

A FDDI network contains two rings, one as a secondary backup in case the primary ring fails. The primary ring offers up to 100 Mbit/s capacity. When a network has no requirement for the secondary ring to do backup, it can also carry data, extending capacity to 200 Mbit/s. The single ring can extend the maximum distance; a dual ring can extend 100 km (62 miles). FDDI has a larger maximum-frame size (4,352 bytes) than standard 100 Mbit/s Ethernet which only supports a maximum-frame size of 1,500 bytes, allowing better throughput.



Designers normally construct FDDI rings in the form of a "dual ring of trees" (see network topology). A small number of devices (typically infrastructure devices such as routers and concentrators rather than host computers) connect to both rings - hence the term "dual-attached". Host computers then connect as single-attached devices to the routers or concentrators. The dual ring in its most degenerate form simply collapses into a single device. Typically, a computer-room contains the whole dual ring, although some implementations have deployed FDDI as a Metropolitan area network.

1.8.1 Standards

FDDI standards include:

- ANSI X3.139-1987, Media Access Control (MAC) — also ISO 9314-2
- ANSI X3.148-1988, Physical Layer Protocol (PHY) — also ISO 9314-1
- ANSI X3.166-1989, Physical Medium Dependent (PMD) — also ISO 9314-3
- ANSI X3.184-1993, Single Mode Fiber Physical Medium Dependent (SMF-PMD) — also ISO 9314-4
- ANSI X3.229-1994, Station Management (SMT) — also ISO 9314-6

1.9 DQDB

In telecommunication, a Distributed-Queue Dual-Bus network (DQDB) is a distributed multi-access network that (a) supports integrated communications using a dual bus and distributed queuing, (b) provides access to local or metropolitan area networks, and (c) supports connectionless data transfer, connection-oriented data transfer, and isochronous communications, such as voice communications.

IEEE 802.6 is an example of a network providing DQDB access methods.

1.9.1 DQDB Concept of Operation

The DQDB Medium Access Control (MAC) algorithm is generally credited to Robert Newman who developed this algorithm in his PhD thesis in the 1980s at the University of Western Australia. To appreciate the innovative value of the DQDB MAC algorithm, it must be seen against the background of LAN protocols at that time, which were based on broadcast (such as ethernet IEEE 802.3) or a ring (like token ring IEEE 802.5 and FDDI). The DQDB may be thought of as two token rings, one carrying data in each direction around the ring. The ring is broken between two of the nodes in the ring. (An advantage of this is that if the ring breaks somewhere else, the broken link can be closed to form a ring with only one break again. This gives reliability which is important in Metropolitan Area Networks (MAN), where repairs may take longer than in a LAN and wifi because the damage may be inaccessible).

The DQDB standard IEEE 802.6 was developed while ATM (Broadband ISDN) was still in early development, but there was strong interaction between the two standards. ATM cells and DQDB frames were harmonized. They both settled on essentially a 48-byte data frame with a 5-byte header. In the DQDB algorithm, a distributed queue was implemented by communicating queue state information via the header. Each

NOTES

node in a DQDB network maintains a pair of state variables which represent its position in the distributed queue and the size of the queue. The headers on the reverse bus communicated requests to be inserted in the distributed queue so that upstream nodes would know that they should allow DQDB cells to pass unused on the forward bus. *The algorithm was remarkable for its extreme simplicity.*

Currently DQDB systems are being installed by many carriers in entire cities, with lengths that reach up to 160 km (100 miles) with speeds of a DS3 line (44.736 Mbit/s). Other implementations use optical fiber for a length of up to 100 km and speeds around 150 Mbit/s

1.10 HIPPI

HIPPI (High Performance Parallel Interface) is a computer bus for the attachment of high speed storage devices to supercomputers. It was popular in the late 1980s and into the mid-to-late 1990s, but has since been replaced by ever-faster standard interfaces like SCSI and Fibre Channel.

The first HIPPI standard defined a 50-wire twisted pair cable, running at 800 Mbit/s (100 MB/s), but was soon upgraded to include a 1600 Mbit/s (200 MB/s) mode running on fibre optic cable. An effort to improve the speed resulted in HIPPI-6400, [2] which was later re-named GSN (for Gigabyte System Network) but saw little use due to competing standards. GSN had a full-duplex bandwidth of 6400 Mbit/s or 800 MB/s in each direction.

To understand why HIPPI is no longer used, consider that Ultra3 SCSI offers rates of 160 MB/s, and is available at almost any corner computer store. Meanwhile Fibre Channel offered simple interconnect with both HIPPI and SCSI (it can run both protocols) and speeds of up to 400 MB/s on fibre and 100 MB/s on a single pair of twisted pair copper wires.

HIPPI was the first “near-gigabit” (0.8 Gbit/s) (ANSI) standard for network data transmission. It was specifically designed for supercomputers and was *never intended* for mass market networks such as Ethernet. Many of the features developed for HIPPI are being integrated into such technologies as InfiniBand. What was remarkable about HIPPI is that it came out when Ethernet was still a 10 Mbit/s data link and SONET at OC-3 (155 Mbit/s) was considered leading edge technology.

1.11 GIGABIT ETHERNET

^{Am 2} Gigabit Ethernet (GbE or 1 GigE) is a term describing various technologies for transmitting Ethernet frames at a rate of a gigabit per second, as defined by the IEEE 802.3-2008 standard. Half-duplex gigabit links connected through hubs are allowed by the specification but in the marketplace full-duplex with switches are normal.

The result of research done at Xerox Corporation in the early 1970s, Ethernet has evolved into the most widely implemented physical and link layer protocol today. Fast Ethernet increased speed from 10 to 100 megabits per second (Mbit/s). Gigabit Ethernet was the next iteration, increasing the speed to 1000 Mbit/s. The initial standard for gigabit Ethernet was produced by the IEEE in June 1998 as IEEE 802.3z, and required optical fiber. 802.3z is commonly referred to as 1000BASE-X, where -X refers to either -CX, -SX, -LX, or (non-standard) -ZX.

Ann 2

IEEE 802.3ab, ratified in 1999, defines gigabit Ethernet transmission over unshielded twisted pair (UTP) category 5, 5e, or 6 cabling and became known as 1000BASE-T. With the ratification of 802.3ab, gigabit Ethernet became a desktop technology as organizations could use their existing copper cabling infrastructure.

IEEE 802.3ah, ratified in 2004 added two more Gigabit fiber standards, 1000BASE-LX10 (which was already widely implemented as vendor specific extension) and 1000BASE-BX10. This was part of a larger group of protocols known as Ethernet in the First Mile.

Initially, gigabit Ethernet was deployed in high-capacity backbone network links (for instance, on a high-capacity campus network). In 2000, Apple's Power Mac G4 and PowerBook G4 were the first mass produced personal computers featuring the 1000BASE-T connection. It quickly became a built-in feature in many other computers. As of 2009, Gigabit NICs (1000BASE-T) are included in almost all desktop and server computer systems.

Faster 10 gigabit Ethernet standards have become available as the IEEE ratified a fiber-based standard in 2002, and a twisted pair standard in 2006. As of 2009, 10Gb Ethernet is replacing 1Gb as the backbone network and has just begun to migrate down to high-end server systems.

There are four different physical layer standards for gigabit Ethernet using optical fiber (1000BASE-X), twisted pair cable (1000BASE-T), or balanced copper cable (1000BASE-CX).

The IEEE 802.3z standard includes 1000BASE-SX for transmission over multi-mode fiber, 1000BASE-LX for transmission over single-mode fiber, and the nearly obsolete 1000BASE-CX for transmission over balanced copper cabling. These standards use 8b/10b encoding, which inflates the line rate by 25%, from 1,000-1,250 Mbit/s to ensure a DC balanced signal. The symbols are then sent using NRZ.

IEEE 802.3ab, which defines the widely used 1000BASE-T interface type, uses a different encoding scheme in order to keep the symbol rate as low as possible, allowing transmission over twisted pair.

Ethernet in the First Mile later added 1000BASE-LX10 and -BX10.

<i>Name</i>	<i>Medium</i>	<i>Specified distance</i>
1000BASE CX	Shielded single twisted-pair cable	25 meters
1000BASE SX	Multi-mode fiber	220 to 550 meters
	dependent on fiber diameter and bandwidth	
1000BASE LX	Multi-mode fiber	550 meters[3]
1000BASE LX	Single-mode fiber	5 km[3]
1000BASE LX10	Single-mode fiber using 1,310 nm wavelength	10 km
1000BASE ZX	Single-mode fiber at 1,550 nm wavelength	~ 70 km
1000BASE BX10	Single-mode fiber, over single-strand fiber: 1,490 nm downstream 1,310 nm upstream	10 km
1000BASE T	Twisted-pair cabling (Cat 5, Cat 5e, Cat 6, or Cat 7)	100 meters
1000BASE TX	Twisted-pair cabling (Cat 6, Cat 7)	100 meters

NOTES

1.11.1 IEEE versions**1.11.1.1 1000BASE-X**

1000BASE-X is used in industry to refer to gigabit Ethernet transmission over fiber, where options include 1000BASE-CX, 1000BASE-LX, and 1000BASE-SX, 1000BASE-LX10, 1000BASE-BX10 or the non-standard -ZX implementations.

NOTES

1.11.1.2 1000BASE-CX

1000BASE-CX is an initial standard for gigabit Ethernet connections over twinaxial copper cabling with maximum distances of 25 meters using balanced shielded twisted pair and either DE-9 or 8P8C connectors. The short segment length is due to very high signal transmission rate. It is still used for specific applications where cabling is not done by general users, for instance the IBM BladeCenter uses 1000BASE-CX for the Ethernet connections between the blade servers and the switch modules. 1000BASE-T succeeded it for general copper wiring use.

1.11.1.3 1000BASE-SX

1000BASE-SX is a fiber optic gigabit Ethernet standard for operation over multi-mode fiber using a 770 to 860 nanometer, near infrared (NIR) light wavelength.

The standard specifies a distance capability between 220 meters (62.5/125 μm fiber with low modal bandwidth) and 550 meters (50/125 μm fiber with high modal bandwidth). In practice, with good quality fibre and terminations, 1000BASE-SX will usually work over significantly longer distances.[citation needed]

This standard is highly popular for intra-building links in large office buildings, co-location facilities and carrier neutral internet exchanges.

Optical power specifications of SX interface: Minimum output power = "9.5 dBm. Minimum receive sensitivity = "17 dBm.

1.11.1.4 1000BASE-LX

1000BASE-LX is a fiber optic gigabit Ethernet standard specified in IEEE 802.3 Clause 38 which uses a long wavelength laser (1,270–1,355 nm), and a maximum RMS spectral width of 4 nm.

1000BASE-LX is specified to work over a distance of up to 5 km over 10 μm single-mode fiber.

1000BASE-LX can also run over all common types of multi-mode fiber with a maximum segment length of 550 m. For link distances greater than 300 m, the use of a special launch conditioning patch cord may be required.[4] This launches the laser at a precise offset from the center of the fiber which causes it to spread across the diameter of the fiber core, reducing the effect known as differential mode delay which occurs when the laser couples onto only a small number of available modes in multi-mode fiber.

1.11.1.5 1000BASE-LX10

1000BASE-LX10 was standardized six years after the initial gigabit fiber versions as part of the Ethernet in the First Mile task group. It is very similar to 1000BASE-LX, but achieves longer distances up to 10 km over a pair of single-mode fiber due to higher quality optics. Before it was standardized 1000BASE-LX10 was essentially

already in widespread use by many vendors as a proprietary extension called either 1000BASE-LX/LH or 1000BASE-LH.

1.11.1.6 1000BASE-BX10

1000BASE-BX10 is capable of up to 10 km over a single strand of single-mode fiber, with a different wavelength going in each direction. The terminals on each side of the fibre are not equal, as the one transmitting downstream (from the center of the network to the outside) uses the 1,490 nm wavelength, and the one transmitting upstream uses the 1,310 nm wavelength.

1.11.2 Non IEEE versions

1.11.2.1 1000BASE-ZX

1000BASE-ZX is a non-standard but industry accepted[citation needed] term to refer to gigabit Ethernet transmission using 1,550 nm wavelength to achieve distances of at least 70 km over single-mode fiber.

1.11.2.2 1000BASE-T

1000BASE-T (also known as IEEE 802.3ab) is a standard for gigabit Ethernet over copper wiring.

Each 1000BASE-T network segment can be a maximum length of 100 meters (328 feet), and must use Category 5 cable at a minimum. Category 5e cable or Category 6 cable may also be used.

Autonegotiation is a requirement for using 1000BASE-T according to Section 28D.5 Extensions required for Clause40 (1000BASE-T).[7] At least the clock source has to be negotiated, as one has to be master and the other slave

1000BASE-T requires all four pairs to be present. If two gigabit devices are connected through a non-compliant Cat-5 cable with two pairs only, negotiation takes place on two pairs only, so the devices successfully choose 'gigabit' as the highest common denominator (HCD), but the link never comes up. Most gigabit physical devices have a specific register to diagnose this behaviour. Some drivers offer an "Ethernet@Wirespeed" option where this situation leads to a slower yet functional connection.

1.11.2.3 1000BASE-T details

In a departure from both 10BASE-T and 100BASE-TX, 1000BASE-T uses all four cable pairs for simultaneous transmission in both directions through the use of echo cancellation and a 5-level pulse amplitude modulation (PAM-5) technique. The symbol rate is identical to that of 100BASE-TX (125 Mbaud) and the noise immunity of the 5-level signaling is also identical to that of the 3-level signaling in 100BASE-TX, since 1000BASE-T uses 4-dimensional trellis coded modulation (TCM) to achieve a 6 dB coding gain across the 4 pairs.

The data are transmitted over four copper pairs, eight bits at a time. First, eight bits of data are expanded into four 3-bit symbols through a non-trivial scrambling procedure based on a linear feedback shift register; this is similar to what is done in 100BASE-T2, but uses different parameters. The 3-bit symbols are then mapped to voltage levels which vary continuously during transmission. One example mapping is as follows:

NOTES

STUDENT ACTIVITY 1.2

1. What are OSI physical layer concepts?

2. What are OSI network layer concepts?

NOTES

Symbol	Line signal level
000	0
001	+1
010	+2
011	-1
100	0
101	+1
110	-2
111	-1

1.11.3 Automatic crossover

Automatic MDI/MDI-X Configuration is specified as an optional feature in the 1000BASE-T standard, meaning that straight-through cables will often work between Gigabit capable interfaces. This feature eliminates the need for crossover cables, making obsolete the uplink/normal ports and manual selector switches found on many older hubs and switches and greatly reducing installation errors.

1.11.3.1 1000BASE-TX

The Telecommunications Industry Association (TIA) created and promoted a version of 1000BASE-T that was simpler to implement, calling it 1000BASE-TX (TIA/EIA-854). The simplified design would, in theory, have reduced the cost of the required electronics by only using one pair of wires in each direction. However, this solution required Category 6 cable and has been a commercial failure, likely due to the cabling requirement as well as the rapidly falling cost of 1000BASE-T products. Many 1000BASE-T products are advertised as 1000BASE-TX due to lack of knowledge that 1000BASE-TX is actually a different standard. The most popular form of Fast Ethernet (100 Mbit/s) is known as 100BASE-TX.

1.12 WIRELESS ETHERNET

IEEE/802.11 is a set of standards carrying out wireless local area network (WLAN) computer communication in the 2.4, 3.6 and 5 GHz frequency bands. They are created and maintained by the IEEE LAN/MAN Standards Committee (IEEE 802). The base current version of the standard is IEEE 802.11-2007.

1.12.1 General description

The 802.11 family includes over-the-air modulation techniques that use the same basic protocol. The most popular are those defined by the 802.11b and 802.11g protocols, which are amendments to the original standard. 802.11-1997 was the first wireless networking standard, but 802.11b was the first widely accepted one, followed by 802.11g and 802.11n. Security was originally purposefully weak due to export requirements of some governments, and was later enhanced via the 802.11i amendment after governmental and legislative changes. 802.11n is a new multi-streaming modulation technique. Other standards in the family (c-f, h, j) are service amendments and extensions or corrections to the previous specifications.

NOTES

802.11b and 802.11g use the 2.4 GHz ISM band, operating in the United States under Part 15 of the US Federal Communications Commission Rules and Regulations. Because of this choice of frequency band, 802.11b and g equipment may occasionally suffer interference from microwave ovens, cordless telephones and Bluetooth devices. 802.11b and 802.11g control their interference and susceptibility to interference by using direct-sequence spread spectrum (DSSS) and orthogonal frequency-division multiplexing (OFDM) signaling methods, respectively. 802.11a uses the 5 GHz U-NII band, which, for much of the world, offers at least 19 non-overlapping channels rather than the 3 offered in the 2.4 GHz ISM frequency band. Better or worse performance with higher or lower frequencies (channels) may be realized, depending on the environment.

The segment of the radio frequency spectrum used by 802.11 varies between countries. In the US, 802.11a and 802.11g devices may be operated without a license, as allowed in Part 15 of the FCC Rules and Regulations. Frequencies used by channels one through six of 802.11b and 802.11g fall within the 2.4 GHz amateur radio band. Licensed amateur radio operators may operate 802.11b/g devices under Part 97 of the FCC Rules and Regulations, allowing increased power output but not commercial content or encryption.

1.12.2 History

802.11 technology has its origins in a 1985 ruling by the U.S. Federal Communications Commission that released the ISM band for unlicensed use.

In 1991 NCR Corporation/AT&T (later Lucent Technologies and Agere Systems) invented the precursor to 802.11 in Nieuwegein, The Netherlands. The inventors initially intended to use the technology for cashier systems; the first wireless products were brought on the market under the name WaveLAN with raw data rates of 1 Mbit/s and 2 Mbit/s.

Vic Hayes, who held the chair of IEEE 802.11 for 10 years and has been called the "father of Wi-Fi" was involved in designing the initial 802.11b and 802.11a standards within the IEEE.

In 1992, the Commonwealth Scientific and Industrial Research Organisation (CSIRO) obtained a patent in Australia for wireless data transfer technology. In 1996, they obtained a patent for the same technology in the US. Wi-Fi uses the mathematical formula in the patents. In April 2009, 14 tech companies including Intel, Microsoft, HP, Dell, agreed to pay CSIRO \$250 million for their Wi-Fi patent infringements.

1.12.3 Protocols

802.11 network standards v • d • e

802.11 Protocol	Release	Freq. (GHz)	Band width (MHz)	Data rate per stream	Allowable MIMO streams	Modulation	Approx. indoor range	Approx. Outdoor range
-	Jun 1997	2.4	20	1, 2	1	DSSS, FHSS	20 - 66	100 - 330
a	Sep 1999	5 3.7[y]	20	6, 9, 12, 18, 24, 36, 48, 54	1	OFDM	35 - --	115 - 5,000 120 - 16,000
b	Sep 1999	2.4	20	5.5, 11	1	DSSS	38 - 125	140 - 460

Contd...

g	Jun 2003	2.4	20	6, 9, 12, 18, 24, 36, 48, 54	1	OFDM, DSSS	38	125	140	460
n	Oct 2009	2.4/5	20	7.2, 14.4, 21.7, 28.9, 43.3, 57.8, 65, 72.2	4	OFDM	70	230	250	820[9]
			40	15, 30, 45, 60, 90, 120, 135, 150			70	230	250	820[9]

NOTES

1.12.4 IEEE 802.11

The original version of the standard IEEE 802.11 was released in 1997 and clarified in 1999, but is today obsolete. It specified two net bit rates of 1 or 2 megabits per second (Mbit/s), plus forward error correction code. It specified three alternative physical layer technologies: diffuse infrared operating at 1 Mbit/s; frequency-hopping spread spectrum operating at 1 Mbit/s or 2 Mbit/s; and direct-sequence spread spectrum operating at 1 Mbit/s or 2 Mbit/s. The latter two radio technologies used microwave transmission over the Industrial Scientific Medical frequency band at 2.4 GHz. Some earlier WLAN technologies used lower frequencies, such as the U.S. 900 MHz ISM band.

Legacy 802.11 with direct-sequence spread spectrum was rapidly supplanted and popularized by 802.11b.

1.12.5 IEEE 802.11a-1999

The 802.11a standard uses the same data link layer protocol and frame format as the original standard, but an OFDM based air interface (physical layer). It operates in the 5 GHz band with a maximum net data rate of 54 Mbit/s, plus error correction code, which yields realistic net achievable throughput in the mid-20 Mbit/s.

Since the 2.4 GHz band is heavily used to the point of being crowded, using the relatively unused 5 GHz band gives 802.11a a significant advantage. However, this high carrier frequency also brings a disadvantage: the effective overall range of 802.11a is less than that of 802.11b/g. In theory, 802.11a signals are absorbed more readily by walls and other solid objects in their path due to their smaller wavelength and, as a result, cannot penetrate as far as those of 802.11b. In practice, 802.11b typically has a higher range at low speeds (802.11b will reduce speed to 5 Mbit/s or even 1 Mbit/s at low signal strengths). However, at higher speeds, 802.11a often has the same or greater range due to less interference.

1.12.6 IEEE 802.11b-1999

802.11b has a maximum raw data rate of 11 Mbit/s and uses the same media access method defined in the original standard. 802.11b products appeared on the market in early 2000, since 802.11b is a direct extension of the modulation technique defined in the original standard. The dramatic increase in throughput of 802.11b (compared to the original standard) along with simultaneous substantial price reductions led to the rapid acceptance of 802.11b as the definitive wireless LAN technology.

802.11b devices suffer interference from other products operating in the 2.4 GHz band. Devices operating in the 2.4 GHz range include: microwave ovens, Bluetooth devices, baby monitors and cordless telephones.

1.12.7 IEEE 802.11g-2003

In June 2003, a third modulation standard was ratified: 802.11g. This works in the 2.4 GHz band (like 802.11b), but uses the same OFDM based transmission scheme as 802.11a. It operates at a maximum physical layer bit rate of 54 Mbit/s exclusive of forward error correction codes, or about 22 Mbit/s average throughput. 802.11g hardware is fully backwards compatible with 802.11b hardware and therefore is encumbered with legacy issues that reduce throughput when compared to 802.11a by ~21%.

The then-proposed 802.11g standard was rapidly adopted by consumers starting in January 2003, well before ratification, due to the desire for higher data rates as well as to reductions in manufacturing costs. By summer 2003, most dual-band 802.11a/b products became dual-band/tri-mode, supporting a and b/g in a single mobile adapter card or access point. Details of making b and g work well together occupied much of the lingering technical process; in an 802.11g network, however, activity of an 802.11b participant will reduce the data rate of the overall 802.11g network .

Like 802.11b, 802.11g devices suffer interference from other products operating in the 2.4 GHz band, for example wireless keyboards.

1.12.8 802.11-2007

In 2003, task group TGma was authorized to "roll up" many of the amendments to the 1999 version of the 802.11 standard. REVma or 802.11ma, as it was called, created a single document that merged 8 amendments (802.11a, b, d, e, g, h, i, j) with the base standard. Upon approval on March 8, 2007, 802.11REVma was renamed to the current base standard IEEE 802.11-2007.

1.12.9 IEEE 802.11n-2009

802.11n is a recent amendment which improves upon the previous 802.11 standards by adding multiple-input multiple-output antennas (MIMO) and many other newer features. The IEEE has approved the amendment and it was published in October 2009.[12][13] Prior to the final ratification, enterprises were already migrating to 802.11n networks based on the Wi-Fi Alliance's certification of products conforming to a 2007 draft of the 802.11n proposal.

1.12.10 Channels and international compatibility

802.11 divides each of the above-described bands into channels, analogously to how radio and TV broadcast bands are sub-divided but with greater channel width and overlap. For example the 2.4000–2.4835 GHz band is divided into 13 channels each of width 22 MHz but spaced only 5 MHz apart, with channel 1 centered on 2.412 GHz and 13 on 2.472 GHz to which Japan adds a 14th channel 12 MHz above channel 13.

Availability of channels is regulated by country, constrained in part by how each country allocates radio spectrum to various services. At one extreme, Japan permits the use of all 14 channels (with the exclusion of 802.11g/n from channel 14), while at the other Spain initially allowed only channels 10 and 11 and France allowed only 10, 11, 12 and 13 (now both countries follow the European model of allowing channels 1 through 13[14][15]). Most other European countries are almost as liberal as Japan, disallowing only channel 14, while North America and some Central and South

NOTES

American countries further disallow 12 and 13. For more details on this topic, see List of WLAN channels.

Besides specifying the centre frequency of each channel, 802.11 also specifies (in Clause 17) a spectral mask defining the permitted distribution of power across each channel. The mask requires that the signal be attenuated by at least 30 dB from its peak energy at ± 11 MHz from the centre frequency, the sense in which channels are effectively 22 MHz wide. One consequence is that stations can only use every fourth or fifth channel without overlap, typically 1, 6 and 11 in the Americas, and in theory 1, 5, 9 and 13 in Europe although 1, 6, and 11 is typical there too. Another is that channels 1-13 effectively require the band 2.401–2.483 GHz, the actual allocation being, for example, 2.400–2.4835 GHz in the UK, 2.402–2.4735 GHz in the US, etc.

1.12.11 Spectral masks for 802.11g channels 1-14 in the 2.4 GHz band

Since the spectral mask only defines power output restrictions up to ± 11 MHz from the center frequency to be attenuated by -50 dB, it is often assumed that the energy of the channel extends no further than these limits. It is more correct to say that given the separation between channels 1, 6, and 11, the signal on any channel should be sufficiently attenuated to minimally interfere with a transmitter on any other channel. Due to the near-far problem a transmitter can impact a receiver on a "non-overlapping" channel, but only if it is close to the victim receiver (within a meter) or operating above allowed power levels.

Although the statement that channels 1, 6, and 11 are "non-overlapping" is limited to spacing or product density, the 1-6-11 guideline has merit. If transmitters are close together than channels 1, 6, and 11 (for example, 1, 4, 7, and 10), overlap between the channels may cause unacceptable degradation of signal quality and throughput. However, overlapping channels may be used under certain circumstances. This way more channels are available.

1.12.12 Frames

Current 802.11 standards define "frame" types for use in transmission of data as well as management and control of wireless links:

Frames are divided into very specific and standardized sections. Each frame has a MAC header, payload and FCS. Some frames may not have payload portion. First 2 bytes of MAC header is a frame control field that provides detailed information about the frame. The sub fields of the frame control field is presented in order.

Protocol Version: It is two bits in size and represents the protocol version. Currently used protocol version is zero. Other values are reserved for future use.

Type: It is two bits in size and helps to identify the type of WLAN frame. Control, Data and Management are various frame types defined in IEEE 802.11.

Sub Type: It is four bits in size. Type and Sub type are combined together to identify the exact frame.

ToDS and FromDS: Each is one bit in size. They indicate whether a data frame is headed for a distributed system. Control and management frames set these values to zero. All the data frames will have one of these bits set. However communication within an IBSS network always set these bits to zero.

More Fragment: The More Fragmentation bit is set most notably when higher level

packets have been partitioned and will be set for all non-final sections. Some management frames may require partitioning as well.

Retry: Sometimes frames require retransmission, and for this there is a Retry bit which is set to one when a frame is resent. This aids in the elimination of duplicate frames.

Power Management: The Power Management bit indicates the power management state of the sender after the completion of a frame exchange. Access points are required to manage the connection and will never set the power saver bit.

More Data: The More Data bit is used to buffer frames received in a distributed system. The access point uses this bit to facilitate stations in power saver mode. It indicates that at least one frame is available and addresses all stations connected.

WEP: The WEP bit is modified after processing a frame. It is toggled to one after a frame has been decrypted or if no encryption is set it will have already been one.

Order: This bit is only set when the "strict ordering" delivery method is employed. Frames and fragments are not always sent in order as it causes a transmission performance penalty.

The next two bytes are reserved for the Duration ID field. This field can take one of three forms: Duration, Contention-Free Period (CFP), and Association ID (AID).

An 802.11 frame can have up to four address fields. Each field can carry a MAC address. Address 1 is the receiver, Address 2 is the transmitter, Address 3 is used for filtering purposes by the receiver.

The Sequence Control field is a two-byte section used for identifying message order as well as eliminating duplicate frames. The first 4 bits are used for the fragmentation number and the last 12 bits are the sequence number.

An optional two-byte Quality of Service control field which was added with 802.11e.

The Frame Body field is variable in size, from 0 to 2304 bytes plus any overhead from security encapsulation and contains information from higher layers.

The Frame Check Sequence (FCS) is the last four bytes in the standard 802.11 frame. Often referred to as the Cyclic Redundancy Check (CRC), it allows for integrity check of retrieved frames. As frames are about to be sent the FCS is calculated and appended. When a station receives a frame it can calculate the FCS of the frame and compare it to the one received. If they match, it is assumed that the frame was not distorted during transmission.[18]

Management Frames allow for the maintenance of communication. Some common 802.11 subtypes include:

Authentication frame: 802.11 authentication begins with the WNIC sending an authentication frame to the access point containing its identity. With an open system authentication the WNIC only sends a single authentication frame and the access point responds with an authentication frame of its own indicating acceptance or rejection. With shared key authentication, after the WNIC sends its initial authentication request it will receive an authentication frame from the access point containing challenge text. The WNIC sends an authentication frame containing the encrypted version of the challenge text to the access point. The access point ensures the text was encrypted with the correct key by decrypting it with its own key. The result of this process determines the WNIC's authentication status.

NOTES

NOTES

Association request frame: sent from a station it enables the access point to allocate resources and synchronize. The frame carries information about the WNIC including supported data rates and the SSID of the network the station wishes to associate with. If the request is accepted, the access point reserves memory and establishes an association ID for the WNIC.

Association response frame: sent from an access point to a station containing the acceptance or rejection to an association request. If it is an acceptance, the frame will contain information such as an association ID and supported data rates.

Beacon frame: Sent periodically from an access point to announce its presence and provide the SSID, and other parameters for WNICs within range.

Deauthentication frame: Sent from a station wishing to terminate connection from another station.

Disassociation frame: Sent from a station wishing to terminate connection. It's an elegant way to allow the access point to relinquish memory allocation and remove the WNIC from the association table.

Probe request frame: Sent from a station when it requires information from another station.

Probe response frame: Sent from an access point containing capability information, supported data rates, etc., after receiving a probe request frame.

Reassociation request frame: A WNIC sends a reassociation request when it drops from range of the currently associated access point and finds another access point with a stronger signal. The new access point coordinates the forwarding of any information that may still be contained in the buffer of the previous access point.

Reassociation response frame: Sent from an access point containing the acceptance or rejection to a WNIC reassociation request frame. The frame includes information required for association such as the association ID and supported data rates.

Control frames facilitate in the exchange of data frames between stations. Some common 802.11 control frames include:

Acknowledgement (ACK) frame: After receiving a data frame, the receiving station will send an ACK frame to the sending station if no errors are found. If the sending station doesn't receive an ACK frame within a predetermined period of time, the sending station will resend the frame.

Request to Send (RTS) frame: The RTS and CTS frames provide an optional collision reduction scheme for access point with hidden stations. A station sends a RTS frame to as the first step in a two-way handshake required before sending data frames.

Clear to Send (CTS) frame: A station responds to an RTS frame with a CTS frame. It provides clearance for the requesting station to send a data frame. The CTS provides collision control management by including a time value for which all other stations are to hold off transmission while the requesting stations transmits.

1.13 FAST ACCESS TECHNOLOGIES

Speed is the key to accessing data in the case of modems. Some of the fast accessing technologies based modems are discussed below.

1.14 ADSL

Asymmetric Digital Subscriber Line (ADSL) is one form of the Digital Subscriber Line technology, a data communications technology that enables faster data transmission over copper telephone lines than a conventional voiceband modem can provide. It does this by utilizing frequencies that are not used by a voice telephone call. A splitter - or microfilter - allows a single telephone connection to be used for both ADSL service and voice calls at the same time. ADSL can generally only be distributed over short distances from the central office, typically less than 4 kilometres (2 mi), but has been known to exceed 8 kilometres (5 mi) if the originally laid wire gauge allows for farther distribution. In 2005, the ability to transmit copper ADSL/DSL services over a fiber optic cable became possible by utilizing the RLLH ADSL/DSL fiber optic link, providing distances from one point to the opposite end of the system of more than 30 miles.

At the telephone exchange the line generally terminates at a DSLAM where another frequency splitter separates the voice band signal for the conventional phone network. Data carried by the ADSL is typically routed over the telephone company's data network and eventually reaches a conventional Internet Protocol network.

1.14.1 Operation

Currently, most ADSL communication is full-duplex. Full-duplex ADSL communication is usually achieved on a wire pair by either frequency-division duplex (FDD), echo-cancelling duplex (ECD), or time-division duplex (TDD). FDD uses two separate frequency bands, referred to as the upstream and downstream bands. The upstream band is used for communication from the end user to the telephone central office. The downstream band is used for communicating from the central office to the end user.

With standard ADSL, the band from 26.000 kHz to 137.825 kHz is used for upstream communication, while 138 kHz - 1104 kHz is used for downstream communication. Each of these is further divided into smaller frequency channels of 4.3125 kHz. These frequency channels are sometimes termed bins. During initial training, the ADSL modem tests each of the bins to establish the signal-to-noise ratio at each bin's frequency. The distance from the telephone exchange and the characteristics of the cable mean that some frequencies may not propagate well, and noise on the copper wire, interference from AM radio stations and local interference and electrical noise at the customer end mean that relatively high levels of noise are present at some frequencies both effects the signal-to-noise ratio in some bins (at some frequencies) may be good or completely inadequate. A bad signal-to-noise ratio measured at certain frequencies will mean that those bins will not be used, resulting in a reduced maximum link capacity but with an otherwise functional ADSL connection.

The DSL modem will make a plan on how to exploit each of the bins sometimes termed "bits per bin" allocation. Those bins that have a good signal-to-noise ratio (SNR) will be chosen to transmit signals chosen from a greater number of possible encoded values (this range of possibilities equating to more bits of data sent) in each main clock cycle. The number of possibilities must not be so large that the receiver might mishear which one was intended in the presence of noise. Noisy bins may only be required to carry as few as two bits, a choice from only one of four possible

NOTES

NOTES

patterns, or only one bit per bin in the case of ADSL2+, and really noisy bins are not used at all. If the pattern of noise versus frequencies heard in the bins changes, the DSL modem can alter the bits-per-bin allocations, in a process called "bitswap", where bins that have become more noisy are only required to carry fewer bits and other channels will be chosen to be given a higher burden. The data transfer capacity the DSL modem therefore reports is determined by the total of the bits-per-bin allocations of all the bins combined. Higher signal-to-noise ratios and more bins being in use gives a higher total link capacity, while lower signal-to-noise ratios or fewer bins being used gives a low link capacity.

The total maximum capacity derived from summing the bits-per-bins is reported by DSL modems and is sometimes termed sync rate. This will always be rather misleading as the true maximum link capacity for user data transfer rate will be significantly lower because extra data is transmitted that is termed protocol overhead, a reduced figure of around 84-87% at most for PPPoA connections being a common example. In addition some ISPs will have traffic policies that limit maximum transfer rates further in the networks beyond the exchange, and traffic congestion on the Internet, heavy loading on servers and slowness or inefficiency in customers' computers may all contribute to reductions below the maximum attainable.

The choices the DSL modem make can also be either conservative, where the modem chooses to allocate fewer bits per bin than it possibly could, a choice which makes for a slower connection, or less conservative in which more bits per bin are chosen in which case there is a greater risk case of error should future signal-to-noise ratios deteriorate to the point where the bits-per-bin allocations chosen are too high to cope with the greater noise present. This conservatism involving a choice to using fewer bits per bin as a safeguard against future noise increases is reported as the signal-to-noise ratio margin or SNR margin. The telephone exchange can indicate a suggested SNR margin to the customer's DSL modem when it initially connects, and the modem may make its bits-per-bin allocation plan accordingly. A high SNR margin will mean a reduced maximum throughput but greater reliability and stability of the connection. A low SNR margin will mean high speeds provided the noise level does not increase too much, otherwise the connection will have to be dropped and renegotiated (resynced). ADSL2+ can better accommodate such circumstances, offering a feature termed seamless rate adaptation (SRA), which can accommodate changes in total link capacity with less disruption to communications.

Vendors may support usage of higher frequencies as a proprietary extension to the standard. However, this requires matching vendor-supplied equipment on both ends of the line, and will likely result in crosstalk problems that affect other lines in the same bundle.

There is a direct relationship between the number of channels available and the throughput capacity of the ADSL connection. The exact data capacity per channel depends on the modulation method used.

ADSL initially existed in two flavours (similar to VDSL), namely CAP and DMT. CAP was the de facto standard for ADSL deployments up until 1996, deployed in 90 percent of ADSL installs at the time. However, DMT was chosen for the first ITU-T ADSL standards, G.992.1 and G.992.2 (also called G.dmt and G.lite respectively). Therefore all modern installations of ADSL are based on the DMT modulation scheme.

1.15 CABLE MODEM

NOTES

A cable modem is a type of network bridge and modem that provides bi-directional data communication via radio frequency channels on a cable television (CATV) infrastructure. Cable modems are primarily used to deliver broadband Internet access in the form of cable Internet, taking advantage of the high bandwidth of a cable television network. They are commonly deployed in Australia, Europe, and North and South America. In the United States alone, there were 38,005,172 cable modem users by mid-2009.

1.15.1 Hybrid Networks

Hybrid Networks developed, demonstrated and patented the first high-speed, asymmetrical cable modem system in 1990. A key Hybrid Networks insight was that highly asymmetrical communications would be sufficient to satisfy consumers connected remotely to an otherwise completely symmetric high-speed data communications network. This was important because it was very expensive to provide high speed in the upstream direction, while the CATV systems already had substantial broadband capacity in the downstream direction. Also key was that it saw that the upstream and downstream communications could be on the same or different communications media using different protocols working in each direction to establish a closed loop communications system. The speeds and protocols used in each direction would be very different. The earliest systems used the public switched telephone network (PSTN) for the return path since very few cable systems were bi-directional. Later systems used CATV for the upstream as well as the downstream path. Hybrid's system architecture is used for most cable modem systems today.

1.15.1.1 LANcity

LANcity was an early pioneer in cable modems, developing a proprietary system that was widely deployed in the US. LANcity was sold to Bay Networks which was then acquired by Nortel, which eventually spun the cable modem business off as ARRIS. ARRIS continues to make cable modems and CMTS equipment compliant with the DOCSIS standard.

1.15.1.2 Zenith Homeworks

Zenith offered a cable modem technology using its own protocol which it introduced in 1993, being one of the first cable modem providers. The cable modem technology was used by several cable television systems in the USA and other countries, including GTE's Americast service.

1.15.1.3 Com21

Com21 was another early pioneer in cable modems, and quite successful until proprietary systems were made obsolete by the DOCSIS standardization. The Com21 system used a ComController as central bridge in CATV network head-ends, the ComPort cable modem in various models and the NMAPS management system using HP OpenView as platform. Later they also introduced a return path multiplexer to overcome noise problems when combining return path signals from multiple areas. The proprietary protocol was based on Asynchronous Transfer Mode (ATM). The central ComController switch was a modular system offering one downstream channel (transmitter) and one management module. The remaining slots could be used for

upstream receivers (2 per card), dual Ethernet 10BaseT and later also Fast-Ethernet and ATM interfaces. The ATM interface became the most popular, as it supported the increasing bandwidth demands and also supported VLANs. Com21 developed a DOCSIS modem, but the company filed for bankruptcy in 2003 and closed. The DOCSIS CMTS assets of COM21 were acquired by ARRIS.

NOTES

1.15.1.4 CDLP

CDLP was a proprietary system manufactured by Motorola. CDLP customer premises equipment (CPE) was capable of both PSTN (telephone network) and radio frequency (cable network) return paths. The PSTN-based service was considered 'one way cable' and had many of the same drawbacks as satellite Internet service and, as a result, it quickly gave way to two-way cable. Cable modems that used the RF cable network for the return path were considered 'two-way cable', and were better able to compete with the bi-directional digital subscriber line (DSL) service. The standard is in little use now while new providers use, and existing providers having changed to, the DOCSIS standard. The Motorola CDLP proprietary CyberSURFER is an example of a device that was built to the CDLP standard, capable of a peak 10 Mbit/s downstream and 1.532 Mbit/s upstream. CDLP supported a maximum downstream bandwidth of 30 Mbit/s which could be reached by using several cable modems.

The Australian ISP BigPond employed this system when it started cable modem tests in 1996. For a number of years cable Internet access was only available in Sydney, Melbourne and Brisbane via CDLP. This network ran parallel to the newer DOCSIS system for several years. In 2004, the CDLP network was terminated and replaced by DOCSIS.

1.15.1.5 IEEE 802.14

In the mid-1990s the IEEE 802 committee formed a subcommittee (802.14) to develop a standard for cable modem systems. While significant progress was made, the group was disbanded when North American multi system operators instead backed the then-fledgling DOCSIS specification.

1.15.1.6 DOCSIS

In the late 1990s, a consortium of US cable operators, known as "MCNS" formed to quickly develop an open and interoperable cable modem specification. The group essentially combined technologies from the two dominant proprietary systems at the time, taking the physical layer from the Motorola CDLP system and the MAC layer from the LANcity system. When the initial specification had been drafted, the MCNS consortium handed over control of it to CableLabs which maintained the specification, promoted it in various standards organizations (notably SCTE and ITU), developed a certification testing program for cable modem equipment, and has since drafted multiple extensions to the original specification. Virtually all cable modems operating in the field today are compliant with one of the DOCSIS versions. Because of the differences in the European PAL and USA's NTSC systems two main versions of DOCSIS exist, DOCSIS and EuroDOCSIS. The main differences are found in the width of RF-channels: 6 MHz for the USA and 8 MHz for Europe. Nearly all current cable modem systems use a version of this standard with the exception of those in Japan.

1.15.2 Cable modems and VoIP

With the advent of Voice over Internet Protocol (VoIP) telephony, cable modems

have been extended to provide telephone service. Some companies which offer cable TV service are still offering VOIP phone, allowing customers who already purchased cable TV to eliminate their plain old telephone service (POTS). Because many telephone companies do not offer naked DSL (DSL service without POTS line service), VoIP use is higher amongst cable modem users. [citation needed] Any high-speed Internet service subscriber can use VoIP telephony by subscribing to a third-party service (e.g., Vonage, Skype).

Many cable operators offer their own VoIP service, based on PacketCable. PacketCable allows multiple system operators (MSOs) to offer both high-speed Internet and VoIP through the same cable transmission system. PacketCable service has a significant technical advantage over third-party providers in that voice packets are given guaranteed quality of service across their entire transmission path, so call quality can be assured.

When using cable operator VoIP, a combined customer premises equipment device known as an embedded multimedia terminal adapter (E-MTA) will often be used. An E-MTA is a cable modem and a VoIP adapter (MTA, multimedia terminal adapter) bundled into a single device.

1.15.3 Network architectural functions

In network topology, a cable modem is a network bridge that conforms to IEEE 802.1D for Ethernet networking (with some modifications). The cable modem bridges Ethernet frames between a customer LAN and the coax cable network. Technically, it is a modem because it must modulate data to transmit it over the cable network, and it must demodulate data from the cable network to receive it.

With respect to the OSI model of network design, a cable modem is both Physical Layer (Layer 1) device and a Data Link Layer (Layer 2) forwarder. As an IP addressable network node, cable modems support functionalities at other layers.

Layer 1 is implemented in the Ethernet PHY on its LAN interface, and a DOCSIS defined cable-specific PHY on its HFC cable interface. The term cable modem refers to this cable-specific PHY. The Network Layer (Layer 3) is implemented as a IP host in that it has its own IP address used by the network operator to maintain the device. In the Transport Layer (Layer 4) the cable modem supports UDP in association with its own IP address, and it supports filtering based on TCP and UDP port numbers to, for example, block forwarding of NetBIOS traffic out of the customer's LAN. In the Application Layer (Layer 7), the cable modem supports certain protocols that are used for management and maintenance, notably DHCP, SNMP, and TFTP.

Some cable modems may incorporate a router and a DHCP server to provide the LAN with IP network addressing. From a data forwarding and network topology perspective, this router functionality is typically kept distinct from the cable modem functionality (at least logically) even though the two may share a single enclosure and appear as one unit, sometimes called a residential gateway. So, the cable modem function will have its own IP address and MAC address as will the router.

NOTES

SUMMARY

1. When more than two computers are connected to each other and sharing information, *resources and remote systems* then this is called Networking.
2. PAN is used for communication among the personal devices (intrapersonal communication), or for connecting to a higher level network and Internet. The reach of a PAN is typically a few meters.
3. A network covering a small geographic area, like a home, office, or building is called Local Area Network.
4. WAN technologies generally function at the lower three layers of the OSI reference model: the physical layer, the data link layer, and the network layer.
5. If connected to the Internet, the intranet or extranet is normally protected from being accessed from the Internet without proper authorization
6. All networks are made up of basic hardware building blocks to interconnect network nodes, such as Network Interface Cards (NICs), Bridges, Hubs, Switches, and Routers.
7. A network card, network adapter or NIC (network interface card) is a piece of computer hardware designed to allow computers to communicate over a computer network.
8. A network bridge connects multiple network segments at the data link layer (layer 2) of the OSI model.
9. A hub is a piece of hardware which provides the connectivity of a segment of a network by directing traffic through the network.
10. Switches are the device of networking that directs traffic to the correct node by filtering and forwarding packets between Nodes.
11. Routers are the networking device that forwards data packets along networks by using headers and forwarding tables to determine the best path to forward the packets.
12. Closely related to the concept of a model is that of an *architecture*.
13. A MAN (metropolitan area network) is a larger network that usually spans several buildings in the same city or town.
14. A WAN (wide area network), in comparison to a MAN, is not restricted to a geographical location, although it might be confined within the bounds of a state or country.
15. The length of the cable connecting a computer to a LAN also varies depending on the LAN.
16. One of the major benefits of implementation of LAN is sharing expensive resources such as storage devices, printers, etc.
17. Three major categories of services used in LANs are: File Server; Printer Server and Modem Server.
18. Some popular LAN operating systems are : Novel Netware; Ethernet; Corvus; ArcNet LAN; Server Omni Net; PC Net; IBM PC LAN and Etherlink Plus, etc.
19. Features of LAN are: Typically connects computer in a single building or campus; Developed in 1970s; Medium : optical fibres, coaxial cables, twisted pair, wireless; Low latency (except in high traffic periods); High speed networks (0.2 to 100 Mb/sec); Speeds adequate for most distributed systems; Problems : Multi media based applications; Typically buses or rings and Ethernet, Token Ring.
20. Routers is a special type of device that can be used to connect networks that may not be similar.
21. Features of Wide Area Networks are: Developed in 1960s; Generally covers large distances (states, countries, continents); Medium : communication circuits connected by routers; Routers forwards packets from one to another following a route from the sender to the receiver. Store-and-Forward; Hosts are typically connected (or close to) the routers; Typical latencies : 100ms - 500ms; Problems with delays if using satellites; Typical speed : 20 - 2000 Kbits/s; Not (yet) suitable for distributed computing; and New standards are changing the landscape.

22. Features of MAN are: Generally covers towns and cities (50 kms); Developed in 1980s; Medium: optical fibres, cables; Data rates adequate for distributed computing applications; A typical standard is DQDB (Distributed Queue Dual Bus); Typical latencies : < 1 msec and Message routing is fast.
23. The **Open Systems Interconnection Basic Reference Model** (*OSI Reference Model* or *OSI Model* for short) is a layered, abstract description for communications and computer network protocol design, developed as part of Open Systems Interconnection (OSI) initiative.
24. In 1977, the International Organization for Standardization (ISO), began to develop its OSI networking suite.
25. The application layer is the seventh level of the seven-layer OSI model. It interfaces directly to and performs common application services for the application processes; it also issues requests to the presentation layer.
26. The Presentation layer transforms the data to provide a standard interface for the Application layer.
27. The Session layer controls the dialogues/connections (sessions) between computers.
28. The Transport layer provides transparent transfer of data between end users, providing reliable data transfer services to the upper layers.
29. The Network layer provides the functional and procedural means of transferring variable length data sequences from a source to a destination via one or more networks while maintaining the quality of service requested by the Transport layer.
30. The Data Link layer provides the functional and procedural means to transfer data between network entities and to detect and possibly correct errors that may occur in the Physical layer.
31. The Physical layer defines all the electrical and physical specifications for devices.
32. The **physical layer** is level one in the seven level OSI model. It performs services requested by the data link layer.
33. The **data link layer** is layer two of the seven-layer OSI model as well as of the five-layer TCP/IP reference model.
34. The uppermost sublayer is *Logical Link Control (LLC)*.
35. The **network layer** is level three of the seven level OSI model as well as of the five layer TCP/IP model.
36. In computing and telecommunications, the **transport layer** is the second highest layer in the four and five layer TCP/IP reference models, where it responds to service requests from the application layer and issues service requests to the network layer.
37. The amount of memory on a computer is limited, and without flow control a larger computer might flood a computer with so much information that it can't hold it all before dealing with it.
38. Network congestion occurs when a queue buffer of a network node is full and starts to drop packets.
39. Ports are essentially ways to address multiple entities in the same location.
40. The **presentation layer** is the sixth level of the seven layer OSI model. It responds to service requests from the application layer and issues service requests to the session layer.
41. The **application layer** is the seventh level of the seven-layer OSI model. It interfaces directly to and performs common application services for the application processes; it also issues requests to the presentation layer.
42. The Internet Protocol (IP), in combination with Transmission Control Protocol (TCP), forms the TCP/IP suite, which is the de facto protocol (i.e., universal computer language) that connects the network of networks – that is, the Internet.
43. The number of valid networks and hosts available is always $2^N - 2$ (where N is the number of bits used, and the 2 adjusts for the invalidity of the first and last addresses). Thus, for a class C address with 8 bits available for hosts, the number of hosts is 254.

NOTES

NOTES

44. An IP v4 address is a 32 bit value which allows 4,294,967,296 unique addresses.
45. The terms Netmask and Subnetmask simply refers to a way of separating the Network Address and the Host Address parts of the IP address.
46. Subnets are simply an administrative convenience and allow a method of splitting the host part of the address into smaller pieces by changing the value of the Class Netmask.
47. Classless Inter-Domain routing (CIDR) essentially removes the idea of Class from IP addresses and allows administrations to allocate and route any valid subnet from any convenient base IP irrespective of its Class.

SELF ASSESSMENT QUESTIONS

1. What is a computer network?
2. Define the various types of networks.
3. Which are the various hardware components required for networking?
4. Describe the various models of network computing.
5. What is Local Area Networks?
6. Describe Wide Area Networks.
7. Describe the various network services.
8. What is OSI model?
9. Describe the various layers.
10. Describe OSI Reference Model.
11. What are OSI physical layer concepts?
12. Describe Data-link layer concepts.
13. What are OSI network layer concepts?
14. Describe Transport layer concepts.
15. What are OSI Session layer concepts?
16. Describe OSI Presentation layer concepts.
17. What are OSI Application layer concepts?
18. What is the concept of physical and logical addressing?
19. What are classes?
20. What are sub netting and super netting?
21. Describe loop back concept.
22. What is IP Packet Format?
23. Write short notes on following:

Packet processing	Packet header
Static IP address	Dynamic IP address
Virtual Internet Protocol Network Interface	Network Address Translation
CIDR	NAT
Multicasting Protocols	Subnets
Bitwise AND	Netmask
Subnetmask	IP Address Classes

Multiple Choice Questions

1. LAN is :

(a) Local Area Network	(b) Local Around Network
(c) Love Area Network.	
2. WAN is :

(a) Windows Area Network	(b) Wide Area Network
(c) Well Area Network.	

3. MAN is :
 - (a) My Area Network
 - (b) Metro Area Network
 - (c) Metropolitan Area Network.
4. CAN is :
 - (a) Campus Area Nework
 - (b) Clear Area Network
 - (c) Clean Area Network.
5. PAN is :
 - (a) Power Area Network
 - (b) Personal Area Network
 - (c) Pop Area Network.
6. OSI is :
 - (a) Open System Interconnection
 - (b) Only System Interconnection
 - (c) Open Side Interconnection.
7. FTAM is :
 - (a) File Transfer Mode
 - (b) File Transfer Model
 - (c) Fill Transfer Model.
8. VTAM is :
 - (a) Value Transfer Model
 - (b) Value Terminal Model
 - (c) Virtual Terminal Model.
9. LLC is :
 - (a) Logical Link Control
 - (b) Logical Line Control
 - (c) Low Links Control.
10. TCP is :
 - (a) Transmission Control Protocol
 - (b) Transmission Cover Protocol
 - (c) Transmission Called Protocol.
11. IP is :
 - (a) Internet Protocols
 - (b) Indian Protocol
 - (c) Internet Player.
12. TCP is :
 - (a) Transmission Colour Protocol
 - (b) Transmission Control Protocol
 - (c) Transmission Control Power.

NOTES

True/False Questions

1. A network covering a small geographic area, like a home, office, or building is called Local Area Network.
2. If connected to the Internet, the intranet or extranet is not protected from being accessed from the Internet without proper authorization.
3. A network bridge connects multiple network segments at the data link layer (layer 2) of the OSI model.
4. Routers are the networking device that forwards data packets along networks by using headers and forwarding tables to determine the best path to forward the packets.
5. Closely related to the concept of a model is that of an architecture.
6. The length of the cable connecting a computer to a LAN also varies depending on the LAN.
7. Three major categories of services used in LANs are: File Server; Printer Server and Modem Server.
8. Routers is a special type of device that can be used to connect networks that have to be similar.
9. The term host is most commonly used when discussing TCP/IP related services and functions.

NOTES

10. The purpose of networking is to share resources.
11. In 1977, the International Organization for Standardization (ISO), began to develop its OSI networking suite.
12. The Presentation layer transforms the data to provide a standard interface for the Application layer.
13. The Session layer does not control the dialogues/connections (sessions) between computers.
14. The Data Link layer provides the functional and procedural means to transfer data between network entities and to detect and possibly correct errors that may occur in the Physical layer.
15. The Physical layer defines all the electrical and physical specifications for devices.
16. The lowermost sublayer is Logical Link Control (LLC).
17. The network layer is level three of the seven level OSI model as well as of the five layer TCP/IP model.
18. Network congestion occurs when a queue buffer of a network node is full and starts to drop packets.
19. Ports are essentially ways to address multiple entities in the same location.
20. The presentation layer is the sixth level of the seven layer OSI model. It responds to service requests from the application layer and issues service requests to the session layer.
21. An IPv4 address is a 32 bit value which allows 4,294,967,296 unique addresses.
22. The terms Netmask and Subnetmask simply refer to a way of separating the Network Address and the Host Address parts of the IP address.
23. The IETF has defined a series of IP ranges which may be freely used by any user/organization on their PRIVATE network.
24. IP addresses may be either statically allocated or dynamically allocated by a service provider.
25. A parallel communications transceiver can use loopback for testing its functionality.
26. Internet Protocol (IP) specifies a loopforward network.
27. Networks using the TCP/IP protocol route messages based on the IP address of the destination.

Short Questions with Answers

1. Which are the popular networks?
Ans. Personal Area Network (PAN)
Local Area Network (LAN)
Campus Area Network (CAN)
Metropolitan Area Network (MAN)
Wide Area Network (WAN)
Internetwork
Intranet
Extranet.
2. What are the hardware components of LAN?
Ans. The following are the major hardware components/devices for establishing LAN:
 1. Transmission Channel
 2. Network Interface Unit or NIU
 3. Servers
 4. Workstations.
3. What are the various features of LAN?
Ans. • Typically connects computer in a single building or campus.
• Developed in 1970s.

NOTES

- Medium : optical fibres, coaxial cables, twisted pair, wireless.
- Low latency (except in high traffic periods).
- High speed networks (0.2 to 100 Mb/sec).
- Speeds adequate for most distributed systems
- Problems : Multi media based applications
- Typically buses or rings.
- Ethernet, Token Ring.

4. What are the features of WAN?

- Ans.**
- Developed in 1960s.
 - Generally covers large distances (states, countries, continents).
 - Medium : communication circuits connected by routers.
 - Routers forwards packets from one to another following a route from the sender to the receiver. Store-and-Forward
 - Hosts are typically connected (or close to) the routers.
 - Typical latencies : 100ms - 500ms.
 - Problems with delays if using satellites.
 - Typical speed : 20 - 2000 Kbits/s.
 - Not (yet) suitable for distributed computing.
 - New standards are changing the landscape.

5. What are the features of MAN?

- Ans.**
- Generally covers towns and cities (50 kms)
 - Developed in 1980s.
 - Medium : optical fibres, cables.
 - Data rates adequate for distributed computing applications.
 - A typical standard is DQDB (Distributed Queue Dual Bus).
 - Typical latencies : < 1 msec.
 - Message routing is fast.

6. Which are the seven layers?

- Ans.**
- Layer 7: Application layer
 - Layer 6: Presentation layer
 - Layer 5: Session layer
 - Layer 4: Transport layer
 - Layer 3: Network layer
 - Layer 2: Data Link layer
 - Layer 1: Physical layer.

7. Which are the major functions and services performed by the physical layer?

- Ans.** The major functions and services performed by the physical layer are:
- Establishment and termination of a connection to a communications medium.
 - Participation in the process whereby the communication resources are effectively shared among multiple users. For example, contention resolution and flow control.
 - Modulation, or conversion between the representation of digital data in user equipment and the corresponding signals transmitted over a communications channel. These are signals operating over the physical cabling (such as copper and fiber optic) or over a radio link.

8. What does presentation layer do?

- Ans.** The Presentation layer transforms the data to provide a standard interface for the Application layer.

9. What does the session layer do?

- Ans.** The Session layer controls the dialogues/connections (sessions) between computers.

NOTES

10. What does transport layer provide?

Ans. The Transport layer provides transparent transfer of data between end users, providing reliable data transfer services to the upper layers.

11. What does network layer provide?

Ans. The Network layer provides the functional and procedural means of transferring variable length data sequences from a source to a destination via one or more networks while maintaining the quality of service requested by the Transport layer.

12. What does data link layer provide?

Ans. The Data Link layer provides the functional and procedural means to transfer data between network entities and to detect and possibly correct errors that may occur in the Physical layer.

13. What does physical layer define?

Ans. The Physical layer defines all the electrical and physical specifications for devices.

14. What is application layer?

Ans. The **application layer** is the seventh level of the seven-layer OSI model. It interfaces directly to and performs common application services for the application processes; it also issues requests to the presentation layer.

The common application layer services provide semantic conversion between associated application processes. Examples of common application services of general interest include the virtual file, virtual terminal, and job transfer and manipulation protocols.

The application layer of the four layer and five layer TCP/IP models corresponds to the application layer, the presentation layer and session layer in the seven layer OSI model.

ANSWERS

Multiple Choice Questions

- | | | | |
|-------|-------|-------|-------|
| 1. a | 2. b | 3. c | 4. a |
| 5. b. | 6. a | 7. b | 8. c |
| 9. a | 10. b | 11. a | 12. b |

True False Questions

- | | | | |
|-------|-------|--------|-------|
| 1. T | 2. F | 3. T | 4. T |
| 5. T | 6. T | 7. T | 8. F |
| 9. T | 10. T | 11. T | 12. T |
| 13. F | 14. T | 15. T | 16. F |
| 17. T | 18. T | 19. T | 20. T |
| 21. T | 22. T | 23. T | 24. T |
| 25. F | 26. F | 27. T. | |

Further Readings

1. **Computer Networks:** Ajit Kumar Singh, Firewall Media.
2. **Data and Computer Network Communication:** Prof. Shashi Banzai, Firewall Media.
3. **TCP / IP and Distributed System:** Vivek Archarya, Firewall Media.
4. **Networking:** Balvir Singh, Firewall Media.

UNIT 2

INTRODUCTION TO IPV6

NOTES

STRUCTURE

- 2.1 Why IPv6
- 2.2 Basic Protocol
- 2.3 Extension and Option Support for QoS
- 2.4 Security
- 2.5 Neighbor Discovery
- 2.6 Auto-Configuration
- 2.7 Routing
- 2.8 Changing to other protocols
- 2.9 Application programming interface for IPv6
- 2.10 6bone
 - Summary
 - Self Assessment Questions
 - Further Readings

Learning Objectives

After going through this unit, students will be able to:

- know about IPv6 in general.
- learn about the extension and option supports for QoS.
- know about Routing.
- learn about Auto-Configuration.
- learn about 6Bone.
- know about changing to other protocols.

2.1 WHY IPV6

Ane 2 (A)

NOTES

While CIDR may buy a few more years' time, everyone realizes that the days of IP in its current form (IPv4) are numbered. In addition to these technical problems, there is another issue looming in the background. Up until recently, the Internet has been used largely by universities, high-tech industry, and the government (especially the Dept. of Defense). With the explosion of interest in the Internet starting in the mid 1990s, it is likely that in the next millenium, it will be used by a much larger group of people, especially people with different requirements.

For one thing, millions of people with wireless portables may use it to keep in contact with their home bases. For another, with the impending convergence of the computer, communication, and entertainment industries, it may not be long before every television set in the world is an Internet node, producing a billion machines being used for video on demand. Under these circumstances, it became apparent that IP had to evolve and become more flexible.

Seeing these problems on the horizon, in 1990, IETF started work on a new version of IP, one which would never run out of addresses, would solve a variety of other problems, and be more flexible and efficient as well. Its major goals were to

1. Support billions of hosts, even with inefficient address space allocation.
2. Reduce the size of the routing tables.
3. Simplify the protocol, to allow routers to process packets faster.
4. Provide better security (authentication and privacy) than curreint IP.
5. Pay more attention to type of service, particularly for real-time data.
6. Aid multicasting by allowing scopes to be specified.
7. Make it possible for a host to roam without changing its address.
8. Allow the protocol to evolve in the future.
9. Permit the old and new protocols to coexist for years.

To find a protocol that met all these requirements, IETF issued a call for proposals and discussion in RFC 1550. Twenty-one responses were received, not all of them full proposals. By December 1992, seven serious proposals were on the table. They ranged from making minor patches to IP, to throwing it out altogether and replacing with a completely different protocol.

One proposal was to run TCP over CLNP, which, with its 160-bit addresses would have provided enough address space forever and would have unified two major network layer protocols. However, many people felt that this would have been an admission that something in the OSI world was actually done right, a statement considered Politically Incorrect in Internet circles. CLNP was patterned closely on IP, so the two are not really that different. In fact, the protocol ultimately chosen differs from IP far more than CLNP does. Another strike against CLNP was its poor support for service types, something required to transmit multimedia efficiently.

Three of the better proposals were published in IEEE Network. After much discussion, revision, and jockeying for position, a modified combined version of the Deering and Francis proposals, by now called SIPP (Simple Internet Protocol Plus)

was selected and given the designation IPv6 (IPv5 was already in use for an experimental real-time stream protocol).

IPv6 meets the goals fairly well. It maintains the good features of IP, discards or deemphasizes the bad ones, and adds new ones where needed. In general, IPv6 is not compatible with IPv4, but it is compatible with all the other Internet protocols, including TCP, UDP, ICMP, IGMP, OSPF, BGP, and DNS, sometimes with small modifications being required (mostly to deal with longer addresses). The main features of IPv6 are discussed below. More information about it can be found in RFC 1883 through RFC 1887.

NOTES

First and foremost, IPv6 has longer addresses than IPv4. They are 16 bytes long, which solves the problem that IPv6 was set out to solve: provide an effectively unlimited supply of Internet addresses. We will have more to say about addresses shortly.

The second major improvement of IPv6 is the simplification of the header. It contains only 7 fields (versus 13 in IPv4). This change allows routers to process packets faster and thus improve throughput. We will discuss the header shortly, too.

The third major improvement was better support for options. This change was essential with the new header because fields that previously were required are now optional. In addition, the way options are represented is different, making it simple for routers to skip over options not intended for them. This feature speeds up packet processing time.

A fourth area in which IPv6 represents a big advance is in security. IETF had its fill of newspaper stories about precocious 12-year-olds using their personal computers to break into banks and military bases all over the Internet. There was a strong feeling that something had to be done to improve security. Authentication and privacy are key features of the new IP.

Finally, more attention has been paid to (type of service) than in the past. IPv4 actually has an 8-bit field devoted to this matter, but with the expected growth in multimedia traffic in the future, much more is needed.

Internet Protocol version 6 (IPv6) is an Internet Protocol version which will succeed IPv4, the first implementation which is still in dominant use currently. It is an Internet Layer protocol for packet-switched internetworks. The main driving force for the redesign of Internet Protocol is the foreseeable IPv4 address exhaustion. IPv6 was defined in December 1998 by the Internet Engineering Task Force (IETF) with the publication of an Internet standard specification, RFC 2460.

IPv6 has a vastly larger address space than IPv4. This results from the use of a 128-bit address, whereas IPv4 uses only 32 bits. The new address space thus supports 2^{128} (about 3.4×10^{38}) addresses. This expansion provides flexibility in allocating addresses and routing traffic and eliminates the primary need for network address translation (NAT), which gained widespread deployment as an effort to alleviate IPv4 address exhaustion.

IPv6 also implements new features that simplify aspects of address assignment (stateless address autoconfiguration) and network renumbering (prefix and router announcements) when changing Internet connectivity providers. The IPv6 subnet size has been standardized by fixing the size of the host identifier portion of an address to 64 bits to facilitate an automatic mechanism for forming the host identifier

NOTES

The Internet Protocol Suite

Application Layer

BGP · DHCP · DNS · FTP · GTP · HTTP · IMAP · IRC · LDAP · Megaco ·
 MGCP · NNTP · NTP · POP · RIP · RPC · RTP · RTSP · SDP · SIP ·
 SMTP · SNMP · SOAP · SSH · Telnet · TLS/SSL · XMPP · (more)

Transport Layer

TCP · UDP · DCCP · SCTP · RSVP · ECN · (more)

Internet Layer

IP (IPv4, IPv6) · ICMP · ICMPv6 · IGMP · IPsec · (more)

Link Layer

ARP/InARP · NDP · OSPF · Tunnels (L2TP) · PPP · Media Access
 Control (Ethernet, DSL, ISDN, FDDI) · (more)

from Link Layer media addressing information (MAC address). Network security is integrated into the design of the IPv6 architecture. Internet Protocol Security (IPsec) was originally developed for IPv6, but found widespread optional deployment first in IPv4 (into which it was back-engineered). The IPv6 specifications mandate IPsec implementation as a fundamental interoperability requirement.

In December 2008, despite marking its 10th anniversary as a Standards Track protocol, IPv6 was only in its infancy in terms of general worldwide deployment. A 2008 study by Google Inc. indicated that penetration was still less than one percent of Internet-enabled hosts in any country. IPv6 has been implemented on all major operating systems in use in commercial, business, and home consumer environments

2.2 BASIC PROTOCOL

The first publicly used version of the Internet Protocol, Version 4 (IPv4), provides an addressing capability of about 4 billion addresses (2^{32}). This was deemed sufficient in the early design stages of the Internet when the explosive growth and worldwide proliferation of networks was not anticipated.

During the first decade of operation of the TCP/IP-based Internet, by the late 1980s, it became apparent that methods had to be developed to conserve address space. In the early 1990s, even after the introduction of classless network redesign, it became clear that this would not suffice to prevent IPv4 address exhaustion and that further

changes to the Internet infrastructure were needed. By the beginning of 1992, several proposed systems were being circulated, and by the end of 1992, the IETF announced a call for white papers (RFC 1550) and the creation of the "IP Next Generation" (IPng) area of working groups.

The Internet Engineering Task Force adopted IPng on July 25, 1994, with the formation of several IPng working groups. By 1996, a series of RFCs were released defining Internet Protocol Version 6 (IPv6), starting with RFC 1883.

Incidentally, the IPng architects could not use version number 5 as a successor to IPv4, because it had been assigned to an experimental flow-oriented streaming protocol (Internet Stream Protocol), similar to IPv4, intended to support video and audio.

It is widely expected that IPv4 will be supported alongside IPv6 for the foreseeable future. IPv4-only nodes are not able to communicate directly with IPv6 nodes, and will need assistance from an intermediary.

2.2.1 IPv4 exhaustion

Estimates of the time frame until complete exhaustion of IPv4 addresses varied widely. In 2003, Paul Wilson (director of APNIC) stated that, based on then-current rates of deployment, the available space would last for one or two decades. In September 2005, a report by Cisco Systems suggested that the pool of available addresses would dry up in as little as 4 to 5 years. As of May 2009, a daily updated report projected that the IANA pool of unallocated addresses would be exhausted in June 2011, with the various Regional Internet Registries using up their allocations from IANA in March 2012. There is now consensus among Regional Internet Registries that final milestones of the exhaustion process will be passed in 2010 or 2011 at the latest, and a policy process has started for the end-game and post-exhaustion era.

Even though consumers are most likely to suffer when their equipment has to be replaced they tend to look at networking devices like household appliances that only rarely need repairs and never have to be configured or updated. Commercial grade equipment is more likely to support IPv6, so it is the small consumer with his cost-effective disposable networking technology who will be most affected by the eventual change from IPv4 to IPv6.

Smart equipment that contains software needs explicit IPv6 support. Lower-level equipment like cables, network adapters, and switches may not be affected by the change. In general, layer-1 and layer-2 equipment won't require updates.

As of 2010, IPv6 readiness is not considered in most consumer purchasing decisions. We could be only years from a universal upgrade to IPv6 driven internet where devices without IPv6 support will not function.

IPv6 compatibility is mainly a software/firmware issue like the year-2000. Unlike the year-2000 issue, there is little interest in ensuring compatibility of older equipment and software by manufacturers. The realization that IPv4 exhaustion is imminent is recent and manufacturers haven't shown much initiative in updating equipment. There is hope that a combined IPv4/IPv6 internet will streamline the transition. The internet community is divided on the issue of whether the transition should be a quick switch or a longer process. It has been suggested that all internet servers be prepared to serve IPv6-only clients by 2012. Universal access to IPv6-only servers will be even more of a challenge.

NOTES

NOTES

Most equipment would be fully IPv6 capable with a software or firmware update if the device has sufficient storage and memory space for the new IPv6 stack. However, as with 64-bit Windows and Wi-Fi Protected Access support, manufacturers are unlikely to spend on development costs for hardware they have already sold when they are poised to make more sales from “IPv6-ready” equipment.

The CableLabs consortium published the 160 Mbit/s DOCSIS 3.0 IPv6-ready specification for cable modems in August 2006. The widely used DOCSIS 2.0 does not support IPv6. The new ‘DOCSIS 2.0 + IPv6’ standard also supports IPv6, which may on the cable modem side only require a firmware upgrade. It is expected that only 60% of cable modems’ servers and 40% of cable modems will be DOCSIS 3.0 by 2011.

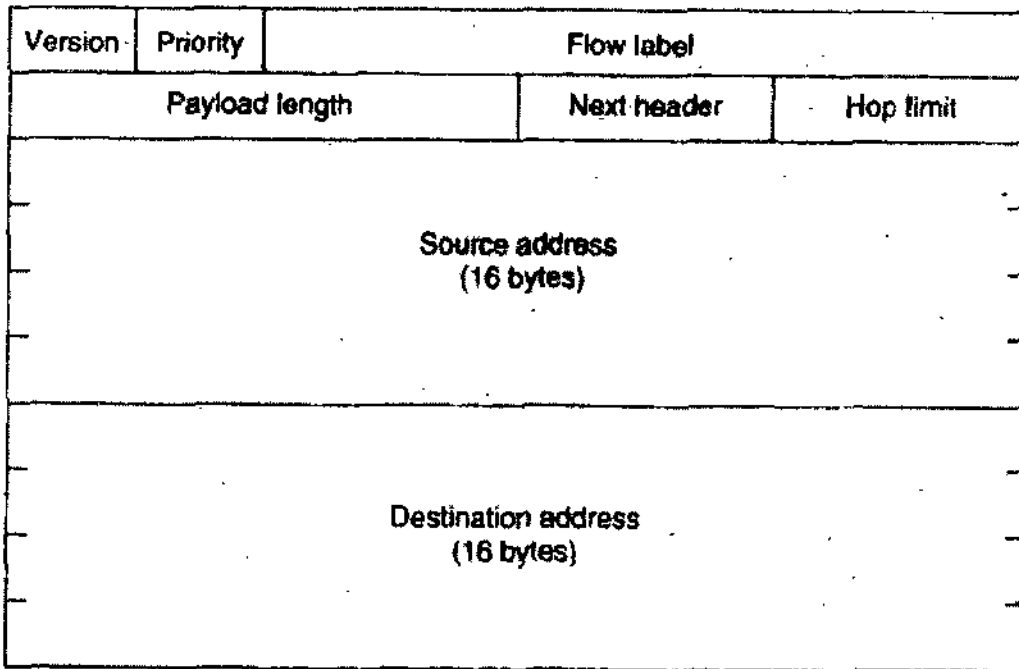

Other equipment which is typically not IPv6-ready ranges from Skype and SIP phones to oscilloscopes and printers. Professional network routers in use should be IPv6-ready. Most personal computers should also be IPv6-ready because the network stack resides in the operating system. Most applications with network capabilities are not ready but could be upgraded with support from the developers. Since Java 1.4 (February 2002) all applications that are 100% Java compatible have support for IPv6 addresses.

2.2.2 The Main IPv6 Header

The IPv6 header is shown in next figure. The Version field is always 6 for IPv6 (and 4 for IPv4). During the transition period from IPv4, which will probably take a decade, routers will be able to examine this field to tell what kind of packet they have. As an aside, making this test wastes a few instructions in the critical path, so many implementations are likely to try to avoid it by using some field in the data link header to distinguish IPv4 packets from IPv6 packets. In this way, packets can be passed to the correct network layer handler directly. However, having the data link layer be aware of network packet types completely violates the design principle that each layer should not be aware of the meaning of the bits given to it from the layer above. The discussions between the “Do it right” and “Make it fast” camps will no doubt be lengthy and vigorous.

The Priority field is used to distinguish between packets whose sources can be flow controlled and those that cannot. Values 0 through 7 are for transmissions that are capable of slowing down in the event of congestion. Values 8 through 15 are for real-time traffic whose sending rate is constant, even if all the packets are being lost. Audio and video fall into the latter category. This distinction allows routers to deal with packets better in the event of congestion. Within each group, lower-numbered packets are less important than higher-numbered ones. The IPv6 standard suggests, for example, to use 1 for news, 4 for FTP, and 6 for Telnet connections, since delaying a news packet for a few seconds is not noticeable, but delaying a Telnet packet certainly is.

The Flow label field is still experimental but will be used to allow a source and destination to set up a pseudoconnection with particular properties and requirements. For example, a stream of packets from one process on a certain source host to a certain process on a certain destination host might have stringent delay requirements and thus need reserved bandwidth. The flow can be set up in advance and given an identifier. When a packet with a nonzero Flow label shows up, all the routers can look it up in internal tables to see what kind of special treatment it requires. In effect,



NOTES

flows are an attempt to have it both ways: the flexibility of a datagram subnet and the guarantees of a virtual circuit subnet.

Each flow is designated by the source address, destination address, and flow number, so many flows may be active at the same time between a given pair of IP addresses. Also, in this way, even if two flows coming from different hosts but with the same flow number pass through the same router, the router will be able to tell them apart using the source and destination addresses. It is expected that flow numbers will be chosen randomly, rather than assigned sequentially starting at 1, to make it easy for routers to hash them.

The Payload length field tells how many bytes follow the 40-byte header of the figure. The name was changed from the IPv4 Total length field because the meaning was changed slightly: the 40 header bytes are no longer counted as part of the length as they used to be.

The Next header field lets the cat out of the bag. The reason the header could be simplified is that there can be additional (optional) extension headers. This field tells which of the (currently) six extension headers, if any, follows this one. If this header is the last IP header, the Next header field tells which transport protocol handler (e.g., TCP, UDP) to pass the packet to.

The Hop limit field is used to keep packets from living forever. It is, in practice, the same as the Time to live field in IPv4, namely, a field that is decremented on each hop. In theory, in IPv4 it was a time in seconds, but no router used it that way, so the name was changed to reflect the way it is actually used.

Next come the Source address and Destination address fields. Deering's original proposal, SIP, used 8-byte addresses, but during the review process many people felt

STUDENT ACTIVITY 2.1

- 1. What is IPv6?

- 2. Describe the various security factors.

3. Which are the basic protocols of IPv6?

4. Describe the Extension and Option Support for QoS.

NOTES

that with 8-byte addresses IPv6 would run out of addresses within a few decades, whereas with 16-byte addresses it would never run out. Other people argued that 16 bytes was overkill, whereas still others favored using 20-byte addresses to be compatible with the OSI datagram protocol. Another faction wanted variable-sized addresses. After much discussion, it was decided that fixed-length 16-byte addresses were the best compromise.

The IPv6 address space is divided up as shown in next figure. Addresses beginning with 80 zeros are reserved for IPv4 addresses. Two variants are supported, distinguished by the next 16 bits. These variants relate to how IPv6 packets will be tunneled over the existing IPv4 infrastructure.

Prefix(binary)	Usage	Fraction
00000000	Reserved(includingIPv4)	1/256
00000001	Unassigned	1/256
0000001	OSINSAPaddresses	1/128
0000010	NovellNetWareIPXaddresses	1/128
0000011	Unassigned	1/128
00001	Unassigned	1/32
0001	Unassigned	1/16
001	Unassigned	1/8
010	Provider-basedaddresses	1/8
011	Unassigned	1/8
100	Geographic-basedaddresses	1/8
101	Unassigned	1/8
110	Unassigned	1/8
1110	Unassigned	1/16
11110	Unassigned	1/32
111110	Unassigned	1/64
1111110	Unassigned	1/128
111111100	Unassigned	1/512
1111111010	Linklocaluseaddresses	1/1024
1111111011	Sitelocal useaddresses	1/1024
11111111	Multicast	1/256

Figure: IPV6 addressed

The use of separate prefixes for provider-based and geographic-based addresses is a compromise between two different visions of the future of the Internet. Provider-based addresses make sense if you think that in the future there will be some number of companies providing Internet service to customers, analogous to AT&T, MCI, Sprint, British Telecom, and so on providing telephone service now. Each of these

companies will be given some fraction of the address space. The first 5 bits following the 010 prefix are used to indicate which registry to look the provider up in. Currently three registries are operating, for North America, Europe, and Asia. Up to 29 new registries can be added later.

Each registry is free to divide up the remaining 15 bytes as it sees fit. It is expected that many of them will use a 3-byte provider number, giving about 16 million providers, in order to allow large companies to act as their own provider. Another possibility is to use 1 byte to indicate national providers and let them do further allocation. In this manner, additional levels of hierarchy can be introduced as needed.

The geographic model is the same as the current Internet, in which providers do not play a large role. In this way, IPv6 can handle both kinds of addresses.

The link and site local addresses have only a local significance. They can be reused at each organization without conflict. They cannot be propagated outside organizational boundaries, making them well suited to organizations that currently use firewalls to wall themselves off from the rest of the Internet.

Multicast addresses have a 4-bit flag field and a 4-bit scope field following the prefix, then a 112-bit group identifier. One of the flag bits distinguishes permanent from transient groups. The scope field allows a multicast to be limited to the current link, site, organization, or planet. These four scopes are spread out over the 16 values to allow new scopes to be added later. For example, the planetary scope is 14, so code 15 is available to allow future expansion of the Internet to other planets, solar systems, and galaxies.

In addition to supporting the standard unicast (point-to-point) and multicast addresses, IPv6 also supports a new kind of addressing: anycast. Anycasting is like multicasting in that the destination is a group of addresses, but instead of trying to deliver the packet to all of them, it tries to deliver it to just one, usually the nearest one. For example, when contacting a group of cooperating file servers, a client can use anycast to reach the nearest one, without having to know which one that is. Anycasting uses regular unicast addresses. It is up to the routing system to choose the lucky host that gets the packet.

A new notation has been devised for writing 16-byte addresses. They are written as eight groups of four hexadecimal digits with colons between the groups, like this:

```
8000:0000:0000:0000:0123:4567:89AB:CDEF
```

Since many addresses will have many zeros inside them, three optimizations have been authorized. First, leading zeros within a group can be omitted, so 0123 can be written as 123. Second, one or more groups of 16 zeros can be replaced by a pair of colons. Thus the above address now becomes

```
8000::123:4567:89AB:CDEF
```

Finally, IPv4 addresses can be written as a pair of colons and an old dotted decimal number, for example

```
::192.31.20.46
```

Perhaps it is unnecessary to be so explicit about it, but there are a lot of 16-byte addresses. Specifically, there are 2128 of them, which is approximately 3×10^{38} . If the entire earth, land and water, were covered with computers, IPv6 would allow 7×10^{23} IP addresses per square meter. Students of chemistry will notice that this number

NOTES

NOTES

is larger than Avogadro's number. While it was not the intention to give every molecule on the surface of the earth its own IP address, we are not that far off.

In practice, the address space will not be used efficiently, just as the telephone number address space is not (the area code for Manhattan, 212, is nearly full, but that for Wyoming, 307, is nearly empty). In RFC 1715, Huitema calculated that using the allocation of telephone numbers as a guide, even in the most pessimistic scenario, there will still be well over 1000 IP addresses per square meter of the earth's surface (land and water). In any likely scenario, there will be trillions of them per square meter. In short, it seems unlikely that we will run out in the foreseeable future. It is also worth noting that only 28 percent of the address space has been allocated so far. The other 72 percent is available for future purposes not yet thought of.

It is instructive to compare the IPv4 header with the IPv6 header to see what has been left out in IPv6. The IHL field is gone because the IPv6 header has a fixed length. The Protocol field was taken out because the Next header field tells what follows the last IP header (e.g., a UDP or TCP segment).

All the fields relating to fragmentation were removed because IPv6 takes a different approach to fragmentation. To start with, all IPv6 conformant hosts and routers must support packets of 576 bytes. This rule makes fragmentation less likely to occur in the first place. In addition, when a host sends an IPv6 packet that is too large, instead of fragmenting it, the router that is unable to forward it sends back an error message. This message tells the host to break up all future packets to that destination. Having the host send packets that are the right size in the first place is ultimately much more efficient than having the routers fragment them on the fly.

Finally, the Checksum field is gone because calculating it greatly reduces performance. With the reliable networks now used, combined with the fact that the data link layer and transport layers normally have their own checksums, the value of yet another checksum was not worth the performance price it extracted. Removing all these features has resulted in a lean and mean network layer protocol. Thus the goal of IPv6—a fast, yet flexible, protocol with plenty of address space—has been met by this design.

2.2.3 Extension Headers

Nevertheless, some of the missing fields are occasionally still needed so IPv6 has introduced the concept of an (optional) extension header. These headers can be supplied to provide extra information, as listed next. Each one is optional, but if more than one is present, they must appear directly after the fixed header, and preferably in the order listed.

Extension header	Description
Hop-by-hop options	Miscellaneous information for routers
Routing	Full or partial route to follow
Fragmentation	Management of datagram fragments
Authentication	Verification of the sender's identity
Encrypted security payload	Information about the encrypted contents
Destination options	Additional information for the destination

Fig.: IPv6 extension headers

Some of the headers have a fixed format; others contain a variable number of variable-length fields. For these, each item is encoded as a (Type, Length, Value) tuple. The Type is a 1-byte field telling which option this is. The Type values have been chosen so that the first 2 bits tell routers that do not know how to process the option what to do. The choices are: skip the option, discard the packet, discard the packet and send back an ICMP packet, and the same as the previous one, except do not send ICMP packets for multicast addresses (to prevent one bad multicast packet from generating millions of ICMP reports).

The Length is also a 1-byte field. It tells how long the value is (0 to 255 bytes). The Value is any information required, up to 255 bytes.

The hop-by-hop header is used for information that all routers along the path must examine. So far, one option has been defined: support of datagrams exceeding 64K. The format of this header is shown in next figure.

Next header	0	194	0
Jumbo payload length			

Figure : The hop-by-hop extension header for large datagrams (jumbograms).

As with all extension headers, this one starts out with a byte telling what kind of header comes next. This byte is followed by one telling how long the hop-by-hop header is in bytes, excluding the first 8 bytes, which are mandatory. The next 2 bytes indicate that this option defines the datagram size (code 194) as a 4-byte number. The last 4 bytes give the size of the datagram. Sizes less than 65,536 are not permitted and will result in the first router discarding the packet and sending back an ICMP error message. Datagrams using this header extension are called jumbograms. The use of jumbograms is important for supercomputer applications that must transfer gigabytes of data efficiently across the Internet.

The routing header lists one or more routers that must be visited on the way to the destination. Both strict routing (the full path is supplied) and loose routing (only selected routers are supplied) are available, but they are combined. The format of the routing header is shown in the next figure.

Next header	0	Number of addresses	Next address
Bit map			
1 - 24 Adresses			

Figure : The extension header for routing

The first 4 bytes of the routing extension header contain four 1-byte integers. the next header type, the routing type (currently 0), the number of addresses present in this header (1 to 24), and the index of the next address to visit. The latter field starts at 0 and is incremented as each address is visited. Then comes a reserved byte followed by a bit map with bits for each of the 24 potential IPv6 addresses following it. These bits tell whether each address must be visited directly after the one before it (strict source routing), or whether other routers may come in between (loose source routing).

The fragment header deals with fragmentation similarly to the way IPv4 does. The header holds the datagram identifier, fragment number, and a bit telling whether more fragments will follow. In IPv6, unlike in IPv4, only the source host can fragment a packet. Routers along the way may not do this. Although this change is a major

philosophical break with the past, it simplifies the routers' work and makes routing go faster. As mentioned above, if a router is confronted with a packet that is too big, it discards the packet and sends an ICMP packet back to the source. This information allows the source host to fragment the packet into smaller pieces using this header and try again.

NOTES

The authentication header provides a mechanism by which the receiver of a packet can be sure of who sent it. With IPv4, no such guarantee is present. The encrypted security payload makes it possible to encrypt the contents of a packet so that only the intended recipient can read it. These headers use cryptographic techniques to accomplish their missions. We will give brief descriptions below, but readers not already familiar with modern cryptography may not understand the full description now.

When a sender and receiver wish to communicate securely, they must first agree on one or more secret keys that only they know. How they do this is outside the scope of IPv6. Each of these keys is assigned a unique 32-bit key number. The key numbers are global, so that if Alice is using key 4 to talk to Bob, she cannot also have a key 4 to talk to Carol. Associated with each key number are other parameters, such as key lifetime, and so on.

To send an authenticated message, the sender first constructs a packet consisting of all the IP headers and the payload and then replaces the fields that change underway (e.g., Hop limit) with zeros. The packet is then padded out with zeros to a multiple of 16 bytes. Similarly, the secret key to be used is also padded out with zeros to a multiple of 16 bytes. Now a cryptographic checksum is computed on the concatenation of the padded secret key, the padded packet, and the padded secret key again. Users may define their own cryptographic checksum algorithms, but cryptographically unsophisticated users should use the default algorithm, MD5.

Now we come to the role of the authentication header. Basically, it contains three parts. The first part consists of 4 bytes holding the next header number, the length of the authentication header, and 16 zero bits. Then comes the 32-bit key number. Finally, the MD5 (or other) checksum is included.

The receiver then uses the key number to find the secret key. The padded version of it is then prepended and appended to the padded payload, the variable header fields are zeroed out, and the checksum computed. If it agrees with the checksum included in the authentication header, the receiver can be sure that the packet came from the sender with whom the secret key is shared and also be sure that the packet was not tampered with underway. The properties of MD5 make it computationally infeasible for an intruder to forge the sender's identity or modify the packet in a way that escapes detection.

It is important to note that the payload of an authenticated packet is sent unencrypted. Any router along the way can read what it says. For many applications, secrecy is not really important, just authentication. For example, if a user instructs his bank to pay his telephone bill, there is probably no real need for secrecy, but there is a very real need for the bank to be absolutely sure it knows who sent the packet containing the payment order.

For packets that must be sent secretly, the encrypted security payload extension header is used. It starts out with a 32-bit key number, followed by the encrypted payload. The encryption algorithm is up to the sender and receiver, but DES in cipher

block chaining mode is the default. When DES-CBC is used, the payload field starts out with the initialization vector (a multiple of 4 bytes), then the payload, then padding out to multiple of 8 bytes. If both encryption and authentication are desired, both headers are needed.

The destination options header is intended for fields that need only be interpreted at the destination host. In the initial version of IPv6, the only options defined are null options for padding this header out to a multiple of 8 bytes, so initially it will not be used. It was included to make sure that new routing and host software can handle it, in case someone thinks of a destination option some day.

NOTES

2.3 EXTENSION AND OPTION SUPPORT FOR QOS

IPv4 has a fixed size (40 octets) of option parameters. In IPv6, options are implemented as additional extension headers after the IPv6 header, which limits their size only by the size of an entire packet. The extension header mechanism allows IPv6 to be easily 'extended' to support future services for QoS, security, mobility, etc. without a redesign of the basic protocol.

2.4 SECURITY

Internet Protocol Security (IPsec), the protocol for IP encryption and authentication, forms an integral part of the base protocol suite in IPv6. IPsec support is mandatory in IPv6; this is unlike IPv4, where it is optional (but usually implemented). IPsec, however, is not widely used at present except for securing traffic between IPv6 Border Gateway Protocol routers.

2.5 NEIGHBOR DISCOVERY

Although IPv4 address exhaustion has been slowed by the introduction of classless inter-domain routing (CIDR) and the extensive use of network address translation (NAT), address uptake has accelerated again in recent years. Most forecasts expect complete depletion between 2010 and 2012.

As of 2008, IPv6 accounts for a minuscule fraction of the used addresses and the traffic in the publicly-accessible Internet which is still dominated by IPv4.

The 2008 Summer Olympic Games were a notable event in terms of IPv6 deployment, being the first time a major world event has had a presence on the IPv6 Internet at <http://ipv6.beijing2008.cn/en> (IP addresses 2001:252:0:1::2008:6 and 2001:252:0:1::2008:8) and all network operations of the Games were conducted using IPv6. It is believed that the Olympics provided the largest showcase of IPv6 technology since the inception of IPv6.

Cellular telephone systems present a large deployment field for Internet Protocol devices as mobile telephone service is being transitioned from 3G systems to next generation (4G) technologies in which voice is provisioned as a Voice over Internet Protocol (VoIP) service. This mandates the use of IPv6 for such networks due to the impending IPv4 address exhaustion. In the U.S., cellular operator Verizon has released technical specifications for devices operating on its future networks. The specification

mandates IPv6 operation according to the 3GPP Release 8 Specifications (March 2009) and deprecates IPv4 as an optional capability.

Some implementations of the BitTorrent peer-to-peer file transfer protocol make extensive use of IPv6 to avoid NAT issues.

NOTES

2.5.1 Major announcements and availability

Year Announcements and availability

- 1996 Alpha quality IPv6 support in Linux kernel development version 2.1.8.
6bone (an IPv6 virtual network for testing) is started.
- 1997 By the end of 1997 IBM's AIX 4.3 is the first commercial platform supporting IPv6.
Also in 1997, Early Adopter Kits for DEC's operating systems, Tru64 and OpenVMS, are made available.
- 1998 Microsoft Research[34] releases its first experimental IPv6 stack. This support is not intended for use in a production environment.
- 1999 The Freenet6 tunnel broker service is launched.
- 2000 Production-quality BSD support for IPv6 becomes generally available in early to mid-2000 in FreeBSD, OpenBSD, and NetBSD via the KAME project.
Microsoft releases an IPv6 technology preview version for Windows 2000 in March 2000.
Sun Solaris supports IPv6 in Solaris 8 in February.
Compaq ships IPv6 with Tru64.
- 2001 In January, Compaq ships IPv6 with OpenVMS.
Cisco Systems introduces IPv6 support on Cisco IOS routers and L3 switches.
HP introduces IPv6 with HP-UX 11i v1.
- 2002 Microsoft Windows NT 4.0 and Windows 2000 SP1 have limited IPv6 support for research and testing since at least 2002.
Microsoft Windows XP (2001) supports IPv6 for developmental purposes. In Windows XP SP1 (2002) and Windows Server 2003, IPv6 is included as a core networking technology, suitable for commercial deployment.
IBM z/OS supports IPv6 since version 1.4 (generally availability in September 2002).
- 2003 Apple Mac OS X v10.3 "Panther" (2003) supports IPv6 which is enabled by default.
In July, ICANN announces that IPv6 address records for the Japan (jp) and Korea (kr) country code top-level domain nameservers are visible in the DNS root server zone files with serial number 2004072000. The IPv6 records for France (fr) are added later. This makes IPv6 DNS publicly operational.
- 2005 Linux 2.6.12 removes experimental status from its IPv6 implementation.
- 2007 Microsoft Windows Vista (2007) supports IPv6 which is enabled by default.

Apple's AirPort Extreme 802.11n base station includes an IPv6 gateway in its default configuration. It uses 6to4 tunneling and manually configured static tunnels. (Note: 6to4 was disabled by default in later firmware revisions.)

2008 On February 4, 2008, IANA adds AAAA records for the IPv6 addresses of six root name servers. With this transition, it is now possible for two Internet hosts to fully communicate without using IPv4.

On March 12, 2008, Google launches a public IPv6 web interface to its popular search engine at the URL <http://ipv6.google.com>.

2009 In January 2009, Google extends its IPv6 initiative with Google over IPv6, which offers IPv6 support for Google services to compatible networks.

2010 In January 2010, Comcast announces public trials of IPv6 on its production network. In April 2010, XS4ALL announced public trials & Verizon announced testing on its FiOS network.

In May/June 2010, Facebook became accessible on IPv6 via <http://www.v6.facebook.com/>

NOTES

2.6 AUTO-CONFIGURATION

IPv6 hosts can configure themselves automatically when connected to a routed IPv6 network using ICMPv6 router discovery messages. When first connected to a network, a host sends a link-local multicast router solicitation request for its configuration parameters; if configured suitably, routers respond to such a request with a router advertisement packet that contains network-layer configuration parameters.

If IPv6 stateless address autoconfiguration is unsuitable for an application, a network may use stateful configuration with the Dynamic Host Configuration Protocol for IPv6 (DHCPv6) or hosts may be configured statically.

Routers present a special case of requirements for address configuration, as they often are sources for autoconfiguration information, such as router and prefix advertisements. Stateless configuration for routers can be achieved with a special router renumbering protocol.

2.7 ROUTING

AM. 3 (C)

A number of simplifications have been made to the packet header, and the process of packet forwarding has been simplified, in order to make packet processing by routers simpler and hence more efficient. Specifically:

- The packet header in IPv6 is simpler than that used in IPv4, with many rarely used fields moved to separate options; as a result, although the addresses in IPv6 are four times larger, the option-less IPv6 header is only twice the size of the option-less IPv4 header.
- IPv6 routers do not perform fragmentation. IPv6 hosts are required to either perform PMTU discovery, perform end-to-end fragmentation, or to send packets smaller than the IPv6 minimum MTU size of 1280 octets.

- Ann 3(LC)
- The IPv6 header is not protected by a checksum; integrity protection is assumed to be assured by both a link layer checksum and a higher layer (TCP, UDP, etc.) checksum. (UDP/IPv4 may actually have a checksum of 0 indicating no checksum; IPv6 requires UDP must have its own checksum.) In effect, IPv6 routers do not need to recompute a checksum when header fields (such as the TTL or Hop Count) change. This improvement may have been made less necessary by the development of routers that perform checksum computation at link speed using dedicated hardware, but it is still relevant for software based routers.
 - The Time-to-Live field of IPv4 has been renamed to Hop Limit, reflecting the fact that routers are no longer expected to compute the time a packet has spent in a queue.

2.8 CHANGING TO OTHER PROTOCOLS

After the Regional Internet Registries have exhausted their pools of available IPv4 addresses, it is likely that hosts newly added to the Internet might only have IPv6 connectivity. For these clients to have backward-compatible connectivity to existing IPv4-only resources, suitable translation mechanisms must be deployed.

One form of translation is the use of a dual-stack application-layer proxy; for example a web proxy.

NAT-like techniques for application-agnostic translation at the lower layers have also been proposed. Most have been found to be too unreliable in practice because of the wide range of functionality required by common application-layer protocols, and are considered by many to be obsolete.

2.9 APPLICATION PROGRAMMING INTERFACE FOR IPV6

The increased length of network addresses emphasizes a most important change when moving from IPv4 to IPv6. IPv6 addresses are 128 bits long, whereas IPv4 addresses are 32 bits; where the IPv4 address space contains roughly 4.3×10^9 (4.3 billion) addresses, IPv6 has enough room for 3.4×10^{38} (340 trillion trillion trillion) unique addresses.

IPv6 addresses are normally written with hexadecimal digits and colon separators like 2001:db8:85a3::8a2c:370:7334, as opposed to the dot-decimal notation of the 32 bit IPv4 addresses. IPv6 addresses are typically composed of two logical parts: a 64-bit (sub-)network prefix, and a 64-bit host part.

IPv6 addresses are classified into three types: unicast addresses which uniquely identify network interfaces, anycast addresses which identify a group of interfaces—mostly at different locations—for which traffic flows to the nearest one, and multicast addresses which are used to deliver one packet to many interfaces. Broadcast addresses are not used in IPv6. Each IPv6 address also has a 'scope', which specifies in which part of the network it is valid and unique. Some addresses have node scope or link scope; most addresses have global scope (i.e. they are unique globally).

Some IPv6 addresses are used for special purposes, like the loopback address. Also, some address ranges are considered special, like link-local addresses (for use in the local network only) and solicited-node multicast addresses (used in the Neighbor Discovery Protocol).

2. ~~10~~ 6BONE

AW 3 (B)

NOTES

The 6bone was a testbed for Internet Protocol version 6; it was an outgrowth of the IETF IPng project that created the IPv6 protocols intended to eventually replace the current Internet network layer protocols known as IPv4. The 6bone was started outside the official IETF process at the March 1996 IETF meetings, and became a worldwide informal collaborative project, with eventual oversight from the "NGtrans" (IPv6 Transition) Working Group of the IETF.

The original mission of the 6bone was to establish a network to foster the development, testing, and deployment of IPv6 using a model to be based upon the experiences from the Mbone, hence the name "6bone".

The 6bone started as a virtual network (using IPv6 over IPv4 tunneling/encapsulation) operating over the IPv4-based Internet to support IPv6 transport, and slowly added native links specifically for IPv6 transport. Although the initial 6bone focus was on testing of standards and implementations, the eventual focus became more on testing of transition and operational procedures, as well as actual IPv6 network usage.

The 6bone operated under the IPv6 Testing Address Allocation, which specified the 3FFE::/16 IPv6 prefix for 6bone testing purposes.

At its peak in mid-2003, over 150 6bone top level 3FFE::/16 network prefixes were routed, interconnecting over 1000 sites in more than 50 countries. When it became obvious that the availability of IPv6 top level production prefixes was assured, and that commercial and private IPv6 networks were being operated outside the 6bone using these prefixes, a plan was developed to phase out the 6bone (see RFC 3701).

The phaseout plan called for a halt to new 6bone prefix allocations on 1 January 2004 and the complete cessation of 6bone operation and routing over the 6bone testing prefixes on 6 June 2006. Addresses within the 6bone testing prefix have now reverted to the IANA, and should no longer be used.

SUMMARY

1. Internet Protocol version 6 (IPv6) is an Internet Protocol version which will succeed IPv4, the first implementation which is still in dominant use currently.
2. The first publicly used version of the Internet Protocol, Version 4 (IPv4), provides an addressing capability of about 4 billion addresses (2^{32}).
3. Estimates of the time frame until complete exhaustion of IPv4 addresses varied widely.
4. Most equipment would be fully IPv6 capable with a software or firmware update if the device has sufficient storage and memory space for the new IPv6 stack.
5. In IPv6, options are implemented as additional extension headers after the IPv6 header, which limits their size only by the size of an entire packet.
6. IPv6 hosts can configure themselves automatically when connected to a routed IPv6 network using ICMPv6 router discovery messages.

NOTES

7. IPv6 routers do not perform fragmentation. IPv6 hosts are required to either perform PMTU discovery, perform end-to-end fragmentation, or to send packets smaller than the IPv6 minimum MTU size of 1280 octets.
8. After the Regional Internet Registries have exhausted their pools of available IPv4 addresses, it is likely that hosts newly added to the Internet might only have IPv6 connectivity.
9. IPv6 addresses are 128 bits long, whereas IPv4 addresses are 32 bits; where the IPv4 address space contains roughly 4.3×10^9 (4.3 billion) addresses, IPv6 has enough room for 3.4×10^{38} (340 trillion trillion trillion) unique addresses.
10. The 6bone was a testbed for Internet Protocol version 6; it was an outgrowth of the IETF IPng project that created the IPv6 protocols intended to eventually replace the current Internet network layer protocols known as IPv4.

SELFASSESSMENT QUESTIONS

1. What is IPv6?
2. Which are the basic protocols of IPv6?
3. Describe the Extension and Option Support for QoS.
4. Describe the various security factors.
5. What do you understand by Neighbor Discovery?
6. Describe the details the various announcements and availabilities.
7. What do you understand by Auto-Configuration?
8. What is Routing?
9. How would you change to other Protocols?
10. Describe the Application Programming Interface for IPv6.
11. What is 6bone?

Multiple Choice Questions

1. IPv6 is:

(a) Internet Protocol version 6	(b) Intranet Protocol version 6
(c) Internet Provider version 6	
2. CIDR is:

(a) Classless Inter-domain Route	(b) Classless Inter-domain Routing
(c) Class Inter-domain Routing	
3. VoIP is:

(a) Voice over Internet Port	(b) Vocal over Internet Protocol
(c) Voice over Internet Protocol	
4. IPsec is:

(a) Internet Protocol Security	(b) Internet Port Security
(c) Intranet Protocol Survey	
5. DHCPv6 is:

(a) Data Host Configuration Protocol for IPv6	
(b) Dynamic Host Configuration Protocol for IPv6	
(c) Dynamic Host Configuration Port for IPv6	

True/False Questions

1. The first publicly used version of the Internet Protocol, Version 4 (IPv4), provides an addressing capability of about 4 billion addresses (2^{32}).
2. Most equipment would be fully IPv6 capable with a software or firmware update if the device has sufficient storage and memory space for the new IPv6 stack.
3. **IPv6 hosts cannot** configure themselves automatically when connected to a routed IPv6 network using ICMPv6 router discovery messages.
4. After the Regional Internet Registries have exhausted their pools of available IPv4 addresses, it is likely that hosts newly added to the Internet might only have IPv6 connectivity.
5. **The 6bone was not a testbed** for Internet Protocol version 6; it was an outgrowth of the IETF IPng project that created the IPv6 protocols intended to eventually replace the current Internet network layer protocols known as IPv4.

NOTES

Short Questions with Answers

1. What is IPv6?

Ans. Internet Protocol version 6 (IPv6) is an Internet Protocol version which will succeed IPv4, the first implementation which is still in dominant use currently. It is an Internet Layer protocol for packet-switched internetworks. The main driving force for the redesign of Internet Protocol is the foreseeable IPv4 address exhaustion. IPv6 was defined in December 1998 by the Internet Engineering Task Force (IETF) with the publication of an Internet standard specification, RFC 2460.

2. Which are the basic protocols used for IPv6?

Ans. The first publicly used version of the Internet Protocol, Version 4 (IPv4), provides an addressing capability of about 4 billion addresses (2^{32}). This was deemed sufficient in the early design stages of the Internet when the explosive growth and worldwide proliferation of networks was not anticipated. The Internet Engineering Task Force adopted IPng on July 25, 1994, with the formation of several IPng working groups.

By 1996, a series of RFCs were released defining Internet Protocol Version 6 (IPv6), starting with RFC 1883. Incidentally, the IPng architects could not use version number 5 as a successor to IPv4, because it had been assigned to an experimental flow-oriented streaming protocol (Internet Stream Protocol), similar to IPv4, intended to support video and audio.

It is widely expected that IPv4 will be supported alongside IPv6 for the foreseeable future. IPv4-only nodes are not able to communicate directly with IPv6 nodes, and will need assistance from an intermediary.

3. What is the advantage of IPv6?

Ans. IPv6 compatibility is mainly a software/firmware issue like the year-2000. Unlike the year-2000 issue, there is little interest in ensuring compatibility of older equipment and software by manufacturers. The realization that IPv4 exhaustion is imminent is recent and manufacturers haven't shown much initiative in updating equipment. There is hope that a combined IPv4/IPv6 internet will streamline the transition. The internet community is divided on the issue of whether the transition should be a quick switch or a longer process. It has been suggested that all internet servers be prepared to serve IPv6-only clients by 2012. Universal access to IPv6-only servers will be even more of a challenge.

4. What is the major difference between IPv4 and IPv6?

Ans. IPv4 has a fixed size (40 octets) of option parameters. In IPv6, options are implemented as additional extension headers after the IPv6 header, which limits their size only by the size of an entire packet. The extension header mechanism allows IPv6 to be easily 'extended' to support future services for QoS, security, mobility, etc. without a redesign of the basic protocol.

NOTES

5. What is VoIP?
- Ans. Cellular telephone systems present a large deployment field for Internet Protocol devices as mobile telephone service is being transitioned from 3G systems to next generation (4G) technologies in which voice is provisioned as a Voice over Internet Protocol (VoIP) service. This mandates the use of IPv6 for such networks due to the impending IPv4 address exhaustion. In the U.S., cellular operator Verizon has released technical specifications for devices operating on its future networks. The specification mandates IPv6 operation according to the 3GPP Release 8 Specifications (March 2009) and deprecates IPv4 as an optional capability.
6. What is Auto-configuration?
- Ans. IPv6 hosts can configure themselves automatically when connected to a routed IPv6 network using ICMPv6 router discovery messages. When first connected to a network, a host sends a link-local multicast router solicitation request for its configuration parameters; if configured suitably, routers respond to such a request with a router advertisement packet that contains network-layer configuration parameters.
7. What is IPv6 addressing system?
- Ans. IPv6 addresses are classified into three types: unicast addresses which uniquely identify network interfaces, anycast addresses which identify a group of interfaces—mostly at different locations—for which traffic flows to the nearest one, and multicast addresses which are used to deliver one packet to many interfaces. Broadcast addresses are not used in IPv6. Each IPv6 address also has a 'scope', which specifies in which part of the network it is valid and unique. Some addresses have node scope or link scope; most addresses have global scope (i.e. they are unique globally).
8. What is 6bone?
- Ans. The original mission of the 6bone was to establish a network to foster the development, testing, and deployment of IPv6 using a model to be based upon the experiences from the Mbone, hence the name "6bone". The 6bone started as a virtual network (using IPv6 over IPv4 tunneling/encapsulation) operating over the IPv4-based Internet to support IPv6 transport, and slowly added native links specifically for IPv6 transport. Although the initial 6bone focus was on testing of standards and implementations, the eventual focus became more on testing of transition and operational procedures, as well as actual IPv6 network usage.

ANSWERS

Multiple Choice Questions

1. a 2. b 3. c 4. a
5. b

True False Questions

1. T 2. T 3. F 4. T
5. F

Further Readings

1. **Elements of Data Communication and Networks:** S. A. Amatha, ~~University~~ University Science Press.
2. **Wide Area Networks:** Navneet Sharma, Firewall Media.
3. **Data Communication System:** Monika Khwaran, Firewall Media.
4. **Computer Network:** Bharat Bhushan Agarwal and Sumit Prakash Tayal, University Science Press.
5. **Computer Network:** Rachna Sharma, University Science Press.
6. **Computer Network:** Sumit Raj Chauhan and Er. Punit Soni, University Science Press.

This model allowed users of personal computers to view and manipulate data stored at a central repository, while not forcing them to share a central CPU for all of the processing of that data. Users did share the CPU running the server process, but to the extent that data processing was performed by the clients, the burden on the central CPU hosting the server process was lessened.

The client/server architecture was soon extended to include more than two processes. The original client/server model began to be called 2-tier client/server, to indicate two processes: one client and one server. More elaborate architectures were called 3-tier, to indicate three processes, 4-tier, to indicate four processes, or N-tier, to indicate people were getting tired of counting processes. Eventually, as more processes became involved, the distinction between client and server blurred, and people just started using the term *distributed processing* to encompass all of these schemes.

The distributed processing model leveraged the network and the proliferation of processors by dividing processing work loads among many processors while allowing those processors to share data. Although this model had many advantages...

STRUCTURE

MOBILITY IN NETWORK

UNIT 3

NOTES

Mobility in Network

Additional Material 77

3.1 MOBILITY IN NETWORK

NOTES

AWA

Prior to the advent of the personal computer, the dominant computing model was the large mainframe computer serving multiple users. By time-sharing, a mainframe computer divided its attention among several users, who logged onto the mainframe at dumb terminals. Software applications were stored on disks attached to the mainframe computer, allowing multiple users to share the same applications while they shared the same CPU. A drawback of this model was that if one user ran a CPU-intensive job, all other users would experience degraded performance.

The appearance of the microprocessor led to the proliferation of the personal computer. This change in the hardware status quo changed the software paradigm as well. Rather than sharing software applications stored at a mainframe computer, individual users had individual copies of software applications stored at each personal computer. Because each user ran software on a dedicated CPU, this new model of computing addressed the difficulties of dividing CPU-time among many users attempting to share one mainframe CPU.

Initially, personal computers operated as unconnected islands of computing. The dominant software model was of isolated executables running on isolated personal computers. But soon, personal computers began to be connected to networks. Because a personal computer gave its user undivided attention, it addressed the CPU-time sharing difficulties of mainframes. But unless personal computers were connected to a network, they couldn't replicate the mainframe's ability to let multiple users view and manipulate a central repository of data.

As personal computers connected to networks became the norm, another software model began to increase in importance: client/server. The client/server model divides work between two processes running on two different computers: a client process runs on the end-user's personal computer, and a server process runs on some other computer hooked to the same network. The client and server processes communicated with one another by sending data back and forth across the network. The server process often simply accepted data query commands from clients across the network, retrieved the requested data from a central database, and sent the retrieved data back across the network to the client. Upon receiving the data, the client processed it and allowed the user to manipulate it.

mainframe model, there was one notable disadvantage: distributed processing systems were more difficult to administer than mainframe systems. On mainframe systems, software applications were stored on a disk attached to the mainframe. Even though an application could serve many users, it only needed to be installed and maintained in one place. When an application was upgraded, all users got the new version the next time they logged on and started the application. By contrast, the software executables for different components of a distributed processing system were usually stored on many different disks. In a client/server architecture, for example, each computer that hosted a client process usually had its own copy of the client software stored on its local disk. As a result, a system administrator had to install and maintain the various components of a distributed software system in many different places. When a software component was upgraded, the system administrator had to physically upgrade each copy of the component on each computer that hosted it. As a result, system administration was more difficult for the distributed processing model than for the mainframe model.

NOTES

The arrival of Java, with an architecture that enabled the network-mobility of software, heralded yet another model for computing. Building on the prevailing distributed processing model, the new model added the automatic delivery of software across networks to computers that ran the software. This addressed the difficulties involved in system administration of distributed processing systems. For example, in a client/server system, client software could be stored at one central computer attached to the network. Whenever an end-user needed to use the client software, the binary executable would be sent from the central computer across the network to the end-user's computer, where the software would run.

So network-mobility of software represented another step in the evolution of the computing model. In particular, it addressed the difficulty of administering a distributed processing system. It simplified the job of distributing any software that was to be used on more than one CPU. It allowed data to be delivered together with the software that knows how to manipulate or display the data. Because code was sent along with data, end-users would always have the most up-to-date version of the code. Thus, because of network-mobility, software can be administered from a central computer, reminiscent of the mainframe model, but processing can still be distributed among many CPUs.

Java's architectural support for network-mobility begins with its support for platform independence and security. Although they are not strictly required for network-mobility, platform independence and security help make network-mobility practical. Platform independence makes it easier to deliver a program across the network because you don't have to maintain a separate version of the program for different platforms, and you don't have to figure out how to get the right version to each computer. One version of a program can serve all computers. Java's security features help promote network-mobility because they give end-users confidence to download class files from untrusted sources. In practice, therefore, Java's architectural support for platform independence and security facilitate the network-mobility of its class files.

Beyond platform independence and security, Java's architectural support for network-mobility is focused on managing the time it takes to move software across a network. If you store a program on a server and download it across a network when you need it, it will likely take longer for your program to start than if you had started the same program from a local disk. Thus, one of the primary issues of network-mobile software

NOTES

is the time it takes to send a program across a network. Java's architecture addresses this issue by rejecting the traditional monolithic binary executable in favor of small binary pieces: *Java class files*. Class files can travel across networks independently, and because Java programs are dynamically linked and dynamically extensible, an end-user needn't wait until all of a program's class files are downloaded before the program starts. The program starts when the first class file arrives. Class files themselves are designed to be compact, so that they fly more quickly across networks. Therefore, the main way Java's architecture facilitates network-mobility directly is by breaking up the monolithic binary executable into compact class files, which can be loaded as needed.

The execution of a Java application begins at a `main()` method of some class, and other classes are loaded and dynamically linked as they are needed by the application. If a class is never actually used during one session, that class won't ever be loaded during that session. For example, if you are using a word processor that has a spelling checker, but during one session you never invoke the spelling checker, the class files for the spelling checker will not be loaded during that session.

In addition to dynamic linking, Java's architecture also enables dynamic extension. Dynamic extension is another way the loading of class files (and the downloading of them across a network) can be delayed in a Java application. Using user-defined class loaders or the `forname()` method of class `Class`, a Java program can load extra classes at run-time, which then become a part of the running program. Therefore, dynamic linking and dynamic extension give a Java programmer some flexibility in designing when class files for a program are loaded, and as a result, how much time an end-user must spend waiting for class files to come across the network.

Besides dynamic linking and dynamic extension, another way Java's architecture directly supports network mobility is through the class file format itself. To reduce the time it takes to send them across networks, class files are designed to be compact. In particular, the bytecode streams they contain are designed to be compact. They are called "bytecodes" because each instruction occupies only one byte. With only two exceptions, all opcodes and their ensuing operands are byte aligned to make the bytecode streams smaller. The two exceptions are opcodes that may have one to three bytes of padding after the opcode and before the start of the operands, so that the operands are aligned on word boundaries.

One of the implications of the compactness goal for class files is that Java compilers are not likely to do much local optimization. Because of binary compatibility rules, Java compilers can't perform global optimizations such as inlining the invocation of another class's method. (Inlining means replacing the method invocation with the code performed by the method, which saves the time it takes to invoke and return from the method as the code executes.) Binary compatibility requires that a method's implementation can be changed without breaking compatibility with pre-existing class files that depend on the method. Inlining could be performed in some circumstances on methods within a single class, but in general that kind of optimization is not done by Java compilers, partly because it goes against the grain of class file compactness. Optimizations are often a tradeoff between execution speed and code size. Therefore, Java compilers generally leave optimization up to the Java virtual machine, which can optimize code as it loads classes for interpreting, just-in-time compiling, or adaptive optimization.

Beyond the architectural features of dynamic linking, dynamic extension and cla

file compactness, there are some strategies that, although they are really not necessarily part of the architecture, help manage the time it takes to move class files across a network. Because HTTP protocols require that each class file of Java applet be requested individually, it turns out that often a large percentage of applet download time is due not to the actual transmission of class files across the network, but to the network handshaking of each class file request. The overhead for a file request is multiplied by the number of class files being requested. To address this problem, Java 1.1 included support for JAR (Java ARchive) files. JAR files enable many class files to be sent in one network transaction, which greatly reduces the overhead time required to move class files across a network compared with sending one class file at a time. Moreover, the data inside a JAR file can be compressed, which results in an even shorter download time. So sometimes it pays to send software across a network in one big chunk. If a set of class files is definitely needed by a program before that program can start, those class files can be more speedily transmitted if they are sent together in a JAR file.

One other strategy to minimize an end-user's wait time is to not download class files on-demand. Through various techniques, such as the subscription model used by Marimba Castanet, class files can be downloaded before they are needed, resulting in a program that starts up faster.

Therefore, other than platform independence and security, which help make network-mobility practical, the main focus of Java's architectural support for network-mobility is managing the time it takes to send class files across a network. Dynamic linking and dynamic extension allow Java programs to be designed in small functional units that are downloaded as needed by the end-user. Class file compactness helps reduce the time it takes to move a Java program across the network. The JAR file enables compression and the sending of multiple class files across the network in a single network file-transfer transaction.

Java is a network-oriented technology that first appeared at a time when the network was looking increasingly like the next revolution in computing. The reason Java was adopted so rapidly and so widely, however, was not simply because it was a timely technology, but because it had timely marketing. Java was not the only network-oriented technology being developed in the early to mid 1990s. And although it was a good technology, it wasn't the necessarily the best technology—but it probably had the best marketing. Java was the one technology to hit a slim market window in early 1995, resulting in such a strong response that many companies developing similar technologies canceled their projects. Companies that carried on with their technologies, such as AT&T did with a network-oriented technology named Inferno, saw Java steal much of their potential thunder.

There were several important factors in how Java was initially unleashed on the world that contributed to its successful marketing. First, it had a cool name—one that could be appreciated by programmers and non-programmers alike. Second, it was, for all practical purposes, free—always a strong selling point among prospective buyers. But the most critical factor contributing to the successful marketing of Java, however, was that Sun's engineers hooked Java technology to the World Wide Web at the precise moment Netscape was looking to transform their web browser from a graphical hypertext viewer to a full-fledged computing platform. As the World Wide Web swept through the software industry (and the global consciousness) like an ever-increasing tidal wave, Java rode with it. Therefore, in a sense Java became a success

NOTES

because Java “surfed the web.” It caught the wave at just the right time and kept riding it as one by one, its potential competitors dropped uneventfully into the cold, dark sea. The way the engineers at Sun hooked Java technology to the World Wide Web—and therefore, the key way Java was successfully marketed—was by creating a special flavor of Java program that ran inside a web browser: the Java applet.

The Java applet showed off all of Java’s network-oriented features: platform independence, network- mobility, and security. Platform independence was one of the main tenets of the World Wide Web, and Java applets fit right in. Java applets can run on any platform so long as there is a Java-capable browser for that platform. Java applets also demonstrated Java’s security capabilities, because they ran inside a strict sandbox. But most significantly, Java applets demonstrated the promise of network-mobility.

3.2 MOBILE SECURITY RELATED ISSUE

AM

Computer or network security has been violated when unauthorized access by any party occurs. So it becomes your duty as the Network Administrator to work on the security to guard against any such incidence taking place.

3.2.1 Need of Network Security

Computer security is required because most organizations can be damaged by hostile software or intruders. There may be several forms of damage which are obviously interrelated. These include:

- Damage or destruction of computer systems.
- Damage or destruction of internal data.
- Loss of sensitive information to hostile parties.
- Use of sensitive information to steal items of monetary value.
- Use of sensitive information against the organization’s customers which may result in legal action by customers against the organization and loss of customers.
- Damage to the reputation of an organization.
- Monetary damage due to loss of sensitive information, destruction of data, hostile use of sensitive data, or damage to the organization’s reputation.

The methods used to accomplish these unscrupulous objectives are many and varied depending on the circumstances. This guide will help administrators understand some of these methods and explain some countermeasures.

3.2.1.1 Security Issues

Computer security can be very complex and may be very confusing to many people. It can even be a controversial subject. Network administrators like to believe that their network is secure and those who break into networks may like to believe that they can break into any network.

Overconfidence plays an important rôle in allowing networks to be intruded upon.

There are many fallacies that network administrators may fall victim to. These fallacies may allow administrators to wrongfully believe that their network is more secure than it really is. This guide will attempt to clarify many issues related to security by doing the following:

- Help you determine what you are protecting.
- Break computer security into categories.
- Explain security terms and methods.
- Point out some common fallacies that may allow administrators to be overconfident.
- Categorize many common attacks against networks and computers.
- Explain some attack methods.
- Describe tools that can be used to help make a network more secure.

NOTES

3.2.1.2 Security Interdependence

There are many different aspects to computer and network security as you will read in this unit. These different areas of computer security are interdependent on each other in order for a network to be secure. If one or more areas of computer security are ignored, then the entire security integrity of the organization's network may be compromised. A clear example of this is in the area of computer virus or worm protection.

Computer virus protection programs can only filter known viruses or worms. There are viruses or worms that are not yet recognized as virus programs immediately after their release.

The best way to make unrecognized virus or worm programs less effective is by quickly removing the vulnerabilities that they use.

Some of these vulnerabilities are operating system and application program errors. When security patches are created for software, they should be quickly applied. In this way the vulnerability to viruses is minimized but not eliminated. There are other steps which may further reduce this vulnerability, but it can never be completely eliminated.

3.2.1.3 Security Limitations and Applications

If you are reading this document and are thinking that you can get all the information required to make your network completely secure, then you are sadly mistaken. In many ways, computer security is almost a statistical game.

You can reduce but not eliminate the chance that you may be penetrated by an intruder or virus. This is mainly for one reason.

No one can ever know all the software vulnerabilities of all software used on a system.

This is why even those who consider themselves hackers will say that the number one computer security threat is the lack of quality in the applications and operating systems. At this point, I could talk about the various corporate entities that write

software and why software lacks the quality that many of us believe that it should possess, but that subject is not only way beyond the scope of this document, but also way beyond the scope of this project.

The bottom line here is that unless you can remove all the application and operating system problems that allow viruses and intruders to penetrate networks, you can never secure your network.

Additionally the users on your network are potentially a greater security risk than any programs. Obviously removing all vulnerabilities is impossible and will not secure your network against user errors. I have even considered the possibility that an operating system without a network interface can be completely secure, but even this cannot be guaranteed. Unknown viruses or trojan programs can creep in with applications on CDs or floppies. This has been known to happen. Although an attacker may not be able to get data from the system, they can damage or destroy data.

3.2.1.4 Layered Security

The fact that complete security is impossible is the reason security experts recommend "layered security". The idea is to have multiple ways of preventing an intrusion to decrease the chance that intrusions will be successful. For example, you should have virus protection on your client computers. To help layer this security you should also filter viruses at your email server.

To help even more, you should block the most dangerous types of email attachments to prevent unrecognized viruses and other hostile software from entering your network.

Another good defense layer would also include educating your users about viruses, how they spread, and how to avoid them.

3.2.1.5 Hackers

There are many documents that attempt to define the term hacker. I believe that the term hacker is a connotative term. This means that it is more defined by people's beliefs rather than by a dictionary.

Some believe that a hacker is a very skilled computer person.

Others believe that hackers are those that perform unauthorized break-ins to computer systems. The media and many sources have caused many uninformed people to believe that a hacker is a threat to computer and network security while this is not the case. A hacker is no more likely to break the law than anyone else. It would be wise to use the more accurate descriptive term, "intruder" to describe those who intrude into networks or systems without authorization.

3.2.1.6 Physical Security

This unit will not talk about physical computer security beyond this paragraph. Your organization should be aware how physically secure every aspect of its network is because if an intruder gets physical access, they can get your data.

Be sure that your organization properly secures locations and consider the following:

- **Servers** - Contain your data and information about how to access that data.

- **Workstations** - May contain some sensitive data and can be used to attack other computers.
- **Routers, switches, bridges, hubs** and any other network equipment may be used as an access point to your network.
- **Network wiring and media** and where they pass through may be used to access your network or place a wireless access point to your network.
- **External media** which may be used between organizational sites or to other sites the organization does business with.
- **Locations of staff** who may have information that a hostile party can use.
- **Some employees** may take data home or may take laptops home or use laptops on the internet from home then bring them to work. Any information on these laptops should be considered to be at risk and these laptops should be secure according to proper policy when connected externally on the network (more on this later).

NOTES

3.2.1.7 Some Technical Terms

This paragraph describes some commonly used computer security terms.

- **Protocol** - Well defined specification allowing computer communication.
- **Confidentiality** - Information is available only to people with rightful access.
- **Integrity** - Information can only be changed by authorized personnel. Integrity - The receiver of the message should be able to tell that the message was not modified. Requires key exchange.
- **Availability** - Information is available to only those who need it.
- **Verification - nonrepudiation** - There is proof that the sender sent the message.
- **Authentication** - The receiver of the message should be able to be sure of the origin of the message. Requires a digital signature (One way hash, public key algorithm, and symmetric algorithm) or a public key algorithm.
- **Spyware** - A computer program whose purpose is to spy on your internet activities usually for marketing purposes and usually done by a shady corporate entity.
- **Malware** - A computer program with some evil intent. It may on the surface have a good or useful intent, but may be a trojan (with a hidden purpose) which can be used to gain unauthorized access to your computer.

3.3 IP MULTICASTING

IP multicast is a method of sending Internet Protocol (IP) datagrams to a group of interested receivers in a single transmission. It is often employed for streaming media applications on the Internet and private networks. The method is the IP-specific version of the general concept of multicast networking. It uses specially reserved multicast address blocks in IPv4 and IPv6.

STUDENT ACTIVITY 3.1

1. What do you understand by Mobility in Network?

2. What do you understand by IP Multicasting?

3. Describe multicasting routing protocols.

4. What is the need for security in Mobile?

3.3.1 Implementations

Pay-TV operators and some educational institutions with significant on-campus student housing have deployed IP multicast to deliver one-way streaming media such as high-speed video to large groups of receivers. Additionally, there have been some uses of audio and video conferencing using multicast technologies. These are far less prevalent and are most often relegated to research and education institutions, which often have a greater degree of network capacity to handle the demands. Some technical conferences and meetings are transmitted using IP multicast. Until recently many of the sessions at the IETF meetings were delivered using multicast.

Another use of multicast within campus and commercial networks is for file distribution, particularly to deliver operating system images and updates to remote hosts. The key advantage of multicast boot images over unicasting boot images is significantly lower network bandwidth usage.

IP multicast has also seen deployment within the financial sector for applications such as stock tickers and hoot-n-holler systems.

While IP multicast has seen some success in each of these areas, IP multicast is not widely deployed and is generally not available as a service for the average end user. There are at least two primary factors for the lack of widespread deployment, both somewhat related to the other. On the one hand, forwarding multicast traffic, particularly for two-way communication, requires a great deal of protocol complexity. On the other hand, there are a number of additional operational concerns in being able to run a multicast network successfully, largely stemming from the complexity of a widely-deployed implementation, not the least of which is the enabling of additional avenues of failure, particularly from denial-of-service attacks. Many of these issues are covered in further detail below.

RFC 3170 (IP Multicast Applications: Challenges & Solutions) provides an overview of deployment issues.

3.3.2 Addressing

There are four forms of IP addressing, each with its own unique properties.

- **Unicast:** The most common concept of an IP address is in unicast addressing, available in both IPv4 and IPv6. It normally refers to a single sender or a single receiver, and can be used for both sending and receiving. Usually, a unicast address is associated with a single device or host, but it is not a one-to-one correspondence. Some individual PCs have several distinct unicast addresses, each for its own distinct purpose. Sending the same data to multiple unicast addresses requires the sender to send all the data many times over, once for each recipient.
- **Broadcast:** In IPv4 it is possible to send data to all possible destinations ("all-hosts broadcast"), which permits the sender to send the data only once, and all receivers receive a copy of it. In the IPv4 protocol, the address 255.255.255.255 is used for local broadcast. In addition, a directed (limited) broadcast can be made by combining the network prefix with a host suffix composed entirely of binary 1s. For example, the destination address used for a directed broadcast to devices on the 192.0.2.0/24 network is 192.0.2.255.
- **Multicast:** A multicast address is associated with a group of interested

receivers. According to RFC 3171, addresses 224.0.0.0 to 239.255.255.255, the former Class D addresses, are designated as multicast addresses in IPv4. The sender sends a single datagram (from the sender's unicast address) to the multicast address, and the intermediary routers take care of making copies and sending them to all receivers that have registered their interest in data from that sender.

- Anycast: Like broadcast and multicast, anycast is a one-to-many routing topology. However, the data stream is not transmitted to all receivers, just the one which the router decides is the "closest" in the network. Anycast is useful for global load balancing and is commonly used in DNS communications.

3.3.3 Protocols and applications

Since multicast is a different transmission mode from unicast, only protocols designed for multicast can be sensibly used with multicast.

Most of the existing application protocols that use multicast run on top of the User Datagram Protocol (UDP). In many applications, the Real-time Transport Protocol (RTP) is used for framing of multimedia content over multicast; the Resource Reservation Protocol (RSVP) may be used for bandwidth reservation in a network supporting multicast distribution.

On the local network, multicast delivery is controlled by IGMP (on IPv4 network) and MLD (on IPv6 network); inside a routing domain, PIM or MOSPF are used; between routing domains, one uses inter-domain multicast routing protocols, such as MBGP.

A number of errors can happen if packets intended for unicast are accidentally sent to a multicast address; in particular, sending ICMP packets to a multicast address has been used in the context of DoS attacks as a way of achieving packet amplification.

3.3.4 Routing

Each host (and in fact each application on the host) that wants to be a receiving member of a multicast group (i.e. receive data corresponding to a particular multicast address) must use the Internet Group Management Protocol (IGMP) to join. Adjacent routers also use this protocol to communicate.

In unicast routing, each router examines the destination address of an incoming packet and looks up the destination in a table to determine which interface to use in order for that packet to get closer to its destination. The source address is irrelevant to the router.

However, in multicast routing, the source address (which is a simple unicast address) is used to determine data stream direction. The source of the multicast traffic is considered upstream. The router determines which downstream interfaces are destinations for this multicast group (the destination address), and sends the packet out through the appropriate interfaces. The term reverse path forwarding is used to describe this concept of routing packets away from the source, rather than towards the destination.

3.3.5 Layer 2 delivery

Unicast packets are delivered to a specific recipient on an Ethernet or IEEE 802.3

NOTES

NOTES

subnet by setting a specific layer 2 MAC address on the Ethernet packet address. Broadcast packets make use of a broadcast MAC address (FF:FF:FF:FF:FF:FF), which includes setting the broadcast/multicast bit in the address. The IANA owns the OUI MAC address 01:00:5e, therefore multicast packets are delivered by using the Ethernet MAC address range 01:00:5e:00:00:00 - 01:00:5e:7f:ff:ff. This is 23 bits of available address space. The first octet (01) includes the broadcast/multicast bit. The lower 23 bits of the 28-bit multicast IP address are mapped into the 23 bits of available Ethernet address space. This means that there is ambiguity in delivering packets. If two hosts on the same subnet each subscribe to a different multicast group whose address differs only in the first 5 bits, Ethernet packets for both multicast groups will be delivered to both hosts, requiring the network software in the hosts to discard the unrequired packets.

For IPv6 Multicast addresses, the Ethernet MAC is derived by the four low-order octets OR'ed with the MAC 33:33:00:00:00:00, so for example the IPv6 address FF02:DEAD:BEEF::1:3 would map to the Ethernet MAC address 33:33:00:01:00:03.

Switches that cannot understand multicast addresses will broadcast traffic sent to a multicast group to all the members of a LAN; in this case the system's network card (and operating system) has to filter the packets sent to multicast groups they are not subscribed to.

There are switches that listen to IGMP traffic and maintain a state table of which network systems are subscribed to a given multicast group. This table is then used to forward traffic destined to a given group only to a limited set of hosts (ports). This is done through the use of IGMP snooping.

Additionally, some switches with layer 3 capabilities can act as an IGMP querier. In networks where there is no router present (or enabled) to act as a multicast router a switch might be used to generate the needed IGMP messages to get users to subscribe to multicast traffic.

3.3.6 Reliable multicast

Multicast, by its very nature, is not a connection-oriented mechanism, so protocols such as TCP, which allows for retransmission of missing packets, are not appropriate. For applications such as streaming audio and video, the occasional dropped packet is not a problem. But for distribution of critical data, a mechanism is required for requesting retransmission.

One such scheme, proposed by Cisco, is PGM (originally Pretty Good Multicasting, but changed for trademark reasons to Pragmatic General Multicast), documented in RFC 3208. In this scheme, multicast packets have sequence numbers and when a packet is missed a recipient can request that the packet be re-multicast with other members of the Multicast group ignoring the replacement data if not needed. An expanded version, PGM-CC, has attempted to make IP Multicasting more "TCP friendly" by stepping the entire group down to the bandwidth available by the worst receiver.

Two other schemes documented by the Internet Engineering Task Force (IETF) are NACK-Oriented Reliable Multicast (NORM), documented in RFC 3940 and RFC 5401, and File Delivery over Unidirectional Transport (FLUTE), documented in RFC 3926. Open-source, in addition to proprietary, implementations exist for these. Other such protocols exist, such as Scalable Reliable Multicast, and are defined by a variety

of sources. Such protocols vary in the means of error detection, the mechanisms used in error recovery, the scalability of such recovery and the underlying ideas involved in what it means to be reliable. A list of reliable multicast protocols from the ACM SIGCOMM Multicast Workshop, August 27, 1996, documents a number of approaches to the problem.

Independent groups like the Internet Protocol Multicast Standards Initiative (IPMSI) have claimed that the lack of a truly scalable Secure Reliable IP Multicast protocol like the proposed Secure Multicast for Advanced Repeating of Television (SMART) have hampered the adoption of IP Multicast in inter-domain routing. The lack of a widely adopted system that has AES level security and scalable reliability have kept mass media transmissions of sporting events (like the Super Bowl) and/or breaking news events from being transmitted on the Public Internet.

Reliable IP Multicasting protocols, such as PGM and SMART, are experimental; the only standards-track protocol is NORM (the standards-track revision of RFC 3941 is specified in RFC 5401, the standards-track revision of RFC 3940 is specified in RFC 5740).

3.3.7 Wireless (802.11) considerations

An 802.11 wireless network will handle multicast traffic differently, depending on the configuration of 802.11 power-save mode, DTIM (Delivery Traffic Indication Message), and beacon interval settings. If power-save mode is disabled, then access points will deliver multicast traffic after each beacon (default interval = 100ms, but it can be adjusted). If power-save mode is enabled, the access point will deliver multicast traffic after each DTIM, which by default is every 1,2, or 3 beacon intervals in most access points. As a result, the DTIM and beacon interval settings should be adjusted for optimum performance when implementing multicast in wireless networks.

3.4 MULTICASTING ROUTING PROTOCOLS

Protocol-Independent Multicast (PIM) is a family of multicast routing protocols for Internet Protocol (IP) networks that provide one-to-many and many-to-many distribution of data over a LAN, WAN or the Internet. It is termed protocol-independent because PIM does not include its own topology discovery mechanism, but instead uses routing information supplied by other traditional routing protocols such as the Border Gateway Protocol (BGP).

There are four variants of PIM:

- PIM Sparse Mode (PIM-SM) explicitly builds unidirectional shared trees rooted at a rendezvous point (RP) per group, and optionally creates shortest-path trees per source. PIM-SM generally scales fairly well for wide-area usage.
- PIM Dense Mode (PIM-DM) uses dense multicast routing. It implicitly builds shortest-path trees by flooding multicast traffic domain wide, and then pruning back branches of the tree where no receivers are present. PIM-DM generally has poor scaling properties. Often, DVMRP is used in dense mode.
- Bidirectional-PIM explicitly builds shared bi-directional trees. It never builds a shortest path tree, so may have longer end-to-end delays than PIM-SM, but scales well because it needs no source-specific state.

NOTES

Two kinds of group addresses are supported: permanent addresses and temporary ones. A permanent group is always there and does not have to be set up. Each permanent group has a permanent group address. Some examples of permanent group addresses are

NOTES

- 224.0.0.1 All systems on a LAN
- 224.0.0.2 All routers on a LAN
- 224.0.0.5 All OSPF routers on a LAN
- 224.0.0.6 All designated OSPF routers on a LAN

Temporary groups must be created before they can be used. A process can ask its host to join a specific group. It can also ask its host to leave the group. When the last process on a host leaves a group, that group is no longer present on the host. Each host keeps track of which groups its processes currently belong to. Multicasting is implemented by special multicast routers, which may or may not be colocated with the standard routers. About once a minute, each multicast router sends a hardware (i.e., data link layer) multicast to the hosts on its LAN (address 224.0.0.1) asking them to report back on the groups their processes currently belong to. Each host sends back responses for all the class D addresses it is interested in.

These query and response packets use a protocol called IGMP (Internet Group Management Protocol), which is vaguely analogous to ICMP. It has only two kinds of packets: query and response, each with a simple fixed format containing some control information in the first word of the payload field and a class D address in the second word. It is described in RFC 1112. Multicast routing is done using spanning trees. Each multicast router exchanges information with its neighbors using a modified distance vector protocol in order for each one to construct a spanning tree per group covering all group members. Various optimizations are used to prune the tree to eliminate routers and networks not interested in particular groups. The protocol makes heavy use of tunneling to avoid bothering nodes not in a spanning tree.

3.4.3 Mobile IP

Many users of the Internet have portable computers and want to stay connected to the Internet when they visit a distant Internet site and even on the road in between. Unfortunately, the IP addressing system makes working far from home easier said than done. In this section we will examine the problem and the solution. A more detailed description is given in (Johnson, 1995).

The real villain is the addressing scheme itself. Every IP address contains three fields: the class, the network number, and the host number. For example, consider the machine with IP address 160.80.40.20. The 160.80 gives the class (B) and network number (8272); the 40.20 is the host number (10260). Routers all over the world have routing tables telling which line to use to get to network 160.80. Whenever a packet comes in with a destination IP address of the form 160.80.xxx.yyy, it goes out on that line.

If all of a sudden, the machine with that address is carted off to some distant site, the packets for it will continue to be routed to its home LAN (or router). The owner will no longer get email, and so on. Giving the machine a new IP address corresponding to its new location is unattractive because large numbers of people, programs, and databases would have to be informed of the change.

Another approach is to have the routers use complete IP addresses for routing, instead

of just the class and network. However, this strategy would require each router to have millions of table entries, at astronomical cost to the Internet.

When people began demanding the ability to have mobile hosts, the IETF set up a Working Group to find a solution. The Working Group quickly formulated a number of goals considered desirable in any solution. The major ones were

1. Each mobile host must be able to use its home IP address anywhere.
2. Software changes to the fixed hosts were not permitted.
3. Changes to the router software and tables were not permitted.
4. Most packets for mobile hosts should not make detours on the way.
5. No overhead should be incurred when a mobile host is at home.

NOTES

3.5 ADDRESS ASSIGNMENTS

There are several subtle points that often deserve consideration when assigning multicast addresses. We've collected these as advice and rationale here.

- Avoid 224.0.0.x—Traffic to addresses of the form 224.0.0.x is often flooded to all switch ports. This address range is reserved for link-local uses. Many routing protocols assume that all traffic within this range will be received by all routers on the network. Hence (at least all Cisco) switches flood traffic within this range. The flooding behavior overrides the normal selective forwarding behavior of a multicast-aware switch (e.g. IGMP snooping, CGMP, etc.).
- Watch for 32:1 overlap—32 non-contiguous IP multicast addresses are mapped onto each Ethernet multicast address. A receiver that joins a single IP multicast group implicitly joins 31 others due to this overlap. Of course, filtering in the operating system discards undesired multicast traffic from applications, but NIC bandwidth and CPU resources are nonetheless consumed discarding it. The overlap occurs in the 5 high-order bits, so it's best to use the 23 low-order bits to make distinct multicast streams unique. For example, IP multicast addresses in the range 239.0.0.0 to 239.127.255.255 all map to unique Ethernet multicast addresses. However, IP multicast address 239.128.0.0 maps to the same Ethernet multicast address as 239.0.0.0, 239.128.0.1 maps to the same Ethernet multicast address as 239.0.0.1, etc.
- Avoid x.0.0.y and x.128.0.y—Combining the above two considerations, it's best to avoid using IP multicast addresses of the form x.0.0.y and x.128.0.y since they all map onto the range of Ethernet multicast addresses that are flooded to all switch ports.
- Watch for address assignment conflicts—IANA administers Internet multicast addresses. Potential conflicts with Internet multicast address assignments can be avoided by using GLOP addressing or administratively scoped addresses. Such addresses can be safely used on a network connected to the Internet without fear of conflict with multicast sources originating on the Internet. Administratively scoped addresses are roughly analogous to the unicast address space for private internets. Site-local multicast addresses are of the form

239.255.x.y, but can grow down to 239.252.x.y if needed. Organization-local multicast addresses are of the form 239.192-251.x.y, but can grow down to 239.x.y.z if needed.

NOTES

3.6 SESSION DISCOVERY

The invention provides a method of accessing data relating to announced media sessions that are to take place over a communications network such as the Multicast Internet. The data may include scheduling, title and content information, for example.

The method comprises the steps of:-

- i) establishing a communications channel between a user terminal and a database system, where the database system comprises session specific data relating to respective media sessions available over a communications network;
- ii) determining the identity of the user by unique identification code or otherwise;
 - ii) retrieving user specific data for that user;
- iii) selecting from the media sessions, in accordance with said user specific data at least one media session relevant to the user; and
- iv) returning session specific data identifying the relevant media session or sessions to the user terminal. In another aspect the invention provides a method of configuring a media session database

SUMMARY

1. Prior to the advent of the personal computer, the dominant computing model was the large mainframe computer serving multiple users.
2. The distributed processing model leveraged the network and the proliferation of processors by dividing processing work loads among many processors while allowing those processors to share data.
3. In addition to dynamic linking, Java's architecture also enables dynamic extension.
4. So network-mobility of software represented another step in the evolution of the computing model.
5. Java's architectural support for network-mobility begins with its support for platform independence and security.
6. Java is a network-oriented technology that first appeared at a time when the network was looking increasingly like the next revolution in computing.
7. Computer security is required because most organizations can be damaged by hostile software or intruders.
8. Network administrators like to believe that their network is secure and those who break into networks may like to believe that they can break into any network.
9. Computer virus protection programs can only filter known viruses or worms.
10. Unknown viruses or trojan programs can creep in with applications on CDs or floppies.
11. A hacker is no more likely to break the law than anyone else.
12. IP multicast is a method of sending Internet Protocol (IP) datagrams to a group of interested receivers in a single transmission.

13. Pay-TV operators and some educational institutions with significant on-campus student housing have deployed IP multicast to deliver one-way streaming media such as high-speed video to large groups of receivers.
14. The most common concept of an IP address is in unicast addressing, available in both IPv4 and IPv6.
15. In the IPv4 protocol, the address 255.255.255.255 is used for local broadcast.
16. A multicast address is associated with a group of interested receivers.
17. Most of the existing application protocols that use multicast run on top of the User Datagram Protocol (UDP).
18. In unicast routing, each router examines the destination address of an incoming packet and looks up the destination in a table to determine which interface to use in order for that packet to get closer to its destination.
19. Independent groups like the Internet Protocol Multicast Standards Initiative (IPMSI) have claimed that the lack of a truly scalable Secure Reliable IP Multicast protocol like the proposed Secure Multicast for Advanced Repeating of Television (SMART) have hampered the adoption of IP Multicast in inter-domain routing.
20. Protocol-Independent Multicast (PIM) is a family of multicast routing protocols for Internet Protocol (IP) networks that provide one-to-many and many-to-many distribution of data over a LAN, WAN or the Internet.

NOTES

SELF ASSESSMENT QUESTIONS

1. What do you understand by Mobility in Network?
2. What is the need for security in Mobile?
3. What do you understand by IP Multicasting?
4. Describe multicasting routing protocols.
5. How would you address assignments?
6. What do you understand by Session discovery?

Multiple Choice Questions

1. UDP is:
 - (a) User Datagram Protocol
 - (b) Use Datagram Protocol
 - (c) User Datagram Port
2. RSVP is:
 - (a) Reserve Reservation Protocol
 - (b) Resource Reservation Protocol
 - (c) Resource Reservation Port
3. IGMP is:
 - (a) Internet Grand Management Protocol
 - (b) Internet Group Management Port
 - (c) Internet Group Management Protocol
4. RTP is:
 - (a) Real-time Transport Protocol
 - (b) Real-time Transfer Protocol
 - (c) Real-time Transport Port

NOTES

5. PGM is:
 - (a) Pragmatic General Market
 - (b) Pragmatic General Multicast
 - (c) People General Multicast
6. SMART is:
 - (a) Secure Multicast for Aware Repeating of Television
 - (b) Sealed Multicast for Advanced Repeating of Television
 - (c) Secure Multicast for Advanced Repeating of Television
7. IETF is :
 - (a) Internet Engineering Task Force
 - (b) Intranet Engineering Task Force
 - (c) Internet Engineering Tell Force
8. DTIM is:
 - (a) Delivery Through Indication Message
 - (b) Delivery Traffic Indication Message
 - (c) Delivery Traffic Indicating Message
9. IPMSI is:
 - (a) Internet Port Multicast Standards Initiative
 - (b) Internet Protocol Multicast Standards Interest
 - (c) Internet Protocol Multicast Standards Initiative
10. PIM is:
 - (a) Protocol-Independent Multicast
 - (b) Port-Independent Multicast
 - (c) Protocol-Independent Market
11. PIMBGP is:
 - (a) Protocol-Independent Market Border Gateway Protocol
 - (b) Protocol-Independent Multicast Border Gateway Protocol
 - (c) Port-Independent Multicast Border Gateway Protocol
12. PIMSSM is:
 - (a) Protocol-Independent Market Source-specific Multicast
 - (b) Port-Independent Multicast Source-specific Multicast
 - (c) Protocol-Independent Multicast Source-specific Multicast

True/False Questions

1. The distributed processing model leveraged the network and the proliferation of processors by dividing processing work loads among many processors while allowing those processors to share data.
2. Network-mobility of software represent another step in the evolution of the computing model.

3. Java is a network-oriented technology that first appeared at a time when the network was looking increasingly like the next revolution in computing.
4. Network administrators like to believe that their network is secure and those who break into networks may like to believe that they can break into any network.
5. Unknown viruses or trojan programs cannot creep in with applications on CDs or floppies.
6. IP multicast is a method of sending Internet Protocol (IP) datagrams to a group of interested receivers in a single transmission.
7. The most common concept of an IP address is in unicast addressing, available in both IPv4 and IPv6.
8. A multicast address is not associated with a group of interested receivers.
9. In unicast routing, each router examines the destination address of an incoming packet and looks up the destination in a table to determine which interface to use in order for that packet to get closer to its destination.
10. Protocol-Independent Multicast (PIM) is a family of multicast routing protocols for Internet Protocol (IP) networks that provide one-to-many and many-to-many distribution of data over a LAN, WAN or the Internet.

NOTES

Short Questions with Answers

1. How did Java effected the Mobility of Networks?

Ans. The arrival of Java, with an architecture that enabled the network-mobility of software, heralded yet another model for computing. Building on the prevailing distributed processing model, the new model added the automatic delivery of software across networks to computers that ran the software. This addressed the difficulties involved in system administration of distributed processing systems. For example, in a client/server system, client software could be stored at one central computer attached to the network. Whenever an end-user needed to use the client software, the binary executable would be sent from the central computer across the network to the end-user's computer, where the software would run.

Java is a network-oriented technology that first appeared at a time when the network was looking increasingly like the next revolution in computing. The reason Java was adopted so rapidly and so widely, however, was not simply because it was a timely technology, but because it had timely marketing. Java was not the only network-oriented technology being developed in the early to mid 1990s. And although it was a good technology, it wasn't the necessarily the best technology—but it probably had the best marketing. Java was the one technology to hit a slim market window in early 1995, resulting in such a strong response that many companies developing similar technologies canceled their projects. Companies that carried on with their technologies, such as AT&T did with a network-oriented technology named Inferno, saw Java steal much of their potential thunder.

NOTES

2. How does Java architecture help in setting up the networks?

Ans. In addition to dynamic linking, Java's architecture also enables dynamic extension. Dynamic extension is another way the loading of class files (and the downloading of them across a network) can be delayed in a Java application. Using user-defined class loaders or the forname() method of class Class, a Java program can load extra classes at run-time, which then become a part of the running program. Therefore, dynamic linking and dynamic extension give a Java programmer some flexibility in designing when class files for a program are loaded, and as a result, how much time an end-user must spend waiting for class files to come across the network.

3. Why is computer security needed?

Ans. Computer security is required because most organizations can be damaged by hostile software or intruders. There may be several forms of damage which are obviously interrelated. These include:

- Damage or destruction of computer systems.
- Damage or destruction of internal data.
- Loss of sensitive information to hostile parties.
- Use of sensitive information to steal items of monetary value.
- Use of sensitive information against the organization's customers which may result in legal action by customers against the organization and loss of customers.
- Damage to the reputation of an organization.
- Monetary damage due to loss of sensitive information, destruction of data, hostile use of sensitive data, or damage to the organization's reputation.

4. Describe some commonly used security terms

Ans. This paragraph describes some commonly used computer security terms.

- **Protocol** - Well defined specification allowing computer communication.
- **Confidentiality** - Information is available only to people with rightful access.
- **Integrity** - Information can only be changed by authorized personnel. Integrity - The receiver of the message should be able to tell that the message was not modified. Requires key exchange.
- **Availability** - Information is available to only those who need it.
- **Verification - nonrepudiation** - There is proof that the sender sent the message.
- **Authentication** - The receiver of the message should be able to be sure of the origin of the message. Requires a digital signature (One way hash, public key algorithm, and symmetric algorithm) or a public key algorithm.
- **Spyware** - A computer program whose purpose is to spy on your internet activities usually for marketing purposes and usually done by a shady corporate entity.
- **Malware** - A computer program with some evil intent. It may on the surface have a good or useful intent, but may be a trojan (with a hidden purpose) which can be used to gain unauthorized access to your computer.

5. Describe the various forms of IP addressing.

Ans. There are four forms of IP addressing, each with its own unique properties.

- **Unicast:** The most common concept of an IP address is in unicast addressing, available in both IPv4 and IPv6. It normally refers to a single sender or a single receiver, and can be used for both sending and receiving. Usually, a unicast address is associated with a single device or host; but it is not a one-to-one correspondence. Some individual PCs have several distinct unicast addresses, each for its own distinct purpose. Sending the same data to multiple unicast addresses requires the sender to send all the data many times over, once for each recipient.
- **Broadcast:** In IPv4 it is possible to send data to all possible destinations (“all-hosts broadcast”); which permits the sender to send the data only once, and all receivers receive a copy of it. In the IPv4 protocol, the address 255.255.255.255 is used for local broadcast. In addition, a directed (limited) broadcast can be made by combining the network prefix with a host suffix composed entirely of binary 1s. For example, the destination address used for a directed broadcast to devices on the 192.0.2.0/24 network is 192.0.2.255.
- **Multicast:** A multicast address is associated with a group of interested receivers. According to RFC 3171, addresses 224.0.0.0 to 239.255.255.255, the former Class D addresses, are designated as multicast addresses in IPv4. The sender sends a single datagram (from the sender’s unicast address) to the multicast address, and the intermediary routers take care of making copies and sending them to all receivers that have registered their interest in data from that sender.
- **Anycast:** Like broadcast and multicast, anycast is a one-to-many routing topology. However, the data stream is not transmitted to all receivers, just the one which the router decides is the “closest” in the network. Anycast is useful for global load balancing and is commonly used in DNS communications.

6. What is Multicast?

Ans. Multicast, by its very nature, is not a connection-oriented mechanism, so protocols such as TCP, which allows for retransmission of missing packets, are not appropriate. For applications such as streaming audio and video, the occasional dropped packet is not a problem. But for distribution of critical data, a mechanism is required for requesting retransmission. One such scheme, proposed by Cisco, is PGM (originally Pretty Good Multicasting, but changed for trademark reasons to Pragmatic General Multicast), documented in RFC 3208. In this scheme; multicast packets have sequence numbers and when a packet is missed a recipient can request that the packet be re-multicast with other members of the Multicast group ignoring the replacement data if not needed. An expanded version, PGM-CC, has attempted to make IP Multicasting more “TCP friendly” by stepping the entire group down to the bandwidth available by the worst receiver.

7. Describe the various types of Protocol Independent Multicasts.

Ans. There are four variants of PIM:

- **PIM Sparse Mode (PIM-SM)** explicitly builds unidirectional shared trees rooted at a rendezvous point (RP) per group, and optionally creates shortest-path trees per source. PIM-SM generally scales fairly well for wide-area usage.
- **PIM Dense Mode (PIM-DM)** uses dense multicast routing. It implicitly builds shortest-path trees by flooding multicast traffic domain wide, and then pruning back branches of the tree where no receivers are present. PIM-DM generally has poor scaling properties. Often, DVMRP is used in dense mode.

NOTES

- Bidirectional PIM explicitly builds shared bi-directional trees. It never builds a shortest path tree, so may have longer end-to-end delays than PIM-SM, but scales well because it needs no source-specific state.
- PIM source-specific multicast (PIM-SSM) builds trees that are rooted in just one source, offering a more secure and scalable model for a limited amount of applications (mostly broadcasting of content). In SSM, an IP datagram is transmitted by a source S to an SSM destination address G, and receivers can receive this datagram by subscribing to channel (S,G).

ANSWERS

Multiple Choice Questions

- | | | | |
|------|-------|-------|-------|
| 1. a | 2. b | 3. c | 4. a |
| 5. b | 6. c | 7. a | 8. b |
| 9. c | 10. a | 11. b | 12. c |

True False Questions

- | | | | |
|------|-------|------|------|
| 1. T | 2. T | 3. T | 4. T |
| 5. F | 6. T | 7. T | 8. F |
| 9. T | 10. T | | |

Further Readings

1. **Computer Networks:** Ajit Kumar Singh, Firewal Media.
2. **TCP / IP and Distributed System:** Vivek Archarya, Firewal Media.
3. **Elements of Data Communication and Networks:** S. A. Amutha Jeevakumari, University Science Press.
4. **Data Communication System:** Monika Khuaran, Firewal Media.
5. **Computer Network:** Rachna-Sharma, University Science Press.

UNIT 4

TCP EXTENSIONS

NOTES

STRUCTURE

- 4.1 TCP Extensions for high-speed networks
- 4.2 Transaction-oriented application
- 4.3 Other new option in TCP
 - Summary
 - Self Assessment Questions
 - Further Readings

Learning Objectives

After going through this unit, students will be able to:

- know about TCP extensions for high-speed networks.
- learn about Transaction-oriented application including network including advantages and disadvantages.
- know about other options in TCP including, Apple Talk, and other protocols like Printer Access Protocol, Zone Information Protocol, Routing Table Maintenance Protocol.

4.1 TCP EXTENSIONS FOR HIGH-SPEED NETWORKS

NOTES

IPX/SPX stands for **Internetwork Packet Exchange/Sequenced Packet Exchange**. It is a networking protocol used by the Novell NetWare operating systems. Like UDP, IPX is a datagram protocol used for connectionless communications. IPX and SPX are derived from Xerox Network Services' IDP and SPP protocols.

SPX is a transport layer protocol (layer 4 of the OSI Model) used in Novell Netware networks. The SPX layer sits on top of the IPX layer (layer 3 - the network layer) and provides connection-oriented services between two nodes on the network. SPX is used primarily by client/server applications.

IPX and SPX both provide connection services similar to TCP/IP, with the IPX protocol having similarities to IP, and SPX having similarities to TCP. IPX/SPX was primarily designed for local area networks (LANs), and is a very efficient protocol for this purpose (typically its performance exceeds that of TCP/IP on a LAN). TCP/IP has, however, become the *de facto* protocol. This is in part due to its superior performance over wide area networks and the Internet (which uses TCP/IP exclusively), and also because TCP/IP is a more mature protocol, designed specifically with this purpose in mind.

Novell is largely responsible for the use of IPX as a popular computer networking protocol due to their dominance in the network operating system software market (with Novell Netware) from the late 1980s through to the mid-1990s.

Because of IPX/SPX's prevalence in LANs in the 1990's, Microsoft added support for the protocols into Windows' networking stack, starting with Windows for Workgroups and Windows NT. Microsoft even named their implementation "NWLink", implying that the inclusion of the layer 3/4 transports provided NetWare connectivity. In reality, the protocols were supported as a native transport for Windows' SMB/NetBIOS, and NetWare connectivity required additional installation of an NCP client (Microsoft provided a basic NetWare client with Windows 95 and later, but it was not automatically installed, and initially only supported NetWare bindery mode). NWLink is still provided with Windows (up to and including Windows 2003). However, it is neither included on nor supported in Windows Vista, Novell will be starting a closed beta in March 2007, but its use is strongly discouraged, because it cannot be used for Windows networking except as a transport for NetBIOS, which is deprecated.

For the most part, Novell's 32-bit Windows client software have eschewed NWLink for an alternative developed by Novell, although some versions permit use of Microsoft's IPX/SPX implementation (with warnings about potential incompatibility).

IPX usage has declined in recent years as the rise of the Internet has made TCP/IP ubiquitous. Novell's initial attempt to support TCP/IP as a client protocol, called NetWare/IP, simply "tunnelled" IPX within IP packets, allowing NetWare clients and servers to communicate over pure TCP/IP networks. However, due to complex implementation, and a significant loss in performance due to the tunnelling overhead, NetWare/IP was largely ignored except as a mechanism to route IPX through TCP/IP-only routers and WAN links. NetWare 5.x introduced native support for NCP over TCP/IP, which is now the preferred configuration.

Both Microsoft and Novell have provided support (through Proxy Server/ISA Server and Border Manager, respectively) for IPX/SPX as an intranet protocol to communicate through a firewall.

NOTES

This allows a machine using client software to access the Internet without having TCP/IP installed locally; the client software emulates a native TCP/IP stack and provides WinSock support for local applications (e.g. web browsers), but actually communicates with the firewall over IPX/SPX. In addition to simplifying migration for legacy IPX LANs, this provides a measure of security, as the use of the IPX protocol on the internal network provides a natural barrier against intruders, should the firewall be compromised.

One area where IPX remains useful is to sidestep VPNs that force all TCP/IP traffic to traverse the VPN, preventing any access to local resources such as printers and shared disks.

4.1.1 TCP Connection Management

Connections are established in TCP using the three-way handshake. To establish a connection, one side, say the server, passively waits for an incoming connection by executing the LISTEN and ACCEPT primitives, either specifying a specific source or nobody in particular.

The other side, say the client, executes a CONNECT primitive, specifying the IP address and port to which it wants to connect, the maximum TCP segment size it is willing to accept, and optionally some user data (e.g., a password). The connect primitive sends a TCP segment with the SYN bit on and ACK bit off and waits for a response.

When this segment arrives at the destination, the TCP entity there checks to see if there is a process that has done a LISTEN on the port given in the Destination port field. If not, it sends a reply with the RST bit on to reject the connection.

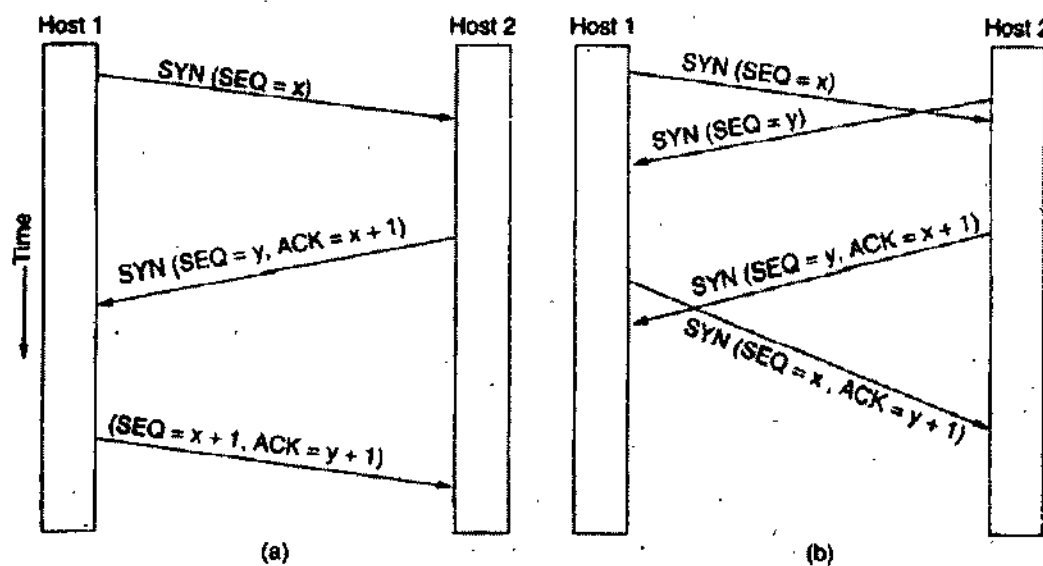


Figure : (a) TCP connection establishment in the normal case, (b) Call collision.

If some process is listening to the port, that process is given the incoming TCP

segment. It can then either accept or reject the connection. If it accepts, an acknowledgement segment is sent back. The sequence of TCP segments sent in the normal case is shown in figure (a). Note that a SYN segment consumes 1 byte of sequence space so it can be acknowledged unambiguously.

NOTES

In the event that two hosts simultaneously attempt to establish a connection between the same two sockets, the sequence of events is as illustrated in figure (b). The result of these events is that just one connection is established, not two because connections are identified by their end points. If the first setup results in a connection identified by (x, y) and the second one does too, only one table entry is made, namely, for (x, y).

The initial sequence number on a connection is not 0 for the reasons we discussed earlier. A clock-based scheme is used, with a clock tick every 4 fisec. For additional safety, when a host crashes, it may not reboot for the maximum packet lifetime (120 sec) to make sure that no packets from previous connections are still roaming around the Internet somewhere.

Although TCP connections are full duplex, to understand how connections are released it is best to think of them as a pair of simplex connections. Each simplex connection is released independently of its sibling. To release a connection, either party can send a TCP segment with the FIN bit set, which means that it has no more data to transmit. When the FIN is acknowledged, that direction is shut down for new data. Data may continue to flow indefinitely in the other direction, how ever. When both directions have been shut down, the connection is released. Normally, four TCP segments are needed to release a connection, one FIN and one ACK for each direction. However, it is possible for the first ACK and the second FIN to be contained in the same segment, reducing the total count to three.

Just as with telephone calls in which both people say goodbye and hang up the phone simultaneously, both ends of a TCP connection may send PIN segments at the same time. These are each acknowledged in the usual way, and the connection shut down. There is, in fact, no essential difference between the two hosts releasing sequentially or simultaneously.

To avoid the two-army problem, timers are used. If a response to a FIN is not forthcoming within two maximum packet lifetimes, the sender of the FIN releases the connection. The other side will eventually notice that nobody seems to be listening to it any more, and time out as well. While this solution is not perfect, given the fact that a perfect solution is theoretically impossible, it will have to do. In practice, problems rarely arise.

The steps required to establish and release connections can be represented in a finite state machine with the 11 states listed in next figure. In each state, certain events are legal. When a legal event happens, some action may be taken. If some other event happens, an error is reported.

State	Description
CLOSED	No connection is active or pending
LISTEN	The server is waiting for an incoming call
SYNRCVD	A connection request has arrived; wait for ACK
SYNSENT	The application has started to open a connection

ESTABLISHED	The normal data transfer state
FINWAIT1	The application has said it is finished
FINWAIT2	The other side has agreed to release
TIMEDWAIT	Wait for all packets to die off
CLOSING	Both sides have tried to close simultaneously
CLOSEWAIT	The other side has initiated a release
LASTACK	Wait for all packets to die off

NOTES

Figure : The states used in the TCP connection management finite state machine

Each connection starts in the CLOSED state. It leaves that state when it does either a passive open (listen), or an active open (CONNECT). If the other side does the opposite one, a connection is established and the state becomes ESTABLISHED. Connection release can be initiated by either side. When it is complete, the state returns to CLOSED.

The finite state machine itself is shown in the figure. The common case of a client actively connecting to a passive server is shown with heavy lines—solid for the client, dotted for the server. The lightface lines are unusual event sequences. Each line in the figure is marked by an event/action pair. The event can either be a user-initiated system call (CONNECT, LISTEN, SEND, or close), a segment arrival (SYN, FIN, ACK, or RST), or in one case, a timeout of twice the maximum packet lifetime. The action is the sending of a control segment (SYN, FIN, or RST) or nothing, indicated by —. Comments are shown in parentheses.

The diagram can best be understood by first following the path of a client (the heavy solid line) then later the path of a server (the heavy dashed line). When an application on the client machine issues a CONNECT request, the local TCP entity creates a connection record, marks it as being in the SYN SENT state, and sends a SYN segment. Note that many connections may be open (or being opened) at the same time on behalf of multiple applications, so the state is per connection and recorded in the connection record. When the SYN+ACK arrives, TCP sends the final ACK of the three-way handshake and switches into the ESTABLISHED state. Data can now be sent and received.

When an application is finished, it executes a CLOSE primitive, which causes the local TCP entity to send a FIN segment and wait for the corresponding ACK (dashed box marked active close). When the ACK arrives, a transition is made to state FIN WAFT 2 and one direction of the connection is now closed. When the other side closes, too, a FIN comes in, which is acknowledged. Now both sides are closed, but TCP waits a time equal to the maximum packet lifetime to guarantee that all packets from the connection have died off, just in case the acknowledgement was lost. When the timer goes off, TCP deletes the connection record.

Now let us examine connection management from the server's viewpoint. The server does a listen and settles; down to see who turns up. When a SYN comes in, it is acknowledged and the server goes to the SYN RCVD state. When the server's SYN is itself acknowledged, the three-way handshake is complete and the server goes to the ESTABLISHED state. Data transfer can now occur.

When the client has had enough, it does a CLOSE, which causes a FIN to arrive at

NOTES

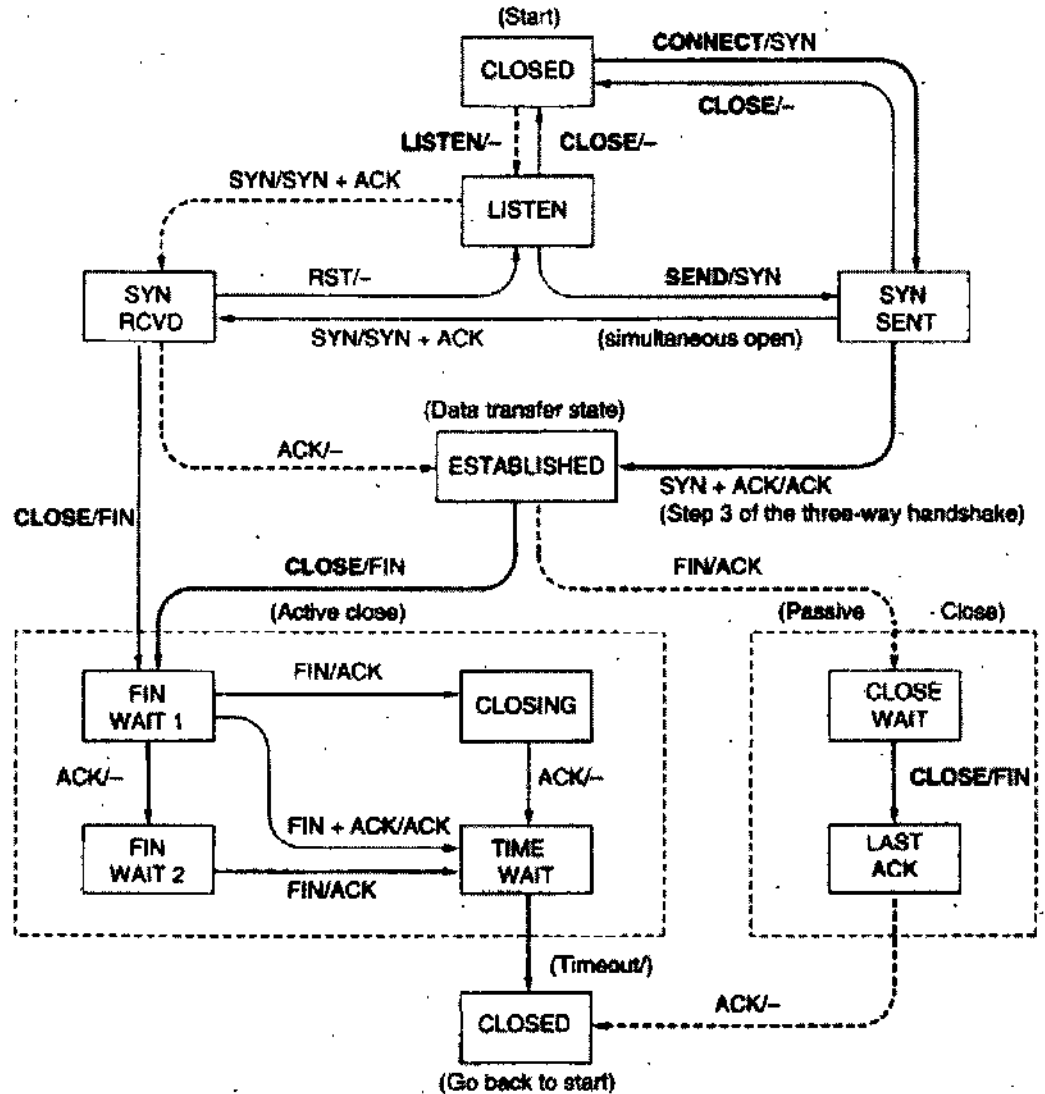


Figure : TCP, connection management finite state machine. The heavy solid line is the normal path for a client. The heavy dashed line is the normal path for a server. The light lines are unusual events.

the server (dashed box marked passive close). The server is then signaled. When it, too, does a CLOSE, a FIN is sent to the client. When the client's acknowledgement shows up, the server releases the connection and deletes the connection record.

4.1.2 TCP Transmission Policy

Window management in TCP is not directly tied to acknowledgements as it is in most data link protocols. For example, suppose the receiver has a 4096-byte buffer as shown in next figure. If the sender transmits a 2048-byte segment that is correctly received, the receiver will acknowledge the segment. However, since it now has only 2048 of buffer space (until the application removes some data from the buffer), it will advertise a window of 2048 starting at the next byte expected.

Now the sender transmits another 2048 bytes, which are acknowledged, but the advertised window is 0. The sender must stop until the application process on the receiving host has removed some data from the buffer, at which time TCP can advertise a larger window.

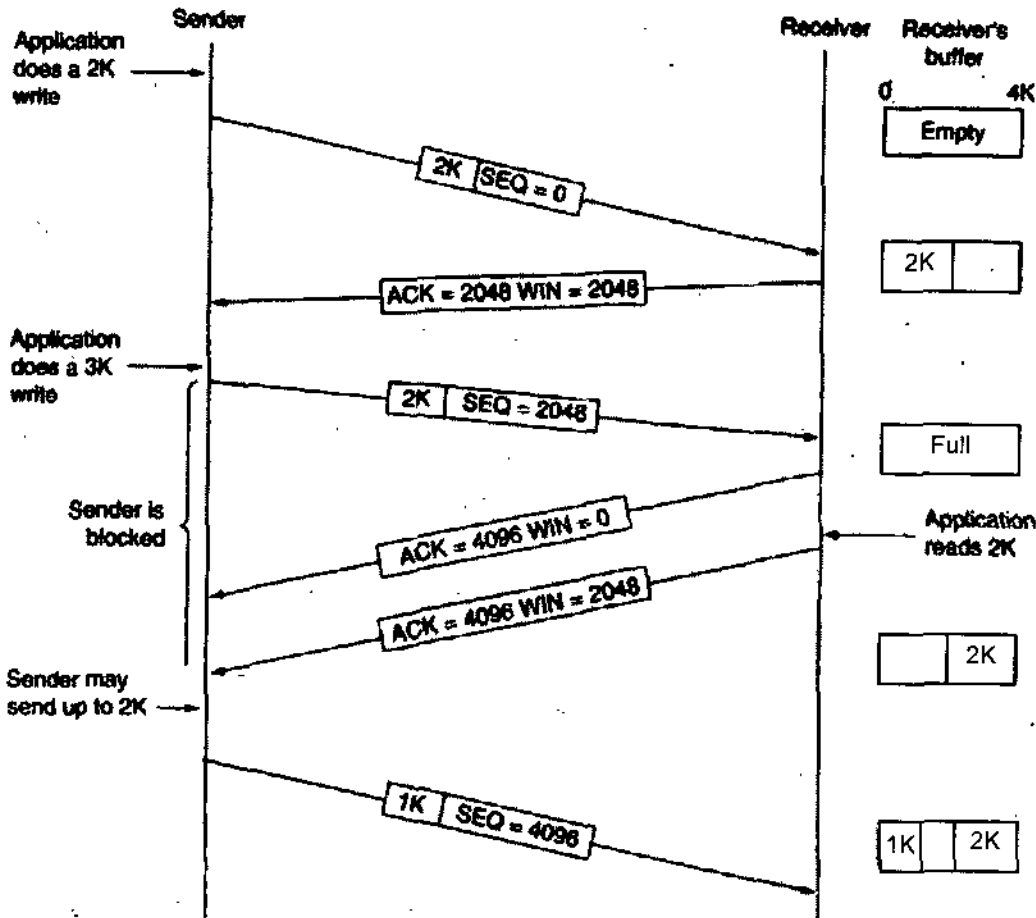


Figure : Window management in TCP

When the window is 0, the sender may not normally send segments, with two exceptions. First, urgent data may be sent, for example, to allow the user to kill the process running on the remote machine. Second, the sender may send a 1-byte segment to make the receiver reannounce the next byte expected and window size. The TCP standard explicitly provides this option to prevent deadlock if a window announcement ever gets lost.

Senders are not required to transmit data as soon as they come in from the application. Neither are receivers required to send acknowledgements as soon as possible. For example, in Fig. 6-29, when the first 2 KB of data came in, TCP, knowing that it had a 4-KB window available, would have been completely correct in just buffering the data until another 2 KB came in, to be able to transmit a segment with a 4-KB payload. This freedom can be exploited to improve performance.

Consider a TELNET connection to an interactive editor that reacts on every keystroke. In the worst case, when a character arrives at the sending TCP entity, TCP creates a 21-byte TCP segment, which it gives to IP to send as a 41-byte IP datagram. At the receiving side, TCP immediately sends a 40-byte acknowledgement (20 bytes of TCP header and 20 bytes of IP header). Later, when the editor has read the byte, TCP sends a window update, moving the window 1 byte to the right. This packet is also 40 bytes. Finally, when the editor has processed the character, it echoes it as a 41-byte packet. In all, 162 bytes of bandwidth are used and four segments are sent for each character typed. When bandwidth is scarce, this method of doing business is not desirable.

NOTES

One approach that many TCP implementations use to optimize this situation is to delay acknowledgements and window updates for 500 msec in the hope of acquiring some data on which to hitch a free ride. Assuming the editor echoes within 500 msec, only one 41-byte packet now need be sent back to the remote user, cutting the packet count and bandwidth usage in half.

NOTES

Although this rule reduces the load placed on the network by the receiver, the sender is still operating inefficiently by sending 41-byte packets containing 1 byte of data. A way to reduce this usage is known as Nagle's algorithm (Nagle, 1984). What Nagle suggested is simple: when data come into the sender one byte at a time, just send the first byte and buffer all the rest until the outstanding byte is acknowledged. Then send all the buffered characters in one TCP segment and start buffering again until they are all acknowledged. If the user is typing quickly and the network is slow, a substantial number of characters may go in each segment, greatly reducing the bandwidth used. The algorithm additionally allows a new packet to be sent if enough data have trickled in to fill half the window or a maximum segment.

Nagle's algorithm is widely used by TCP implementations, but there are times when it is better to disable it. In particular, when an X-Windows application is being run over the Internet, mouse movements have to be sent to the remote computer. Gathering them up to send in bursts makes the mouse cursor move erratically, which makes for unhappy users.

Another problem that can ruin TCP performance is the silly window syndrome. This problem occurs when data are passed to the sending TCP entity in large blocks, but an interactive application on the receiving side reads data 1 byte at a time. To see the problem, look at the next figure. Initially, the TCP buffer on the receiving side is full and the sender knows this (i.e., has a window of size 0). Then the interactive application reads one character from the TCP stream. This action makes the receiving TCP happy, so it sends a window update to the sender saying that it is all right to send 1 byte. The sender obliges and sends 1 byte. The buffer is now full, so the receiver acknowledges the 1-byte segment but sets the window to 0. This behavior can go on forever.

Clark's solution is to prevent the receiver from sending a window update for 1 byte. Instead it is forced to wait until it has a decent amount of space available and advertise that instead. Specifically, the receiver should not send a window update until it can handle the maximum segment size it advertised when the connection was established, or its buffer is half empty, whichever is smaller. Furthermore, the sender can also help by not sending tiny segments. Instead, it should try to wait until it has accumulated enough space in the window to send a full segment or at least one containing half of the receiver's buffer size (which it must estimate from the pattern of window updates it has received in the past).

Nagle's algorithm and Clark's solution to the silly window syndrome are complementary. Nagle was trying to solve the problem caused by the sending application delivering data to TCP a byte at a time. Clark was trying to solve the problem of the receiving application sucking the data up from TCP a byte at a time. Both solutions are valid and can work together. The goal is for the sender not to send small segments and the receiver not to ask for them.

The receiving TCP can go further in improving performance than just doing window updates in large units. Like the sending TCP, it also has the ability to buffer data, so

NOTES

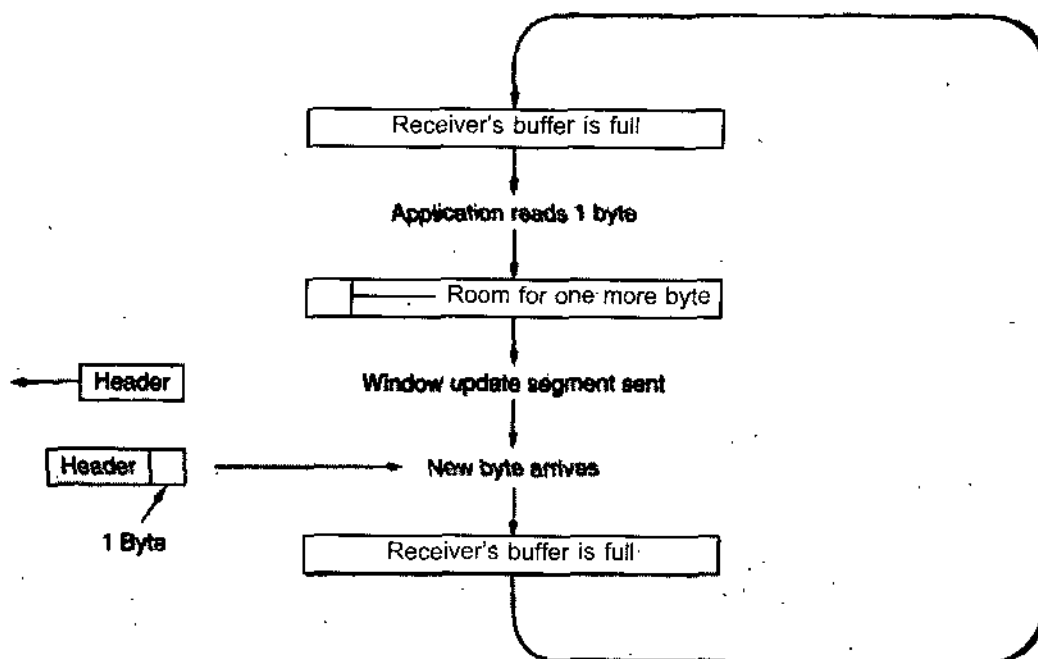


Figure : Silly window syndrome

it can block a READ request from the application until it has a large chunk of data to provide. Doing this reduces the number of calls to TCP, and hence the overhead. Of course, it also increases the response time, but for noninteractive applications like file transfer, efficiency may outweigh response time to individual requests.

Another receiver issue is what to do with out of order segments. They can be kept or discarded, at the receiver's discretion. Of course, acknowledgements can be sent only when all the data up to the byte acknowledged have been received. If the receiver gets segments 0, 1, 2, 4, 5, 6, and 7, it can acknowledge everything up to and including the last byte in segment 2. When the sender times out, it then retransmits segment 3. If the receiver has buffered segments 4 through 7, upon receipt of segment 3 it can acknowledge all bytes up to the end of segment 7.

4.1.3 TCP Congestion Control

When the load offered to any network is more than it can handle, congestion builds up. The internet is no exception. In this section we will discuss algorithms that have been developed over the past decade to deal with congestion. Although the network layer also tries to manage congestion, most of the heavy lifting is done by TCP because the real solution to congestion is to slow down the data rate.

In theory, congestion can be dealt with by employing a principle borrowed from physics: the law of conservation of packets. The idea is not to inject a new packet into the network until an old one leaves (i.e., is delivered). TCP attempts to achieve this goal by dynamically, manipulating the window size.

The first step in managing congestion is detecting it. In the old days, detecting congestion was difficult. A timeout caused by a lost packet could have been caused by either (1) noise on a transmission line or (2) packet discard at a congested router. Telling the difference was difficult.

Nowadays, packet loss due to transmission errors is relatively rare because most

NOTES

long-haul trunks are fiber (although wireless networks are a different story). Consequently, most transmission timeouts on the Internet are due to congestion. All the Internet TCP algorithms assume that timeouts are caused by congestion and monitor timeouts for signs of trouble the way miners watch their canaries. Before discussing how TCP reacts to congestion, let us first describe what it does to try to prevent it from occurring in the first place. When a connection is established, a suitable window size has to be chosen. The receiver can specify a window based on its buffer size. If the sender sticks to this window size, problems will not occur due to buffer overflow at the receiving end, but they may still occur due to internal congestion within the network.

In next figure, we see this problem illustrated hydraulically. In (a), we see a thick pipe leading to a small-capacity receiver. As long as the sender does not send more water than the bucket can contain, no water will be lost. In (b), the limiting factor is not the bucket capacity, but the internal carrying capacity of the network. If too much water comes in too fast, it will back up and some will be lost (in this case by overflowing the funnel).




Figure : (a) A fast network feeding a low-capacity receiver, (b) A slow network feeding a high-capacity receiver.

The Internet solution is to realize that two potential problems exist—network capacity and receiver capacity—and to deal with each of them separately. To do so, each sender maintains two windows: the window the receiver has granted and a second window, the congestion window. Each reflects the number of bytes the sender may transmit. The number of bytes that may be sent is the minimum of the two windows. Thus the effective window is the minimum of what the sender thinks is all right and what the receiver thinks is all right. If the receiver says “Send 8K” but the sender

knows that bursts of more than 4K clog the network up, it sends 4K. On the other hand, if the receiver says "Send 8K" and the sender knows that bursts of up to 32K get through effortlessly, it sends the full 8K requested.

When a connection is established, the sender initializes the congestion window to the size of the maximum segment in use on the connection. It then sends one maximum segment. If this segment is acknowledged before the timer goes off, it adds one segment's worth of bytes to the congestion window to make it two maximum size segments and sends two segments. As each of these segments is acknowledged, the congestion window is increased by one maximum segment size. When the congestion window is n segments, if all n are acknowledged on time, the congestion window is increased by the byte count corresponding to n segments. In effect, each burst successfully acknowledged doubles the congestion window.

The congestion window keeps growing exponentially until either a timeout occurs or the receiver's window is reached. The idea is that if bursts of size, say, 1024, 2048, and 4096 bytes work fine, but a burst of 8192 bytes gives a timeout, the congestion window should be set to 4096 to avoid congestion. As long as the congestion window remains at 4096, no bursts longer than that will be sent, no matter how much window space the receiver grants. This algorithm is called slow start, but it is not slow at all (Jacobson, 1988). It is exponential. All TCP implementations are required to support it.

Now let us look at the Internet congestion control algorithm. It uses a third parameter, the threshold, initially 64K, in addition to the receiver and congestion windows. When a timeout occurs, the threshold is set to half of the current congestion window, and the congestion window is reset to one maximum segment. Slow start is then used to determine what the network can handle, except that exponential growth stops when the threshold is hit. From that point on, successful transmissions grow the congestion window linearly (by one maximum segment for each burst) instead of one per segment. In effect, this algorithm is guessing that it is probably acceptable to cut the congestion window in half, and then it gradually works its way up from there.

As an illustration of how the congestion algorithm works, see the next figure. The maximum segment size here is 1024 bytes. Initially the congestion window was 64K, but a timeout occurred, so the threshold is set to 32K and the congestion window to 1K for transmission 0 here. The congestion window then grows exponentially until it hits the threshold (32K). Starting then it grows linearly.

Transmission 13 is unlucky (it should have known) and a timeout occurs. The threshold is set to half the current window (by now 40K, so half is 20K) and slow start initiated all over again. When the acknowledgements from transmission 14 start coming in, the first four each double the congestion window, but after that, growth becomes linear again.

If no more timeouts occur, the congestion window will continue to grow up to the size of the receiver's window. At that point, it will stop growing and remain constant as long as there are no more timeouts and the receiver's window does not change size. As an aside, if an ICMP SOURCE QUENCH packet comes in and is passed to TCP, this event is treated the same way as a timeout.

Work on improving the congestion control mechanism is continuing. For example, Brakmo et al. (1994) have reported improving TCP throughput by 40 percent to 70

NOTES

NOTES




Figure : An example of the Internet congestion algorithm.

percent by managing the clock more accurately, predicting congestion before timeouts occur, and using this early warning system to improve the slow start algorithm.

4.1.4 TCP Timer Management

TCP uses multiple timers (at least conceptually) to do its work. The most important of these is the retransmission timer. When a segment is sent, a retransmission timer is started. If the segment is acknowledged before the timer expires, the timer is stopped. If, on the other hand, the timer goes off before the acknowledgement comes in, the segment is retransmitted (and the timer started again). The question that arises is: How long should the timeout interval be?

This problem is much more difficult in the Internet transport layer than in the generic data link protocols. In the latter case, the expected delay is highly predictable (i.e., has a low variance), so the timer can be set to go off just slightly after the acknowledgement is expected, as shown in next figure. Since acknowledgements are rarely delayed in the data link layer, the absence of an acknowledgement at the expected time generally means the frame or the acknowledgement has been lost.

TCP is faced with a radically different environment. The probability density function for the time it takes for a TCP acknowledgement to come back looks more like (b) than (a). Determining the round-trip time to the destination is tricky. Even when it is known, deciding on the timeout interval is also difficult. If the timeout is set too short, say T_1 in (b), unnecessary retransmissions will occur, clogging the Internet with useless packets. If it is set too long, (T_2), performance will suffer due to the long retransmission delay when ever a packet is lost. Furthermore, the mean and variance of the acknowledgement arrival distribution can change rapidly within a few seconds as congestion builds up or is resolved.

The solution is to use a highly dynamic algorithm that constantly adjusts the timeout




Figure : (a) Probability density of acknowledgement arrival times in the data link layer, (b) Probability density of acknowledgement arrival times for TCP.

NOTES

interval, based on continuous measurements of network performance. The algorithm generally used by TCP is due to Jacobson (1988) and works as follows. For each connection, TCP maintains a variable, RTT, that is the best current estimate of the round-trip time to the destination in question. When a segment is sent, a timer is started, both to see how long the acknowledgement takes and to trigger a retransmission if it takes too long. If the acknowledgement gets back before the timer expires, TCP measures how long the acknowledgement took, say, M . It then updates RTT according to the formula

$$RTT = \alpha RTT + (1 - \alpha) M$$

where α is a smoothing factor that determines how much weight is given to the old value. Typically $\alpha = 7/8$.

Even given a good value of RTT, choosing a suitable retransmission timeout is a nontrivial matter. Normally, TCP uses βRTT , but the trick is choosing β . In the initial implementations, β was always 2, but experience showed that a constant value was inflexible because it failed to respond when the variance went up.

In 1988, Jacobson proposed making β roughly proportional to the standard deviation of the acknowledgement arrival time probability density function so a large variance means a large β and vice versa. In particular, he suggested using the mean deviation as a cheap estimator of the standard deviation. His algorithm requires keeping track of another smoothed variable, D , the deviation. Whenever an acknowledgement comes in, the difference between the expected and observed values, $|RTT - M|$ is computed. A smoothed value of this is maintained in D by the formula

$$D = \alpha D + (1 - \alpha) |RTT - M|$$

where α may or may not be the same value used to smooth RTT. While D is not exactly the same as the standard deviation, it is good enough and Jacobson showed how it could be computed using only integer adds, subtracts, and shifts, a big plus. Most TCP implementations now use this algorithm and set the timeout interval to

$$\text{Timeout} = RTT + 4 * D$$

The choice of the factor 4 is somewhat arbitrary, but it has two advantages. First,

STUDENT ACTIVITY 4.1

1. What are protocols?

2. Describe Network IPX/SPX.

3. What is AppleTalk?

4. Describe Digital Network Architecture.

NOTES

multiplication by 4 can be done with a single shift. Second, it minimizes unnecessary timeouts and retransmissions because less than one percent of all packets come in more than four standard deviations late. (Actually, Jacobson initially said to use 2, but later work has shown that 4 gives better performance.)

One problem that occurs with the dynamic estimation of RTT is what to do when a segment times out and is sent again. When the acknowledgement comes in, it is unclear whether the acknowledgement refers to the first transmission or a later one. Guessing wrong can seriously contaminate the estimate of RTT. Phil Karn discovered this problem the hard way. He is an amateur radio enthusiast interested in transmitting TCP/IP packets by ham radio, a notoriously unreliable medium (on a good day, half the packets get through). He made a simple proposal: do not update RTT on any segments that have been retransmitted. Instead, the timeout is doubled on each failure until the segments get through the first time. This fix is called Karn's algorithm. Most TCP implementations use it.

The retransmission timer is not the only one TCP uses. A second timer is the persistence timer. It is designed to prevent the following deadlock. The receiver sends an acknowledgement with a window size of 0, telling the sender to wait. Later, the receiver updates the window, but the packet with the update is lost. Now both the sender and the receiver are waiting for each other to do something. When the persistence timer goes off, the sender transmits a probe to the receiver. The response to the probe gives the window size. If it is still zero, the persistence timer is set again and the cycle repeats. If it is nonzero, data can now be sent.

A third timer that some implementations use is the keepalive timer. When a connection has been idle for a long time, the keepalive timer may go off to cause one side to check if the other side is still there. If it fails to respond, the connection is terminated. This feature is controversial because it adds overhead and may terminate an otherwise healthy connection due to a transient network partition.

The last timer used on each TCP connection is the one used in the TIMED WAIT state while closing. It runs for twice the maximum packet lifetime to make sure that when a connection is closed, all packets created by it have died off.

4.2 TRANSACTION-ORIENTED APPLICATIONS

An **intranet** is a private computer network that uses Internet protocols, network connectivity to securely share part of an organization's information or operations with its employees. Sometimes the term refers only to the most visible service, the internal website. The same concepts and technologies of the Internet such as clients and servers running on the Internet protocol suite are used to build an intranet. HTTP and other Internet protocols are commonly used as well, FTP. There is often an attempt to use Internet technologies to provide new interfaces with corporate 'legacy' data and information systems.

Briefly, an **intranet** can be understood as "a private version of the Internet," or as a version of the internet confined to an organization.

4.2.1 Extranets

Intranets differ from "Extranets" in that the former is generally restricted to employees of the organization while extranets can generally be accessed by customers, suppliers,

or other approved parties.

There does not necessarily have to be any access from the organization's internal network to the internet itself. When such access is provided it is usually through a gateway with a firewall, along with user authentication, encryption of messages, and often make use of virtual private networks (VPN's). Through such devices and systems off-site employees can access company information, computing resources and internal communications.

Increasingly, intranets are being used to deliver tools and applications, e.g., collaboration (to facilitate working in groups and teleconferencing) or sophisticated corporate directories, sales and CRM tools, project management etc., to advance productivity. Intranets are also being used as culture change platforms. For example, large numbers of employees discussing key issues in online forums could lead to new ideas.

Intranet traffic, like public-facing web site traffic, is better understood by using web metrics software to track overall activity, as well as through surveys of users. Intranet "User Experience", "Editorial", and "Technology" teams work together to produce in-house sites. Most commonly, intranets are owned by the communications, HR or CIO areas of large organizations, or some combination of the three.

4.2.2 Advantages

1. **Workforce productivity:** Intranets can help users to locate and view information faster and use applications relevant to their roles and responsibilities. With the help of a web browser interface such as Internet Explorer or Firefox, users can access data held in any database the organization wants to make available, anytime and - subject to security provisions - from anywhere within the company workstations, increasing employees' ability to perform their jobs faster, more accurately, and with confidence that they have the right information. It also helps to improve the services provided to the users.
2. **Time:** With intranets, organizations can make more information available to employees on a "pull" basis (ie: employees can link to relevant information at a time which suits them) rather than being deluged indiscriminately by emails.
3. **Communication:** Intranets can serve as powerful tools for communication within an organization, vertically and horizontally. From a communications standpoint, intranets are useful to communicate strategic initiatives that have a global reach throughout the organization. The type of information that can easily be conveyed is the purpose of the initiative and what the initiative is aiming to achieve, who is driving the initiative, results achieved to date, and who to speak to for more information. By providing this information on the intranet, staff have the opportunity to keep up-to-date with the strategic focus of the organisation.
4. **Web publishing** allows 'cumbersome' corporate knowledge to be maintained and easily accessed throughout the company using hypermedia and Web technologies. Examples include: employee manuals, benefits documents, company policies, business standards, newsfeeds, and even training, can be accessed using common Internet standards (Acrobat files, Flash files, CGI applications). Because each business unit can update the online copy of a document, the most recent version is always available to employees using the intranet.

NOTES

NOTES

5. Business operations and management: Intranets are also being used as a platform for developing and deploying applications to support business operations and decisions across the internetworked enterprise.
6. Cost-effective: Users can view information and data via web-browser rather than maintaining physical documents such as procedure manuals, internal phone list and requisition forms.
7. Promote common corporate culture: Every user is viewing the same information within the Intranet.
8. Enhance Collaboration: With information easily accessible by all authorised users, teamwork is enabled.
9. Cross-platform Capability: Web browsers that support Java under Windows, Mac, Unix are supported.

4.2.3 Disadvantages

1. Inappropriate or incorrect information can be posted on an intranet which can reduce its credibility and effectiveness.
2. In a devolved and highly interactive intranet there is freedom to post abusive and possibly illegal materials. There is a balance to be struck between taking advantage of this freedom to achieve corporate goals and having appropriate controls in place to meet an organization's legal or moral responsibilities.
3. Training is required to educate people of what intranet can do.
4. Need expertise in field to administer and develop Intranet information within the organization.
5. Security of the intranet becomes an issue. Other users may post sensitive information which may appear to another user. Furthermore, in an industry with high turnover there is the potential for an employee to acquire sensitive information which may significantly benefit their new position at a competing company.
6. As information can be posted by any user, information overload may occur during the cause if it is not controlled well.

4.2.4 Planning and creating an intranet

Most organizations devote considerable resources into the planning and implementation of their intranet as it is of strategic importance to the organization's success. Some of the planning would include topics such as:

- What they hope to achieve from the intranet
- Which person or department would "own" (take control of) the technology and the implementation
- How and when existing systems would be phased out/replaced
- How they intend to make the intranet secure
- How they'll ensure to keep it within legislative and other constraints
- Level of interactivity or interactivity (eg on-line forms) desired.

- Is the input of new data and updating of existing data to be centrally controlled or devolved.

These are in addition to the hardware and software decisions (like Content Management Systems), participation issues (like good taste, harassment, confidentiality), and features to be supported.

The actual implementation would include steps such as

1. User involvement to identify users' information needs.
2. Setting up a web server with the correct hardware and software.
3. Setting up web server access using a TCP/IP network.
4. Installing the user programs on all required computers.
5. Creating a homepage for the content to be hosted.
6. User involvement in testing and promoting use of intranet.

4.2.5 Industry examples

- KPMG moved all of its information assets to an intranet called KWorld.
- Calor Gas uses an Intranet called Interact Intranet to share documents and communicate with staff. They also use it for *electronic forms and workflow*.
- "The success of Cisco Systems has been largely attributed to its innovative corporate intranet".
- The People's Republic of China (PRC) is attempting to build a national intranet while limiting access to information forbidden by Chinese Internet regulations. If they are successful in their attempt, it will be the largest intranet.
- North Korea controls access to the internet among its people by using a country-wide intranet called "Kwangbyong" ('Bright'), while those with actual access to the internet are controlled by a government agency called the Korean Computer Centre. They are employed in searching the web for useful information, which is then copied and moved onto the Kwangbyong.
- Ford Motor Co has more than 175,000 employees in 950 locations worldwide, each of whom had access to the company's intranet. The intranet gave employees information about benefits, demographics, salary history, general company news and human resources forms.
- ShoreBank's branch, affiliate, and consulting service employees around the world communicate and collaborate using SIREN. SIREN is an intranet, extranet, and knowledge management solution implemented in 2006 using Intranet DASHBOARD.
- The Australian National University uses an Intranet called Claromentis to maintain one of its external sites.

4.2.6 Physical implementation

The initial default hardware implementation for AppleTalk was a high-speed serial protocol known as **LocalTalk** that used the Macintosh's built-in RS-422 ports at 230.4 kbit/s. LocalTalk used a splitter box in the RS-422 port to provide an upstream and downstream cable from a single port. The system was slow by today's standards,

NOTES

NOTES

but at the time the additional cost and complexity of networking on PC machines was such that it was common that Macs were the only networked machines in the office.

Other physical implementations were also available. One common replacement for LocalTalk was **PhoneNet**, a 3rd party solution (from a company called Farallon) that also used the RS-422 port and was indistinguishable from LocalTalk as far as Apple's LocalTalk port drivers were concerned, but ran over two unused wires in existing phone cabling. PhoneNet was considerably less expensive to install and maintain. Ethernet and TokenRing was also supported, known as **EtherTalk** and **TokenTalk** respectively. EtherTalk in particular gradually became the dominant implementation method for AppleTalk as Ethernet became generally popular in the PC industry throughout the 1990s.

4.2.6.1 Networking model

OSI Model	Corresponding AppleTalk layers
Application	Apple Filing Protocol (AFP)
Presentation	Apple Filing Protocol (AFP)
Session	Zone Information Protocol (ZIP) AppleTalk Session Protocol (ASP)
Transport	AppleTalk Data Stream Protocol (ADSP) AppleTalk Transaction Protocol (ATP) AppleTalk Echo Protocol (AEP) Name Binding Protocol (NBP) Routing Table Maintenance Protocol (RTMP)
Network	Datagram Delivery Protocol (DDP)
Data link	EtherTalk Link Access Protocol (ELAP) LocalTalk Link Access Protocol (LLAP) TokenTalk Link Access Protocol (TLAP) Fiber Distributed Data Interface (FDDI)
Physical	LocalTalk driver Ethernet driver Token Ring driver FDDI driver

4.2.7 Cross Platform Solutions

The BSD and Linux operating systems support AppleTalk through an open source project called Netatalk, which implements the complete protocol suite and allows them to both act as native file or print servers for Macintoshes, and print to LocalTalk printers over the network.

In addition, Columbia University released the Columbia AppleTalk Package (CAP) which implemented the protocol suite for various Unix flavours including Ultrix, SunOS, *BSD and IRIX. This package is no longer actively maintained.

4.3 OTHER NEW OPTION IN TCP

NOTES

DECnet is a proprietary suite of network protocols created by Digital Equipment Corporation, originally released in 1975 in order to connect two PDP-11 minicomputers. It evolved into one of the first peer-to-peer network architectures, thus transforming DEC into a networking powerhouse in the 1980s.

Initially built with four layers, it later (1992) evolved into a seven layer OSI compliant networking protocol, around the time when open systems (POSIX compliant, i.e. Unix-like) were grabbing marketshare from the proprietary OSes like VAX/VMS and AlphaVMS.

DECnet was built right into the DEC flagship operating system VAX/VMS since its inception. Digital ported it to its own Ultrix variant of UNIX, as well as Apple Macintosh computers and PCs running both DOS and Windows under the name **DEC Pathworks**, transforming these systems into DECnet end-nodes in a network of VAX machines. More recently, an open-source version has been developed for the Linux OS: see Linux-DECnet on Sourceforge.

Some of the new options are listed next.

4.3.1 Apple Talk

AppleTalk is a proprietary suite of protocols developed by Apple Inc for computer networking. It was included in the original Macintosh (1984) and is now deprecated by Apple in favour of TCP/IP networking.

4.3.1.1 Design

The design fairly rigorously followed the OSI model of protocol layering. Unlike most other early LAN systems, AppleTalk was not built on the archetypal Xerox XNS system, as the intended target was not Ethernet and did not have 48-bit addresses to route. Nevertheless many portions of the AppleTalk system have direct analogs in XNS.

One key differentiator for AppleTalk was that the system contained two protocols aimed at making the system completely self-configuring. The **AppleTalk address resolution protocol (AARP)** allowed AppleTalk hosts to automatically generate their own network addresses, and the **Name Binding Protocol (NBP)** was essentially a dynamic Domain Name System (DNS), system which mapped network addresses to user-readable names. Although systems similar to AARP existed in other systems, Banyan VINES for instance, nothing like NBP has existed until recently.

Both AARP and NBP had defined ways to allow "controller" devices to override the default mechanisms.

The concept here was to allow routers to provide all of this information, or additionally "hardwire" the system to known addresses and names. On larger networks where AARP could cause problems as new nodes searched for free addresses, the addition of a router could dramatically reduce "chattiness".

NOTES

Together AARP and NBP made AppleTalk perhaps the easiest to use networking system yet developed. New machines were added to the network simply by plugging them in, and optionally giving them a name. The NBP lists were examined and displayed by a program known as the **Chooser** (originally because it allowed you to choose your default printer) which would display a list of machines on the local network, divided into classes such as file servers and printers. All of this was completely automated.

One problem for AppleTalk was that it was originally intended to be part of a project known as **Macintosh Office**, which would consist of a host machine providing routing, printer sharing and file sharing. However, this project was cancelled in 1986, although the LaserWriter included AppleTalk built-in. Apple did eventually release a File and Print Server known as the AppleShare File and Print Servers.

For some the old AppleTalk protocol was considered clunky and often called 'chatty', notably on larger networks and Wide area networks (WAN) where the naming services generated considerable unwanted traffic. AppleTalk Phase 2, included with System 7, reduced the chattiness significantly.

Today AppleTalk support is provided for backward compatibility in many products, but the default networking on the Mac is TCP/IP. Starting with Mac OS X v10.2, Bonjour (originally named *Rendezvous*) provides similar discovery and configuration services for TCP/IP-based networks. Bonjour is Apple's implementation of ZeroConf, which was written specifically to bring NBP's ease-of-use to the TCP/IP world.

4.3.1.2 Addressing

An AppleTalk address was a 4-byte quantity. This consisted of a two-byte network number, a one-byte node number, and a one-byte socket number. Of these, only the network number required any configuration, being obtained from a router. Each node dynamically chose its own node number, according to a protocol which handled contention between different nodes accidentally choosing the same number. For socket numbers, a few well-known numbers were reserved for special purposes specific to the AppleTalk protocol itself. Apart from these, all application-level protocols were expected to use dynamically-assigned socket numbers at both the client and server end.

Because of this dynamism, users could not be expected to access services by specifying their address. Instead, all services had *names* which, being chosen by humans, could be expected to be meaningful to users, and also could be sufficiently long enough to minimize the chance of conflicts.

Note that, because a name translated to an address which included a socket number as well as a node number, a name in AppleTalk mapped directly to a *service* being provided by a machine, which was entirely separate from the name of the machine itself. Thus, services could be moved to a different machine and, so long as they kept the same service name, there was no need for users to do anything different to continue accessing the service. And the same machine could host any number of instances of services of the same type, without any network connection conflicts.

Contrast this with *A records* in the DNS, where a name translates only to a machine address, not including the port number that might be providing a service. Thus, if people are accustomed to using a particular machine name to access a particular service, their access will break when the service is moved to a different machine. This

can be mitigated somewhat by insistence on using *CNAME records* indicating service rather than actual machine names to refer to the service, but there is no way of guaranteeing that users will follow such a convention. Some newer protocols, such as Kerberos and Active Directory use DNS *SRV records* to identify services by name, which is much closer to the AppleTalk model.

4.3.1.3 AppleTalk Address Resolution Protocol

AARP resolves AppleTalk addresses to physical layer, usually MAC, addresses. It is functionally equivalent to ARP.

AARP is a fairly simple system. When powered on, an AppleTalk machine broadcasts an **AARP probe packet** asking for a network address, intending to hear back from controllers such as routers. If no address is provided, one is picked at random from the "base subnet", 0. It then broadcasts another packet saying "I am selecting this address", and then waits to see if anyone else on the network complains. If another machine has that address, it will pick another address, and keep trying until it finds a free one. On a network with many machines it may take several tries before a free address is found, so for performance purposes the successful address is "written down" in NVRAM and used as the default address in the future. This means that in most real-world setups where machines are added a few at a time, only one or two tries are needed before the address effectively becomes constant.

4.3.1.4 AppleTalk Data Stream Protocol

This was a comparatively late addition to the AppleTalk protocol suite, done when it became clear that a TCP-style reliable connection-oriented transport was needed. Significant differences from TCP were:

- a connection attempt could be rejected
- there were no "half-open" connections; once one end initiated a tear-down of the connection, the whole connection would be closed (*i.e.*, ADSP is full-duplex, not dual simplex).

4.3.1.5 Apple Filing Protocol

The Apple Filing Protocol (AFP), formerly AppleTalk Filing Protocol, is the protocol for communicating with AppleShare file servers. Built on top of AppleTalk Session Protocol, it provides services for authenticating users (extensible to different authentication methods including two-way random-number exchange) and for performing operations specific to the Macintosh HFS filesystem. AFP is still in use in Mac OS X, even though most other AppleTalk protocols have been deprecated.

4.3.1.6 AppleTalk Session Protocol

ASP was an intermediate protocol, built on top of ATP, which in turn was the foundation of AFP. It provided basic services for requesting responses to arbitrary *commands* and performing out-of-band status queries. It also allowed the server to send asynchronous *attention* messages to the client.

4.3.1.7 AppleTalk Transaction Protocol

ATP was the original reliable transport-level protocol for AppleTalk, built on top of DDP. At the time it was being developed, a full, reliable connection-oriented protocol like TCP was considered to be too expensive to implement for most of the intended

NOTES

uses of AppleTalk. Thus, ATP was a simple request/response exchange, with no need to set up or tear down connections.

An ATP *request* packet could be answered by up to eight *response* packets. The requestor then sent an *acknowledgement* packet containing a bit mask indicating which of the response packets it received, so the responder could retransmit the remainder.

ATP could operate in either “at-least-once” mode or “exactly-once” mode. Exactly-once mode was essential for operations which were not idempotent; in this mode, the responder kept a copy of the response buffers in memory until successful receipt of a *release* packet from the requestor, or until a timeout elapsed. This way, it could respond to duplicate requests with the same transaction ID by resending the same response data, without performing the actual operation again.

4.3.1.8 Datagram Delivery Protocol

DDP was the lowest-level data-link-independent transport protocol. It provided a datagram service with no guarantees of delivery. All application-level protocols, including the infrastructure protocols NBP, RTMP and ZIP, were built on top of DDP.

4.3.1.9 Name Binding Protocol

NBP was a dynamic, distributed system for managing AppleTalk names. When a service started up on a machine, it registered a name for itself on that machine, as chosen by a human administrator. At this point, NBP provided a system for checking that no other machine had already registered the same name. Then later, when a client wanted to access that service, it used NBP to query machines to find that service. NBP provided browseability (“what are the names of all the services available?”) as well as the ability to find a service with a particular name.

As would be expected from Apple, names were truly human readable, containing spaces, upper and lower case letters, and including support for searching.

4.3.1.10 Printer Access Protocol

PAP was the standard way of communicating with PostScript printers. It was built on top of ATP. When a PAP connection was opened, each end sent the other an ATP request which basically meant “send me more data”. The client’s response to the server was to send a block of PostScript code, while the server could respond with any diagnostic messages that might be generated as a result, after which another “send-more-data” request was sent. This use of ATP provided automatic flow control; each end could only send data to the other end if there was an outstanding ATP request to respond to.

PAP also provided for out-of-band status queries, handled by separate ATP transactions. Even while it was busy servicing a print job from one client, a PAP server could continue to respond to status requests from any number of other clients. This allowed other Macintoshes on the LAN that were waiting to print to display status messages indicating that the printer was busy, and what the job was that it was busy with.

4.3.1.11 Routing Table Maintenance Protocol

RTMP was the protocol by which routers kept each other informed about the topology of the network. This was the only part of AppleTalk that required periodic unsolicited broadcasts: every 10 seconds, each router had to send out a list of all the network numbers it knew about and how far away it thought they were.

4.3.1.12 Zone Information Protocol

ZIP was the protocol by which AppleTalk network numbers were associated with zone names. A *zone* was a subdivision of the network that made sense to humans (for example, "Accounting Department"); but while a network number had to be assigned to a topologically-contiguous section of the network, a zone could include several different discontinuous portions of the network.

NOTES

SUMMARY

1. In computing, a protocol is a convention or standard that controls or enables the connection, communication, and data transfer between two computing endpoints.
2. The pair of Internet Protocol (or IP) and Transmission Control Protocol (or TCP) are the most important of these, and the term TCP/IP refers to a collection (or protocol suite) of its most used protocols.
3. IPX/SPX stands for Internetwork Packet Exchange/Sequenced Packet Exchange.
4. IPX and SPX both provide connection services similar to TCP/IP, with the IPX protocol having similarities to IP, and SPX having similarities to TCP.
5. An intranet is a private computer network that uses Internet protocols, network connectivity to securely share part of an organization's information or operations with its employees.
6. Intranets differ from "Extranets" in that the former is generally restricted to employees of the organization while extranets can generally be accessed by customers, suppliers, or other approved parties.
7. Most organizations devote considerable resources into the planning and implementation of their intranet as it is of strategic importance to the organization's success.
8. AppleTalk is a proprietary suite of protocols developed by Apple Inc for computer networking.
9. Today AppleTalk support is provided for backward compatibility in many products, but the default networking on the Mac is TCP/IP.
10. AARP resolves AppleTalk addresses to physical layer, usually MAC, addresses. It is functionally equivalent to ARP.
11. The Apple Filing Protocol (AFP), formerly AppleTalk Filing Protocol, is the protocol for communicating with AppleShare file servers.
12. Datagram Delivery Protocol (DDP) was the lowest-level data-link-independent transport protocol.
13. Name Binding Protocol (NBP) was a dynamic, distributed system for managing AppleTalk names.
14. Routing Table Maintenance Protocol (RTMP) was the protocol by which routers kept each other informed about the topology of the network.
15. Zone Information Protocol (ZIP) was the protocol by which AppleTalk network numbers were associated with zone names.
16. Digital Network Architecture (DECnet) is a proprietary suite of network protocols created by Digital Equipment Corporation, originally released in 1975 in order to connect two PDP-11 minicomputers.

SELF ASSESSMENT QUESTIONS

1. What are protocols?
2. Which are common protocols?
3. Describe Network IPX/SPX.

NOTES

4. Describe the various Intranet Protocols.
5. What is AppleTalk?
6. Describe Digital Network Architecture.
7. Write short note on following:

DECnet Phase IV protocol suite
AppleTalk Data Stream Protocol
AppleTalk Session Protocol
Datagram Delivery Protocol
Printer Access Protocol
Extranets.

AppleTalk Address Resolution Protocol
Apple Filing Protocol
AppleTalk Transaction Protocol
Name Binding Protocol
Routing Table Maintenance Protocol

Multiple Choice Questions

1. IETF is :
(a) Internet Engineering Task Force (b) Indian Engineering Task Force
(c) Internet Engineering Travel Force.
2. HTTP is :
(a) Hypertext Travel Protocol (b) Hypertext Transfer Protocol
(c) Hypertext Transfer Port.
3. AARP is :
(a) AppleTalk added resolution protocol
(b) AppleTalk address resolution port
(c) AppleTalk address resolution protocol.
4. AFP is :
(a) Apple Filing Protocol (b) Apple Find Protocol
(c) Apple Filing Port.
5. ELAP is :
(a) EtherTalk Link Added Protocol (b) EtherTalk Link Access Protocol
(c) EtherTalk Link Access Port.

True/False Questions

1. IPX/SPX stands for Internetwork Packet Exchange/Sequenced Packet Exchange.
2. An intranet is a private computer network that uses Internet protocols, network connectivity to securely share part of an organization's information or operations with its employees.
3. AppleTalk is not a proprietary suite of protocols developed by Apple Inc for computer networking.
4. AARP resolves AppleTalk addresses to physical layer, usually MAC, addresses. It is functionally equivalent to ARP.
5. Datagram Delivery Protocol (DDP) was the highest-level data-link-independent transport protocol.
6. Name Binding Protocol (NBP) was a dynamic, distributed system for managing AppleTalk names.
7. Zone Information Protocol (ZIP) was the protocol by which AppleTalk network numbers were associated with zone names.

Short Questions with Answers

1. What properties do most protocols follow?

Ans. Most protocols specify one or more of the following properties:

- Detection of the underlying physical connection (wired or wireless), or the existence of the other endpoint or node

- Handshaking
 - Negotiation of various connection characteristics
 - How to start and end a message
 - How to format a message
 - What to do with corrupted or improperly formatted messages (error correction)
 - How to detect unexpected loss of the connection, and what to do next
 - Termination of the session or connection.
2. Which are the most common protocols?

Ans. The common protocols are:

- IP (Internet Protocol)
- DHCP (Dynamic Host Configuration Protocol)
- TCP (Transmission Control Protocol)
- HTTP (Hypertext Transfer Protocol)
- FTP (File Transfer Protocol)
- POP3 (Post Office Protocol 3)
- SMTP (Simple Mail Transfer Protocol)
- IMAP (Internet Message Access Protocol)

3. What is Intranet?

Ans. An **intranet** is a private computer network that uses Internet protocols, network connectivity to securely share part of an organization's information or operations with its employees. Sometimes the term refers only to the most visible service, the internal website.

4. Which topics would you consider while planning for an Intranet?

Ans. Some of the planning would include topics such as:

- What they hope to achieve from the intranet
- Which person or department would "own" (take control of) the technology and the implementation
- How and when existing systems would be phased out/replaced
- How they intend to make the intranet secure
- How they'll ensure to keep it within legislative and other constraints
- Level of interactivity or interactivity (eg on-line forms) desired.
- Is the input of new data and updating of existing data to be centrally controlled or devolved.

5. What does AARP do?

Ans. AARP is a fairly simple system. When powered on, an AppleTalk machine broadcasts an AARP probe packet asking for a network address, intending to hear back from controllers such as routers. If no address is provided, one is picked at random from the "base subnet", 0.

6. What is the difference between TCP and AppleTalk Data Stream Protocol?

Ans. Significant differences from TCP were:

- a connection attempt could be rejected
- there were no "half-open" connections; once one end initiated a tear-down of the connection, the whole connection would be closed (*i.e.*, ADSP is full-duplex, not dual simplex).

7. What is AFP?

Ans. The Apple Filing Protocol (AFP), formerly AppleTalk Filing Protocol, is the protocol for communicating with AppleShare file servers.

8. What is DDP?

Ans. Datagram Delivery Protocol (DDP) was the lowest-level data-link-independent transport protocol.

9. What is NBP?

Ans. Name Binding Protocol (NBP) was a dynamic, distributed system for managing AppleTalk names.

NOTES

NOTES

10. What is RTMP?

Ans. Routing Table Maintenance Protocol (RTMP) was the protocol by which routers kept each other informed about the topology of the network.

11. What is ZIP?

Ans. Zone Information Protocol (ZIP) was the protocol by which AppleTalk network numbers were associated with zone names.

12. What is DECnet?

Ans. Digital Network Architecture (DECnet) is a proprietary suite of network protocols created by Digital Equipment Corporation, originally released in 1975 in order to connect two PDP-11 minicomputers.

ANSWERS

Multiple Choice Questions

- | | | | |
|-------|------|------|------|
| 1. a | 2. b | 3. c | 4. a |
| 5. b. | | | |

True/False Questions

- | | | | |
|------|------|-------|------|
| 1. T | 2. T | 3. F | 4. T |
| 5. F | 6. F | 7. T. | |

Further Readings

1. **Data and Computer Network Communication:** Prof. Shashi Banzai, Firewal Media.
2. **Networking:** Balvir Singh, Firewal Media.
3. **Wide Area Networks:** Navneet Sharma, Firewal Media.
4. **Computer Network:** Bharat Bhushan Agarwal and Sumit Prakash Tayal, University Science Press.

UNIT 5

NETWORK SECURITY

NOTES

STRUCTURE

- 5.1 Network security at various levels
- 5.2 Secure HTTP
- 5.3 SSL
- 5.4 ESP
- 5.5 Authentication header
- 5.6 Key distribution protocols
- 5.7 Digital signatures
- 5.8 Digital certificates
- 5.9 Firewall
 - Summary
 - Self Assessment Questions
 - Further Readings

Learning Objectives

After going through this unit, students will be able to:

- know about networking security at various levels of transmission.
- learn about Secure HTTP.
- know about topics like SSL, and ESP.
- learn about authentication headers and its usage.
- learn about digital signatures.
- know about firewall and its usage.

5.1 NETWORK SECURITY AT VARIOUS LEVELS

NOTES

Computer or network security has been violated when unauthorized access by any party occurs. So it becomes your duty as the Network Administrator to work on the security to guard against any such incidence taking place.

5.1.1 Need of Network Security

Computer security is required because most organizations can be damaged by hostile software or intruders. There may be several forms of damage which are obviously interrelated. These include:

- Damage or destruction of computer systems.
- Damage or destruction of internal data.
- Loss of sensitive information to hostile parties.
- Use of sensitive information to steal items of monetary value.
- Use of sensitive information against the organization's customers which may result in legal action by customers against the organization and loss of customers.
- Damage to the reputation of an organization.
- Monetary damage due to loss of sensitive information, destruction of data, hostile use of sensitive data, or damage to the organization's reputation.

The methods used to accomplish these unscrupulous objectives are many and varied depending on the circumstances. This guide will help administrators understand some of these methods and explain some countermeasures.

5.1.1.1 Security Issues

Computer security can be very complex and may be very confusing to many people. It can even be a controversial subject. Network administrators like to believe that their network is secure and those who break into networks may like to believe that they can break into any network.

Overconfidence plays an important role in allowing networks to be intruded upon.

There are many fallacies that network administrators may fall victim to. These fallacies may allow administrators to wrongfully believe that their network is more secure than it really is. This guide will attempt to clarify many issues related to security by doing the following:

- Help you determine what you are protecting.
- Break computer security into categories.
- Explain security terms and methods.
- Point out some common fallacies that may allow administrators to be overconfident.
- Categorize many common attacks against networks and computers.

- Explain some attack methods.
- Describe tools that can be used to help make a network more secure.

5.1.1.2 Security Interdependence

There are many different aspects to computer and network security as you will read in this book. These different areas of computer security are interdependent on each other in order for a network to be secure. If one or more areas of computer security are ignored, then the entire security integrity of the organization's network may be compromised. A clear example of this is in the area of computer virus or worm protection.

Computer virus protection programs can only filter known viruses or worms. There are viruses or worms that are not yet recognized as virus programs immediately after their release.

The best way to make unrecognized virus or worm programs less effective is by quickly removing the vulnerabilities that they use.

Some of these vulnerabilities are operating system and application program errors. When security patches are created for software, they should be quickly applied. In this way the vulnerability to viruses is minimized but not eliminated. There are other steps which may further reduce this vulnerability, but it can never be completely eliminated.

5.1.1.3 Security Limitations and Applications

If you are reading this document and are thinking that you can get all the information required to make your network completely secure, then you are sadly mistaken. In many ways, computer security is almost a statistical game.

You can reduce but not eliminate the chance that you may be penetrated by an intruder or virus. This is mainly for one reason.

No one can ever know all the software vulnerabilities of all software used on a system.

This is why even those who consider themselves hackers will say that the number one computer security threat is the lack of quality in the applications and operating systems. At this point, I could talk about the various corporate entities that write software and why software lacks the quality that many of us believe that it should possess, but that subject is not only way beyond the scope of this document, but also way beyond the scope of this project.

The bottom line here is that unless you can remove all the application and operating system problems that allow viruses and intruders to penetrate networks, you can never secure your network.

Additionally the users on your network are potentially a greater security risk than any programs. Obviously removing all vulnerabilities is impossible and will not secure your network against user errors. I have even considered the possibility that an operating system without a network interface can be completely secure, but even this cannot be guaranteed. Unknown viruses or trojan programs can creep in with

NOTES

applications on CDs or floppies. This has been known to happen. Although an attacker may not be able to get data from the system, they can damage or destroy data.

5.1.1.4 Layered Security

The fact that complete security is impossible is the reason security experts recommend "layered security". The idea is to have multiple ways of preventing an intrusion to decrease the chance that intrusions will be successful. For example, you should have virus protection on your client computers. To help layer this security you should also filter viruses at your email server.

To help even more, you should block the most dangerous types of email attachments to prevent unrecognized viruses and other hostile software from entering your network.

Another good defense layer would also include educating your users about viruses, how they spread, and how to avoid them.

5.1.1.5 Hackers

There are many documents that attempt to define the term hacker. I believe that the term hacker is a connotative term. This means that it is more defined by people's beliefs rather than by a dictionary.

Some believe that a hacker is a very skilled computer person.

Others believe that hackers are those that perform unauthorized break-ins to computer systems. The media and many sources have caused many uninformed people to believe that a hacker is a threat to computer and network security while this is not the case. A hacker is no more likely to break the law than anyone else. It would be wise to use the more accurate descriptive term, "intruder" to describe those who intrude into networks or systems without authorization.

5.1.1.6 Physical Security

This book will not talk about physical computer security beyond this paragraph. Your organization should be aware how physically secure every aspect of its network is because if an intruder gets physical access, they can get your data.

Be sure that your organization properly secures locations and consider the following:

- **Servers** - Contain your data and information about how to access that data.
- **Workstations** - May contain some sensitive data and can be used to attack other computers.
- **Routers, switches, bridges, hubs** and any other network equipment may be used as an access point to your network.
- **Network wiring and media** and where they pass through may be used to access your network or place a wireless access point to your network.
- **External media** which may be used between organizational sites or to other sites the organization does business with.
- **Locations of staff** who may have information that a hostile party can use.
- **Some employees** may take data home or may take laptops home or use laptops

on the internet from home then bring them to work. Any information on these laptops should be considered to be at risk and these laptops should be secure according to proper policy when connected externally on the network (more on this later).

5.1.1.7 Some Technical Terms

This paragraph describes some commonly used computer security terms.

- **Protocol** - Well defined specification allowing computer communication.
- **Confidentiality** - Information is available only to people with rightful access.
- **Integrity** - Information can only be changed by authorized personnel. Integrity - The receiver of the message should be able to tell that the message was not modified. Requires key exchange.
- **Availability** - Information is available to only those who need it.
- **Verification - nonrepudiation** - There is proof that the sender sent the message.
- **Authentication** - The receiver of the message should be able to be sure of the origin of the message. Requires a digital signature (One way hash, public key algorithm, and symmetric algorithm) or a public key algorithm.
- **Spyware** - A computer program whose purpose is to spy on your internet activities usually for marketing purposes and usually done by a shady corporate entity.
- **Malware** - A computer program with some evil intent. It may on the surface have a good or useful intent, but may be a trojan (with a hidden purpose) which can be used to gain unauthorized access to your computer.

5.1.2 Attackers vs Hackers: attackers from Within and External

In a security context, a Hacker is someone who is involved in computer security/insecurity, specializing in the discovery of exploits in systems (for exploitation or prevention), or in obtaining or preventing unauthorized access to systems through skills, tactics and detailed knowledge. In the most common general form of this usage, "hacker" refers to a black-hat hacker (a malicious or criminal hacker); many who use the term in any other sense insist on using the term cracker to refer to such hackers.

There are also ethical hackers (more commonly referred to as white hats), and those more ethically ambiguous (grey hats).

The context of computer security hacking forms a subculture which is often referred to as the network hacker subculture or simply the computer underground. According to its adherents, cultural values center around the idea of creative and extraordinary computer usage. Proponents claim to be motivated by artistic and political ends, but are often unconcerned about the use of criminal means to achieve them.

5.1.2.1 Artifacts and Customs

Contrary to the academic hacker subculture, networking hackers have no inherently close connection to the academic world. They have a tendency to work anonymously and in private. It is common among them to use aliases for the purpose of concealing

NOTES

NOTES

identity, rather than revealing their real names. This practice is uncommon within and even frowned upon by the academic hacker subculture. Members of the network hacking scene are often being stereotypically described as crackers by the academic hacker subculture, yet see themselves as hackers and even try to include academic hackers in what they see as one wider hacker culture, a view harshly rejected by the academic hacker subculture itself. Instead of a hacker–cracker dichotomy, they give more emphasis to a spectrum of different categories, such as white hat (“ethical hacking”), grey hat, black hat and script kiddie.

In contrast to the academic hackers, they usually reserve the term cracker to refer to black hat hackers, or more generally hackers with unlawful intentions.

The network hacking subculture is supported by regular gatherings, so called Hacker cons. These have drawn more and more people every year including SummerCon (Summer), DEF CON, HoHoCon (Christmas), PumpCon (Halloween), H.O.P.E. (Hackers on Planet Earth) and HEU (Hacking at the End of the Universe). They have helped expand the definition and solidify the importance of the network hacker subculture. In Germany, members of the subculture are organized mainly around the Chaos Computer Club.

The subculture has given birth to what its members consider to be novel forms of art, most notably ascii art. It has also produced its own slang and various forms of unusual alphabet use, for example leetspeak. Both things are usually seen as an especially silly aspect by the academic hacker subculture. In part due to this, the slangs of the two subcultures differ substantially.

Political attitude usually includes views for freedom of information, freedom of speech, a right for anonymity and most have a strong opposition against copyright. Writing programs and performing other activities to support these views is referred to as hacktivism by the subculture. Some go as far as seeing illegal computer cracking ethically justified for this goal; the most common form is website defacement.

The security hackers have also edited some publications, most notably

- “2600: The Hacker Quarterly”
- “Hakin9”
- “Binary Revolution Magazine 2006”
- “Blacklisted 411”

5.1.2.2 Academic hackers

In the academic hacker culture, a computer hacker is a person who enjoys designing software and building programs with a sense for aesthetics and playful cleverness.

According to Eric S. Raymond, the academic hacker subculture developed in the 1960s among hackers working on early minicomputers in academic computer science environments. After 1969 it fused with the technical culture of the pioneers of the Internet. One PDP-10 machine at MIT connected to the Internet provided an early hacker meeting point. It was called AI and ran ITS. After 1980 the subculture coalesced with the culture of Unix, and after 1987 with elements of the early microcomputer hobbyists that themselves had connections to radio amateurs in the 1920s. Since the mid-1990s, it has been largely coincident with what is now called the free software movement and the open source movement.

Many programmers have been labeled “great hackers,” but the specifics of who that label applies to is a matter of opinion. Certainly major contributors to computer science such as Edsger Dijkstra and Donald Knuth, as well as the inventors of popular software such as Linus Torvalds (Linux), and Dennis Ritchie and Ken Thompson (the C programming language) are likely to be included in any such list. People primarily known for their contributions to the consciousness of the academic hacker culture include Richard Stallman, the founder of the free software movement and the GNU project, president of the Free Software Foundation and author of the famous Emacs text editor as well as the GNU Compiler Collection (GCC), and Eric S. Raymond, one of the founders of the Open Source Initiative and writer of the famous text *The Cathedral and the Bazaar* and many other essays, maintainer of the Jargon File (which was previously maintained by Guy L. Steele, Jr.).

NOTES

Within the academic hacker culture, the term hacker is also used for a programmer who reaches a goal by employing a series of modifications to extend existing code or resources. In this sense, it can have a negative connotation of using kludges to accomplish programming tasks that are ugly, inelegant, and inefficient.

This derogatory form of the noun “hack” is even used among users of the positive sense of “hacker” (some argue that it should not be, due to this negative meaning; others argue that some kludges can, for all their ugliness and imperfection, still have “hack value”).

In a very universal sense, a hacker also means someone who makes things work beyond perceived limits in a clever way in general, for example reality hackers.

5.1.3 Hobby Hackers

The hobby hacking subculture relates to hobbyist home computing of the late 1970s, beginning with the availability of MITS Altair. An influential organization was the Homebrew Computer Club.

The parts that didn't fuse with the academic hacker subculture focus mainly on commercial computer and video games, software cracking and exceptional computer programming (demo scene), but also to the modification of computer hardware and other electronic devices.

5.1.4 Overlaps and differences

The main basic difference between academic and computer security hackers is their separate historical origin and development. The Jargon File reports that although considerable overlap existed for the early phreaking at the beginning of the 1970s, it quickly started to break when people joined in the activity who did it in a less responsible way.

Academic hackers usually work openly and use their real name, while computer security hackers prefer secretive groups and identity-concealing aliases.

Also, their activities in practice are largely distinct. The former focus on creating new and improving existing infrastructure (especially the software environment they work with), while the latter primarily and strongly emphasize the general act of circumvention of security measures, with the effective use of the knowledge (which can be to report and help fixing the security bugs, or exploitation for criminal purpose)

NOTES

being only rather secondary. The most visible difference in these views was in the design of the MIT hackers' Incompatible Timesharing System, which deliberately didn't have any security measures.

There are some subtle overlaps, however, since basic knowledge about computer security is also common within the academic hacker community.

It sees secondary circumvention of security mechanisms as legitimate if it is done to get practical barriers out of the way for doing actual work. In special forms, that can even be an expression of playful cleverness. However, the systematic and primary engagement in such activities is not one of the actual interests of the academic hacker subculture and it doesn't have significance in its actual activities, either.

A further difference is that, historically, academic hackers were working at academic institutions and used the computing environment there. In contrast, the prototypical computer security hacker had access exclusively to a home computer and a modem.

Since the mid-1980s, there are some overlaps in ideas and members with the computer security hacking community. The most prominent case is Robert T. Morris, who was a user of MIT-AI, yet wrote the Morris worm. The Jargon File hence calls him "a true hacker who blundered". Nevertheless, members of the academic subculture have a tendency to look down and disassociate from these overlaps. They commonly refer disparagingly to people in the computer security subculture as crackers, and refuse to accept any definition of hacker that encompasses such activities.

The computer security hacking subculture on the other hand tends not to distinguish between the two subcultures as harshly, instead acknowledging that they have much in common including many members, political and social goals, and a love of learning about technology. They restrict the use of the term cracker to their categories of script kiddies and black hat hackers instead. There is also overlap into the other direction. Since the mid-1990s, with home computers that could run Unix-like operating systems and with inexpensive internet home access being available for the first time, many people from outside of the academic world started to take part in the academic hacking subculture.

All three subcultures have relations to hardware modifications. In the early days of network hacking, phreaks were building blue boxes and various variants. The academic hacker culture has stories about several hardware hacks in its folklore, such as a mysterious 'magic' switch attached to a PDP-10 computer in MIT's AI lab, that, when turned off, crashed the computer.

The early hobbyist hackers built their home computers themselves, from construction kits.

However, all these activities have died out during the 1980s, when the phone network switched to digitally controlled switchboards, causing network hacking to shift to dialling remote computers with modems, when preassembled inexpensive home computers were available, and when academic institutions started to give individual mass-produced workstation computers to scientists instead of using a central timesharing system. The only kind of widespread hardware modification nowadays is case modding.

An encounter of the academic and the computer security hacker subculture occurred at the end of the 1980s, when a group of hackers, sympathizing with the Chaos Computer Club (who disclaimed any knowledge in these activities), broke into computers of American military organizations and academic institutions. They sold data from these machines to the Soviet secret service, one of them in order to fund his drug addiction. The case could be solved when scientists from the environment of the academic hacker subculture found ways to log the attacks and to trace them back.

NOTES

5.2 SECURE HTTP

Hypertext Transfer Protocol Secure (HTTPS) is a combination of the Hypertext Transfer Protocol with the SSL/TLS protocol to provide encryption and secure (website security testing) identification of the server. It uses port 443. HTTPS connections are often used for payment transactions on the World Wide Web and for sensitive transactions in corporate information systems. HTTPS should not be confused with Secure HTTP (S-HTTP) specified in RFC 2660.

The main idea of HTTPS is to create a secure channel over an insecure network. This ensures reasonable protection from eavesdroppers and man-in-the-middle attacks, provided that adequate cipher suites are used and that the server certificate is verified and trusted.

The trust inherent in HTTPS is based on major certificate authorities which come pre-installed in browser software (this is equivalent to saying "I trust certificate authority (e.g. VeriSign/Microsoft/etc.) to tell me who I should trust"). Therefore an HTTPS connection to a website can be trusted if and only if all of the following are true:

- The user trusts that their browser software correctly implements HTTPS with correctly pre-installed certificate authorities.
- The user trusts the certificate authority to vouch only for legitimate websites without misleading names.
- The website provides a valid certificate (an invalid certificate shows a warning in most browsers), which means it was signed by a trusted authority.
- The certificate correctly identifies the website (e.g. visiting <https://example.com> and receiving a certificate for "Example Inc." and not anything else).
- Either the intervening hops on the Internet are trustworthy, or the user trusts the protocol's encryption layer (TLS or SSL) is unbreakable by an eavesdropper.

HTTP
Persistence · Compression · HTTP Secure
Headers
Etag · Cookie · Referrer · Location
Status codes
301 Moved permanently
302 Found
303 See Other
403 Forbidden
404 Not Found

5.2.1 Difference from HTTP

As opposed to HTTP URLs which begin with "http://" and use port 80 by default, HTTPS URLs begin with "https://" and use port 443 by default.

HTTP is insecure and is subject to man-in-the-middle and eavesdropping attacks which can let attackers gain access to website accounts and sensitive information. HTTPS is designed to withstand such attacks and is considered secure (with the exception of older deprecated versions of SSL).

5.2.2 Network layers

HTTP operates at the highest layer of the OSI Model, the Application layer; but the security protocol operates at a lower sublayer, encrypting an HTTP message prior to transmission and decrypting a message upon arrival. Strictly speaking, HTTPS is not a separate protocol, but refers to use of ordinary HTTP over an encrypted Secure Sockets Layer (SSL) or Transport Layer Security (TLS) connection.

5.2.3 Server setup

To prepare a web server to accept HTTPS connections, the administrator must create a public key certificate for the web server. This certificate must be signed by a trusted certificate authority for the web browser to accept it. The authority certifies that the certificate holder is indeed the entity it claims to be. Web browsers are generally distributed with the signing certificates of major certificate authorities so that they can verify certificates signed by them.

5.2.4 Acquiring certificates

Organizations may also run their own certificate authority, particularly if they are responsible for setting up browsers to access their own sites (for example, sites on a company intranet, or major universities). They can easily add copies of their own signing certificate to the trusted certificates distributed with the browser.

Peer-to-peer certificate authorities also exist.

5.2.5 Use as access control

The system can also be used for client authentication in order to limit access to a web server to authorized users. To do this, the site administrator typically creates a certificate for each user, a certificate that is loaded into his/her browser. Normally, that contains the name and e-mail address of the authorized user and is automatically checked by the server on each reconnect to verify the user's identity, potentially without even entering a password.

5.2.6 In case of compromised private key

A certificate may be revoked before it expires, for example because the secrecy of the private key has been compromised. Newer versions of popular browsers such as Firefox, Opera, and Internet Explorer on Windows Vista implement the Online Certificate Status Protocol (OCSP) to verify that this is not the case. The browser sends the certificate's serial number to the certificate authority or its delegate via OCSP and the authority responds, telling the browser whether or not the certificate is still valid.

5.3 SSL

SSL comes in two options; simple and mutual.

The mutual flavor is more secure but requires the user to install a personal certificate in their browser in order to authenticate themselves.

Whatever strategy is used (simple or mutual) the level of protection strongly depends on the correctness of the implementation of the web browser and the server software and the actual cryptographic algorithms supported.

SSL doesn't prevent the entire site from being indexed using a web crawler, and the URI of the encrypted resource can be inferred by knowing only the intercepted request/response size. This allows an attacker to have access to the plaintext (the publicly-available static content), and the encrypted text (the encrypted version of the static content), permitting a cryptographic attack.

Because SSL operates below HTTP and has no knowledge of higher-level protocols, SSL servers can only strictly present one certificate for a particular IP/port combination. This means that, in most cases, it is not feasible to use name-based virtual hosting with HTTPS.

A solution called Server Name Indication (SNI) exists which sends the hostname to the server before encrypting the connection, although many older browsers don't support this extension. Support for SNI is available since Firefox 2, Opera 8, and Internet Explorer 7 on Windows Vista.

If parental controls are enabled on Mac OS X, HTTPS sites must be explicitly allowed using the Always Allow list.

From a architectural point of view:

1. An SSL connection is managed by the first front machine which initiate the SSL connection. If for any reasons (routing, traffic optimization, etc.) this front machine is not the application server and it has to decipher data, solutions have to be found to propagate user authentication informations or certificate to the application server which needs to know who is going to be connected.
2. For SSL with mutual authentication, the SSL session is managed by the first server which initiates the connection. In situations where encryption has to be propagated along chained servers, session timeOut management becomes extremely tricky to implement.
3. With mutual SSL, security is maximal, but on the client-side, there is no way to properly end the SSL connection and disconnect the user except by waiting for the SSL server session to expire or closing all related client applications.
4. For performance reasons static contents are usually delivered through a non-encrypted front server or separate server instance with no SSL, as a consequence these contents are usually not protected.

5.4 ESP

Encapsulating Security Payload (ESP) is a member of the IPsec protocol suite. In

NOTES

NOTES

IPsec it provides origin authenticity, integrity, and confidentiality protection of packets. ESP also supports encryption-only and authentication-only configurations, but using encryption without authentication is strongly discouraged because it is insecure. Unlike Authentication Header (AH), ESP does not protect the IP packet header. However, in Tunnel Mode, where the entire original IP packet is encapsulated with a new packet header added, ESP protection is afforded to the whole inner IP packet (including the inner header) while the outer header remains unprotected. ESP operates directly on top of IP, using IP protocol number 50.

5.4.1 Security Parameters Index (32 bits)

Arbitrary value which is used (together with the source IP address) to identify the security association of the sending party.

5.4.2 Sequence Number (32 bits)

A monotonically increasing sequence number (incremented by 1 for every packet sent) to protect against replay attacks. There is a separate counter kept for every security association.

5.4.3 Payload data (variable)

The protected contents of the original IP packet, including any data used to protect the contents (e.g. an Initialisation Vector for the cryptographic algorithm). The type of content that was protected is indicated by the Next Header field.

5.4.4 Padding (0-255 octets)

Padding for encryption, to extend the payload data to a size that fits the encryption's cypher block size, and to align the next field.

5.4.5 Pad Length (8 bits)

Size of the padding in octets.

5.4.6 Next Header (8 bits)

Type of the next header. The value is taken from the list of IP protocol numbers.

5.4.7 Integrity Check Value (multiple of 32 bits)

Variable length check value. It may contain padding to align the field to an 8-octet boundary for IPv6, or a 4-octet boundary for IPv4.

5.5 AUTHENTICATION HEADER

Authentication Header (AH) is a member of the IPsec protocol suite. AH guarantees connectionless integrity and data origin authentication of IP packets. Further, it can optionally protect against replay attacks by using the sliding window technique and discarding old packets.

- In IPv4, the AH protects the IP payload and all header fields of an IP datagram except for mutable fields (i.e. those that might be altered in transit). Mutable (and therefore unauthenticated) IP header fields are DSCP/TOS, ECN, Flags, Fragment Offset, TTL and Header Checksum.

- In IPv6, the AH protects the AH itself, the Destination Options extension header after the AH, and the IP payload. It also protects the fixed IPv6 header and all extension headers before the AH, except for the mutable fields: DSCP, ECN, Flow Label, and Hop Limit.

AH operates directly on top of IP, using IP protocol number 51.

NOTES

5.6 KEY DISTRIBUTION PROTOCOLS

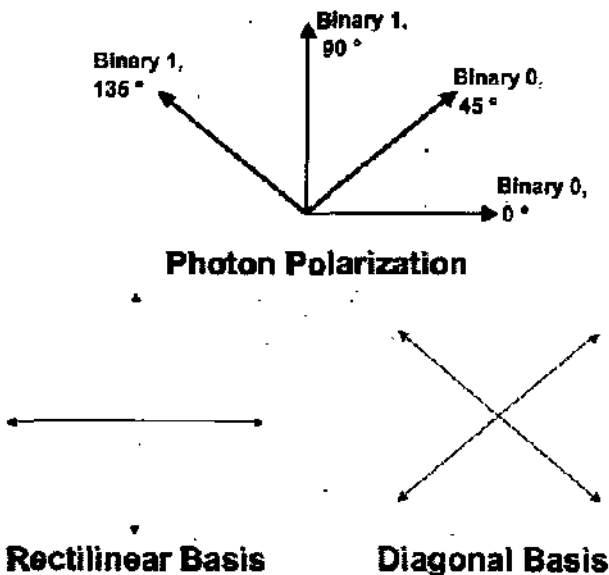
In 1984 Charles Bennett and Gilles Brassard published the first QKD protocol. It was based on Heisenberg's Uncertainty Principle and is simply known as the BB84 protocol after the authors names and the year in which it was published. It is still one of the most prominent protocols and one could argue that all of the other HUP based protocols are essentially variants of the BB84 idea. The basic idea for all of these protocols then is that Alice can transmit a random secret key to Bob by sending a string of photons where the secret key's bits are encoded in the polarization of the photons. Heisenberg's Uncertainty Principle can be used to guarantee that an Eavesdropper cannot measure these photons and transmit them on to Bob without disturbing the photon's state in a detectable way thus revealing her presence.

5.6.1 BB84 Protocol

The figure here shows how a bit can be encoded in the polarization state of a photon in BB84. We define a binary 0 as a polarization of 0 degrees in the rectilinear bases or 45 degrees in the diagonal bases. Similarly a binary 1 can be 90 degrees in the rectilinear bases or 135 in diagonal bases. Thus a bit can be represented by polarizing the photon in either one of two bases.

In the first phase, Alice will communicate to Bob over a quantum channel. Alice begins by choosing a random string of bits and for each bit, Alice will randomly choose a basis, rectilinear or diagonal, by which to encode the bit. She will transmit a photon for each bit with the corresponding polarization, as just described, to Bob. For every photon Bob receives, he will measure the photon's polarization by a randomly chosen basis. If, for a particular photon, Bob chose the same basis as Alice, then in principle, Bob should measure the same polarization and thus he can correctly infer the bit that Alice intended to send. If he chose the wrong basis, his result, and thus the bit he reads, will be random.

In the second phase, Bob will notify Alice over any insecure channel what basis he used to measure each photon. Alice will report back to Bob whether he chose the correct basis for each photon. At this point Alice and Bob will discard the bits corresponding to the photons which Bob measured with a different basis. Provided no errors



STUDENT ACTIVITY 5.1

1. Describe network security.

2. Who is an Attacker?

3. Who are Hobby Hackers?

4. What is the need for network security?

NOTES

occurred or no one manipulated the photons, Bob and Alice should now both have an identical string of bits which is called a sifted key. The example below shows the bits Alice chose, the bases she encoded them in, the bases Bob used for measurement, and the resulting sifted key after Bob and Alice discarded their bits as just mentioned.

Before they are finished however, Alice and Bob agree upon a random subset of the bits to compare to ensure consistency. If the bits agree, they are discarded and the remaining bits form the shared secret key.


In the absence of noise or any other measurement error, a disagreement in any of the bits compared would indicate the presence of an eavesdropper on the quantum channel.

Alice's bit	0	1	1	0	1	0	0	1
Alice's basis	+	+	X	+	X	X	X	+
Alice's polarization	↑	→	↖	↑	↖	↗	↗	→
Bob's basis	+	X	X	X	+	X	+	+
Bob's measurement	↑	↗	↖	↗	→	↗	→	→
Public discussion								
Shared Secret key	0		1			0		1

This is because the eavesdropper, Eve, were attempting to determine the key, she would have no choice but to measure the photons sent by Alice before sending them on to Bob. This is true because the no cloning theorem assures that she cannot replicate a particle of unknown state. Since Eve will not know what bases Alice used to encoded the bit until after Alice and Bob discuss their measurements, Eve will be forced to guess. If she measures on the incorrect bases, the Heisenberg Uncertainty Principle ensures that the information encoded on the other bases is now lost. Thus when the photon reaches Bob, his measurement will now be random and he will read a bit incorrectly 50% of the time. Given that Eve will choose the measurement basis incorrectly on average 50% of the time, 25% of Bob's measured bits will differ from Alice. If Eve has eavesdropped on all the bits then after n bit comparisons by Alice and Bob, they will reduce the probability that Eve will go undetected to $\frac{3}{4}^n$. The chance that an eavesdropper learned the secret is thus negligible if a sufficiently long sequence of the bits are compared.

5.6.2 B92 Protocol

In 1992, Charles Bennett proposed what is essentially a simplified version of BB84 in his paper, "Quantum cryptography using any two non-orthogonal states". The key difference in B92 is that only two states are necessary rather than the possible 4 polarization states in BB84. As shown in figure here, can be encoded as 0 degrees in the rectilinear basis and 1 can be encoded by 45 degrees in the diagonal basis. Like the BB84, Alice transmits to Bob a string of photons encoded with randomly chosen bits but this time the bits Alice chooses dictates which bases she must use. Bob still randomly chooses a basis by which to measure but if he chooses the wrong basis, he will not measure anything; a condition in quantum mechanics which is known as an erasure. Bob can simply tell Alice after each bit she sends whether or not he measured it correctly.



5.6.3 Other Uncertainty Based Protocols

Another variant of BB84 is the Six-State Protocol (SSP) proposed by Pasquinucci and Gisin in 1999. SSP is identical to BB84 except, as its name implies, rather than using two or four states, SSP uses six states on three orthogonal bases by which to encode the bits sent. This means that an eavesdropper would have to choose the right basis from among 3 possibilities. This extra choice causes the eavesdropper to produce a higher rate of error thus becoming easier to detect. Brus and Micchiavello proved in 2002 that such higher-dimensional systems offer increased security.

While there are a number of other BB84 variants, one of the more recent was proposed in 2004 by Scarani, Acin, Ribordy, and Gisin. The SARG04 protocol shares the exact same first phase as BB84. In the second phase, when Alice and Bob determine for which bits their bases matched, Alice does not directly announce her bases. Rather she announces a pair of non-orthogonal states, one of which she used to encode her bit. If Bob used the correct basis, he will measure the correct state. If he chose incorrectly, he will not measure either of Alice's states and he will not be able to determine the bit. This protocol has a specific advantage when used in practical equipment as will be discussed in Section 5.

BB84 was the first proposed QKD protocol and it was based on Heisenberg's Uncertainty Principle. A whole series of protocols followed which built on the ideas of BB84. Some of the most notable of these were B92, SSP, and Sarg04. The next section describes the alternate approach to QKD which is based on the principle of quantum entanglement.

NOTES

5.7 DIGITAL SIGNATURES

A digital signature or digital signature scheme is a mathematical scheme for demonstrating the authenticity of a digital message or document. A valid digital signature gives a recipient reason to believe that the message was created by a known sender, and that it was not altered in transit. Digital signatures are commonly used for software distribution, financial transactions, and in other cases where it is important to detect forgery and tampering.

Digital signatures are often used to implement electronic signatures, a broader term that refers to any electronic data that carries the intent of a signature, but not all electronic signatures use digital signatures. In some countries, including the United States, India, and members of the European Union, electronic signatures have legal significance. However, laws concerning electronic signatures do not always make clear whether they are digital cryptographic signatures in the sense used here, leaving the legal definition, and so their importance, somewhat confused.

Digital signatures employ a type of asymmetric cryptography. For messages sent through an insecure channel, a properly implemented digital signature gives the receiver reason to believe the message was sent by the claimed sender. Digital signatures are equivalent to traditional handwritten signatures in many respects; properly implemented digital signatures are more difficult to forge than the handwritten type. Digital signature schemes in the sense used here are cryptographically based, and must be implemented properly to be effective. Digital signatures can also provide non-repudiation, meaning that the signer cannot successfully claim they did not sign a message, while also claiming their private key remains secret; further, some non-

NOTES

repudiation schemes offer a time stamp for the digital signature, so that even if the private key is exposed, the signature is valid nonetheless. Digitally signed messages may be anything representable as a bitstring: examples include electronic mail, contracts, or a message sent via some other cryptographic protocol.

A digital signature scheme typically consists of three algorithms:

- A key generation algorithm that selects a private key uniformly at random from a set of possible private keys. The algorithm outputs the private key and a corresponding public key.
- A signing algorithm which, given a message and a private key, produces a signature.
- A signature verifying algorithm which given a message, public key and a signature, either accepts or rejects the message's claim to authenticity.

Two main properties are required. First, a signature generated from a fixed message and fixed private key should verify the authenticity of that message by using the corresponding public key. Secondly, it should be computationally infeasible to generate a valid signature for a party who does not possess the private key.

5.7.1 Notions of security

In their foundational paper, Goldwasser, Micali, and Rivest lay out a hierarchy of attack models against digital signatures:

- In a key-only attack, the attacker is only given the public verification key.
- In a known message attack, the attacker is given valid signatures for a variety of messages known by the attacker but not chosen by the attacker.
- In an adaptive chosen message attack, the attacker first learns signatures on arbitrary messages of the attacker's choice.

They also describe a hierarchy of attack results:

- A total break results in the recovery of the signing key.
- A universal forgery attack results in the ability to forge signatures for any message.
- A selective forgery attack results in a signature on a message of the adversary's choice.
- An existential forgery merely results in some valid message/signature pair not already known to the adversary.
- The strongest notion of security, therefore, is security against existential forgery under an adaptive chosen message attack.

5.7.2 Uses of digital signatures

As organizations move away from paper documents with ink signatures or authenticity stamps, digital signatures can provide added assurances of the evidence to provenance, identity, and status of an electronic document as well as acknowledging informed consent and approval by a signatory.

The United States Government Printing Office (GPO) publishes electronic versions of the budget, public and private laws, and congressional bills with digital signatures.

Universities including Penn State, University of Chicago, and Stanford are publishing electronic student transcripts with digital signatures.

Below are some common reasons for applying a digital signature to communications:

5.7.3 Authentication

Although messages may often include information about the entity sending a message, that information may not be accurate. Digital signatures can be used to authenticate the source of messages. When ownership of a digital signature secret key is bound to a specific user, a valid signature shows that the message was sent by that user. The importance of high confidence in sender authenticity is especially obvious in a financial context. For example, suppose a bank's branch office sends instructions to the central office requesting a change in the balance of an account. If the central office is not convinced that such a message is truly sent from an authorized source, acting on such a request could be a grave mistake.

5.7.4 Integrity

In many scenarios, the sender and receiver of a message may have a need for confidence that the message has not been altered during transmission. Although encryption hides the contents of a message, it may be possible to change an encrypted message without understanding it. (Some encryption algorithms, known as nonmalleable ones, prevent this, but others do not.) However, if a message is digitally signed, any change in the message after signature will invalidate the signature. Furthermore, there is no efficient way to modify a message and its signature to produce a new message with a valid signature, because this is still considered to be computationally infeasible by most cryptographic hash functions (see collision resistance).

5.7.5 Additional security precautions

5.7.5.1 Using Private/Public Key

All public key/private key cryptosystems depend entirely on keeping the private key secret. A private key can be stored on a user's computer, and protected by a local password, but this has two disadvantages:

- the user can only sign documents on that particular computer
- the security of the private key depends entirely on the security of the computer

A more secure alternative is to store the private key on a smart card. Many smart cards are designed to be tamper-resistant (although some designs have been broken, notably by Ross Anderson and his students). In a typical digital signature implementation, the hash calculated from the document is sent to the smart card, whose CPU encrypts the hash using the stored private key of the user, and then returns the encrypted hash. Typically, a user must activate his smart card by entering a personal identification number or PIN code (thus providing two-factor authentication). It can be arranged that the private key never leaves the smart card, although this is not always implemented. If the smart card is stolen, the thief will still need the PIN code to generate a digital signature. This reduces the security of the scheme to that of the PIN system, although it still requires an attacker to possess the card. A mitigating factor is that private keys, if generated and stored on smart cards, are usually regarded as difficult to copy, and are assumed to exist in exactly one copy.

NOTES

Thus, the loss of the smart card may be detected by the owner and the corresponding certificate can be immediately revoked. Private keys that are protected by software only may be easier to copy, and such compromises are far more difficult to detect.

5.7.5.2 Using smart card readers with a separate keyboard

NOTES

Entering a PIN code to activate the smart card commonly requires a numeric keypad. Some card readers have their own numeric keypad. This is safer than using a card reader integrated into a PC, and then entering the PIN using that computer's keyboard. Readers with a numeric keypad are meant to circumvent the eavesdropping threat where the computer might be running a keystroke logger, potentially compromising the PIN code. Specialized card readers are also less vulnerable to tampering with their software or hardware and are often EAL3 certified.

5.7.5.3 Other smart card designs

Smart card design is an active field, and there are smart card schemes which are intended to avoid these particular problems, though so far with little security proofs.

5.7.6 Using digital signatures only with trusted applications

One of the main differences between a digital signature and a written signature is that the user does not "see" what he signs. The user application presents a hash code to be encrypted by the digital signing algorithm using the private key. An attacker who gains control of the user's PC can possibly replace the user application with a foreign substitute, in effect replacing the user's own communications with those of the attacker. This could allow a malicious application to trick a user into signing any document by displaying the user's original on-screen, but presenting the attacker's own documents to the signing application.

To protect against this scenario, an authentication system can be set up between the user's application (word processor, email client, etc.) and the signing application. The general idea is to provide some means for both the user app and signing app to verify each other's integrity. For example, the signing application may require all requests to come from digitally-signed binaries.

5.7.6.1 WYSIWYS

Technically speaking, a digital signature applies to a string of bits, whereas humans and applications "believe" that they sign the semantic interpretation of those bits. In order to be semantically interpreted the bit string must be transformed into a form that is meaningful for humans and applications, and this is done through a combination of hardware and software based processes on a computer system.

The problem is that the semantic interpretation of bits can change as a function of the processes used to transform the bits into semantic content. It is relatively easy to change the interpretation of a digital document by implementing changes on the computer system where the document is being processed. From a semantic perspective this creates uncertainty about what exactly has been signed. WYSIWYS (What You See Is What You Sign) means that the semantic interpretation of a signed message cannot be changed.

In particular this also means that a message cannot contain hidden info that the signer is unaware of, and that can be revealed after the signature has been applied. WYSIWYS is a desirable property of digital signatures that is difficult to guarantee because of the increasing complexity of modern computer systems.

5.7.6.2 Digital signatures vs. ink on paper signatures

An ink signature can be easily replicated from one document to another by copying the image manually or digitally. Digital signatures cryptographically bind an electronic identity to an electronic document and the digital signature cannot be copied to another document. Paper contracts often have the ink signature block on the last page, and the previous pages may be replaced after a signature is applied. Digital signatures can be applied to an entire document, such that the digital signature on the last page will indicate tampering if any data on any of the pages have been altered.

NOTES

5.8 DIGITAL CERTIFICATES

In cryptography, a public key certificate (also known as a digital certificate or identity certificate) is an electronic document which uses a digital signature to bind together a public key with an identity — information such as the name of a person or an organization, their address, and so forth. The certificate can be used to verify that a public key belongs to an individual. In a typical public key infrastructure (PKI) scheme, the signature will be of a certificate authority (CA). In a web of trust scheme, the signature is of either the user (a self-signed certificate) or other users (“endorsements”). In either case, the signatures on a certificate are attestations by the certificate signer that the identity information and the public key belong together.

For provable security this reliance on something external to the system has the consequence that any public key certification scheme has to rely on some special setup assumption, such as the existence of a certificate authority. Certificates can be created for Unix-based servers with tools such as OpenSSL’s `ssl-ca` or SuSE’s `gensslcert`. These may be used to issue unmanaged certificates, Certification Authority (CA) certificates for managing other certificates, and user and/or computer certificate requests to be signed by the CA, as well as a number of other certificate related functions.

Similarly, Microsoft Windows 2000 Server and Windows Server 2003 contain a Certification Authority (CA) as part of Certificate Services for the creation of digital certificates. In Windows Server 2008 the CA may be installed as part of Active Directory Certificate Services. The CA is used to manage and centrally issue certificates to users and/or computers. Microsoft also provides a number of different certificate utilities, such as `SelfSSL.exe` for creating unmanaged certificates, and `Certreq.exe` for creating and submitting certificate requests to be signed by the CA, and `certutil.exe` for a number of other certificate related functions.

5.8.1 Contents of a typical digital certificate

Serial Number: Used to uniquely identify the certificate.

Subject: The person, or entity identified.

Signature Algorithm: The algorithm used to create the signature.

Issuer: The entity that verified the information and issued the certificate.

Valid-From: The date the certificate is first valid from.

Valid-To: The expiration date.

Key-Usage: Purpose of the public key (e.g. encipherment, signature, certificate signing...).

Public Key: the purpose of SSL when used with HTTP is not just to encrypt the traffic, but also to authenticate who the owner of the website is, and that someone's been willing to invest time and money into proving the authenticity and ownership of their domain.

Thumbprint Algorithm: The algorithm used to hash the certificate.

Thumbprint: The hash itself to ensure that the certificate has not been tampered with.

5.8.2 Optional Certificate Fields

Subject Alternative Name (SAN): Provide multiple identities which the certificate can authenticate. Systems that may operate under multiple identities, such as server farms, and some software platforms, such as Microsoft Exchange, may use SAN certificates to simplify the support of the environment.

5.8.3 Classification

VeriSign introduced the concept of classes of digital certificates[citation needed]:

- Class 1 for individuals, intended for email
- Class 2 for organizations, for which proof of identity is required
- Class 3 for servers and software signing, for which independent verification and checking of identity and authority is done by the issuing certificate authority
- Class 4 for online business transactions between companies
- Class 5 for private organizations or governmental security.

5.8.4 Certificates and web site security

The most common use of certificates is for HTTPS-based web sites. A web browser validates that an SSL (Transport Layer Security) web server is authentic, so that the user can feel secure that their interaction with the web site has no eavesdroppers and that the web site is who it claims to be. This security is important for electronic commerce. In practice, a web site operator obtains a certificate by applying to a certificate provider with a certificate signing request.

The certificate request is an electronic document that contains the web site name, contact email address, and company information. The certificate provider signs the request, thus producing a public certificate. This public certificate is served to any web browser that connects to the web site and proves to the web browser that the provider believes it has issued a certificate to the owner of the web site. Before issuing a certificate, the certificate provider will request the contact email address for the web site from a public domain name registrar, and check that published address against the email address supplied in the certificate request. Therefore, an https web site is only secure to the extent that the end user can be sure that the web site is operated by someone in contact with the person that registered the domain name.

As an example, when a user connects to `https://www.example.com/` with their browser, if the browser gives no certificate warning message, then the user can be theoretically sure that interacting with `https://www.example.com/` is equivalent to

interacting with the entity in contact with the email address listed in the public registrar under "example.com", even though that email address may not be displayed anywhere on the web site. No other surety of any kind is implied. Further, the relationship between the purchaser of the certificate, the operator of the web site, and the generator of the web site content may be tenuous and is not guaranteed. At best, the certificate guarantees uniqueness of the web site, provided that the web site itself has not been compromised (hacked) or the certificate issuing process subverted.

NOTES

5.9 FIREWALL

A firewall's basic task is to transfer traffic between computer networks of different trust levels. Typical examples are the Internet which is a zone with no trust and an internal network which is a zone of higher trust. A zone with an intermediate trust level, situated between the Internet and a trusted internal network, is often referred to as a "perimeter network" or Demilitarized zone (DMZ).

A firewall's function within a network is similar to firewalls with fire door in building construction. In former case, it is used to prevent network intrusion to the private network. In later case, it is intended to contain and delay structural fire from spreading to adjacent structures.

An analogy of network firewall is a fire-resistance rated wall with a fire-resistance rated, self-closing, solid-core, inside unlock able, outside key-lockable door between a house and its attached garage.

Without proper configuration, a firewall can often become worthless. Standard security practices dictate a "default-deny" firewall rule set, in which the only network connections which are allowed are the ones that have been explicitly allowed. Unfortunately, such a configuration requires detailed understanding of the network applications and endpoints required for the organization's day-to-day operation. Many businesses lack such understanding, and therefore, implement a "default-allow" rule set, in which all traffic is allowed unless it has been specifically blocked. This configuration makes inadvertent network connections and system compromise much more likely.

5.9.1 History

Firewall technology emerged in the late 1980s when the Internet was a fairly new technology in terms of its global use and connectivity. The original idea was formed in response to a number of major internet security breaches, which occurred in the late 1980s. In 1988 an employee at the NASA Ames Research Center in California sent a memo by email to his colleagues that read, "We are currently under attack from an Internet VIRUS! It has hit Berkeley, UC San Diego, Lawrence Livermore, Stanford, and NASA Ames." The Morris Worm spread itself through multiple vulnerabilities in the machines of the time. Although it was not malicious in intent, the Morris Worm was the first large scale attack on Internet security; the online community was neither expecting an attack nor prepared to deal with one.

5.9.2 First generation - Packet Filters

The first paper published on firewall technology was in 1988, when engineers from Digital Equipment Corporation (DEC) developed filter systems known as packet

filter firewalls. This fairly basic system was the first generation of what would become a highly evolved and technical internet security feature. At AT&T Bill Cheswick and Steve Bellovin were continuing their research in packet filtering and developed a working model for their own company based upon their original first generation architecture.

NOTES

Packet filters act by inspecting the “packets” which represent the basic unit of data transfer between computers on the Internet. If a packet matches the packet filter’s set of rules, the packet filter will drop (silently discard) the packet, or reject it (discard it, and send “error responses” to the source).

This type of packet filtering pays no attention to whether a packet is part of an existing stream of traffic (it stores no information on connection “state”).

Instead, it filters each packet based only on information contained in the packet itself (most commonly using a combination of the packet’s source and destination address, its protocol, and, for TCP and UDP traffic, which comprises most internet communication, the port number).

Because TCP and UDP traffic by convention uses well known ports for particular types of traffic, a “stateless” packet filter can distinguish between, and thus control, those types of traffic (such as web browsing, remote printing, email transmission, file transfer), unless the machines on each side of the packet filter are both using the same non-standard ports.

5.9.3 Second generation - “stateful” filters

From 1980-1990 three colleagues from AT&T Bell Laboratories, Dave Presetto, Howard Trickey, and Kshitij Nigam developed the second generation of firewalls, calling them circuit level firewalls.

This technology is generally referred to as a ‘stateful firewall’ as it maintains records of all connections passing through the firewall, and is able to determine whether a packet is the start of a new connection, or part of an existing connection.

Though there’s still a set of static rules in such a firewall, the state of a connection can in itself be one of the criteria which trigger specific rules.

This type of firewall can help prevent attacks which exploit existing connections, or certain Denial-of-service attacks, including the SYN flood which sends improper sequences of packets to consume resources on systems behind a firewall.

5.9.4 Third generation - application layer

Publications by Gene Spafford of Purdue University, Bill Cheswick at AT&T Laboratories and Marcus Ranum described a third generation firewall known as application layer firewall, also known as proxy based firewalls. Marcus Ranum’s work on the technology spearheaded the creation of the first commercial product. The product was released by DEC who named it the SEAL product. DEC’s first major sale was on June 13, 1991 to a chemical company based on the East Coast of the USA.

The key benefit of application layer filtering is that it can “understand” certain applications and protocols (such as File Transfer Protocol, DNS or web

browsing), and can detect whether an unwanted protocol is being sneaked through on a non-standard port, or whether a protocol is being abused in a known harmful way.

This type of filtering can be carried out by proxy servers, but if the filtering is done by a standalone firewall appliance, or in a device for traffic shaping, the technology is likely to be referred to as deep packet inspection.

NOTES

5.9.5 Subsequent developments

In 1992, Bob Braden and Annette DeSchon at the University of Southern California (USC) were developing their own fourth generation packet filter firewall system. The product known as "Visas" was the first system to have a visual integration interface with colours and icons, which could be easily implemented to and accessed on a computer operating system such as Microsoft's Windows or Apple's MacOS. In 1994 an Israeli company called Check Point Software Technologies built this into readily available software known as FireWall-1.

A second generation of proxy firewalls was based on Kernel Proxy technology. This design is constantly evolving but its basic features and codes are currently in widespread use in both commercial and domestic computer systems. Cisco, one of the largest internet security companies in the world released their PIX product to the public in 1997.

Some modern firewalls leverage their existing deep packet inspection engine by sharing this functionality with an Intrusion-prevention system (IPS).

Currently, the Middlebox Communication Working Group of the Internet Engineering Task Force (IETF) is working on standardizing protocols for managing firewalls and other middleboxes, a way of transferring policy enforcement.

5.9.6 Types of Firewalls

There are several classifications of firewalls depending on where the communication is taking place, where the communication is intercepted and the state that is being traced.

Application-layer

Application-layer firewalls work on the application level of the TCP/IP stack (i.e., all browser traffic, or all telnet or ftp traffic), and may intercept all packets traveling to or from an application. They block other packets (usually dropping them without acknowledgement to the sender). In principle, application firewalls can prevent all unwanted outside traffic from reaching protected machines.

By inspecting all packets for improper content, firewalls can restrict or prevent outright the spread of networked computer worms and trojans. In practice, however, this becomes so complex and so difficult to attempt (given the variety of applications and the diversity of content each may allow in its packet traffic) that comprehensive firewall design does not generally attempt this approach.

The XML firewall exemplifies a more recent kind of application-layer firewall. It performs mainly three functions i.e; simplest. sees only; a- address b-service protocol. auditing is difficult.

Proxies

A proxy device (running either on dedicated hardware or as software on a general-purpose machine) may act as a firewall by responding to input packets (connection requests, for example) in the manner of an application, whilst blocking other packets.

Proxies make tampering with an internal system from the external network more difficult and misuse of one internal system would not necessarily cause a security breach exploitable from outside the firewall (as long as the application proxy remains intact and properly configured). Conversely, intruders may hijack a publicly-reachable system and use it as a proxy for their own purposes; the proxy then masquerades as that system to other internal machines.

While use of internal address spaces enhances security, crackers may still employ methods such as IP spoofing to attempt to pass packets to a target network.

5.9.7 Network address translation

Firewalls often have network address translation (NAT) functionality, and the hosts protected behind a firewall commonly have addresses in the "private address range", as defined in RFC 1918. Firewalls often have such functionality to hide the true address of protected hosts.

Originally, the NAT function was developed to address the limited amount of IPv4 routable addresses that could be used or assigned to companies or individuals as well as reduce both the amount and therefore, cost of obtaining enough public addresses for every computer in an organization. Hiding the addresses of protected devices has become an increasingly important defense against network reconnaissance.

5.9.8 Firewall (UNIX and NT) Address and Translation

Network layer firewalls, also called packet filters, operate at a relatively low level of the TCP/IP protocol stack, not allowing packets to pass through the firewall unless they match the established ruleset. The firewall administrator may define the rules; or default rules may apply. The term packet filter originated in the context of BSD operating systems.

Network layer firewalls generally fall into two sub-categories, stateful and stateless. Stateful firewalls maintain context about active sessions, and use that "state information" to speed up packet processing. Any existing network connection can be described by several properties, including source and destination IP address, UDP or TCP ports, and the current stage of the connection's lifetime (including session initiation, handshaking, data transfer, or completion connection. If a packet does not match an existing connection, it will be evaluated according to the ruleset for new connections. If a packet matches an existing connection based on comparison with the firewall's state table, it will be allowed to pass without further processing.

Stateless firewalls have packet-filtering capabilities, but cannot make more complex decisions on what stage communications between hosts have reached. Stateless firewalls, therefore, offer less security.

Modern firewalls can filter traffic based on many packet attributes like source IP address, source port, destination IP address or port, destination service like WWW

or FTP. They can filter based on protocols, TTL values, netblock of originator, domain name of the source, and many other attributes.

Commonly used packet filters on various versions of Unix are ipf (various), ipfw (FreeBSD/Mac OS X), pf (OpenBSD, and all other BSDs), iptables/ipchains (Linux).

NOTES

SUMMARY

1. Computer or network security has been violated when unauthorized access by any party occurs.
2. Network administrators like to believe that their network is secure and those who break into networks may like to believe that they can break into any network.
3. If you are reading this document and are thinking that you can get all the information required to make your network completely secure, then you are sadly mistaken.
4. A hacker is no more likely to break the law than anyone else.
5. In a security context, a Hacker is someone who is involved in computer security/insecurity, specializing in the discovery of exploits in systems (for exploitation or prevention), or in obtaining or preventing unauthorized access to systems through skills, tactics and detailed knowledge.
6. Instead of a hacker-cracker dichotomy, they give more emphasis to a spectrum of different categories, such as white hat ("ethical hacking"), grey hat, black hat and script kiddie.
7. Within the academic hacker culture, the term hacker is also used for a programmer who reaches a goal by employing a series of modifications to extend existing code or resources.
8. The hobby hacking subculture relates to hobbyist home computing of the late 1970s, beginning with the availability of MITS Altair.
9. The main basic difference between academic and computer security hackers is their separate historical origin and development.
10. The computer security hacking subculture on the other hand tends not to distinguish between the two subcultures as harshly, instead acknowledging that they have much in common including many members, political and social goals, and a love of learning about technology.
11. Hypertext Transfer Protocol Secure (HTTPS) is a combination of the Hypertext Transfer Protocol with the SSL/TLS protocol to provide encryption and secure (website security testing) identification of the server.
12. A solution called Server Name Indication (SNI) exists which sends the hostname to the server before encrypting the connection, although many older browsers don't support this extension.
13. Encapsulating Security Payload (ESP) is a member of the IPsec protocol suite.
14. Unlike Authentication Header (AH), ESP does not protect the IP packet header.
15. Authentication Header (AH) is a member of the IPsec protocol suite.
16. A digital signature or digital signature scheme is a mathematical scheme for demonstrating the authenticity of a digital message or document.
17. The United States Government Printing Office (GPO) publishes electronic versions of the budget, public and private laws, and congressional bills with digital signatures.
18. All public key/private key cryptosystems depend entirely on keeping the private key secret.
19. In cryptography, a public key certificate (also known as a digital certificate or identity certificate) is an electronic document which uses a digital signature to bind together a

public key with an identity — information such as the name of a person or an organization, their address, and so forth.

20. A web browser validates that an SSL (Transport Layer Security) web server is authentic, so that the user can feel secure that their interaction with the web site has no eavesdroppers and that the web site is who it claims to be.

NOTES

SELF ASSESSMENT QUESTIONS

1. Describe network security.
2. What is the need for network security?
3. Who is an Attacker?
4. What are the Security Limitations and Applications?
5. Who are Hobby Hackers?
6. Describe Secure HTTP.
7. Describe the various network layers.
8. What is SSL?
9. Describe ESP.
10. What do you understand by Authentic Header?
11. Describe the various distribution protocols.
12. What are Digital signatures?
13. What are Digital certificates?

Multiple Choice Questions

1. HTTPS is:
 - (a) **Hypertext Transfer Protocol Secure**
 - (b) Hypertext Transfer Port Secure
 - (c) Hypertext Transfer Protocol Set
2. SSL is:
 - (a) Safe Sockets Layer
 - (b) **Security Sockets Layer**
 - (c) Security Sock Layer
3. TLS is:
 - (a) Transport Lower Security
 - (b) **Transport Level Security**
 - (c) **Transport Layer Security**
4. OCSP is:
 - (a) **Online Certificate Status Protocol**
 - (b) Online Certificate Self Protocol
 - (c) Online Certificate Status Port
5. SSP is:
 - (a) Six-Sense Protocol
 - (b) **Six-State Protocol**
 - (c) Six-State Port
6. SAN is:
 - (a) Subject Alternative Name
 - (b) Subject Alternative Name
 - (c) **Subject Alternative Name**

True/False Questions

1. Network administrators like to believe that their network is secure and those who break into networks may like to believe that they can break into any network.
2. A hacker is no more likely to break the law than anyone else.
3. Instead of a hacker-cracker dichotomy, they give more emphasis to a spectrum of different categories, such as white hat ("ethical hacking"), grey hat, black hat and script kiddie.
4. **The hobby hacking subculture does not relate** to hobbyist home computing of the late 1970s, beginning with the availability of MITS Altair.
5. The computer security hacking subculture on the other hand tends not to distinguish between the two subcultures as harshly, instead acknowledging that they have much in common including many members, political and social goals, and a love of learning about technology.
6. **A solution called Server Name Indication (SNI) does not exist** which sends the hostname to the server before encrypting the connection, although many older browsers don't support this extension.
7. Unlike Authentication Header (AH), ESP does not protect the IP packet header.
8. A digital signature or digital signature scheme is a mathematical scheme for demonstrating the authenticity of a digital message or document.
9. **All public key/private key** cryptosystems do not depend entirely on keeping the private key secret.
10. A web browser validates that an SSL (Transport Layer Security) web server is authentic, so that the user can feel secure that their interaction with the web site has no eavesdroppers and that the web site is who it claims to be.

NOTES

Short Questions with Answers

1. What is computer security needed?

Ans. Computer security is required because most organizations can be damaged by hostile software or intruders. There may be several forms of damage which are obviously interrelated. These include:

- Damage or destruction of computer systems.
- Damage or destruction of internal data.
- Loss of sensitive information to hostile parties.
- Use of sensitive information to steal items of monetary value.
- Use of sensitive information against the organization's customers which may result in legal action by customers against the organization and loss of customers.
- Damage to the reputation of an organization.
- Monetary damage due to loss of sensitive information, destruction of data, hostile use of sensitive data, or damage to the organization's reputation.

2. Describe the various popular terms of computer network security.

Ans. This paragraph describes some commonly used computer security terms.

- **Protocol** - Well defined specification allowing computer communication.
- **Confidentiality** - Information is available only to people with rightful access.
- **Integrity** - Information can only be changed by authorized personnel. Integrity - The receiver of the message should be able to tell that the message was not modified. Requires key exchange.

NOTES

- **Availability** - Information is available to only those who need it.
- **Verification - nonrepudiation** - There is proof that the sender sent the message.
- **Authentication** - The receiver of the message should be able to be sure of the origin of the message. Requires a digital signature (One way hash, public key algorithm, and symmetric algorithm) or a public key algorithm.
- **Spyware** - A computer program whose purpose is to spy on your internet activities usually for marketing purposes and usually done by a shady corporate entity.
- **Malware** - A computer program with some evil intent. It may on the surface have a good or useful intent, but may be a trojan (with a hidden purpose) which can be used to gain unauthorized access to your computer.

3. What is HTTPS?

Ans. Hypertext Transfer Protocol Secure (HTTPS) is a combination of the Hypertext Transfer Protocol with the SSL/TLS protocol to provide encryption and secure (website security testing) identification of the server. It uses port 443. HTTPS connections are often used for payment transactions on the World Wide Web and for sensitive transactions in corporate information systems. HTTPS should not be confused with Secure HTTP (S-HTTP) specified in RFC 2660.

The main idea of HTTPS is to create a secure channel over an insecure network. This ensures reasonable protection from eavesdroppers and man-in-the-middle attacks, provided that adequate cipher suites are used and that the server certificate is verified and trusted.

4. What are the various key points of SSL?

Ans. The points can be defined as:

- i) An SSL connection is managed by the first front machine which initiate the SSL connection. If for any reasons (routing, traffic optimization, etc.) this front machine is not the application server and it has to decipher data, solutions have to be found to propagate user authentication informations or certificate to the application server which needs to know who is going to be connected.
- ii) For SSL with mutual authentication, the SSL session is managed by the first server which initiates the connection. In situations where encryption has to be propagated along chained servers, session timeOut management becomes extremely tricky to implement.
- iii) With mutual SSL, security is maximal, but on the client-side, there is no way to properly end the SSL connection and disconnect the user except by waiting for the SSL server session to expire or closing all related client applications.
- iv) For performance reasons static contents are usually delivered through a non-crypted front server or separate server instance with no SSL, as a consequence these contents are usually not protected.

5. What is ESP?

Ans. Encapsulating Security Payload (ESP) is a member of the IPsec protocol suite. In IPsec it provides origin authenticity, integrity, and confidentiality protection of packets. ESP also supports encryption-only and authentication-only configurations, but using encryption without authentication is strongly discouraged because it is insecure. Unlike Authentication Header (AH), ESP does not protect the IP packet header. However, in Tunnel Mode, where the entire original IP packet is encapsulated with a new packet header added, ESP protection is afforded to the whole inner IP packet (including the inner header) while the outer header remains unprotected. ESP operates directly on top of IP, using IP protocol number 50.

6. What is Authentication Header?

Ans. Authentication Header (AH) is a member of the IPsec protocol suite. AH guarantees connectionless integrity and data origin authentication of IP packets. Further, it can optionally protect against replay attacks by using the sliding window technique and discarding old packets.

7. What are Digital Signatures used for?

Ans. Digital signatures are often used to implement electronic signatures, a broader term that refers to any electronic data that carries the intent of a signature, but not all electronic signatures use digital signatures. In some countries, including the United States, India, and members of the European Union, electronic signatures have legal significance. However, laws concerning electronic signatures do not always make clear whether they are digital cryptographic signatures in the sense used here, leaving the legal definition, and so their importance, somewhat confused.

8. Describe the various uses of Digital Signatures.

Ans. As organizations move away from paper documents with ink signatures or authenticity stamps, digital signatures can provide added assurances of the evidence to provenance, identity, and status of an electronic document as well as acknowledging informed consent and approval by a signatory. The United States Government Printing Office (GPO) publishes electronic versions of the budget, public and private laws, and congressional bills with digital signatures. Universities including Penn State, University of Chicago, and Stanford are publishing electronic student transcripts with digital signatures.

9. What is the difference between digital signature and written signature?

Ans. One of the main differences between a digital signature and a written signature is that the user does not “see” what he signs. The user application presents a hash code to be encrypted by the digital signing algorithm using the private key. An attacker who gains control of the user's PC can possibly replace the user application with a foreign substitute, in effect replacing the user's own communications with those of the attacker. This could allow a malicious application to trick a user into signing any document by displaying the user's original on-screen, but presenting the attacker's own documents to the signing application.

10. What is the difference between ink signature and digital signature?

Ans. An ink signature can be easily replicated from one document to another by copying the image manually or digitally. Digital signatures cryptographically bind an electronic identity to an electronic document and the digital signature cannot be copied to another document. Paper contracts often have the ink signature block on the last page, and the previous pages may be replaced after a signature is applied. Digital signatures can be applied to an entire document, such that the digital signature on the last page will indicate tampering if any data on any of the pages have been altered.

11. What digital certificates?

Ans. In cryptography, a public key certificate (also known as a digital certificate or identity certificate) is an electronic document which uses a digital signature to bind together a public key with an identity — information such as the name of a person or an organization, their address, and so forth. The certificate can be used to verify that a public key belongs to an individual. In a typical public key infrastructure (PKI) scheme, the signature will be of a certificate authority (CA). In a web of trust scheme, the signature is of either the user (a self-signed certificate) or other users (“endorsements”). In either case, the signatures on a certificate are attestations by the certificate signer that the identity information and the public key belong together.

12. What should a digital signature must have:

Ans. Following is the list of items which a digital signature should have:

Serial Number: Used to uniquely identify the certificate.

Subject: The person, or entity identified.

Signature Algorithm: The algorithm used to create the signature.

Issuer: The entity that verified the information and issued the certificate.

Valid-From: The date the certificate is first valid from.

NOTES

Valid-To: The expiration date.

Key-Usage: Purpose of the public key (e.g. encipherment, signature, certificate signing...).

Public Key: the purpose of SSL when used with HTTP is not just to encrypt the traffic, but also to authenticate who the owner of the website is, and that someone's been willing to invest time and money into proving the authenticity and ownership of their domain.

Thumbprint Algorithm: The algorithm used to hash the certificate.

Thumbprint: The hash itself to ensure that the certificate has not been tampered with.

ANSWERS

Multiple Choice Questions

- | | | | |
|------|------|------|------|
| 1. a | 2. b | 3. c | 4. a |
| 5. b | 6. c | | |

True False Questions

- | | | | |
|------|-------|------|------|
| 1. T | 2. T | 3. T | 4. F |
| 5. T | 6. F | 7. T | 8. T |
| 9. F | 10. T | | |

Further Readings

1. **Computer Networks:** Ajit Kumar Singh, Firewall Media.
2. **Data and Computer Network Communication:** Prof. Shashi Banzai, Firewall Media.
3. **TCP / IP and Distributed System:** Vivek Archarya, Firewall Media.
4. **Networking:** Balvir Singh, Firewall Media.