



VSS University of Technology

----- BURLA -----

DEPARTMENT OF COMPUTER APPLICATIONS (MCA)

MCA-301 INTERNET TECHNOLOGY 5th Semester



Veer Surendra Sai University of Technology, Burla

(A UGC & AICTE affiliated Unitary Technical University)

Sambalpur-768018, Odisha

INDIA

www.vssut.ac.in

5th SEMESTER

MCA-301

INTERNET TECHNOLOGY

L-T-P: 3-1-0

Module I (10 hrs)

Internet architecture: Internet overview, evolution of internet. Internet components – Local Area Networks, Access Networks, Core Networks, Routers, Transmission infrastructure, ISPs. Packet switching fundamentals-Packet Switching versus Circuit Switching, Connectionless packet switching (IP). Internet Standards: Standards bodies and the standards process, IETF, ITU, IEEE, ATM Forum.

Module II (10 hrs)

Networking protocols: Network Protocol Overview: What are networking protocols, and what do they do? Key protocol architectures. IP Network Overview: What are the key IP network capabilities? How will these capabilities adapt to future network? IP protocol operation. IP addressing: IP address classes. Why are IP addresses under pressure, and what fixes are in place? TCP Fundamentals: How does TCP shield end users from IP network problems? TCP protocol operation and capabilities. TCP/IP: routing.

Module III (10 hrs)

Access Methods and Internet working: Access Network Architectures: Access network characteristics. Differences between Access Networks, Local Area Networks and Wide Area Networks. Access Technologies: Why there is an upper limit on modem speeds. Voice grade modems, ADSL, Cable Modems, Frame Relay. DNS: Domain Names. Resolving Domain Names to IP addresses (DNS operation). Registering Domain Names and solving Domain name disputes. Routing: How the key IP routing protocols (OSPF and BGP4) operate. Implications of future Internet growth on routing protocol performance.

Module IV (10 hrs)

Internet applications: FTP, Telnet, Email, Chat. World Wide Web: HTTP protocol. Search Engines. E-Commerce and security issues including symmetric and asymmetric key, encryption and digital signature, and authentication. Emerging trends, Internet telephony, virtual reality over the web, etc. Intranet and extranet, firewall.

Books:

1. Data & Computer Communications, By William Stallings
2. Computer Networks, A system approach By Larry L.Peterson, Bruce S. Davie
3. Internetworking with TCP / IP, Principles, Protocols & Architecture, By Douglas E.Comer.
4. TCP / IP – clearly Explained – by Pete Loshin, Morgan Kaufmann Publishers.
5. TCP / IP Network Administration by Craig Hunt, Shroff Publishers & Distributors Pvt.Ltd.
6. The Internet and its protocols – A Comparative Approach, by A.Farrel I Elseviers, (Morgan Kaufmann Publishers).

Course Outcomes:

1. To be able to understand the technologies and protocols used on the Internet,
2. To be able to understand how effectively Internet tools technologies including current web-based applications, e-mail, and social networking tools can be used.
3. To be able to understand the basics of web search strategies.
4. To be able to understand the basics of web authoring.

DISCLAIMER

This document does not claim any originality and cannot be used as a substitute for prescribed textbooks. The information presented here is merely a collection of knowledge base by the committee members for their respective teaching assignments. Various online/offline sources as mentioned at the end of the document as well as freely available material from internet were helpful for preparing this document. The ownership of the information lies with the respective authors/institution/publisher. Further, this study material is not intended to be used for commercial purpose and the committee members make no representations or warranties with respect to the accuracy or completeness of the information contents of this document and specially disclaim any implied warranties of merchantability or fitness for a particular purpose. The committee members shall not be liable for any loss or profit or any other commercial damages, including but not limited to special, incidental, consequential, or other damages.

INTRNET TECHNOLOGY

MODULE - 1

Internet Architecture:

Internet overview:

The Internet system consists of a number of interconnected packet networks supporting communication among host computers using the Internet protocols. These protocols include the Internet Protocol (IP), the Internet Control Message Protocol (ICMP), the Internet Group Management Protocol (IGMP), and a variety of transport and application protocols that depend upon them.

All Internet protocols use IP as the basic data transport mechanism. IP is a datagram, or connectionless, internetwork service and includes provision for addressing, type-of-service specification, fragmentation and reassembly, and security. ICMP and IGMP are considered integral parts of IP, although they are architecturally layered upon IP. ICMP provides error reporting, flow control, first-hop router redirection, and other maintenance and control functions. IGMP provides the mechanisms by which hosts and routers can join and leave IP multicast groups.

Reliable data delivery is provided in the Internet protocol suite by Transport Layer protocols such as the Transmission Control Protocol (TCP), which provides end-end retransmission, re-sequencing and connection control. Transport Layer connectionless service is provided by the User Datagram Protocol (UDP).

Elements of the Architecture:

1. Protocol Layering:

To communicate using the Internet system, a host must implement the layered set of protocols comprising the Internet protocol suite. A host typically must implement at least one protocol from each layer.

Application Layer

The Application Layer is the top layer of the Internet protocol suite. The Internet suite does not further subdivide the Application Layer, although some application layer protocols do contain some internal sub-layering. The application layer of the Internet suite essentially combines the functions of the top two layers - Presentation and Application - of the OSI Reference Model.

Transport Layer

The Transport Layer provides end-to-end communication services. This layer is roughly equivalent to the Transport Layer in the OSI Reference Model, except that it also incorporates some of OSI's Session Layer establishment and destruction functions.

There are two primary Transport Layer protocols at present:

- Transmission Control Protocol (TCP)

· User Datagram Protocol (UDP)

TCP is a reliable connection-oriented transport service that provides end-to-end reliability, resequencing, and flow control. UDP is a connectionless (datagram) transport service.

Internet Layer

All Internet transport protocols use the Internet Protocol (IP) to carry data from source host to destination host. IP is a connectionless or datagram internetwork service, providing no end-to-end delivery guarantees. IP datagrams may arrive at the destination host damaged, duplicated, out of order, or not at all.

Link Layer

To communicate on a directly connected network, a host must implement the communication protocol used to interface to that network. We call this a Link Layer protocol.

Some older Internet documents refer to this layer as the Network Layer, but it is not the same as the Network Layer in the OSI Reference Model

2. Networks

The constituent networks of the Internet system are required to provide only packet (connectionless) transport. According to the IP service specification, datagrams can be delivered out of order, be lost or duplicated, and/or contain errors.

Internet Components.

Local Area Networks:

A datacom system allowing a number of independent devices to communicate directly with each other, within a moderately sized geographic area over a physical communications channel of moderate data rates.

Characteristics of LAN:

- The stations on the network are peers—any station can initiate data exchange with any other station.
- Full connectivity among all stations.
- Fully administered by the owner.
- Runs over a shared transmission medium—often, cabling.
- The network is confined to a small area—a single building or a cluster of buildings.
- The data rate is high—several Mbps (million bits per second)

Early LAN cabling had generally been based on various grades of [coaxial cable](#). Shielded [twisted pair](#) was used in IBM's [Token Ring](#) LAN implementation, but in 1984, [StarLAN](#) showed the potential of simple unshielded [twisted pair](#) by using [Cat3](#) cable—the same simple cable used for telephone systems. This led to the development of [10Base-T](#) (and [its successors](#)) and [structured cabling](#) which is still the basis of most commercial LANs today.

[Fiber-optic](#) cabling is common for links between switches, but [fiber to the desktop](#) is uncommon.

Access Network:

An access network is the part of a [telecommunications network](#) which connects subscribers to their immediate [service provider](#). It is contrasted with the [core network](#), (for example the [Network Switching Subsystem](#) in [GSM](#)) which connects local providers to each other. The access network may be further divided between feeder plant or distribution network, and drop plant or edge network.

- **Access process:**

The process of communicating with a network begins with an access attempt, in which one or more users interact with a [communications system](#) to enable initiation of [user information transfer](#). An access attempt itself begins with an issuance of an access request by an access originator.

An access attempt ends either in successful access or in access failure - an unsuccessful access that results in termination of the attempt in any manner other than initiation of [user information transfer](#) between the intended source and destination ([sink](#)) within the specified maximum access time.

Core Networks:

A core network, or network core, is the central part of a [telecommunications network](#) that provides various [services](#) to customers who are connected by the [access network](#). One of the main functions is to [route telephone calls](#) across the [PSTN](#).

Typically the term refers to the high capacity communication facilities that connect primary nodes. Core/[backbone network](#) provides paths for the exchange of information between different [sub-networks](#). For [enterprise private networks](#) serving one organization, the term backbone is more used, while for [service providers](#), the term core network is more used.

In the United States, [local exchange](#) core networks are linked by several competing [interexchange networks](#); in the rest of the world, the core network has been extended to national boundaries.

Core/backbone network usually has a [mesh topology](#) that provides any-to-any connections among devices on the network. Many main service providers would have their own core/backbone networks, that are interconnected. Some large enterprises have their own core/backbone network, which are typically connected to the public networks.

The devices and facilities in the core / backbone networks are switches and routers. The trend is to push the intelligence and decision making into access and [edge devices](#) and keep the core devices dumb and fast. As a result, switches are more and more often used in the core/backbone network facilities. Technologies used in the core and backbone facilities are [data link layer](#) and [network layer](#) technologies such

as [SONET](#), [DWDM](#), ATM, IP, etc. For enterprise backbone network, [Gigabit Ethernet](#) or [10 Gigabit Ethernet](#) technologies are also often used.

Routers:

In the Internet model, constituent networks are connected together by IP datagram forwarders which are called routers or IP routers. In this document, every use of the term router is equivalent to IP router.

Many older Internet documents refer to routers as gateways. Historically, routers have been realized with packet-switching software executing on a general-purpose CPU. However, as custom hardware development becomes cheaper and as higher throughput is required, special purpose hardware is becoming increasingly common. This specification applies to routers regardless of how they are implemented.

A router connects to two or more logical interfaces, represented by IP subnets or unnumbered point to point lines (discussed in section [2.2.7]). Thus, it has at least one physical interface. Forwarding an IP datagram generally requires the router to choose the address and relevant interface of the next-hop router or (for the final hop) the destination host. This choice, called relaying or forwarding depends upon a route database within the router. The route database is also called a routing table or forwarding table.

The term "router" derives from the process of building this route database; routing protocols and configuration interact in a process called routing. The routing database should be maintained dynamically to reflect the current topology of the Internet system. A router normally accomplishes this by participating in distributed routing and reachability algorithms with other routers. Routers provide datagram transport only, and they seek to minimize the state information necessary to sustain this service in the interest of routing flexibility and robustness.

Packet switching devices may also operate at the Link Layer; such devices are usually called bridges. Network segments that are connected by bridges share the same IP network prefix forming a single IP subnet. These other devices are outside the scope of this document.

ISPs:

An **Internet service provider (ISP)** is an organization that provides services for accessing, using, or participating in the [Internet](#). Internet service providers may be organized in various forms, such as commercial, [community-owned](#), [non-profit](#), or otherwise [privately owned](#).

Internet services typically provided by ISPs include [Internet access](#), [Internet transit](#), [domain name](#) registration, web hosting, [colocation](#).

ISPs provide [Internet access](#), employing a range of technologies to connect users to their network.[22] Available technologies have ranged from computer modems

with [acoustic couplers](#) to telephone lines, to television cable (CATV), wireless Ethernet (wi-fi), and fiber optics.

For users and [small businesses](#), traditional options include copper wires to provide [dial-up](#), DSL (typically [asymmetric digital subscriber line](#), ADSL), [cable modem](#) or [Integrated Services Digital Network](#) (ISDN) (typically [basic rate interface](#)). Using [fiber-optics](#) to end users is called [Fiber To The Home](#) or similar names.

For customers with more demanding requirements (such as medium-to-large businesses, or other ISPs) can use higher-speed DSL (such as [single-pair high-speed digital subscriber line](#)), [Ethernet](#), [metropolitan Ethernet](#), [gigabit Ethernet](#), [Frame Relay](#), ISDN [Primary Rate Interface](#), [ATM \(Asynchronous Transfer Mode\)](#) and [synchronous optical networking](#) (SONET).

[Wireless access](#) is another option, including [satellite Internet access](#).

Packet Switching:

Refers to [protocols](#) in which [messages](#) are divided into [packets](#) before they are sent. Each packet is then transmitted individually and can even follow different routes to its destination. Once all the packets forming a message arrive at the destination, they are recompiled into the original message.

Most modern [Wide Area Network \(WAN\)](#) protocols, including [TCP/IP](#), [X.25](#), and [Frame Relay](#), are based on packet-switching technologies. In contrast, normal telephone service is based on a [circuit-switching](#) technology, in which a dedicated line is allocated for transmission between two parties. Circuit-switching is ideal when data must be transmitted quickly and must arrive in the same order in which it's sent. This is the case with most [real-time](#) data, such as live audio and [video](#). Packet switching is more efficient and robust for data that can withstand some delays in transmission, such as [e-mail](#) messages and [Web pages](#). A new technology, [ATM](#), attempts to combine the best of both worlds -- the guaranteed delivery of circuit-switched networks and the robustness and efficiency of packet-switching networks.

Packet switching VS Circuit switching:

- **Packet Switching:**
 - In packet-based networks, the message gets broken into small data packets.
 - These packets are sent out from the computer and they travel around the network seeking out the most efficient route to travel as circuits become available.
 - This does not necessarily mean that they seek out the shortest route.
 - Each packet may go a different route from the others.
 - Each packet is sent with a 'header address' which tells it where its final destination is, so it knows where to go.

- The header address also describes the sequence for reassembly at the destination computer so that the packets are put back into the correct order.
- One packet also contains details of how many packets should be arriving so that the recipient computer knows if one packet has failed to turn up.
- If a packet fails to arrive, the recipient computer sends a message back to the computer which originally sent the data, asking for the missing packet to be resent.

Advantages

- » Security
- » Bandwidth used to full potential
- » Devices of different speeds can communicate
- » Not affected by line failure (redirects signal)
- » Availability - no waiting for a direct connection to become available
- » During a crisis or disaster, when the public telephone network might stop working, e-mails and texts can still be sent via packet switching

Disadvantages

- » Under heavy use there can be a delay
- » Data packets can get lost or become corrupted
- » Protocols are needed for a reliable transfer
- » Not so good for some types data streams (e.g. real-time video streams can lose frames due to the way packets arrive out of sequence)

• **Circuit switching**

- Circuit switching was designed in 1878 in order to send telephone calls down a dedicated channel.
- This channel remains open and in use throughout the whole call and cannot be used by any other data or phone calls.
- There are three phases in circuit switching:
 - Establish
 - Transfer
 - Disconnect
- The telephone message is sent all together; it is not broken up.
- The message arrives in the same order that it was originally sent.
- In modern circuit-switched networks, electronic signals pass through several switches before a connection is established.
- During a call no other network traffic can use those switches.
- The resources remain dedicated to the circuit during the entire data transfer and the entire message follows the same path.

- Circuit switching can be analog or digital.
- With the expanded use of the Internet for voice and video, analysts predict a gradual shift away from circuit-switched networks.
- A circuit-switched network is excellent for data that needs a constant link from end-to-end, for example, real-time video.

Advantages

- Circuit is dedicated to the call – no interference, no sharing
- Guaranteed the full bandwidth for the duration of the call
- Guaranteed quality of service

Disadvantages

- Inefficient – the equipment may be unused for a lot of the call; if no data is being sent, the dedicated line still remains open.
- It takes a relatively long time to set up the circuit.
- During a crisis or disaster, the network may become unstable or unavailable.
- It was primarily developed for voice traffic rather than data traffic.

Connectionless Packet Switching / Datagram Switching :

Transmission of packets are made on Per-Packet basis.

- Each packet routes individually.
- Packets may route in different paths and sometimes may be out of order.
- Each packet contains the complete addressing of their routed paths sometimes of sequence number of the packet.
- The packets may arrive at the destination machine in an order different from the transmission order.
- Routers in Connectionless switching network maintain a simple but long routing table that contains two columns - Destination Address and Output Port.
- Example : Ethernet , IP , UDP.

Internet standards.

IETF:

The Internet Engineering Task Force (IETF) develops and promotes , in particular the standards that comprise the [Internet protocol suite](#) (TCP/IP). It is an open [standards organization](#), with no formal membership or membership

requirements. All participants and managers are volunteers, though their work is usually funded by their employers or sponsors.

The IETF started out as an activity supported by the US federal government, but since 1993 it has operated as a standards development function under the auspices of the [Internet Society](#), an international membership-based non-profit organization.

Internet Standards R Us most Internet related standards were developed or are maintained by the IETF not including physical network or page display standards does not exist (in a legal sense), no members, no voting. The IETF is “an organized activity of the Internet Society” 1K to 1.5K people at 3/year meetings many many more on mail lists.

IETF standards: not standards “because we say so” they are standards only if people use them formal SDOs can create legally mandated standards no formal recognition for IETF standards by governments or “approved” standards organization but some government standards refer to IETF standards lack of formal government input “a problem” at least to some governments no submitting to “traditional” bodies.

ITU:

ITU (International Telecommunication Union) is the United Nations specialized agency for information and communication technologies.

We allocate global [radio](#) spectrum and satellite orbits, develop the technical standards that ensure networks and technologies seamlessly interconnect, and strive to improve access to ICTs to underserved communities worldwide.

ITU is committed to connecting all the world's people – wherever they live and whatever their means. Through our [work](#), we protect and support everyone's fundamental right to communicate.

[Today](#), ICTs underpin everything we do. They help manage and control emergency [services](#), water supplies, power networks and food distribution chains. They support health care, education, government services, financial markets, transportation systems and environmental management. And they allow people to communicate with colleagues, friends and family anytime, and almost anywhere.

With the help of our membership, ITU brings the benefits of modern communication technologies to people everywhere in an efficient, safe, easy and affordable manner.

ITU membership reads like a Who's Who of the ICT sector. We're unique among UN agencies in having both public and private sector membership. So in addition to our 193 Member [States](#), ITU membership includes ICT regulators, leading academic institutions and some 700 private companies.

An organization based on public-private partnership since its inception, ITU currently has a membership of 193 countries and over 700 private-sector entities and academic institutions. ITU is headquartered in Geneva, Switzerland, and has twelve regional and area offices around the world.

ITU membership represents a cross-section of the global ICT sector, from the world's largest manufacturers and carriers to small, innovative players working with new and emerging technologies, along with leading R&D institutions and academia.

Founded on the principle of international cooperation between governments (Member [States](#)) and the private sector (Sector Members, Associates and Academia), ITU is the premier global forum through which parties [work](#) towards consensus on a wide range of issues affecting the future direction of the ICT industry.

IEEE:

The Institute of Electrical and Electronics Engineers Standards Association (IEEE-SA) is an organization within [IEEE](#) that develops global [standards](#) in a broad range of industries, including: [power](#) and [energy](#), [biomedical](#) and [health care](#), [information technology](#), [telecommunication](#), [transportation](#), [nanotechnology](#), [information assurance](#), and many more.

IEEE-SA has developed standards for over a century, through a program that offers balance, openness, [fair procedures](#), and [consensus](#). Technical experts from all over the world participate in the development of [IEEE](#) standards.

IEEE-SA is not a body formally authorized by any government, but rather a community. Formally recognized international standards organizations (ISO, IEC, ITU, CEN) are federations of national standards bodies (American ANSI, German DIN, Japanese JISC, etc.).

The IEEE standards development process can be broken down into seven basic steps:

1. Securing Sponsorship: An IEEE-approved organization must sponsor a standard. A sponsoring organization is in charge of coordinating and supervising the standard development from inception to completion. The professional societies within IEEE serve as the natural sponsor for many standards.
2. Requesting Project Authorization: To gain authorization for the standard a Project Authorization Request (PAR) is submitted to the IEEE-SA Standards Board. The New Standards Committee (NesCom) of the IEEE-SA Standards Board reviews the PAR and makes a recommendation to the Standards Board about whether to approve the PAR.
3. Assembling a Working Group: After the PAR is approved, a [working group](#) of individuals affected by, or interested in, the standard is organized to develop the standard. IEEE-SA rules ensure that all Working Group meetings are open and that anyone has the right to attend and contribute to the meetings.
4. Drafting the Standard: The Working Group prepares a draft of the proposed standard. Generally, the draft follows the IEEE Standards Style Manual that sets guidelines for the clauses and format of the standards document.

5. **Balloting:** Once a draft of the standard is finalized in the Working Group, the draft is submitted for Balloting approval. The IEEE Standards Department sends an invitation-to-ballot to any individual who has expressed an interest in the subject matter of the standard. Anyone who responds positively to the invitation-to-ballot becomes a member of the balloting group, as long as the individual is an IEEE Standards Association member or has paid a balloting fee. The IEEE requires that a proposed draft of the standard receive a response rate of 75% (i.e., at least 75% of potential ballots are returned) and that, of the responding ballots, at least 75% approve the proposed draft of the standard. If the standard is not approved, the process returns to the drafting of the standard step in order to modify the standard document to gain approval of the balloting group.
6. **Review Committee:** After getting 75% approval, the draft standard, along with the balloting comments, are submitted to the IEEE-SA Standards Board Review Committee (RevCom). The RevCom reviews the proposed draft of the standard against the IEEE-SA Standards Board Bylaws and the stipulations set forth in the IEEE-SA Standards Board Operations Manual. The RevCom then makes a recommendation about whether to approve the submitted draft of the standard document.
7. **Final Vote:** Each member of the IEEE-SA Standards Board places a final vote on the submitted standard document. In some cases external members are invited to vote. It takes a majority vote of the Standards Board to gain final approval of the standard. In general, if the RevCom recommends approval, the Standards Board will vote to approve the standard.

Notable IEEE Standards committees and formats

IEEE 802	LAN/MAN
IEEE 802.1	Standards for LAN/MAN bridging and management and remote media access control (MAC) bridging.
IEEE 802.2	Standards for Logical Link Control (LLC) standards for connectivity.
IEEE 802.3	Ethernet Standards for Carrier Sense Multiple Access with Collision Detection (CSMA/CD).
IEEE 802.4	Standards for token passing bus access.
IEEE 802.4	Standards for Logical Link Control (LLC) standards for connectivity.
IEEE 802.5	Standards for token ring access and for communications between LANs and MANs
IEEE 802.6	Standards for information exchange between systems.
IEEE 802.7	Standards for broadband LAN cabling.
IEEE 802.8	Fiber optic connection.
IEEE 802.9	Standards for integrated services, like voice and data.
IEEE 802.10	Standards for LAN/MAN security implementations.
IEEE 802.11	Wireless Networking - " WiFi ".
IEEE 802.12	Standards for demand priority access method.
IEEE 802.14	Standards for cable television broadband communications.

ATM Forum:

The ATM Forum was founded in 1991 to be the industry consortium to promote [Asynchronous Transfer Mode](#) technology used in telecommunication networks; the founding president and chairman was [Fred Sammartino](#) of [Sun Microsystems](#). It was a non-profit international organization. The ATM Forum created over 200 implementation agreements.

Terminologies

Interworking: The term interworking is used to express interactions between networks, between end systems, or between parts thereof, with the aim of providing a functional entity capable of supporting an end-to-end communication. The interactions required to provide a functional entity rely on functions and on the means to select these functions. Error: Reference source not found.

Interworking Function (IWF): An IWF includes the conversion between protocols and the mapping of one protocol to another. The functionality required between networks can be separated from the functionality, if any, required in end systems. The former functionality is considered to reside in an internetworking network element (INE). Additional details may be found in Error: Reference source not found.

Interworking Network Element (INE): The INE is an entity where user plane, control plane and management plane interworking functions (IWFs) may be implemented. The INE could be a standalone network element, part of the ATM switch or part of an LSR located at the entrance to the MPLS network (LER).

Network interworking: In network interworking, the PCI (Protocol Control Information) of the protocol used in two similar networks and the payload information are transferred, transparently, across an intermediate network by a pair of IWFs.

Bundle: A set of one or more transport LSPs in each direction that provide the appearance of a virtual ATM interface to the ATM control protocols.

Downstream INE: The INE receiving MPLS frames on an LSP.

Upstream INE: The INE sending MPLS frames on an LSP.

Congestion Control Loop

- ABR congestion control is based on the idea that each sender has a current rate, the **Actual Cell Rate (ACR)**, that falls between MCR and PCR. When congestion occurs, ACR is reduced (but not below MCR). When congestion is absent, ACR is increased (but not above PCR).
- After every k data cells, the sender inserts a special **Resource Management (RM)** cell containing the rate at which the sender currently would like to send. This value is called **Explicit Rate (ER)** and must fulfill the condition $MCR \leq ER \leq PCR$.
- As the RM cell passes through the various switches on the way to the receiver (travelling along the same path as the data cells), each switch experiencing congestion can lower the ER value, but not below the minimum cell rate (MCR).
- At the receiver side, the RM cell is looped back to the sender. When the sender gets the RM cell back, it can then see what the minimum acceptable rate is

according to all the switches along the path and can then adjust its Actual Cell Rate (ACR) , to bring it into line with what the slowest switch can handle.

- The rate-based congestion control algorithm is quite robust and fully reliable, Since even a lost RM cell will be noticed by the sender when it fails to return within the expected time interval. The sender then can take action e.g. by lowering its ACR.

The Payload Type (PT) Field

- The payload type field in the ATM cell header consists of three bits. Besides designating various payload types these bits have found a number of additional uses that will be explained in detail in the following slide.

Setup of Switched Virtual Connections based on ATM End-System Addresses

- In order to automatically set up a Switched Virtual Connection (SVC), each ATM end system must be uniquely identified by an ATM address. In a similar way as in an ISDN Q.931 call setup used by the Public Switched Telephone System (PSTN), the addresses of both the calling and called parties are contained in the B-ISDN Q.2931 connection setup messages.

Public ATM Addresses

- These are defined by the ITU-T and are used in public ATM networks

Private ATM Addresses

- These are defined by the ATM Forum and are used in private ATM networks

ATM Public Network Addresses

- ATM Public Network Addresses have the same format as normal public telephone numbers that comply with the ITU-T E.164 International Public Telecommunication Numbering Plan

Private ATM Addresses

- ATM End System Addresses (AESA) are defined by the ATM Forum to have a constant length of 20 octets. They come in many types and they differ mostly in what authority assigns them. Four variants appear in ATM Forum specifications.

AFI - Authority and Format Indicator

- The first octet in an AESA determines the type and the format of the address.

DCC - Data Country Code (AFI = 0x39)

- In this case the IDI is a unique Data Country Code (DCC); these identify particular countries, as specified in ISO 3166. Each country is free to decide the structure and the rules for assignment of the **Domain Specific Part (DSP)**.

ICD - International Code Designator (AFI = 0x47)

- In this case, the IDI is an International Code Designator (ICD); these are allocated by the ISO 6523 registration authorities (the British Standards Institute). ICD codes identify particular international organizations.

HO-DSP - High Order Domain Specific Part

- Used to support flexible, multi-level addressing hierarchies for prefix-based routing protocols (similar to hierarchical IP subnet addresses).

ESI - End System Identifier

- The ESI field is specified to be a 48-bit MAC address, as administered by the IEEE. This facilitates the support of LAN equipment, which is typically hardwired with such addresses.

SEL - Selector

- Used for local multiplexing within end-stations. Has no network significance.

MODULE - 2

NETWORK PROTOCOLS

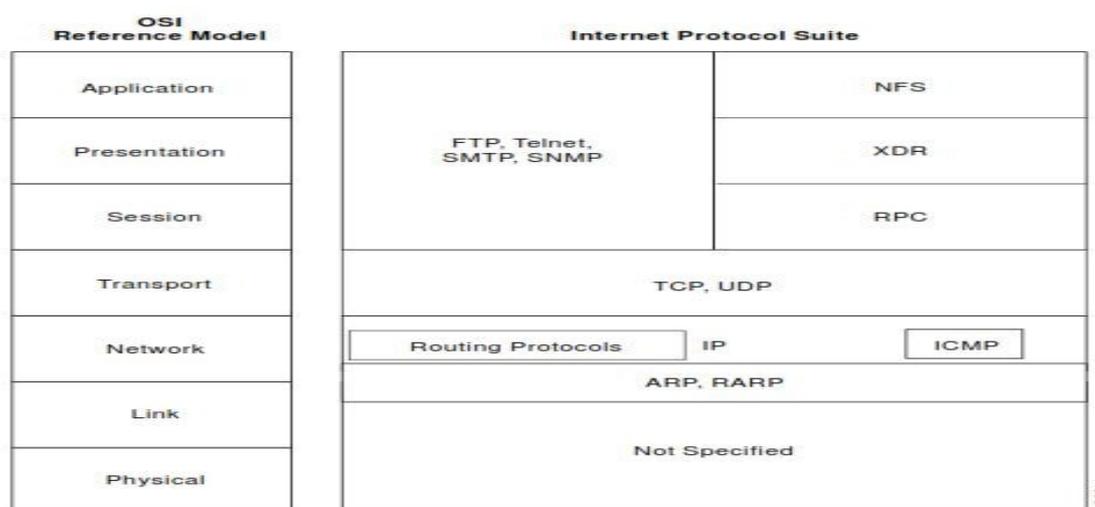
The network protocols are the world's most popular open-system (non-proprietary) protocol suite because they can be used to communicate across any set of interconnected networks and are equally well suited for LAN and WAN communications. The Internet protocols consist of a suite of communication protocols, of which the two best known are the Transmission Control Protocol (TCP) and the Internet Protocol (IP). The Internet protocol suite not only includes lower-layer protocols (such as TCP and IP), but it also specifies common applications such as electronic mail, terminal emulation, and file transfer. This chapter provides a broad introduction to specifications that comprise the Internet protocols. Discussions include IP addressing and key upper-layer protocols used in the Internet. Specific routing protocols are addressed individually in Part 6, Routing Protocols.

Internet protocols were first developed in the mid-1970s, when the Defence Advanced Research Projects Agency (DARPA) became interested in establishing a packet-switched network that would facilitate communication between dissimilar computer systems at research institutions. With the goal of heterogeneous connectivity in mind, DARPA funded research by Stanford University and Bolt, Beranek, and Newman (BBN). The result of this development effort was the Internet protocol suite, completed in the late 1970s.

TCP/IP later was included with Berkeley Software Distribution (BSD) UNIX and has since become the foundation on which the Internet and the World Wide Web (WWW) are based.

Documentation of the Internet protocols (including new or revised protocols) and policies are specified in technical reports called Request for Comments (RFCs), which are published and then reviewed and analysed by the Internet community. Protocol refinements are published in the new RFCs. To illustrate the scope of the Internet protocols, Figure 30-1 maps many of the protocols of the Internet protocol suite and their corresponding OSI layers. This chapter addresses the basic elements and operations of these and other key Internet protocols.

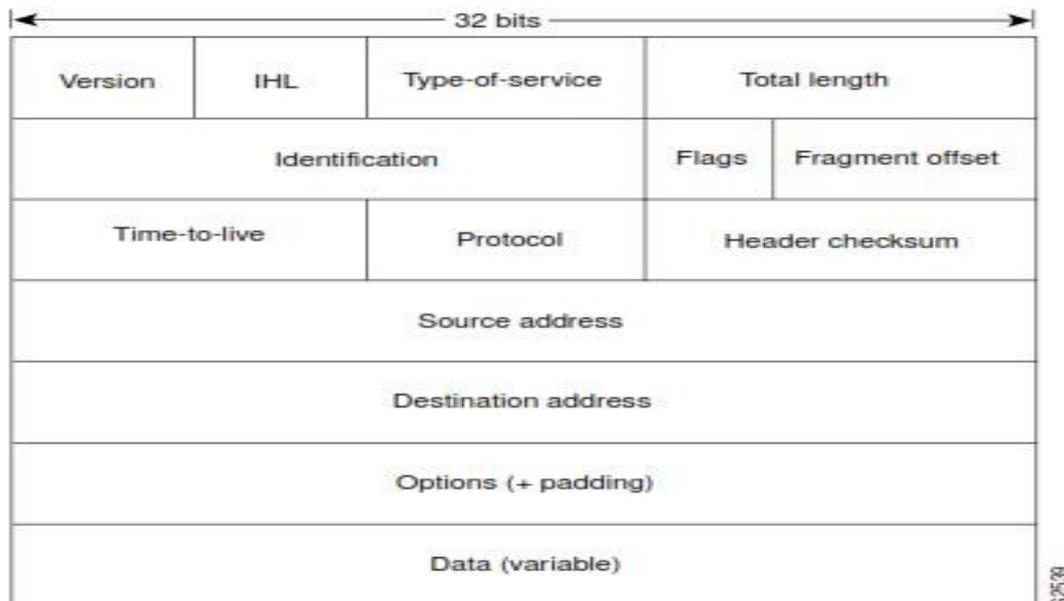
Internet protocols span the complete range of OSI model layers.



Internet Protocol (IP)

The Internet Protocol (IP) is a network-layer (Layer 3) protocol that contains addressing information and some control information that enables packets to be routed. IP is documented in RFC 791 and is the primary network-layer protocol in the Internet protocol suite. Along with the Transmission Control Protocol (TCP), IP represents the heart of the Internet protocols. IP has two primary responsibilities: providing connectionless, best-effort delivery of datagrams through an internetwork; and providing fragmentation and reassembly of datagrams to support data links with different maximum-transmission unit (MTU) sizes.

IP Packet Format



The following discussion describes the IP packet fields illustrated in above Figure.

- Version – Indicates the version of IP currently used.
- IP Header Length (IHL) – Indicates the datagram header length in 32-bit words.
- Type-of-Service – Specifies how an upper-layer protocol would like a current datagram to be handled, and assigns datagrams various levels of importance.
- Total Length – Specifies the length, in bytes, of the entire IP packet, including the data and header.
- Identification – Contains an integer that identifies the current datagram. This field is used to help piece together datagram fragments.
- Flags – Consists of a 3-bit field of which the two low-order (least-significant) bits control fragmentation. The low-order bit specifies whether the packet can be fragmented. The middle bit specifies whether the packet is the last fragment in a series of fragmented packets. The third or high-order bit is not used.
- Fragment Offset – Indicates the position of the fragment’s data relative to the beginning of the data in the original datagram, which allows the destination IP process to properly reconstruct the original datagram.
- Time-to-Live – Maintains a counter that gradually decrements down to zero, at which point the datagram is discarded. This keeps packets from looping endlessly.
- Protocol – Indicates which upper-layer protocol receives incoming packets after IP processing is complete.
- Header Checksum – Helps ensure IP header integrity.
- Source Address – Specifies the sending node.
- Destination Address – Specifies the receiving node.
- Options – Allows IP to support various options, such as security.
- Data – Contains upper-layer information.

IP Addressing

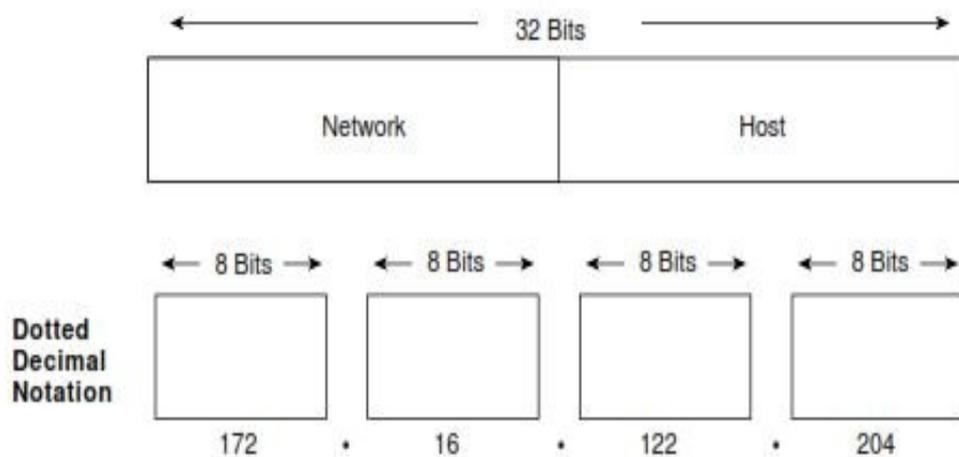
As with any other network-layer protocol, the IP addressing scheme is integral to the process of routing IP datagrams through an internetwork. Each IP address has specific components and follows a basic format. These IP addresses can be

subdivided and used to create addresses for subnetworks, as discussed in more detail later in this chapter. Each host on a TCP/IP network is assigned a unique 32-bit logical address that is divided into two main parts: the network number and the host number. The network number identifies a network and must be assigned by the Internet Network Information Center (InterNIC) if the network is to be part of the Internet. An Internet Service Provider (ISP) can obtain blocks of network addresses from the InterNIC and can itself assign address space as necessary. The host number identifies a host on a network and is assigned by the local network administrator.

IP Address Format

The 32-bit IP address is grouped eight bits at a time, separated by dots, and represented in decimal format (known as dotted decimal notation). Each bit in the octet has a binary weight (128, 64, 32, 16, 8, 4, 2, 1). The minimum value for an octet is 0, and the maximum value for an octet is 255.

Figure 30-3 illustrates the basic format of an IP address.



IP Address Classes

IP addressing supports five different address classes: A, B, C, D, and E. Only classes A, B, and C are available for commercial use. The left-most (high-order) bits indicate the network class. provides reference information about the five IP address classes.

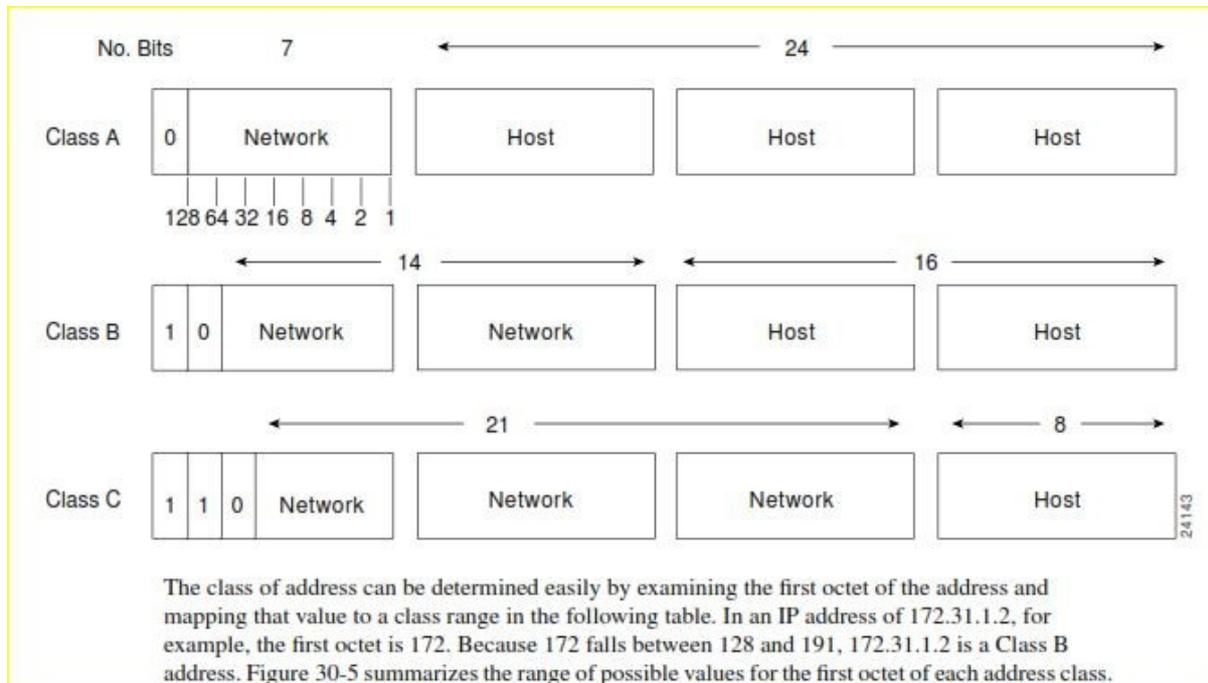
Reference Information About the Five IP Address Classes

IP Address Class	Format	Purpose	High-Order Bit(s)	Address Range	No. Bits Network/Host	Max. Hosts
A	N.H.H.H ¹	Few large organizations	0	1.0.0.0 to 126.0.0.0	7/24	16,777,214 ² ($2^{24} - 2$)
B	N.N.H.H	Medium-size organizations	1, 0	128.1.0.0 to 191.254.0.0	14/16	65,543 ($2^{16} - 2$)
C	N.N.N.H	Relatively small organizations	1, 1, 0	192.0.1.0 to 223.255.254.0	22/8	245 ($2^8 - 2$)
D	N/A	Multicast groups (RFC 1112)	1, 1, 1, 0	224.0.0.0 to 239.255.255.255	N/A (not for commercial use)	N/A
E	N/A	Experimental	1, 1, 1, 1	240.0.0.0 to 254.255.255.255	N/A	N/A

1 N = Network number, H = Host number.

2 One address is reserved for the broadcast address, and one address is reserved for the network.

Reference Information About the Five IP Address Classes



A range of possible values exists for the first octet of each address class.

Address Class	First Octet in Decimal	High-Order Bits
Class A	1 - 126	0
Class B	128 - 191	10
Class C	192 - 223	110
Class D	224 - 239	1110
Class E	240 - 254	1111

IP Subnet Addressing

IP networks can be divided into smaller networks called subnetworks (or subnets). Subnetting provides the network administrator with several benefits, including extra flexibility, more efficient use of network addresses, and the capability to contain broadcast traffic (a broadcast will not cross a router).

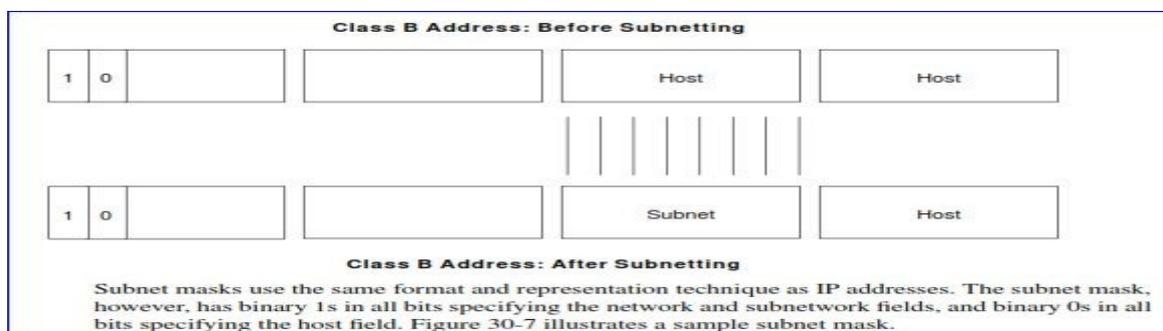
Subnets are under local administration. As such, the outside world sees an organization as a single network and has no detailed knowledge of the organization's internal structure.

A given network address can be broken up into many subnetworks. For example, 172.16.1.0, 172.16.2.0, 172.16.3.0, and 172.16.4.0 are all subnets within network 172.16.0.0. (All 0s in the host portion of an address specifies the entire network.)

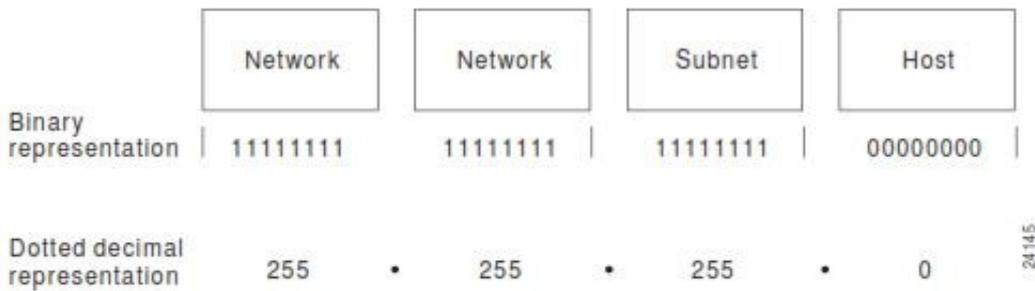
IP Subnet Mask

A subnet address is created by "borrowing" bits from the host field and designating them as the subnet field. The number of borrowed bits varies and is specified by the subnet mask. Figure shows how bits are borrowed from the host address field to create the subnet address field.

Bits are borrowed from the host address field to create the subnet address field.



A sample subnet mask consists of all binary 1s and 0s.



Subnet mask bits should come from the high-order (left-most) bits of the host field, as Figure 30-8 illustrates. Details of Class B and C subnet mask types follow. Class A addresses are not discussed in this chapter because they generally are subnetted on an 8-bit boundary.

Subnet mask bits come from the high-order bits of the host field.

128	64	32	16	8	4	2	1	
↓	↓	↓	↓	↓	↓	↓	↓	
1	0	0	0	0	0	0	0	= 128
1	1	0	0	0	0	0	0	= 192
1	1	1	0	0	0	0	0	= 224
1	1	1	1	0	0	0	0	= 240
1	1	1	1	1	0	0	0	= 248
1	1	1	1	1	1	0	0	= 252
1	1	1	1	1	1	1	0	= 254
1	1	1	1	1	1	1	1	= 255

Various types of subnet masks exist for Class B and C subnets.

The default subnet mask for a Class B address that has no subnetting is 255.255.0.0, while the subnet mask for a Class B address 171.16.0.0 that specifies eight bits of subnetting is 255.255.255.0. The reason for this is that eight bits of subnetting or $2^8 - 2$ (1 for the network address and 1 for the broadcast address) = 254 subnets possible, with $2^8 - 2 = 254$ hosts per subnet. The subnet mask for a Class C address 192.168.2.0 that specifies five bits of subnetting is 255.255.255.248. With five bits available for subnetting, $2^5 - 2 = 30$ subnets possible, with $2^3 - 2 = 6$ hosts per subnet. The reference charts shown in table 30-2 and table 30-3 can be used when planning Class B and C networks to determine the required number of subnets and hosts, and the appropriate subnet mask.

Class B Subnetting Reference Chart

Number of Bits	Subnet Mask	Number of Subnets	Number of Hosts
2	255.255.192.0	2	16382
3	255.255.224.0	6	8190
4	255.255.240.0	14	4094
5	255.255.248.0	30	2046
6	255.255.252.0	62	1022
7	255.255.254.0	126	510
8	255.255.255.0	254	254
9	255.255.255.128	510	126
10	255.255.255.192	1022	62
11	255.255.255.224	2046	30
12	255.255.255.240	4094	14

Class C Subnetting Reference Chart

Number of Bits	Subnet Mask	Number of Subnets	Number of Hosts
2	255.255.255.192	2	62
3	255.255.255.224	6	30
4	255.255.255.240	14	14
5	255.255.255.248	30	6
6	255.255.255.252	62	2

How Subnet Masks are Used to Determine the Network Number?

The router performs a set process to determine the network (or more specifically, the subnetwork) address. First, the router extracts the IP destination address from the incoming packet and retrieves the internal subnet mask. It then performs a logical AND operation to obtain the network number.

This causes the host portion of the IP destination address to be removed, while the destination network number remains. The router then looks up the destination network number and matches it with an outgoing interface. Finally, it forwards the frame to the destination IP address. Specifics regarding the logical AND operation are discussed in the following section.

Logical AND Operation

Three basic rules govern logically “ANDing” two binary numbers. First, 1 “ANDed” with 1 yields 1. Second, 1 “ANDed” with 0 yields 0. Finally, 0 “ANDed” with 0 yields 0. The truth table provided in table 30-4 illustrates the rules for logical AND operations.

Rules for Logical AND Operations

Input	Input	Output
1	1	1
1	0	0
0	1	0
0	0	0

Two simple guidelines exist for remembering logical AND operations: Logically “ANDing” a 1 with a 1 yields the original value, and logically “ANDing” a 0 with any number yields 0. Figure illustrates that when a logical AND of the destination IP address and the subnet mask is performed, the subnetwork number remains, which the router uses to forward the packet.

Applying a logical AND the destination IP address and the subnet mask produces the subnetwork number.

	Network	Subnet	Host
Destination IP Address	171.16.1.2	00000001	00000010
Subnet Mask	255.255.255.0	11111111	00000000
		00000001	00000000
		1	0

Address Resolution Protocol (ARP) Overview

For two machines on a given network to communicate, they must know the other machine’s physical (or MAC) addresses. By broadcasting Address Resolution Protocols (ARPs), a host can dynamically discover the MAC-layer address corresponding to a particular IP network-layer address. After receiving a MAC-layer address, IP devices create an ARP cache to store the recently acquired IP-to-MAC address mapping, thus avoiding having to broadcast ARPs when they want to recontact a device. If the device does not respond within a specified time frame, the cache entry is flushed. In addition to the Reverse Address Resolution Protocol (RARP) is used to map MAC-layer addresses to IP addresses. RARP, which is the logical inverse of ARP, might be used by diskless workstations that do not know their IP addresses when they boot. RARP relies on the presence of a RARP server with table entries of MAC-layer-to-IP address mappings.

Internet Routing

Internet routing devices traditionally have been called gateways. In today’s terminology, however, the term gateway refers specifically to a device that performs application-layer protocol translation between devices. Interior gateways refer to devices that perform these protocol functions between machines or networks under the same administrative control or authority, such as a corporation’s internal network. These are known as autonomous systems. Exterior gateways perform

protocol functions between independent networks. Routers within the Internet are organized hierarchically. Routers used for information exchange within autonomous systems are called interior routers, which use a variety of Interior Gateway Protocols (IGPs) to accomplish this purpose. The Routing Information Protocol (RIP) is an example

of an IGP. Routers that move information between autonomous systems are called exterior routers. These routers use an exterior gateway protocol to exchange information between autonomous systems. The Border Gateway Protocol (BGP) is an example of an exterior gateway protocol.

IP Routing

IP routing protocols are dynamic. Dynamic routing calls for routes to be calculated automatically at regular intervals by software in routing devices. This contrasts with static routing, where routes are established by the network administrator and do not change until the network administrator changes them. An IP routing table, which consists of destination address/next hop pairs, is used to enable dynamic routing. An entry in this table, for example, would be interpreted as follows: to get to network 172.31.0.0, send the packet out Ethernet interface 0 (E0). IP routing specifies that IP datagrams travel through internetworks one hop at a time. The entire route is not known at the onset of the journey, however. Instead, at each stop, the next destination is calculated by matching the destination address within the datagram with an entry in the current node's routing table. Each node's involvement in the routing process is limited to forwarding packets based on internal information. The nodes do not monitor whether the packets get to their final destination, nor does IP provide for error reporting back to the source when routing anomalies occur. This task is left to another Internet protocol, the Internet Control-Message Protocol (ICMP), which is discussed in the following section.

Internet Control Message Protocol (ICMP)

The Internet Control Message Protocol (ICMP) is a network-layer Internet protocol that provides message packets to report errors and other information regarding IP packet processing back to the source. ICMP is documented in RFC 792.

ICMP Messages

ICMPs generate several kinds of useful messages, including Destination Unreachable, Echo Request and Reply, Redirect, Time Exceeded, and Router Advertisement and Router Solicitation. If an ICMP message cannot be delivered, no second one is generated. This is to avoid an endless flood of ICMP messages.

When an ICMP destination-unreachable message is sent by a router, it means that the router is unable to send the package to its final destination. The router then discards the original packet. Two reasons exist for why a destination might be unreachable. Most commonly, the source host has specified a non-existent address. Less frequently, the router does not have a route to the destination. Destination-unreachable messages include four basic types: network unreachable, host unreachable, protocol unreachable, and port unreachable. Network-unreachable messages usually mean that a failure has occurred in the routing or addressing of a packet. Host-unreachable messages usually indicate delivery failure, such as a

wrong subnet mask. Protocol-unreachable messages generally mean that the destination does not support the upper-layer protocol specified in the packet. Port-unreachable messages imply that the TCP socket or port is not available. An ICMP echo-request message, which is generated by the ping command, is sent by any host to test node reachability across an internetwork. The ICMP echo-reply message indicates that the node can be successfully reached. An ICMP Redirect message is sent by the router to the source host to stimulate more efficient routing. The router still forwards the original packet to the destination. ICMP redirects allow host routing tables to remain small because it is necessary to know the address of only one router, even if that router does not provide the best path. Even after receiving an ICMP Redirect message, some devices might continue using the less-efficient route.

Transmission Control Protocol (TCP)

An ICMP Time-exceeded message is sent by the router if an IP packet's Time-to-Live field (expressed in hops or seconds) reaches zero. The Time-to-Live field prevents packets from continuously circulating the internetwork if the internetwork contains a routing loop. The router then discards the original packet.

ICMP Router-Discovery Protocol (IDRP)

IDRP uses Router-Advertisement and Router-Solicitation messages to discover the addresses of routers on directly attached subnets. Each router periodically multicasts Router-Advertisement messages from each of its interfaces. Hosts then discover addresses of routers on directly attached subnets by listening for these messages. Hosts can use Router-Solicitation messages to request immediate advertisements rather than waiting for unsolicited messages. IDRP offers several advantages over other methods of discovering addresses of neighbouring routers. Primarily, it does not require hosts to recognize routing protocols, nor does it require manual Configuration by an administrator. Router-Advertisement messages enable hosts to discover the existence of neighbouring routers, but not which router is best to reach a particular destination. If a host uses a poor first-hop router to reach a particular destination, it receives a Redirect message identifying a better choice.

Transmission Control Protocol (TCP)

The TCP provides reliable transmission of data in an IP environment. TCP corresponds to the transport layer (Layer 4) of the OSI reference model. Among the services TCP provides are stream data transfer, reliability, efficient flow control, full-duplex operation, and multiplexing. With stream data transfer, TCP delivers an unstructured stream of bytes identified by sequence numbers. This service benefits applications because they do not have to chop data into blocks before handing it off to TCP. Instead, TCP groups bytes into segments and passes them to IP for delivery. TCP offers reliability by providing connection-oriented, end-to-end reliable packet delivery through an internetwork. It does this by sequencing bytes with a forwarding acknowledgment number that indicates to the destination the next byte the source expects to receive. Bytes not acknowledged within a specified time period are retransmitted. The reliability mechanism of TCP allows devices to deal with lost, delayed, duplicate, or misread packets. A time-out mechanism allows devices to detect lost packets and request retransmission. TCP offers efficient flow control, which means that, when sending acknowledgments back to the source, the receiving

TCP process indicates the highest sequence number it can receive without overflowing its internal buffers. Full-duplex operation means that TCP processes can both send and receive at the same time. Finally, TCP's multiplexing means that numerous simultaneous upper-layer conversations can be multiplexed over a single connection.

TCP Connection Establishment

To use reliable transport services, TCP hosts must establish a connection-oriented session with one another. Connection establishment is performed by using a "three-way handshake" mechanism. A three-way handshake synchronizes both ends of a connection by allowing both sides to agree upon initial sequence numbers. This mechanism also guarantees that both sides are ready to transmit data and know that the other side is ready to transmit as well. This is necessary so that packets are not transmitted or retransmitted during session establishment or after session termination. Each host randomly chooses a sequence number used to track bytes within the stream it is sending and receiving. Then, the three-way handshake proceeds in the following manner: The first host (Host A) initiates a connection by sending a packet with the initial sequence number (X) and SYN bit set to indicate a connection request. The second host (Host B) receives the SYN, records the sequence number X, and replies by acknowledging the SYN (with an ACK = X + 1). Host B includes its own initial sequence number (SEQ = Y). An ACK = 20 means the host has received bytes 0 through 19 and expects byte 20 next. This technique is called forward acknowledgment. Host A then acknowledges all bytes Host B sent with a forward acknowledgment indicating the next byte Host A expects to receive (ACK = Y + 1). Data transfer then can begin.

Positive Acknowledgment and Retransmission (PAR)

A simple transport protocol might implement a reliability-and-flow-control technique where the source sends one packet, starts a timer, and waits for an acknowledgment before sending a new packet. If the acknowledgment is not received before the timer expires, the source retransmits the packet. Such a technique is called positive acknowledgment and retransmission (PAR). By assigning each packet a sequence number, PAR enables hosts to track lost or duplicate packets caused by network delays that result in premature retransmission. The sequence numbers are sent back in the acknowledgments so that the acknowledgments can be tracked. PAR is an inefficient use of bandwidth, however, because a host must wait for an acknowledgment before sending a new packet, and only one packet can be sent at a time.

TCP Sliding Window

A TCP sliding window provides more efficient use of network bandwidth than PAR because it enables hosts to send multiple bytes or packets before waiting for an acknowledgment. In TCP, the receiver specifies the current window size in every packet. Because TCP provides a byte-stream connection, window sizes are expressed in bytes. This means that a window is the number of data bytes that the sender is allowed to send before waiting for an acknowledgment. Initial window

sizes are indicated at connection setup, but might vary throughout the data transfer to provide flow control. A window size of zero, for instance, means “Send no data.”

In a TCP sliding-window operation, for example, the sender might have a sequence of bytes to send (numbered 1 to 10) to a receiver who has a window size of five. The sender then would place a window around the first five bytes and transmit them together. It would then wait for an acknowledgment. The receiver would respond with an ACK = 6, indicating that it has received bytes 1 to 5 and is expecting byte 6 next. In the same packet, the receiver would indicate that its window size is 5.

The sender then would move the sliding window five bytes to the right and transmit bytes 6 to 10. The receiver would respond with an ACK = 11, indicating that it is expecting sequenced byte 11 next. In this packet, the receiver might indicate that its window size is 0 (because, for example, its internal buffers are full). At this point, the sender cannot send any more bytes until the receiver sends another packet with a window size greater than 0.

TCP Packet Format

illustrates the fields and overall format of a TCP packet.

Twelve fields comprise a TCP packet.

Source port		Destination port	
Sequence number			
Acknowledgment number			
Data offset	Reserved	Flags	Window
Checksum		Urgent pointer	
Options (+ padding)			
Data (variable)			

TCP Packet Field Descriptions

The following descriptions summarize the TCP packet fields illustrated in Figure

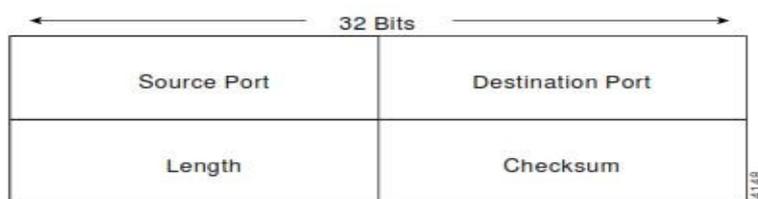
- Source Port and Destination Port—Identifies points at which upper-layer source and destination processes receive TCP services.
- Sequence Number—Usually specifies the number assigned to the first byte of data in the current message. In the connection-establishment phase, this field also can be used to identify an initial sequence number to be used in an upcoming transmission.

- Acknowledgment Number – Contains the sequence number of the next byte of data the sender of the packet expects to receive.
- Data Offset – Indicates the number of 32-bit words in the TCP header.
- Reserved – Remains reserved for future use.
- Flags – Carries a variety of control information, including the SYN and ACK bits used for connection establishment, and the FIN bit used for connection termination.
- Window – Specifies the size of the sender’s receive window (that is, the buffer space available for incoming data).
- Checksum – Indicates whether the header was damaged in transit.
- Urgent Pointer – Points to the first urgent data byte in the packet.
- Options – Specifies various TCP options.
- Data – Contains upper-layer information.

User Datagram Protocol (UDP)

The User Datagram Protocol (UDP) is a connectionless transport-layer protocol (Layer 4) that belongs to the Internet protocol family. UDP is basically an interface between IP and upper-layer processes. UDP protocol ports distinguish multiple applications running on a single device from one another. Unlike the TCP, UDP adds no reliability, flow-control, or error-recovery functions to IP. Because of UDP’s simplicity, UDP headers contain fewer bytes and consume less network overhead than TCP. UDP is useful in situations where the reliability mechanisms of TCP are not necessary, such as in cases where a higher-layer protocol might provide error and flow control. UDP is the transport protocol for several well-known application-layer protocols, including Network File System (NFS), Simple Network Management Protocol (SNMP), Domain Name System (DNS), and Trivial File Transfer Protocol (TFTP). The UDP packet format contains four fields, as shown in Figure. These include source and destination ports, length, and checksum fields.

A UDP packet consists of four fields.



Source and destination ports contain the 16-bit UDP protocol port numbers used to demultiplex datagrams for receiving application-layer processes. A length field specifies the length of the UDP header and data. Checksum provides an (optional) integrity check on the UDP header and data.

Internet Protocols Application-Layer Protocols

The Internet protocol suite includes many application-layer protocols that represent a wide variety of applications, including the following:

- File Transfer Protocol (FTP) – Moves files between devices
- Simple Network-Management Protocol (SNMP) – Primarily reports anomalous network conditions and sets network threshold values
- Telnet – Serves as a terminal emulation protocol
- X Windows – Serves as a distributed windowing and graphics system used for communication between X terminals and UNIX workstations

- Network File System (NFS), External Data Representation (XDR), and Remote Procedure Call (RPC) – Work together to enable transparent access to remote network resources
- Simple Mail Transfer Protocol (SMTP) – Provides electronic mail services
- Domain Name System (DNS) – Translates the names of network nodes into network addresses lists these higher-layer protocols and the applications that they support.

Higher-Layer Protocols and Their Applications

Application	Protocols
File transfer	FTP
Terminal emulation	Telnet
Electronic mail	SMTP
Network management	SNMP
Distributed file services	NFS, XDR, RPC, X Windows