Components, Elements and Concepts of Computer Networking

Klara Oden

Bruno Barbee

First Edition, 2012

ISBN 978-81-323-1265-9

Published by:
College Publishing House
4735/22 Prakashdeep Bldg,
Ansari Road, Darya Ganj,
Delhi - 110002
Email: info@wtbooks.com

Table of Contents

Chapter 1

Computer Network

A **computer network**, often simply referred to as a network, is a collection of computers and devices interconnected by communications channels that facilitate communications among users and allows users to share resources. Networks may be classified according to a wide variety of characteristics. A computer network allows sharing of resources and information among interconnected devices.

History

Early networks of communicating computers included the military radar system Semi-Automatic Ground Environment (SAGE) and its relative the commercial airline reservation system Semi-Automatic Business Research Environment (SABRE), started in the late 1950s. In the 1960s, the Advanced Research Projects Agency (ARPA) started funding the design of the Advanced Research Projects Agency Network (ARPANET) for the United States Department of Defense. Development of the network began in 1969, based on designs developed during the 1960s. The ARPANET evolved into the modern Internet.

Purpose

Computer networks can be used for a variety of purposes:

- *Facilitating communications.* Using a network, people can communicate efficiently and easily via email, instant messaging, chat rooms, telephone, video telephone calls, and video conferencing.
- *Sharing hardware.* In a networked environment, each computer on a network may access and use hardware resources on the network, such as printing a document on a shared network printer.
- *Sharing files, data, and information.* In a network environment, authorized user may access data and information stored on other computers on the network. The capability of providing access to data and information on shared storage devices is an important feature of many networks.
- *Sharing software.* Users connected to a network may run application programs on remote computers.
- *Information preservation.*

- *Security.*
- Easy communication

Network classification

The following list presents categories used for classifying networks.

Connection method

Computer networks can be classified according to the hardware and software technology that is used to interconnect the individual devices in the network, such as optical fiber, Ethernet, wireless LAN, HomePNA, power line communication or G.hn.

Ethernet as it is defined by IEEE 802 utilizes various standards and mediums that enable communication between devices. Frequently deployed devices include hubs, switches, bridges, or routers. Wireless LAN technology is designed to connect devices without wiring. These devices use radio waves or infrared signals as a transmission medium. ITU-T G.hn technology uses existing home wiring (coaxial cable, phone lines and power lines) to create a high-speed (up to 1 Gigabit/s) local area network

Wired technologies

- *Twisted pair wire* is the most widely used medium for telecommunication. Twisted-pair cabling consist of copper wires that are twisted into pairs. Ordinary telephone wires consist of two insulated copper wires twisted into pairs. Computer networking cabling consist of 4 pairs of copper cabling that can be utilized for both voice and data transmission. The use of two wires twisted together helps to reduce crosstalk and electromagnetic induction. The transmission speed ranges from 2 million bits per second to 100 million bits per second. Twisted pair cabling comes in two forms which are Unshielded Twisted Pair (UTP) and Shielded twisted-pair (STP) which are rated in categories which are manufactured in different increments for various scenarios.

- *Coaxial cable* is widely used for cable television systems, office buildings, and other work-sites for local area networks. The cables consist of copper or aluminum wire wrapped with insulating layer typically of a flexible material with a high dielectric constant, all of which are surrounded by a conductive layer. The layers of insulation help minimize interference and distortion. Transmission speed range from 200 million to more than 500 million bits per second.

- *Optical fiber cable* consists of one or more filaments of glass fiber wrapped in protective layers that carries a data by means of pulses of light. It transmits light which can travel over extended distances. Fiber-optic cables are not affected by electromagnetic radiation. Transmission speed may reach trillions of bits per second. The transmission speed of fiber optics is hundreds of times faster than for coaxial cables and thousands of times faster than a twisted-pair wire.A recent innovation in fiber-optic cable is the use of colored light.Instead of carrying one message in a stream of white light impulses, this technology can carry multiple signals in a single strand.

Wireless technologies

- *Terrestrial microwave* – Terrestrial microwaves use Earth-based transmitter and receiver. The equipment looks similar to satellite dishes. Terrestrial microwaves use low-gigahertz range, which limits all communications to line-of-sight. Path between relay stations spaced approx, 30 miles apart. Microwave antennas are usually placed on top of buildings, towers, hills, and mountain peaks.

- *Communications satellites* – The satellites use microwave radio as their telecommunications medium which are not deflected by the Earth's atmosphere. The satellites are stationed in space, typically 22,000 miles (for geosynchronous satellites) above the equator. These Earth-orbiting systems are capable of receiving and relaying voice, data, and TV signals.

- *Cellular and PCS systems* – Use several radio communications technologies. The systems are divided to different geographic areas. Each area has a low-power transmitter or radio relay antenna device to relay calls from one area to the next area.

- *Wireless LANs* – Wireless local area network use a high-frequency radio technology similar to digital cellular and a low-frequency radio technology. Wireless LANs use spread spectrum technology to enable communication between multiple devices in a limited area. An example of open-standards wireless radio-wave technology is IEEE.

- Infrared communication, which can transmit signals between devices within small distances not more than 10 meters peer to peer or (face to face) without any body in the line of transmitting.

Scale

Networks are often classified as local area network (LAN), wide area network (WAN), metropolitan area network (MAN), personal area network (PAN), virtual private network (VPN), campus area network (CAN), storage area network (SAN), and others, depending on their scale, scope and purpose, e.g., controller area network (CAN) usage, trust level, and access right often differ between these types of networks. LANs tend to be designed for internal use by an organization's internal systems and employees in individual physical locations, such as a building, while WANs may connect physically separate parts of an organization and may include connections to third parties.

Functional relationship (network architecture)

Computer networks may be classified according to the functional relationships which exist among the elements of the network, e.g., active networking, client–server, Wireless ad hoc network and peer-to-peer (workgroup) architecture.

Network topology

Computer networks may be classified according to the network topology upon which the network is based, such as bus network, star network, ring network, mesh network. Network topology is the coordination by which devices in the network are arranged in their logical relations to one another, independent of physical arrangement. Even if networked computers are physically placed in a linear arrangement and are connected to a hub, the network has a star topology, rather than a bus topology. In this regard the visual and operational characteristics of a network are distinct. Networks may be classified based on the method of data used to convey the data, these include digital and analog networks.

Types of networks based on physical scope

Common types of computer networks may be identified by their scale.

Local area network

A local area network (LAN) is a network that connects computers and devices in a limited geographical area such as home, school, computer laboratory, office building, or closely positioned group of buildings. Each computer or device on the network is a node. Current wired LANs are most likely to be based on Ethernet technology, although new standards like ITU-T G.hn also provide a way to create a wired LAN using existing home wires (coaxial cables, phone lines and power lines).

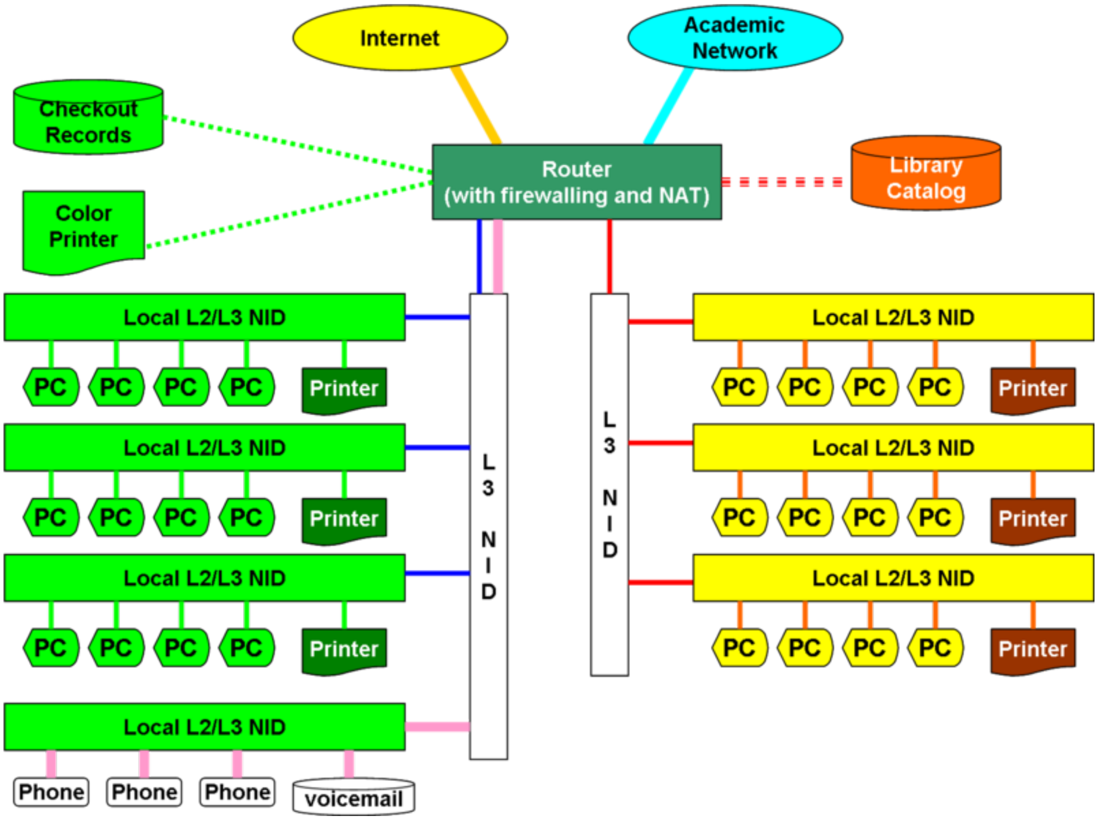

Typical library network, in a branching tree topology and controlled access to resources

All interconnected devices must understand the network layer (layer 3), because they are handling multiple subnets (the different colors). Those inside the library, which have only 10/100 Mbit/s Ethernet connections to the user device and a Gigabit Ethernet connection to the central router, could be called "layer 3 switches" because they only have Ethernet interfaces and must understand IP. It would be more correct to call them access routers, where the router at the top is a distribution router that connects to the Internet and academic networks' customer access routers.

The defining characteristics of LANs, in contrast to WANs (Wide Area Networks), include their higher data transfer rates, smaller geographic range, and no need for leased telecommunication lines. Current Ethernet or other IEEE 802.3 LAN technologies operate at speeds up to 10 Gbit/s. This is the data transfer rate. IEEE has projects investigating the standardization of 40 and 100 Gbit/s.

Personal area network

A personal area network (PAN) is a computer network used for communication among computer and different information technological devices close to one person. Some examples of devices that are used in a PAN are personal computers, printers, fax machines, telephones, PDAs, scanners, and even video game consoles. A PAN may include wired and wireless devices. The reach of a PAN typically extends to 10 meters. A wired PAN is usually constructed with USB and Firewire connections while technologies such as Bluetooth and infrared communication typically form a wireless PAN.

Home area network

A home area network (HAN) is a residential LAN which is used for communication between digital devices typically deployed in the home, usually a small number of personal computers and accessories, such as printers and mobile computing devices. An important function is the sharing of Internet access, often a broadband service through a CATV or Digital Subscriber Line (DSL) provider. It can also be referred to as an office area network (OAN).

Wide area network

A wide area network (WAN) is a computer network that covers a large geographic area such as a city, country, or spans even intercontinental distances, using a communications channel that combines many types of media such as telephone lines, cables, and air waves. A WAN 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.

Campus network

A campus network is a computer network made up of an interconnection of local area networks (LAN's) within a limited geographical area. The networking equipments (switches, routers) and transmission media (optical fiber, copper plant, Cat5 cabling etc.) are almost entirely owned (by the campus tenant / owner: an enterprise, university, government etc.).

In the case of a university campus-based campus network, the network is likely to link a variety of campus buildings including; academic departments, the university library and student residence halls.

Metropolitan area network

A Metropolitan area network is a large computer network that usually spans a city or a large campus.

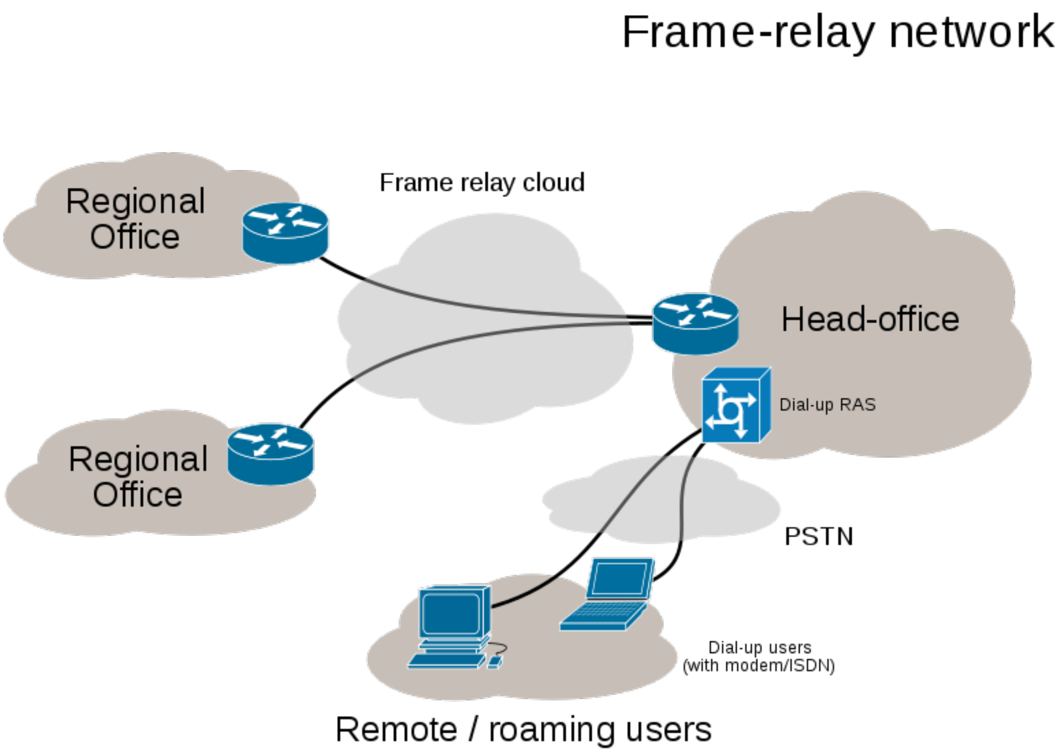

Sample EPN made of Frame relay WAN connections and dialup remote access

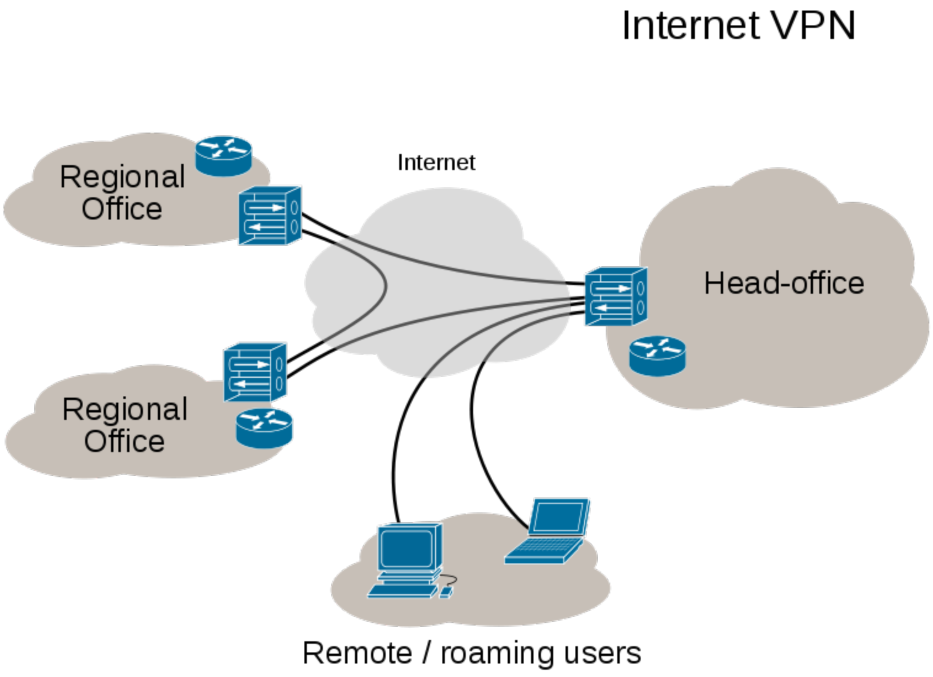

Sample VPN used to interconnect 3 offices and remote users

Enterprise private network

An enterprise private network is a network build by an enterprise to interconnect various company sites, e.g., production sites, head offices, remote offices, shops, in order to share computer resources.

Virtual private network

A virtual private network (VPN) is a computer network in which some of the links between nodes are carried by open connections or virtual circuits in some larger network (e.g., the Internet) instead of by physical wires. The data link layer protocols of the virtual network are said to be tunneled through the larger network when this is the case. One common application is secure communications through the public Internet, but a VPN need not have explicit security features, such as authentication or content encryption. VPNs, for example, can be used to separate the traffic of different user communities over an underlying network with strong security features.

VPN may have best-effort performance, or may have a defined service level agreement (SLA) between the VPN customer and the VPN service provider. Generally, a VPN has a topology more complex than point-to-point.

Internetwork

An internetwork is the connection of two or more private computer networks via a common routing technology (OSI Layer 3) using routers. The Internet is an aggregation of many internetworks, hence its name was shortened to Internet.

Backbone network

A Backbone network (BBN) A backbone network or network backbone is part of a computer network infrastructure that interconnects various pieces of network, providing a path for the exchange of information between different LANs or subnetworks. A backbone can tie together diverse networks in the same building, in different buildings in a campus environment, or over wide areas. Normally, the backbone's capacity is greater than the networks connected to it.

A large corporation that has many locations may have a backbone network that ties all of the locations together, for example, if a server cluster needs to be accessed by different departments of a company that are located at different geographical locations. The pieces of the network connections (for example: ethernet, wireless) that bring these departments together is often mentioned as network backbone. Network congestion is often taken into consideration while designing backbones.

Backbone networks should not be confused with the Internet backbone.

Global area network

A global area network (GAN) is a network used for supporting mobile communications across an arbitrary number of wireless LANs, satellite coverage areas, etc. The key challenge in mobile communications is handing off the user communications from one local coverage area to the next. In IEEE Project 802, this involves a succession of terrestrial wireless LANs.

Internet

The Internet is a global system of interconnected governmental, academic, corporate, public, and private computer networks. It is based on the networking technologies of the Internet Protocol Suite. It is the successor of the Advanced Research Projects Agency Network (ARPANET) developed by DARPA of the United States Department of Defense. The Internet is also the communications backbone underlying the World Wide Web (WWW).

Participants in the Internet use a diverse array of methods of several hundred documented, and often standardized, protocols compatible with the Internet Protocol Suite and an addressing system (IP addresses) administered by the Internet Assigned Numbers Authority and address registries. Service providers and large enterprises exchange information about the reachability of their address spaces through the Border Gateway Protocol (BGP), forming a redundant worldwide mesh of transmission paths.

Intranets and extranets

Intranets and extranets are parts or extensions of a computer network, usually a local area network.

An intranet is a set of networks, using the Internet Protocol and IP-based tools such as web browsers and file transfer applications, that is under the control of a single administrative entity. That administrative entity closes the intranet to all but specific, authorized users. Most commonly, an intranet is the internal network of an organization. A large intranet will typically have at least one web server to provide users with organizational information.

An extranet is a network that is limited in scope to a single organization or entity and also has limited connections to the networks of one or more other usually, but not necessarily, trusted organizations or entities—a company's customers may be given access to some part of its intranet—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 an extranet cannot consist of a single LAN; it must have at least one connection with an external network.

Overlay network

An overlay network is a virtual computer network that is built on top of another network. Nodes in the overlay are connected by virtual or logical links, each of which corresponds to a path, perhaps through many physical links, in the underlying network.

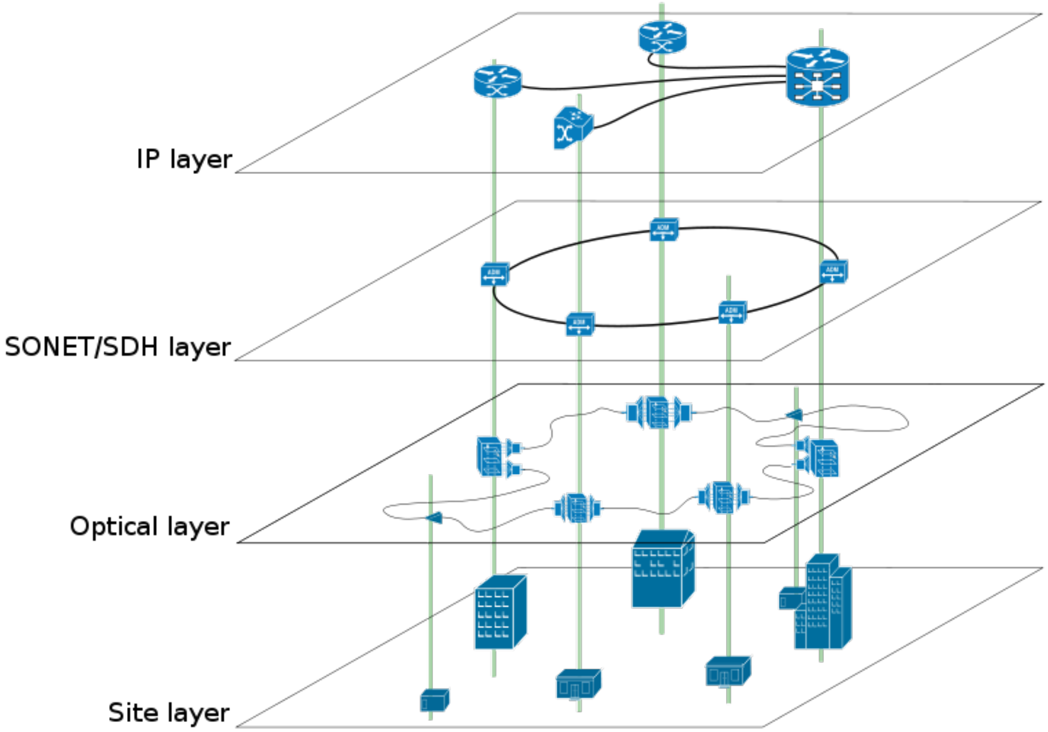

A sample overlay network: IP over SONET over Optical

For example, many peer-to-peer networks are overlay networks because they are organized as nodes of a virtual system of links run on top of the Internet. The Internet was initially built as an overlay on the telephone network .

Overlay networks have been around since the invention of networking when computer systems were connected over telephone lines using modem, before any data network existed.

Nowadays the Internet is the basis for many overlaid networks that can be constructed to permit routing of messages to destinations specified by an IP address. For example, distributed hash tables can be used to route messages to a node having a specific logical address, whose IP address is known in advance.

Overlay networks have also been proposed as a way to improve Internet routing, such as through quality of service guarantees to achieve higher-quality streaming media.

Previous proposals such as IntServ, DiffServ, and IP Multicast have not seen wide acceptance largely because they require modification of all routers in the network. On the other hand, an overlay network can be incrementally deployed on end-hosts running the overlay protocol software, without cooperation from Internet service providers. The overlay has no control over how packets are routed in the underlying network between two overlay nodes, but it can control, for example, the sequence of overlay nodes a message traverses before reaching its destination.

For example, Akamai Technologies manages an overlay network that provides reliable, efficient content delivery (a kind of multicast). Academic research includes End System Multicast and Overcast for multicast; RON (Resilient Overlay Network) for resilient routing; and OverQoS for quality of service guarantees, among others. A backbone network or network backbone is a part of computer network infrastructure that interconnects various pieces of network, providing a path for the exchange of information between different LANs or subnetworks. A backbone can tie together diverse networks in the same building, in different buildings in a campus environment, or over wide areas. Normally, the backbone's capacity is greater than the networks connected to it.

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 802.12) or optical cable ("optical fiber").

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 often provides a low-level addressing system through the use of MAC addresses.

Each network interface card has its unique id. This is written on a chip which is mounted on the card.

Repeaters

A repeater is an electronic device that receives a signal, cleans it of unnecessary noise, regenerates it, and retransmits it at a higher power level, or to the other side of an obstruction, so that the signal can cover longer distances without degradation. In most twisted pair Ethernet configurations, repeaters are required for cable that runs longer than 100 meters. A repeater with multiple ports is known as a hub. Repeaters work on the Physical Layer of the OSI model. Repeaters require a small amount of time to regenerate the signal. This can cause a propagation delay which can affect network communication when there are several repeaters in a row. Many network architectures limit the number of repeaters that can be used in a row (e.g. Ethernet's 5-4-3 rule).

Bridges

A network bridge connects multiple network segments at the data link layer (layer 2) of the OSI model. Bridges broadcast to all ports except the port on which the broadcast was received. However, bridges do not promiscuously copy traffic to all ports, as hubs do, but learn which MAC addresses are reachable through specific ports. Once the bridge associates a port and an address, it will send traffic for that address to that port only.

Bridges learn the association of ports and addresses by examining the source address of frames that it sees on various ports. Once a frame arrives through a port, its source address is stored and the bridge assumes that MAC address is associated with that port. The first time that a previously unknown destination address is seen, the bridge will forward the frame to all ports other than the one on which the frame arrived.

Bridges come in three basic types:

- Local bridges: Directly connect local area networks (LANs)
- Remote bridges: Can be used to create a wide area network (WAN) link between LANs. Remote bridges, where the connecting link is slower than the end networks, largely have been replaced with routers.
- Wireless bridges: Can be used to join LANs or connect remote stations to LANs.

Switches

A network switch is a device that forwards and filters OSI layer 2 datagrams (chunks of data communication) between ports (connected cables) based on the MAC addresses in the packets. A switch is distinct from a hub in that it only forwards the frames to the ports involved in the communication rather than all ports connected. A switch breaks the collision domain but represents itself as a broadcast domain. Switches make forwarding decisions of frames on the basis of MAC addresses. A switch normally has numerous ports, facilitating a star topology for devices, and cascading additional switches. Some switches are capable of routing based on Layer 3 addressing or additional logical levels; these are called multi-layer switches. The term *switch* is used loosely in marketing to encompass devices including routers and bridges, as well as devices that may distribute traffic on load or by application content (e.g., a Web URL identifier).

Routers

A router is an internetworking device that forwards packets between networks by processing information found in the datagram or packet (Internet protocol information from Layer 3 of the OSI Model). In many situations, this information is processed in conjunction with the routing table (also known as forwarding table). Routers use routing tables to determine what interface to forward packets (this can include the "null" also known as the "black hole" interface because data can go into it, however, no further processing is done for said data).

Firewalls

Firewalls are the most important aspect of a network with respect to security. A firewalled system does not need every interaction or data transfer monitored by a human, as automated processes can be set up to assist in rejecting access requests from unsafe sources, and allowing actions from recognized ones. The vital role firewalls play in network security grows in parallel with the constant increase in 'cyber' attacks for the purpose of stealing/corrupting data, planting viruses, etc.

Chapter 2

Computer Networking

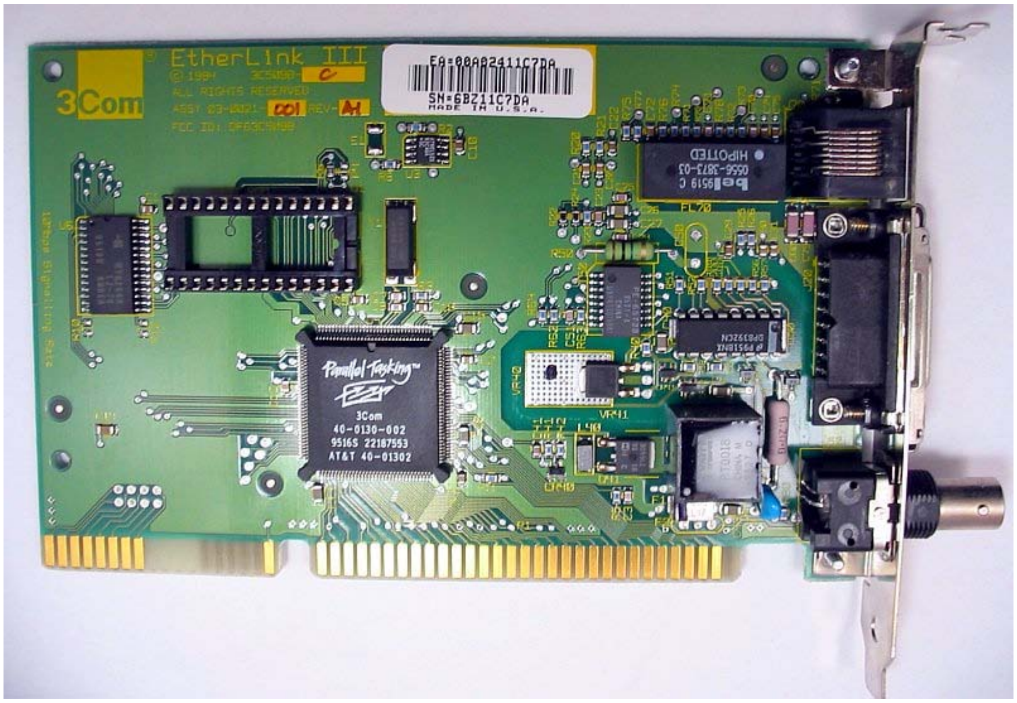

Network cards such as this one can transmit and receive data at high rates over various types of network cables. This card is a 'Combo' card which supports three cabling standards.

Computer networking or **Data communications** (**Datacom**) is the engineering discipline concerned with the communication between computer systems or devices. A computer network is any set of computers or devices connected to each other with the ability to exchange data. Computer networking is sometimes considered a sub-discipline of telecommunications, computer science, information technology and/or computer engineering since it relies heavily upon the theoretical and practical application of these scientific and engineering disciplines. The three types of networks are: the Internet, the intranet, and the extranet. Examples of different network methods are:

- Local area network (LAN), which is usually a small network constrained to a small geographic area. An example of a LAN would be a computer network within a building.
- Metropolitan area network (MAN), which is used for medium size area. examples for a city or a state.
- Wide area network (WAN) that is usually a larger network that covers a large geographic area.
- Wireless LANs and WANs (WLAN & WWAN) are the wireless equivalent of the LAN and WAN.

All networks are interconnected to allow communication with a variety of different kinds of media, including twisted-pair copper wire cable, coaxial cable, optical fiber, power lines and various wireless technologies. The devices can be separated by a few meters (e.g. via Bluetooth) or nearly unlimited distances (e.g. via the interconnections of the Internet). Networking, routers, routing protocols, and networking over the public Internet have their specifications defined in documents called RFCs.

Views of networks

Users and network administrators typically have different views of their networks. Users can share printers and some servers from a workgroup, which usually means they are in the same geographic location and are on the same LAN, whereas a Network Administrator is responsible to keep that network up and running. A community of interest has less of a connection of being in a local area, and should be thought of as a set of arbitrarily located users who share a set of servers , and possibly also communicate via peer-to-peer technologies.

Network administrators can see networks from both physical and logical perspectives. The physical perspective involves geographic locations, physical cabling, and the network elements (e.g., routers, bridges and application layer gateways that interconnect the physical media. Logical networks, called, in the TCP/IP architecture, subnets, map onto one or more physical media. For example, a common practice in a campus of buildings is to make a set of LAN cables in each building appear to be a common subnet, using virtual LAN (VLAN) technology.

Both users and administrators will be aware, to varying extents, of the trust and scope characteristics of a network. Again using TCP/IP architectural terminology, an intranet is a community of interest under private administration usually by an enterprise, and is only accessible by authorized users (e.g. employees). Intranets do not have to be connected to the Internet, but generally have a limited connection. An extranet is an extension of an intranet that allows secure communications to users outside of the intranet (e.g. business partners, customers).

Unofficially, the Internet is the set of users, enterprises, and content providers that are interconnected by Internet Service Providers (ISP). From an engineering viewpoint, the Internet is the set of subnets, and aggregates of subnets, which share the registered

IP address space and exchange information about the reachability of those IP addresses using the Border Gateway Protocol. Typically, the human-readable names of servers are translated to IP addresses, transparently to users, via the directory function of the Domain Name System (DNS).

Over the Internet, there can be business-to-business (B2B), business-to-consumer (B2C) and consumer-to-consumer (C2C) communications. Especially when money or sensitive information is exchanged, the communications are apt to be **secured** by some form of communications security mechanism. Intranets and extranets can be securely superimposed onto the Internet, without any access by general Internet users, using secure Virtual Private Network (VPN) technology.

History of computer networks

Before the advent of computer networks that were based upon some type of telecommunications system, communication between calculation machines and early computers was performed by human users by carrying instructions between them. Many of the social behaviors seen in today's Internet were demonstrably present in the nineteenth century and arguably in even earlier networks using visual signals.

- In September 1940 George Stibitz used a teletype machine to send instructions for a problem set from his Model at Dartmouth College to his Complex Number Calculator in New York and received results back by the same means. Linking output systems like teletypes to computers was an interest at the Advanced Research Projects Agency (ARPA) when, in 1962, J.C.R. Licklider was hired and developed a working group he called the "Intergalactic Network", a precursor to the ARPANET.
- In 1964, researchers at Dartmouth developed the Dartmouth Time Sharing System for distributed users of large computer systems. The same year, at Massachusetts Institute of Technology, a research group supported by General Electric and Bell Labs used a computer to route and manage telephone connections.
- Throughout the 1960s Leonard Kleinrock, Paul Baran and Donald Davies independently conceptualized and developed network systems which used packets that could be used in a network between computer systems.
- 1965 Thomas Merrill and Lawrence G. Roberts created the first wide area network (WAN).
- The first widely used telephone switch that used true computer control was introduced by Western Electric in 1965.
- In 1969 the University of California at Los Angeles, the Stanford Research Institute, University of California at Santa Barbara, and the University of Utah were connected as the beginning of the ARPANET network using 50 kbit/s circuits.
- Commercial services using X.25 were deployed in 1972, and later used as an underlying infrastructure for expanding TCP/IP networks.

Today, computer networks are the core of modern communication. All modern aspects of the Public Switched Telephone Network (PSTN) are computer-controlled, and telephony increasingly runs over the Internet Protocol, although not necessarily the public Internet. The scope of communication has increased significantly in the past decade, and this boom in communications would not have been possible without the progressively advancing computer network. Computer networks, and the technologies needed to connect and communicate through and between them, continue to drive computer hardware, software, and peripherals industries. This expansion is mirrored by growth in the numbers and types of users of networks from the researcher to the home user.

Networking methods

One way to categorize computer networks is by their geographic scope, although many real-world networks interconnect Local Area Networks (LAN) via Wide Area Networks (WAN) and wireless wide area networks (WWAN). These three (broad) types are:

Local area network (LAN)

A local area network is a network that spans a relatively small space and provides services to a small number of people.

A peer-to-peer or client-server method of networking may be used. A peer-to-peer network is where each client shares their resources with other workstations in the network. Examples of peer-to-peer networks are: Small office networks where resource use is minimal and a home network. A client-server network is where every client is connected to the server and each other. Client-server networks use servers in different capacities. These can be classified into two types:

1. Single-service servers
2. Print servers

The server performs one task such as file server, while other servers can not only perform in the capacity of file servers and print servers, but also can conduct calculations and use them to provide information to clients (Web/Intranet Server). Computers may be connected in many different ways, including Ethernet cables, Wireless networks, or other types of wires such as power lines or phone lines.

The ITU-T G.hn standard is an example of a technology that provides high-speed (up to 1 Gbit/s) local area networking over existing home wiring (power lines, phone lines and coaxial cables).

Wide area network (WAN)

A wide area network is a network where a wide variety of resources are deployed across a large domestic area or internationally. An example of this is a multinational business that uses a WAN to interconnect their offices in different countries. The largest and best example of a WAN is the Internet, which is a network composed of many smaller networks. The Internet is considered the largest network in the world. The PSTN (Public Switched Telephone Network) also is an extremely large network that is converging to use Internet technologies, although not necessarily through the public Internet.

A Wide Area Network involves communication through the use of a wide range of different technologies. These technologies include Point-to-Point WANs such as Point-to-Point Protocol (PPP) and High-Level Data Link Control (HDLC), Frame Relay, ATM (Asynchronous Transfer Mode) and Sonet (Synchronous Optical Network). The difference between the WAN technologies is based on the switching capabilities they perform and the speed at which sending and receiving bits of information (data) occur.

Wireless networks (WLAN, WWAN)

A wireless network is basically the same as a LAN or a WAN but there are no wires between hosts and servers. The data is transferred over sets of radio transceivers. These types of networks are beneficial when it is too costly or inconvenient to run the necessary cables.

The most common IEEE 802.11 WLANs cover, depending on antennas, ranges from hundreds of meters to a few kilometers. For larger areas, either communications satellites of various types, cellular radio, or wireless local loop (IEEE 802.16) all have advantages and disadvantages. Depending on the type of mobility needed, the relevant standards may come from the IETF or the ITU.

Network topology

The network topology defines the way in which computers, printers, and other devices are connected, physically and logically. A network topology describes the layout of the wire and devices as well as the paths used by data transmissions.

Network topology has two types:

- Physical
- Logical

Commonly used topologies include:

- Bus

- Star
- Tree (hierarchical)
- Linear
- Ring
- Mesh
    - partially connected
    - fully connected (sometimes known as *fully redundant*)

The network topologies mentioned above are only a general representation of the kinds of topologies used in computer network and are considered basic topologies

Chapter 3

OSI Model

The **Open Systems Interconnection model** (**OSI model**) is a product of the Open Systems Interconnection effort at the International Organization for Standardization. It is a way of sub-dividing a communications system into smaller parts called layers. A layer is a collection of similar functions that provide services to the layer above it and receives services from the layer below it. On each layer, an *instance* provides services to the instances at the layer above and requests service from the layer below.

For example, a layer that provides error-free communications across a network provides the path needed by applications above it, while it calls the next lower layer to send and receive packets that make up the contents of the path. Two instances at one layer are connected by a horizontal connection on that layer.

Most network protocols used in the market today are based on TCP/IP stacks.

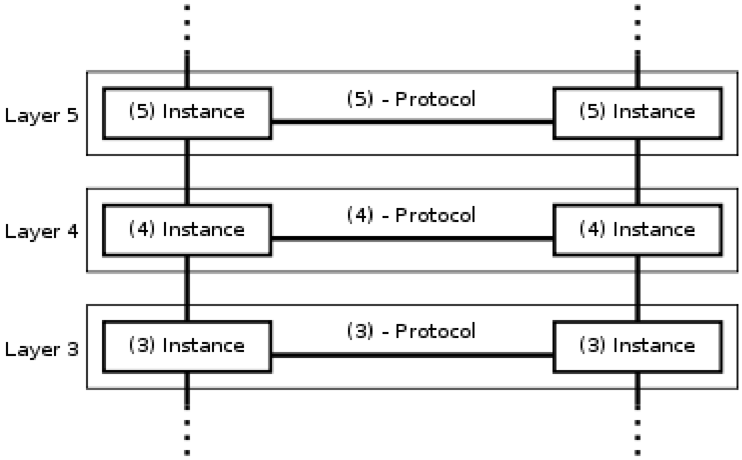

Communication in the OSI-Model (Example with layers 3 to 5)

History

Work on a layered model of network architecture was started and the International Organization for Standardization (ISO) began to develop its OSI framework architecture. OSI has two major components: an abstract model of networking, called the Basic Reference Model or seven-layer model, and a set of specific protocols.

Note: The standard documents that describe the OSI model can be freely downloaded from the ITU-T as the **X.200**-series of recommendations. A number of the protocol specifications are also available as part of the ITU-T X series. The equivalent ISO and ISO/IEC standards for the OSI model are available from ISO, but only some of them at no charge.

The concept of a 7 layer model was provided by the work of Charles Bachman, Honeywell Information Services. Various aspects of OSI design evolved from experiences with the ARPANET, the fledgling Internet, NPLNET, EIN, CYCLADES network and the work in IFIP WG6.1. The new design was 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.

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

Description of OSI layers

According to recommendation X.200, there are seven layers, each generically known as an N layer. An N+1 entity requests services from the N entity.

At each level, two entities (N-entity peers) interact by means of the N protocol by transmitting protocol data units (PDU).

A Service Data Unit (SDU) is a specific unit of data that has been passed down from an OSI layer to a lower layer, and which the lower layer has not yet encapsulated into a protocol data unit (PDU). An SDU is a set of data that is sent by a user of the services of a given layer, and is transmitted semantically unchanged to a peer service user.

The PDU at any given layer, layer N, is the SDU of the layer below, layer N-1. In effect the SDU is the 'payload' of a given PDU. That is, the process of changing a SDU to a PDU, consists of an encapsulation process, performed by the lower layer. All the data contained in the SDU becomes encapsulated within the PDU. The layer N-1 adds headers or footers, or both, to the SDU, transforming it into a PDU of layer N-1. The added headers or footers are part of the process used to make it possible to get data from a source to a destination.

OSI Model

	Data unit	Layer	Function
Host layers	Data	7. Application	Network process to application
		6. Presentation	Data representation, encryption and decryption, convert machine dependent data to machine independent data
		5. Session	Interhost communication
	Segments	4. Transport	End-to-end connections and reliability, flow control
Media layers	Packet/Datagram	3. Network	Path determination and logical addressing
	Frame	2. Data Link	Physical addressing
	Bit	1. Physical	Media, signal and binary transmission

Some orthogonal aspects, such as management and security, involve every layer.

Security services are not related to a specific layer: they can be related by a number of layers, as defined by ITU-T X.800 Recommendation.

These services are aimed to improve the CIA triad of transmitted data. Actually the availability of communication service is determined by network design and/or network management protocols. Appropriate choices for these are needed to protect against denial of service.

Layer 1: Physical Layer

The Physical Layer defines the electrical and physical specifications for devices. In particular, it defines the relationship between a device and a transmission medium, such as a copper or optical cable. This includes the layout of pins, voltages, cable specifications, hubs, repeaters, network adapters, host bus adapters (HBA used in storage area networks) and more.

To understand the function of the Physical Layer, contrast it with 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, whereas the Data Link Layer is concerned more with the interactions of multiple devices (i.e., at least two) with a shared medium. 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 optical fiber) 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, ITU-T G.hn and IEEE 802.11, as well as personal area networks such as Bluetooth and IEEE 802.15.4.

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. Originally, this layer was intended for point-to-point and point-to-multipoint media, characteristic of wide area media in the telephone system. Local area network architecture, which included broadcast-capable multiaccess media, was developed independently of the ISO work in IEEE Project 802. IEEE work assumed sublayering and management functions not required for WAN use. In modern practice, only error detection, not flow control using sliding window, is present in 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 the 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 Layer by protocols such as TCP, but is still used in niches where X.25 offers performance advantages.

The ITU-T G.hn standard, which provides high-speed local area networking over existing wires (power lines, phone lines and coaxial cables), includes a complete Data Link Layer which provides both error correction and flow control by means of a selective repeat Sliding Window Protocol.

Both WAN and LAN service arrange bits, from the Physical Layer, into logical sequences called frames. Not all Physical Layer bits necessarily go into frames, as some of these bits are purely intended for Physical Layer functions. For example, every fifth bit of the FDDI bit stream is not used by the Layer.

WAN Protocol architecture

Connection-oriented WAN data link protocols, in addition to framing, detect and may correct errors. They are also capable of controlling the rate of transmission. A WAN 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.

IEEE 802 LAN architecture

Practical, connectionless LANs began with the pre-IEEE Ethernet specification, which is the ancestor of IEEE 802.3. This layer manages the interaction of devices with a shared medium, which is the function of a Media Access Control (MAC) sublayer. Above this MAC sublayer is the media-independent IEEE 802.2 Logical Link Control (LLC) sublayer, which deals with addressing and multiplexing on multiaccess media.

While IEEE 802.3 is the dominant wired LAN protocol and IEEE 802.11 the wireless LAN protocol, obsolescent MAC layers include Token Ring and FDDI. The MAC sublayer detects but does not correct errors.

Layer 3: Network Layer

The Network Layer provides the functional and procedural means of transferring variable length data sequences from a source host on one network to a destination host on a different network, while maintaining the quality of service requested by the Transport Layer (in contrast to the data link layer which connects hosts within the same network). 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 not hierarchical.

Careful analysis of the Network Layer indicated that the Network Layer could have at least three sublayers:

1. Subnetwork Access - that considers protocols that deal with the interface to networks, such as X.25;
2. Subnetwork Dependent Convergence - when it is necessary to bring the level of a transit network up to the level of networks on either side;
3. Subnetwork Independent Convergence - which handles transfer across multiple networks.

The best example of this latter case is CLNP, or IPv7 ISO 8473. It manages the connectionless transfer of data one hop at a time, from end system to ingress router, router to router, and from egress router to destination end system. It is not responsible for reliable delivery to a next hop, but only for the detection of erroneous packets so they may be discarded. In this scheme, IPv4 and IPv6 would have to be classed with X.25 as subnet access protocols because they carry interface addresses rather than node addresses.

A number of layer management protocols, a function defined in the Management Annex, ISO 7498/4, belong to the Network Layer. These include routing protocols, multicast group management, Network Layer information and error, and Network Layer address assignment. It is the function of the payload that makes these belong to the Network Layer, not the protocol that carries them.

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 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 Transport layer also provides the acknowledgement of the successful data transmission and sends the next data if no errors occurred.

Although not developed under the OSI Reference Model and not strictly conforming to the OSI definition of the Transport Layer, typical examples of Layer 4 are the Transmission Control Protocol (TCP) and User Datagram Protocol (UDP).

Of the actual OSI protocols, there are five classes of connection-mode transport protocols ranging from class 0 (which is also known as TP0 and provides the least features) to class 4 (TP4, designed for less reliable networks, similar to the Internet). Class 0 contains no error recovery, and was designed for use on network layers that provide error-free connections. Class 4 is closest to TCP, although TCP contains functions, such as the graceful close, which OSI assigns to the Session Layer. Also, all OSI TP connection-mode protocol classes provide expedited data and preservation of record boundaries, of both of which TCP is incapable. Detailed characteristics of TP0-4 classes are shown in the following table:

Feature Name	TP0	TP1	TP2	TP3	TP4
Connection oriented network	Yes	Yes	Yes	Yes	Yes
Connectionless network	No	No	No	No	Yes
Concatenation and separation	No	Yes	Yes	Yes	Yes
Segmentation and reassembly	Yes	Yes	Yes	Yes	Yes
Error Recovery	No	Yes	Yes	Yes	Yes

Reinitiate connection (if an excessive number of PDUs are unacknowledged)	No	Yes	No	Yes	No
Multiplexing and demultiplexing over a single virtual circuit	No	No	Yes	Yes	Yes
Explicit flow control	No	No	Yes	Yes	Yes
Retransmission on timeout	No	No	No	No	Yes
Reliable Transport Service	No	Yes	No	Yes	Yes

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 packet.

Layer 5: Session Layer

The Session Layer controls the dialogues (connections) between computers. It establishes, manages and terminates the connections between the local and remote application. It provides for full-duplex, half-duplex, or simplex 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 the Transmission Control Protocol, and also for session checkpointing and recovery, which is not usually used in the Internet Protocol Suite. The Session Layer is commonly implemented explicitly in application environments that use remote procedure calls.

Layer 6: Presentation Layer

The Presentation Layer establishes context between Application Layer entities, in which the higher-layer entities may use different syntax and semantics if the presentation service provides a mapping between them. If a mapping is available, presentation service data units are encapsulated into session protocol data units, and passed down the stack.

This layer provides independence from data representation (e.g., encryption) by translating between application and network formats. The presentation layer transforms data into the form that the application accepts. This layer formats and encrypts data to be sent across a network. It is sometimes called the syntax layer.

The original presentation structure used the basic encoding rules of Abstract Syntax Notation One (ASN.1), with capabilities such as converting an EBCDIC-coded text file to an ASCII-coded file, or serialization of objects and other data structures from and to XML.

Layer 7: Application Layer

The Application Layer is the OSI layer closest to the end user, which means that both the OSI application layer and the user interact directly with the software application. This layer interacts with software applications that implement a communicating component. Such application programs fall outside the scope of the OSI model. Application layer functions typically include identifying communication partners, determining resource availability, and synchronizing communication. When identifying communication partners, the application layer determines the identity and availability of communication partners for an application with data to transmit. When determining resource availability, the application layer must decide whether sufficient network or the requested communication exist. In synchronizing communication, all communication between applications requires cooperation that is managed by the application layer. Some examples of application layer implementations also include:

- On OSI stack:
  - FTAM File Transfer and Access Management Protocol
  - X.400 Mail
  - Common management information protocol (CMIP)
- On TCP/IP stack:
  - Hypertext Transfer Protocol (HTTP),
  - File Transfer Protocol (FTP),
  - Simple Mail Transfer Protocol (SMTP)
  - Simple Network Management Protocol (SNMP).

Cross-layer functions

There are some functions or services that are not tied to a given layer, but they can affect more than one layer. Examples are

- security service (telecommunication) as defined by ITU-T X.800 Recommendation.
- management functions, i.e functions that permit to configure, instantiate, monitor, terminate the communications of two or more entities: there is a specific application layer protocol Common management information protocol (CMIP) and its corresponding service common management information service (CMIS), they need to interact with every layer in order to deal with their instances.
- MPLS operates at an OSI Model layer that is generally considered to lie between traditional definitions of Layer 2 (Data Link Layer) and Layer 3 (Network Layer), and thus is often referred to as a "Layer 2.5" protocol. It was

designed to provide a unified data-carrying service for both circuit-based clients and packet-switching clients which provide a datagram service model. It can be used to carry many different kinds of traffic, including IP packets, as well as native ATM, SONET, and Ethernet frames.

- ARP is used to translate IPv4 addresses (OSI Layer 3) into Ethernet MAC addresses (OSI Layer 2)

Interfaces

Neither the OSI Reference Model nor OSI protocols specify any programming interfaces, other than as deliberately abstract service specifications. Protocol specifications precisely define the interfaces between different computers, but the software interfaces inside computers are implementation-specific.

For example Microsoft Windows' Winsock, and Unix's Berkeley sockets and System V Transport Layer Interface, are interfaces between applications (Layer 5 and above) and the transport (Layer 4). NDIS and ODI are interfaces between the media (Layer 2) and the network protocol (Layer 3).

Interface standards, except for the Physical Layer to media, are approximate implementations of OSI Service Specifications.

Examples

Layer		OSI protocols	TCP/IP protocols	Signaling System 7	AppleTalk	IPX	SNA	UMTS	Misc. examples
#	**Name**								
7	Application	FTAM, X.400, X.500, DAP, ROSE, RTSE, ACSE CMIP	NNTP, SIP, SSI, DNS, FTP, Gopher, HTTP, NFS, NTP, DHCP, SMPP, SMTP, SNMP, Telnet, RIP, BGP	INAP, MAP, TCAP, ISUP, TUP	AFP, ZIP, RTMP, NBP	RIP, SAP	APPC		HL7, Modbus
6	Presentation	ISO/IEC 8823, X.226, ISO/IEC 9576-1, X.236	MIME, SSL, TLS, XDR		AFP				TDI, ASCII, EBCDIC, MIDI, MPEG
5	Session	ISO/IEC 8327, X.225, ISO/IEC 9548-1, X.235	Sockets. Session establishment in TCP, RTP		ASP, ADSP, PAP	NWLink	DLC?		Named pipes, NetBIOS, SAP, half duplex, full duplex, simplex, RPC

4 Transport	ISO/IEC 8073, TP0, TP1, TP2, TP3, TP4 (X.224), ISO/IEC 8602, X.234	TCP, UDP, SCTP, DCCP			DDP, SPX			NBF
3 Network	ISO/IEC 8208, X.25 (PLP), ISO/IEC 8878, X.223, ISO/IEC 8473-1, CLNP X.233.	IP, IPsec, ICMP, IGMP, OSPF	SCCP, MTP	ATP (TokenTalk or EtherTalk)	IPX		RRC (Radio Resource Control) Packet Data Convergence Protocol (PDCP) and BMC (Broadcast/ Multicast Control)	NBF, Q.931, IS-IS Leaky bucket, token bucket
2 Data Link	ISO/IEC 7666, X.25 (LAPB), Token Bus, X.222, ISO/IEC 8802-2 LLC Type 1 and 2	PPP, SLIP, PPTP, L2TP	MTP, Q.710	LocalTalk, AppleTalk Remote Access, PPP	IEEE 802.3 framing, Ethernet II framing	SDLC	LLC (Logical Link Control), MAC (Media Access Control)	802.3 (Ethernet), 802.11a/b/g/n MAC/LLC, 802.1Q (VLAN), ATM, HDP, FDDI, Fibre Channel, Frame Relay, HDLC, ISL, PPP, Q.921, Token Ring, CDP, ARP (maps layer 3 to layer 2 address), ITU-T G.hn DLL CRC, Bit stuffing, ARQ, Data Over Cable Service Interface Specification (DOCSIS)
1 Physical	X.25 (X.21bis, EIA/TIA-232, EIA/TIA-449, EIA-530, G.703)		MTP, Q.710	RS-232, RS-422, STP, PhoneNet		Twinax	UMTS Physical Layer or L1	RS-232, Full duplex, RJ45, V.35, V.34, I.430, I.431, T1, E1, 10BASE-T, 100BASE-TX, POTS, SONET, SDH, DSL, 802.11a/b/g/n PHY, ITU-T G.hn PHY,

Controller Area Network, Data Over Cable Service Interface Specification (DOCSIS)

Comparison with TCP/IP

In the TCP/IP model of the Internet, protocols are deliberately not as rigidly designed into strict layers as the OSI model. RFC 3439 contains a section entitled "Layering considered harmful." However, TCP/IP does recognize four broad layers of functionality which are derived from the operating scope of their contained protocols, namely the scope of the software application, the end-to-end transport connection, the internetworking range, and lastly the scope of the direct links to other nodes on the local network.

Even though the concept is different from the OSI model, these layers are nevertheless often compared with the OSI layering scheme in the following way: The Internet Application Layer includes the OSI Application Layer, Presentation Layer, and most of the Session Layer. Its end-to-end Transport Layer includes the graceful close function of the OSI Session Layer as well as the OSI Transport Layer. The internetworking layer (Internet Layer) is a subset of the OSI Network Layer, while the Link Layer includes the OSI Data Link and Physical Layers, as well as parts of OSI's Network Layer. 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.

The presumably strict peer layering of the OSI model as it is usually described does not present contradictions in TCP/IP, as it is permissible that protocol usage does not follow the hierarchy implied in a layered model. Such examples exist in some routing protocols (e.g., OSPF), or in the description of tunneling protocols, which provide a Link Layer for an application, although the tunnel host protocol may well be a Transport or even an Application Layer protocol in its own right.

Chapter 4

Physical Layer

The **Physical Layer** is the first and lowest layer in the seven-layer OSI model of computer networking. The implementation of this layer is often termed **PHY**.

The Physical Layer consists of the basic hardware transmission technologies of a network. It is a fundamental layer underlying the logical data structures of the higher level functions in a network. Due to the plethora of available hardware technologies with widely varying characteristics, this is perhaps the most complex layer in the OSI architecture.

The Physical Layer defines the means of transmitting raw bits rather than logical data packets over a physical link connecting network nodes. The bit stream may be grouped into code words or symbols and converted to a physical signal that is transmitted over a hardware transmission medium. The Physical Layer provides an electrical, mechanical, and procedural interface to the transmission medium. The shapes and properties of the electrical connectors, the frequencies to broadcast on, the modulation scheme to use and similar low-level parameters, are specified here.

Within the semantics of the OSI network architecture, the Physical Layer translates logical communications requests from the Data Link Layer into hardware-specific operations to affect transmission or reception of electronic signals.

Physical signaling sublayer

In a local area network (LAN) or a metropolitan area network (MAN) using open systems interconnection (OSI) architecture, the *physical signaling sublayer* is the portion of the Physical Layer that:

- interfaces with the Data Link Layer's medium access control (MAC) sublayer.
- performs character encoding, transmission, reception and decoding.
- performs mandatory isolation functions.

List of services

The major functions and services performed by the Physical Layer are:

- Bit-by-bit or symbol-by-symbol delivery
- Providing a standardized interface to physical transmission media, including
    - Mechanical specification of electrical connectors and cables, for example maximum cable length
    - Electrical specification of transmission line signal level and impedance
    - Radio interface, including electromagnetic spectrum frequency allocation and specification of signal strength, analog bandwidth, etc.
    - Specifications for IR over optical fiber or a wireless IR communication link
- Modulation
- Line coding
- Bit synchronization in synchronous serial communication
- Start-stop signalling and flow control in asynchronous serial communication
- Circuit switching
- Multiplexing
    - Establishment and termination of circuit switched connections
- Carrier sense and collision detection utilized by some level 2 multiple access protocols
- Equalization filtering, training sequences, pulse shaping and other signal processing of physical signals
- Forward error correction for example bitwise convolutional coding
- Bit-interleaving and other channel coding

The Physical Layer is also concerned with

- Bit rate
- Point-to-point, multipoint or point-to-multipoint line configuration
- Physical network topology, for example bus, ring, mesh or star network
- Serial or parallel communication
- Simplex, half duplex or full duplex transmission mode
- Autonegotiation

List of protocols

- Telephone network modems- V.92
- IRDA Physical Layer
- USB Physical Layer
- EIA RS-232, EIA-422, EIA-423, RS-449, RS-485
- Ethernet physical layer Including 10BASE-T, 10BASE2, 10BASE5, 100BASE-TX, 100BASE-FX, 100BASE-T, 1000BASE-T, 1000BASE-SX and other varieties
- Varieties of 802.11 Wi-Fi Physical Layers
- DSL
- ISDN
- T1 and other T-carrier links, and E1 and other E-carrier links
- SONET/SDH

- Optical Transport Network (OTN)
- GSM Um radio interface physical layer
- Bluetooth Physical Layer
- ITU Recommendations
- Firewire
- TransferJet Physical Layer
- Etherloop
- ARINC 818 Avionics Digital Video Bus
- G.hn/G.9960 Physical Layer
- Controller Area Network (CAN) Physical Layer

Hardware equipment (network node) examples

- Network adapter
- Repeater
- Network hub
- Modem
- Fiber Media Converter

Relation to TCP/IP model

The TCP/IP model, defined in RFC 1122 and RFC 1123, is a high-level networking description used for the Internet and similar networks. It does not define an equivalent layer that deals exclusively with hardware-level specifications and interfaces, as this model does not concern itself directly with physical interfaces. Several RFCs mention a physical layer and data link layer, but that is in context of IEEE protocols. RFC 1122 and 1123 do not mention any physical layer functionality or physical layer standards.

Chapter 5

Data Link Layer

The **Data Link Layer** is Layer 2 of the seven-layer OSI model of computer networking. It corresponds to, or is part of the link layer of the TCP/IP reference model.

The Data Link Layer is the protocol layer which transfers data between adjacent network nodes in a wide area network or between nodes on the same local area network segment. The Data Link Layer provides the functional and procedural means to transfer data between network entities and might provide the means to detect and possibly correct errors that may occur in the Physical Layer. Examples of data link protocols are Ethernet for local area networks (multi-node), the Point-to-Point Protocol (PPP), HDLC and ADCCP for point-to-point (dual-node) connections.

The Data Link Layer is concerned with local delivery of frames between devices on the same LAN. Data Link frames, as these protocol data units are called, do not cross the boundaries of a local network. Inter-network routing and global addressing are higher layer functions, allowing Data Link protocols to focus on local delivery, addressing, and media arbitration. In this way, the Data Link layer is analogous to a neighborhood traffic cop; it endeavors to arbitrate between parties contending for access to a medium.

When devices attempt to use a medium simultaneously, frame collisions occur. Data Link protocols specify how devices detect and recover from such collisions, and may provide mechanisms to reduce or prevent them.

Delivery of frames by layer 2 devices is affected through the use of unambiguous hardware addresses. A frame's header contains source and destination addresses that indicate which device originated the frame and which device is expected to receive and process it. In contrast to the hierarchical and routable addresses of the network layer, layer 2 addresses are flat, meaning that no part of the address can be used to identify the logical or physical group to which the address belongs.

The data link thus provides data transfer across the physical link. That transfer can be reliable or unreliable; many data link protocols do not have acknowledgments of successful frame reception and acceptance, and some data link protocols might not even have any form of checksum to check for transmission errors. In those cases, higher-level protocols must provide flow control, error checking, and acknowledgments and retransmission.

In some networks, such as IEEE 802 local area networks, the Data Link Layer is described in more detail with Media Access Control (MAC) and Logical Link Control (LLC) sublayers; this means that the IEEE 802.2 LLC protocol can be used with all of the IEEE 802 MAC layers, such as Ethernet, token ring, IEEE 802.11, etc., as well as with some non-802 MAC layers such as FDDI. Other Data Link Layer protocols, such as HDLC, are specified to include both sublayers, although some other protocols, such as Cisco HDLC, use HDLC's low-level framing as a MAC layer in combination with a different LLC layer. In the ITU-T G.hn standard, which provides a way to create a high-speed (up to 1 Gigabit/s) Local area network using existing home wiring (power lines, phone lines and coaxial cables), the Data Link Layer is divided into three sub-layers (Application Protocol Convergence, Logical Link Control and Medium Access Control).

Within the semantics of the OSI network architecture, the Data Link Layer protocols respond to service requests from the Network Layer and they perform their function by issuing service requests to the Physical Layer.

Sublayers of the Data Link Layer

Logical Link Control sublayer

The uppermost sublayer is *Logical Link Control* (LLC). This sublayer multiplexes protocols running atop the Data Link Layer, and optionally provides flow control, acknowledgment, and error notification. The LLC provides addressing and control of the data link. It specifies which mechanisms are to be used for addressing stations over the transmission medium and for controlling the data exchanged between the originator and recipient machines.

Media Access Control sublayer

The sublayer below it is *Media Access Control* (MAC). Sometimes this refers to the sublayer that determines who is allowed to access the media at any one time (usually CSMA/CD). Other times it refers to a frame structure with MAC addresses inside.

There are generally two forms of media access control: distributed and centralized. Both of these may be compared to communication between people. In a network made up of people speaking, i.e. a conversation, we look for clues from our fellow talkers to see if any of them appear to be about to speak. If two people speak at the same time, they will back off and begin a long and elaborate game of saying "no, you first".

The Media Access Control sublayer also determines where one frame of data ends and the next one starts -- frame synchronization. There are four means of frame synchronization: time based, character counting, byte stuffing and bit stuffing.

- The *time based* approach simply puts a specified amount of time between frames. The major drawback of this is that new gaps can be introduced or old gaps can be lost due to external influences.
- *Character counting* simply notes the count of remaining characters in the frame's header. This method, however, is easily disturbed if this field gets faulty in some way, thus making it hard to keep up synchronization.
- *Byte stuffing* precedes the frame with a special byte sequence such as DLE STX and succeeds it with DLE ETX. Appearances of DLE (byte value 0x10) has to be escaped with another DLE. The start and stop marks are detected at the receiver and removed as well as the inserted DLE characters.
- Similarly, *bit stuffing* replaces these start and end marks with flag consisting of a special bit pattern (e.g. a 0, six 1 bits and a 0). Occurrences of this bit pattern in the data to be transmitted is avoided by inserting a bit. To use the example where the flag is 01111110, a 0 is inserted after 5 consecutive 1's in the data stream. The flags and the inserted 0's are removed at the receiving end. This makes for arbitrary long frames and easy synchronization for the recipient. Note that this stuffed bit is added even if the following data bit is 0, which could not be mistaken for a sync sequence, so that the receiver can unambiguously distinguish stuffed bits from normal bits.

List of Data Link Layer services

- Encapsulation of network layer data packets into frames
- Frame synchronization
- Logical link control (LLC) sublayer:
    - Error control (automatic repeat request,ARQ), in addition to ARQ provided by some Transport layer protocols, to forward error correction (FEC) techniques provided on the Physical Layer, and to error-detection and packet canceling provided at all layers, including the network layer. Data link layer error control (i.e. retransmission of erroneous packets) is provided in wireless networks and V.42 telephone network modems, but not in LAN protocols such as Ethernet, since bit errors are so uncommon in short wires. In that case, only error detection and canceling of erroneous packets are provided.
    - Flow control, in addition to the one provided on the Transport layer. Data link layer error control is not used in LAN protocols such as Ethernet, but in modems and wireless networks.

- Media access control (MAC) sublayer:
    - Multiple access protocols for channel-access control, for example CSMA/CD protocols for collision detection and retransmission in

Ethernet bus networks and hub networks, or the CSMA/CA protocol for collision avoidance in wireless networks.

  - Physical addressing (MAC addressing)
  - LAN switching (packet switching) including MAC filtering and spanning tree protocol
  - Data packet queueing or scheduling
  - Store-and-forward switching or cut-through switching
  - Quality of Service (QoS) control
  - Virtual LANs (VLAN)

Protocol examples

- Address Resolution Protocol (ARP)
- ARCnet
- ATM
- Cisco Discovery Protocol (CDP)
- Controller Area Network (CAN)
- Econet
- Ethernet
- Ethernet Automatic Protection Switching (EAPS)
- Fiber Distributed Data Interface (FDDI)
- Frame Relay
- High-Level Data Link Control (HDLC)
- IEEE 802.2 (provides LLC functions to IEEE 802 MAC layers)
- IEEE 802.11 wireless LAN
- Link Access Procedures, D channel (LAPD)
- LocalTalk
- Multiprotocol Label Switching (MPLS)
- Point-to-Point Protocol (PPP)
- Serial Line Internet Protocol (SLIP) (obsolete)
- Spanning tree protocol
- StarLan
- Token ring
- Unidirectional Link Detection (UDLD)
- and most forms of serial communication.

Interfaces

The Data Link Layer is often implemented in software as a "network card driver". The operating system will have a defined software interface between the data link and the network transport stack above. This interface is not a layer itself, but rather a definition for interfacing between layers.

Relation to TCP/IP model

In the frame work of the TCP/IP (Internet Protocol Suite) model, OSI's Data Link Layer, in addition to other components, is contained in TCP/IP's lowest layer, the Link Layer. The Internet Protocol's Link Layer only concerns itself with hardware issues to the point of obtaining hardware addresses for locating hosts on a physical network link and transmitting data frames onto the link. Thus, the Link Layer is broader in scope and encompasses all methods that affect the local link, which is the group of connections that are limited in scope to other nodes on the local access network.

The TCP/IP model is not a top/down comprehensive design reference for networks. It was formulated for the purpose of illustrating the logical groups and scopes of functions needed in the design of the suite of internetworking protocols of TCP/IP, as needed for the operation of the Internet. In general, direct or strict comparisons of the OSI and TCP/IP models should be avoided, because the layering in TCP/IP is not a principal design criterion and in general considered to be "harmful" (RFC 3439). In particular, TCP/IP does not dictate a strict hierarchical sequence of encapsulation requirements, as is attributed to OSI protocols.

Chapter 6

Transport Layer

In computer networking, the **Transport Layer** provides end-to-end communication services for applications within a layered architecture of network components and protocols. The transport layer provides convenient services such as connection-oriented data stream support, reliability, flow control, and multiplexing.

Transport layers are contained in both the TCP/IP model (RFC 1122), which is the foundation of the Internet, and the Open Systems Interconnection (OSI) model of general networking. The definitions of the Transport Layer are slightly different in these two models.

The most well-known transport protocol is the Transmission Control Protocol (TCP). It lent its name to the title of the entire Internet Protocol Suite, *TCP/IP*. It is used for connection-oriented transmissions, whereas the connectionless User Datagram Protocol (UDP) is used for simpler messaging transmissions. TCP is the more complex protocol, due to its stateful design incorporating reliable transmission and data stream services. Other prominent protocols in this group are the Datagram Congestion Control Protocol (DCCP) and the Stream Control Transmission Protocol (SCTP).

Services

There are many services that can be optionally provided by a Transport Layer protocol, and different protocols may or may not implement them.

- Connection-oriented communication: Interpreting the connection as a data stream can provide many benefits to applications. It is normally easier to deal with than the underlying connection-less models, such as the Transmission Control Protocol's underlying Internet Protocol model of datagrams.
- Byte orientation: Rather than processing the messages in the underlying communication system format, it is often easier for an application to process the data stream as a sequence of bytes. This simplification helps applications work with various underlying message formats.
- Same order delivery: The Network layer doesn't generally guarantee that packets of data will arrive in the same order that they were sent, but often this is a desirable feature. This is usually done through the use of segment

numbering, with the receiver passing them to the application in order. This can cause head-of-line blocking.

- Reliability: Packets may be lost during transport due to network congestion and errors. By means of an error detection code, such as a checksum, the transport protocol may check that the data is not corrupted, and verify correct receipt by sending an ACK or NACK message to the sender. Automatic repeat request schemes may be used to retransmit lost or corrupted data.
- Flow control: The rate of data transmission between two nodes must sometimes be managed to prevent a fast sender from transmitting more data than can be supported by the receiving data buffer, causing a buffer overrun. This can also be used to improve efficiency by reducing buffer underrun.
- Congestion avoidance: Congestion control can control traffic entry into a telecommunications network, so as to avoid congestive collapse by attempting to avoid oversubscription of any of the processing or link capabilities of the intermediate nodes and networks and taking resource reducing steps, such as reducing the rate of sending packets. For example, automatic repeat requests may keep the network in a congested state; this situation can be avoided by adding congestion avoidance to the flow control, including slow-start. This keeps the bandwidth consumption at a low level in the beginning of the transmission, or after packet retransmission.
- Multiplexing: Ports can provide multiple endpoints on a single node. For example, the name on a postal address is a kind of multiplexing, and distinguishes between different recipients of the same location. Computer applications will each listen for information on their own ports, which enables the use of more than one network service at the same time. It is part of the Transport Layer in the TCP/IP model, but of the Session Layer in the OSI model.

Analysis

The Transport Layer is responsible for delivering data to the appropriate application process on the host computers. This involves statistical multiplexing of data from different application processes, i.e. forming data packets, and adding source and destination port numbers in the header of each Transport Layer data packet. Together with the source and destination IP address, the port numbers constitutes a network socket, i.e. an identification address of the process-to-process communication. In the OSI model, this function is supported by the Session Layer.

Some Transport Layer protocols, for example TCP, but not UDP, support virtual circuits, i.e. provide connection oriented communication over an underlying packet oriented datagram network. A byte-stream is delivered while hiding the packet mode communication for the application processes. This involves connection establishment, dividing of the data stream into packets called segments, segment numbering and reordering of out-of order data.

Finally, some Transport Layer protocols, for example TCP, but not UDP, provide end-to-end reliable communication, i.e. error recovery by means of error detecting code and automatic repeat request (ARQ) protocol. The ARQ protocol also provides flow control, which may be combined with congestion avoidance.

UDP is a very simple protocol, and does not provide virtual circuits, nor reliable communication, delegating these functions to the application program. UDP packets are called datagrams, rather than segments.

TCP is used for many protocols, including HTTP web browsing and email transfer. UDP may be used for multicasting and broadcasting, since retransmissions are not possible to a large amount of hosts. UDP typically gives higher throughput and shorter latency, and is therefore often used for real-time multimedia communication where packet loss occasionally can be accepted, for example IP-TV and IP-telephony, and for online computer games.

In many non-IP-based networks, for example X.25, Frame Relay and ATM, the connection oriented communication is implemented at network layer or data link layer rather than the Transport Layer. In X.25, in telephone network modems and in wireless communication systems, reliable node-to-node communication is implemented at lower protocol layers.

The OSI model defines five classes of transport protocols: *TP0*, providing the least error recovery, to *TP4*, which is designed for less reliable networks.

Protocols

The exact definition of what qualifies as a transport layer protocol is not firm. The following is a short list:

- ATP, AppleTalk Transaction Protocol
- CUDP, Cyclic UDP
- DCCP, Datagram Congestion Control Protocol
- FCP, Fiber Channel Protocol
- IL, IL Protocol
- NBF, NetBIOS Frames protocol
- RDP, Reliable Datagram Protocol
- SCTP, Stream Control Transmission Protocol
- SPX, Sequenced Packet Exchange
- SST, Structured Stream Transport
- TCP, Transmission Control Protocol
- UDP, User Datagram Protocol
- UDP Lite
- µTP, Micro Transport Protocol

Comparison of Transport Layer protocols

Feature Name	UDP	UDP Lite	TCP	SCTP	DCCP	RUDP
Packet header size	8 Bytes	8 Bytes	20-60 Bytes	12 Bytes + Variable Chunk Header	12 or 16 bytes	
Transport Layer packet entity	Datagram	Datagram	Segment	Datagram	Datagram	Datagram
Connection oriented	No	No	Yes	Yes	Yes	No
Reliable transport	No	No	Yes	Yes	No	Yes
Unreliable transport	Yes	Yes	No	Yes	Yes	Yes
Preserve message boundary	Yes	Yes	No	Yes	Yes	Unsure
Ordered delivery	No	No	Yes	Yes	No	No
Unordered delivery	Yes	Yes	No	Yes	Yes	Yes
Data checksum	Optional	Yes	Yes	Yes	Yes	Unsure
Checksum size (bits)	16	16	16	32	16	Unsure
Partial checksum	No	Yes	No	No	Yes	No
Path MTU	No	No	Yes	Yes	Yes	Unsure
Flow control	No	No	Yes	Yes	No	
Congestion control	No	No	Yes	Yes	Yes	Unsure
ECN support	No	No	Yes	Yes	Yes	
Multiple streams	No	No	No	Yes	No	No

Multi-homing support	No	No	No	Yes	No	No
Bundling / Nagle	No	No	Yes	Yes	No	Unsure
NAT friendly	Yes		Yes	No	Yes	

Comparison of OSI transport protocols

The OSI model defines five classes of connection-mode transport protocols designated class 0 (TP0) to class 4 (TP4). Class 0 contains no error recovery, and was designed for use on network layers that provide error-free connections. Class 4 is closest to TCP, although TCP contains functions, such as the graceful close, which OSI assigns to the Session Layer. All OSI connection-mode protocol classes provide expedited data and preservation of record boundaries. Detailed characteristics of the classes are shown in the following table:

Service	**TP0**	**TP1**	**TP2**	**TP3**	**TP4**
Connection oriented network	Yes	Yes	Yes	Yes	Yes
Connectionless network	No	No	No	No	Yes
Concatenation and separation	No	Yes	Yes	Yes	Yes
Segmentation and reassembly	Yes	Yes	Yes	Yes	Yes
Error Recovery	No	Yes	No	Yes	Yes
Reinitiate connection (if an excessive number of PDUs are unacknowledged)	No	Yes	No	Yes	No
multiplexing and demultiplexing over a single virtual circuit	No	No	Yes	Yes	Yes
Explicit flow control	No	No	Yes	Yes	Yes
Retransmission on timeout	No	No	No	No	Yes
Reliable Transport Service	No	Yes	No	Yes	Yes

Chapter 7

Session Layer, Application Layer & Presentation Layer

Session Layer

The **Session Layer** is Layer 5 of the seven-layer OSI model of computer networking.

The Session Layer provides the mechanism for opening, closing and managing a session between end-user application processes, i.e. a semi-permanent dialogue. Communication sessions consist of requests and responses that occur between applications. Session Layer services are commonly used in application environments that make use of remote procedure calls (RPCs).

An example of a Session Layer protocol is the OSI protocol suite Session Layer Protocol, also known as X.225 or ISO 8327. In case of a connection loss this protocol may try to recover the connection. If a connection is not used for a long period, the Session Layer Protocol may close it and re-open it. It provides for either full duplex or half-duplex operation and provides synchronization points in the stream of exchanged messages.

Other examples of Session Layer implementations include Zone Information Protocol (ZIP) – the AppleTalk protocol that coordinates the name binding process, and Session Control Protocol (SCP) – the DECnet Phase IV Session Layer protocol.

Within the service layering semantics of the OSI network architecture, the Session Layer responds to service requests from the Presentation Layer and issues service requests to the Transport Layer.

Services

- Authentication
- Permissions
- Session restoration (checkpointing and recovery)

The Session Layer of the OSI model is responsible for session checkpointing and recovery. It allows information of different streams, perhaps originating from different sources, to be properly combined or synchronized.

An example application is web conferencing, in which the streams of audio and video must be synchronous to avoid so-called lip synch problems. Floor control ensures that the person displayed on screen is the current speaker.

Another application is in live TV programs, where streams of audio and video need to be seamlessly merged and transitioned from one to the other to avoid silent airtime or excessive overlap.

Protocols

- ADSP, AppleTalk Data Stream Protocol
- ASP, AppleTalk Session Protocol
- H.245, Call Control Protocol for Multimedia Communication
- ISO-SP, OSI Session Layer Protocol (X.225, ISO 8327)
- iSNS, Internet Storage Name Service
- L2F, Layer 2 Forwarding Protocol
- L2TP, Layer 2 Tunneling Protocol
- NetBIOS, Network Basic Input Output System
- PAP, Password Authentication Protocol
- PPTP, Point-to-Point Tunneling Protocol
- RPC, Remote Procedure Call Protocol
- RTCP, Real-time Transport Control Protocol
- SMPP, Short Message Peer-to-Peer
- SCP, Session Control Protocol
- ZIP, Zone Information Protocol
- SDP, Sockets Direct Protocol

Comparison with TCP/IP model

The TCP/IP reference model does not concern itself with the OSI model's details of application or transport protocol semantics and therefore does not consider a Session Layer. OSI's session management in connection with the typical transport protocols (TCP, SCTP), is contained in the Transport Layer protocols, or otherwise considered the realm of the Application Layer protocols. TCP/IP's layers are *descriptions* of operating scopes (application, host-to-host, network, link) and not detailed *prescriptions* of operating procedures or data semantics.

Application Layer

The Internet Protocol Suite (TCP/IP) and the Open Systems Interconnection model (OSI model) of computer networking each specify a group of protocols and methods identified by the name **Application Layer**.

In TCP/IP, the Application Layer contains all protocols and methods that fall into the realm of process-to-process communications across an Internet Protocol (IP) network. Application Layer methods use the underlying Transport Layer protocols to establish host-to-host connections.

In the OSI model, the definition of its Application Layer is narrower in scope, explicitly distinguishing additional functionality above the Transport Layer at two additional levels, the Session Layer and the Presentation Layer. OSI specifies strict modular separation of functionality at these layers and provides protocol implementations for each layer.

TCP/IP protocols

The following protocols are explicitly mentioned in RFC 1123 (1989), describing the Application Layer of the Internet Protocol Suite.

- Remote Login category
    - Telnet
- File Transfer category
    - FTP
    - TFTP
- Electronic Mail category
    - SMTP
- Support Services category
    - DNS
    - RARP
    - BOOTP
    - SNMP
    - CMOT

Other protocol examples

- 9P, Plan 9 from Bell Labs distributed file system protocol
- AFP,
- APPC, Advanced Program-to-Program Communication
- AMQP, Advanced Message Queuing Protocol
- BitTorrent
- Atom Publishing Protocol
- BOOTP, Bootstrap Protocol

- CFDP, Coherent File Distribution Protocol
- DDS, Data Distribution Service
- DHCP, Dynamic Host Configuration Protocol
- DeviceNet
- DNS, Domain Name System (Service) Protocol
- eDonkey
- ENRP, Endpoint Handlespace Redundancy Protocol
- FastTrack (KaZaa, Grokster, iMesh)
- Finger, User Information Protocol
- Freenet
- FTAM, File Transfer Access and Management
- FTP, File Transfer Protocol
- Gopher, Gopher protocol
- HL7, Health Level Seven
- HTTP, HyperText Transfer Protocol
- H.323, Packet-Based Multimedia Communications System
- IMAP, IMAP4, Internet Message Access Protocol (version 4)

- IRCP, Internet Relay Chat Protocol
- Kademlia
- LDAP, Lightweight Directory Access Protocol
- LPD, Line Printer Daemon Protocol
- MIME (S-MIME), Multipurpose Internet Mail Extensions and Secure MIME
- Modbus
- Netconf
- NFS, Network File System
- NIS, Network Information Service
- NNTP, Network News Transfer Protocol
- NTCIP, National Transportation Communications for Intelligent Transportation System Protocol
- NTP, Network Time Protocol
- OSCAR, AOL Instant Messenger Protocol
- PNRP, Peer Name Resolution Protocol
- POP, POP3, Post Office Protocol (version 3)
- RDP, Remote Desktop Protocol
- Rlogin, Remote Login in UNIX Systems
- RPC, Remote Procedure Call
- RTMP Real Time Messaging Protocol
- RTP, Real-time Transport Protocol
- RTPS, Real Time Publish Subscribe
- RTSP, Real Time Streaming Protocol
- SAP, Session Announcement Protocol
- SDP, Session Description Protocol
- SIP, Session Initiation Protocol
- SLP, Service Location Protocol
- SMB, Server Message Block

- SMTP, Simple Mail Transfer Protocol
- SNMP, Simple Network Management Protocol
- SNTP, Simple Network Time Protocol
- SPTP, Secure Parallel Transfer Protocol
- SSH, Secure Shell
- SSMS, Secure SMS Messaging Protocol
- TCAP, Transaction Capabilities Application Part
- TDS, Tabular Data Stream
- TELNET, Terminal Emulation Protocol of TCP/IP
- TFTP, Trivial File Transfer Protocol
- TSP, Time Stamp Protocol
- VTP, Virtual Terminal Protocol
- Waka, an HTTP replacement protocol
- Whois (and RWhois), Remote Directory Access Protocol
- WebDAV
- X.400, Message Handling Service Protocol
- X.500, Directory Access Protocol (DAP)
- XMPP, Extensible Messaging and Presence Protocol

Presentation Layer

The **Presentation Layer** is Layer 6 of the seven-layer OSI model of computer networking.

The Presentation Layer is responsible for the delivery and formatting of information to the application layer for further processing or display. It relieves the application layer of concern regarding syntactical differences in data representation within the end-user systems. *Note:* An example of a presentation service would be the conversion of an EBCDIC-coded text file to an ASCII-coded file.

The Presentation Layer is the lowest layer at which application programmers consider data structure and presentation, instead of simply sending data in form of datagrams or packets between hosts. This layer deals with issues of string representation - whether they use the Pascal method (an integer length field followed by the specified amount of bytes) or the C/C++ method (null-terminated strings, i.e. "`thisisastring\0`"). The idea is that the application layer should be able to point at the data to be moved, and the Presentation Layer will deal with the rest.

Serialization of complex data structures into flat byte-strings (using mechanisms such as TLV or XML) can be thought of as the key functionality of the Presentation Layer.

Encryption is typically done at this level too, although it can be done on the Application, Session, Transport, or Network Layers; each having its own advantages and disadvantages. Another example is representing structure, which is normally standardized at this level, often by using XML. As well as simple pieces of data, like strings, more complicated things are standardized in this layer. Two common examples are 'objects' in object-oriented programming, and the exact way that streaming video is transmitted.

In many widely used applications and protocols, no distinction is made between the presentation and application layers. For example, HTTP, generally regarded as an application layer protocol, has Presentation Layer aspects such as the ability to identify character encoding for proper conversion, which is then done in the Application Layer.

Within the service layering semantics of the OSI network architecture, the Presentation Layer responds to service requests from the Application Layer and issues service requests to the Session Layer.

Services

- Encryption
- Compression

Sublayers

The Presentation Layer is composed of two sublayers:

- **CASE** (Common Application Service Element)
- **SASE** (Specific Application Service Element)

CASE

The **CASE sublayer** provides services for the Application Layer and request services from the Session Layer. It provides support for common application services, such as:

- ACSE (Association Control Service Element)
- ROSE (Remote Operation Service Element)
- CCR (Commitment Concurrency and Recovery)
- RTSE (Reliable Transfer Service Element)

SASE

The **SASE sublayer** provides application specific services (protocols), such as

- FTAM (File Transfer, Access and Manager)
- VT (Virtual Terminal)

- MOTIS (Message Oriented Text Interchange Standard)
- CMIP (Common Management Information Protocol)
- JTM (Job Transfer and Manipulation) a former OSI standard
- MMS (Manufacturing Messaging Service)
- RDA (Remote Database Access)
- DTP (Distributed Transaction Processing)
- Tel Net(a remote terminal access protocol)

Protocols

- AFP, Apple Filing Protocol
- ASCII, American Standard Code for Information Interchange
- EBCDIC, Extended Binary Coded Decimal Interchange Code
- ICA, Independent Computing Architecture, the Citrix system core protocol
- LPP, Lightweight Presentation Protocol
- NCP, NetWare Core Protocol
- NDR, Network Data Representation
- XDR, eXternal Data Representation
- X.25 PAD, Packet Assembler/Disassembler Protocol

Chapter 8

Spanning Tree Protocol

The **Spanning Tree Protocol** (**STP**) is a network protocol that ensures a loop-free topology for any bridged Ethernet local area network. The basic function of STP is to prevent bridge loops and ensuing broadcast radiation. Spanning tree also allows a network design to include spare (redundant) links to provide automatic backup paths if an active link fails, without the danger of bridge loops, or the need for manual enabling/disabling of these backup links.

STP is a Data Link Layer protocol. It is standardized as IEEE 802.1D. As the name suggests, it creates a spanning tree within a mesh network of connected layer-2 bridges (typically Ethernet switches), and disables those links that are not part of the spanning tree, leaving a single active path between any two network nodes.

STP is based on an algorithm invented by Radia Perlman while working for Digital Equipment Corporation.

Protocol operation

The collection of bridges in a local area network (LAN) can be considered a graph whose nodes are bridges and LAN segments (or cables), and whose edges are the interfaces connecting the bridges to the segments. To break loops in the LAN while maintaining access to all LAN segments, the bridges collectively compute a spanning tree. The spanning tree is not necessarily a minimum cost spanning tree. A network administrator can reduce the cost of a spanning tree, if necessary, by altering some of the configuration parameters in such a way as to affect the choice of the root of the spanning tree. The spanning tree that the bridges compute using the Spanning Tree Protocol can be determined using the following rules. The example network at the right, below, will be used to illustrate the rules.

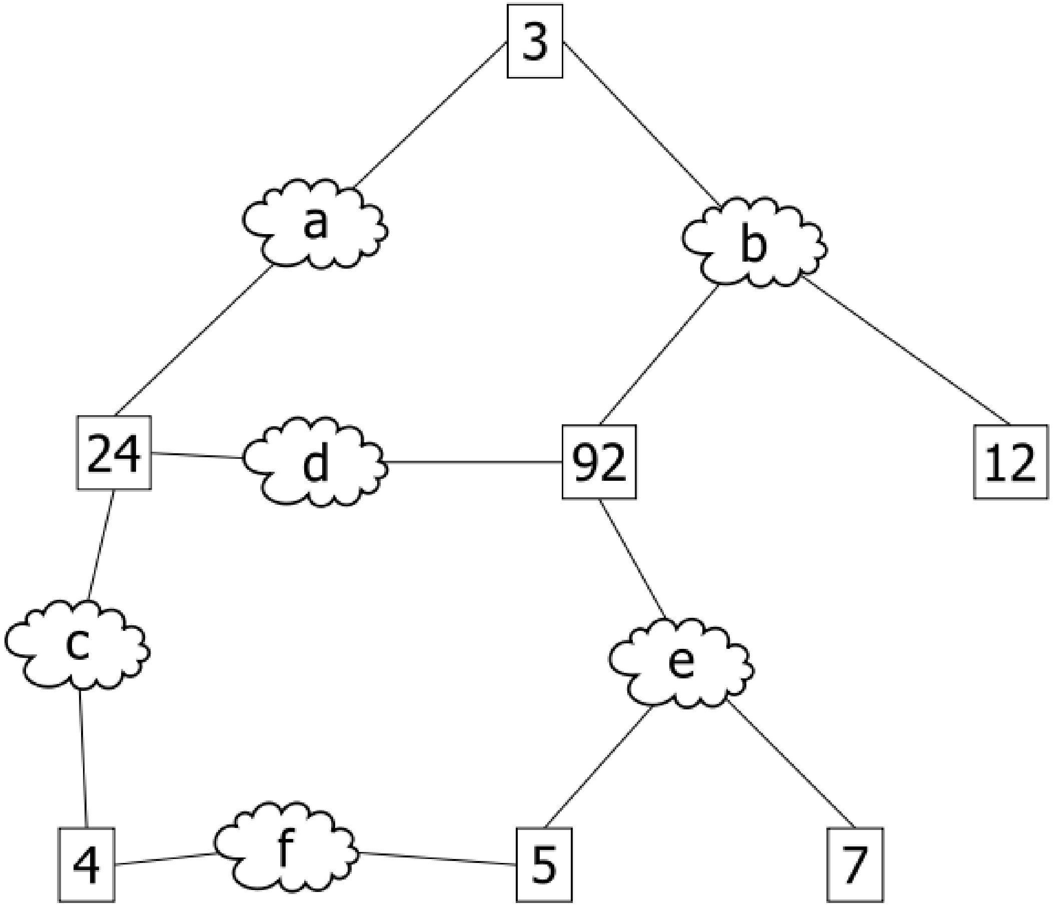

1. An example network. The numbered boxes represent bridges (the number represents the bridge ID). The lettered clouds represent network segments.

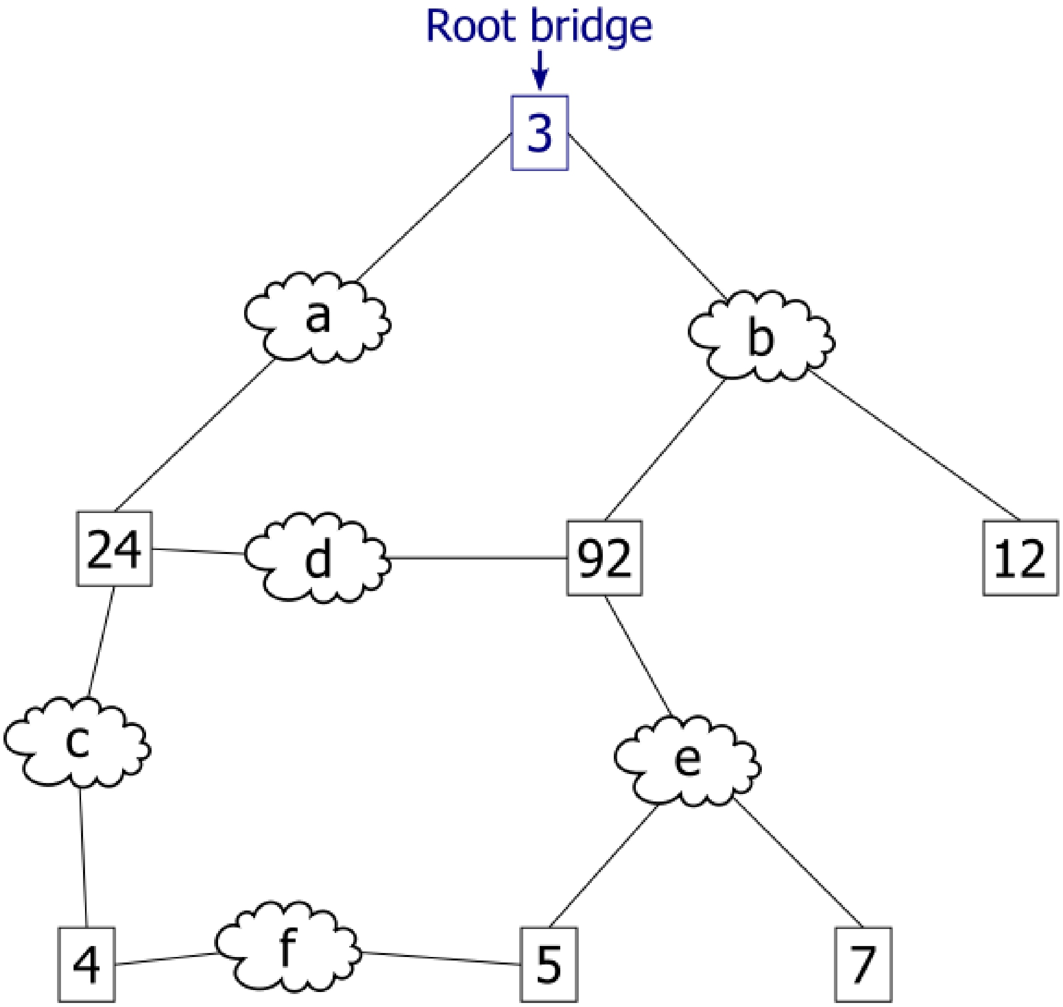

2. The smallest bridge ID is 3. Therefore, bridge 3 is the root bridge

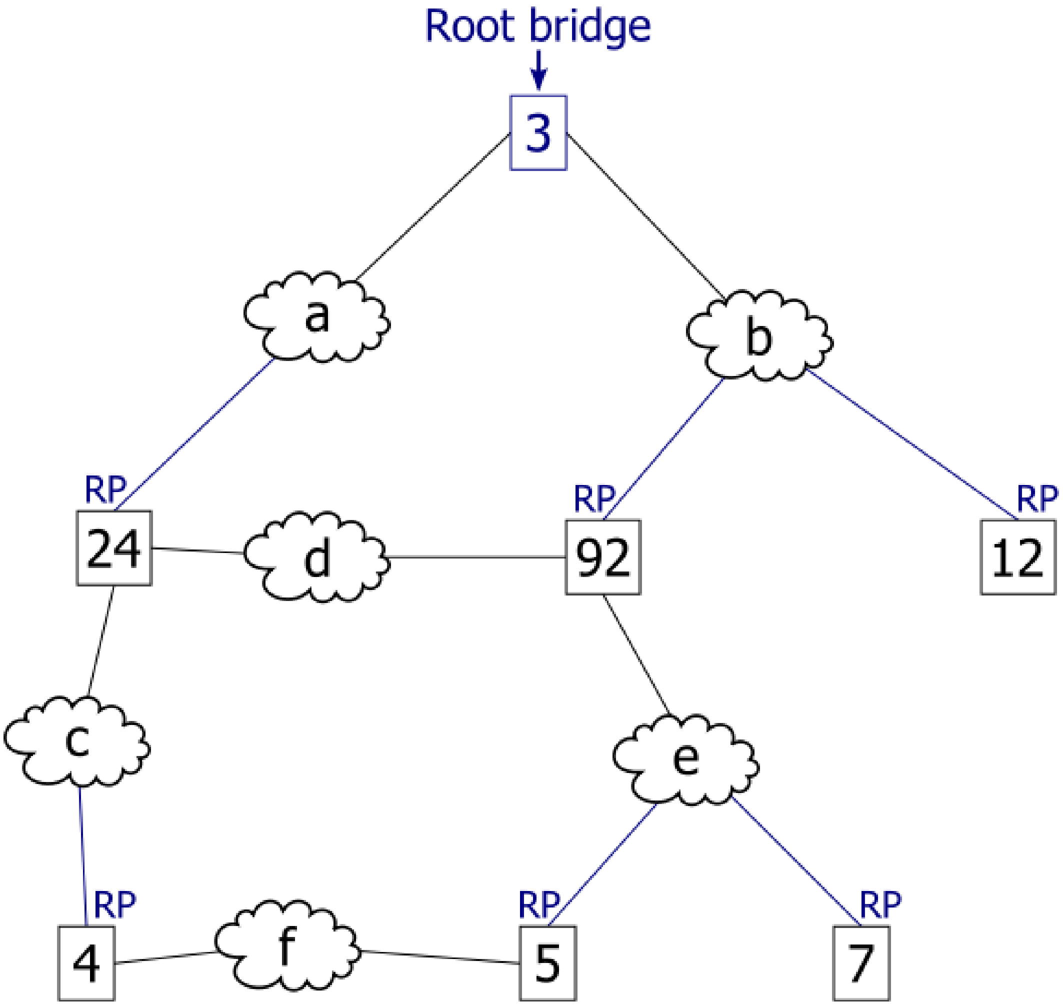

3. Assuming that the cost of traversing any network segment is 1, the least cost path from bridge 4 to the root bridge goes through network segment c. Therefore, the root port for bridge 4 is the one on network segment c.

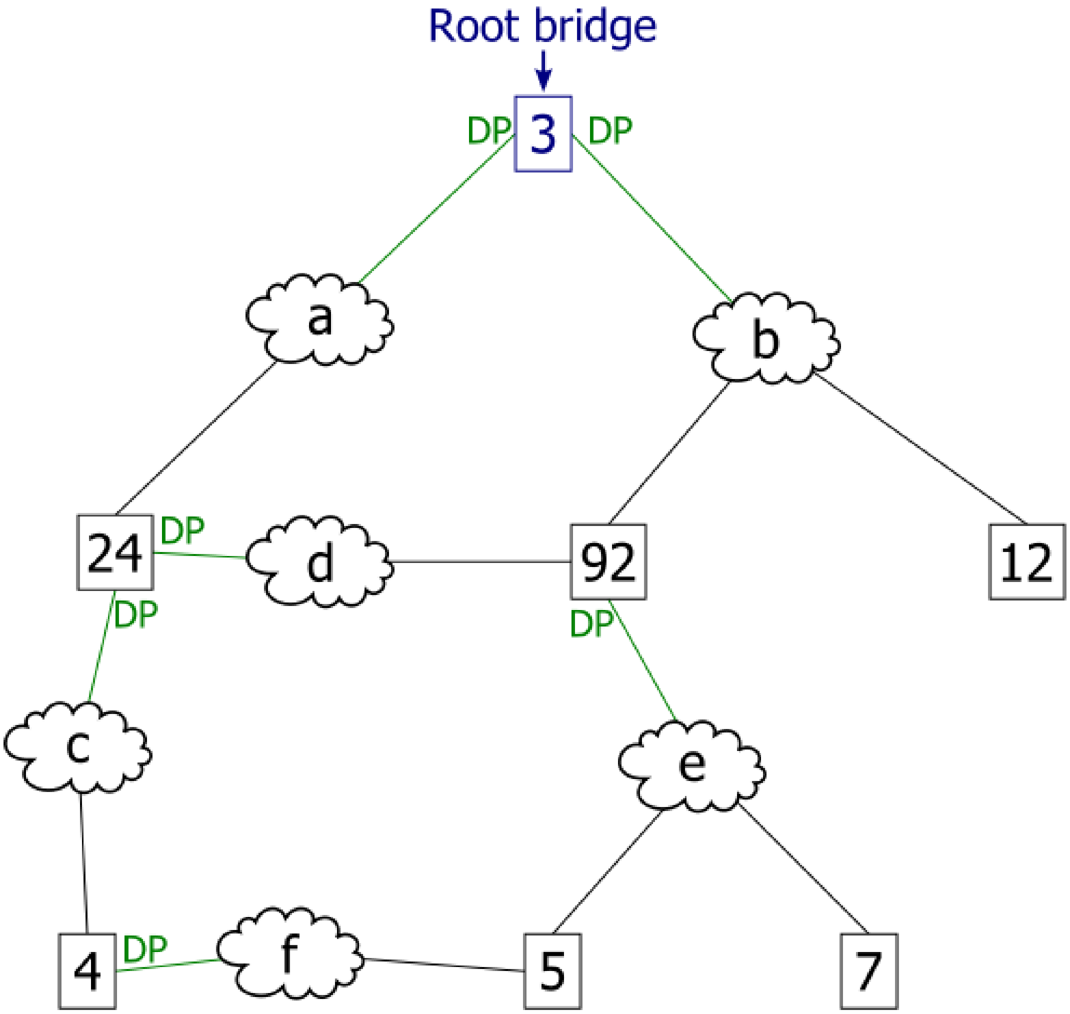

4. The least cost path to the root from network segment e goes through bridge 92. Therefore the designated port for network segment e is the port that connects bridge 92 to network segment e.

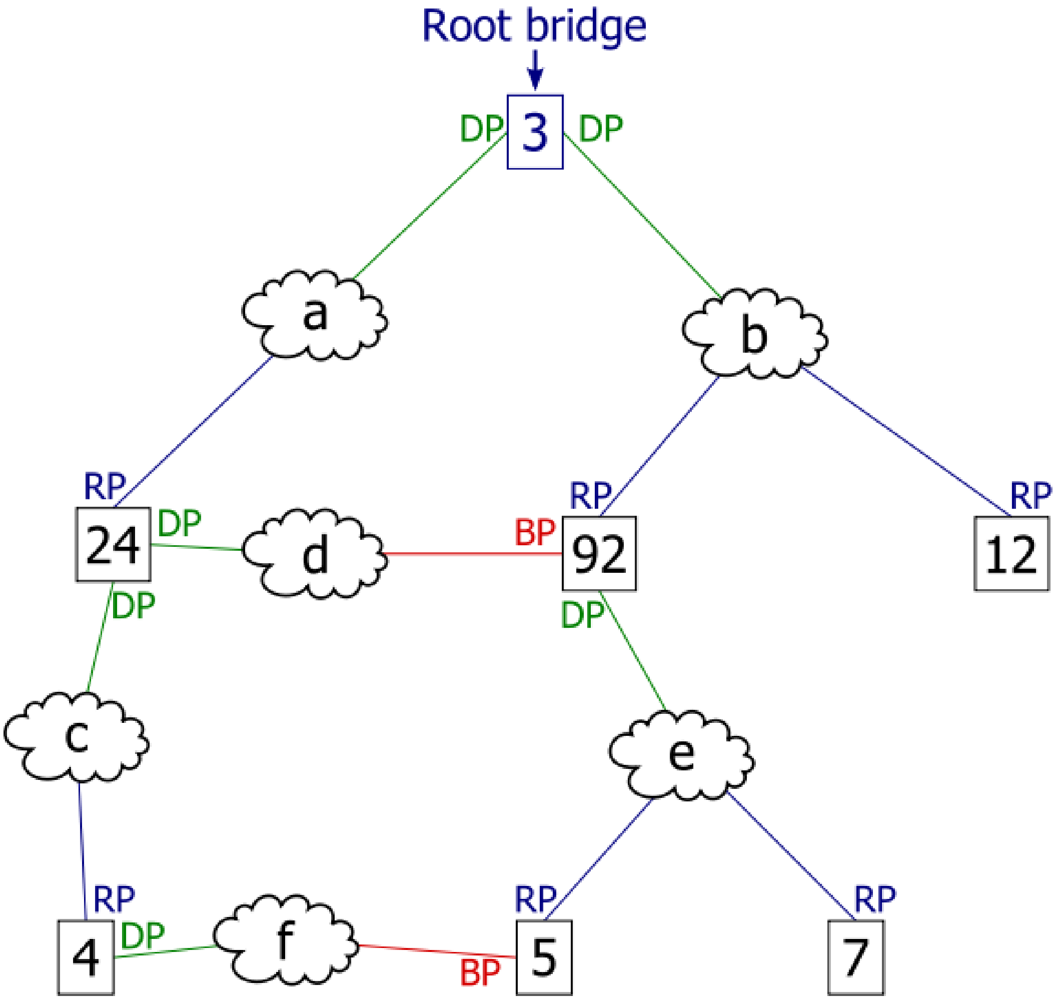

5. This diagram illustrates all port states as computed by the spanning tree algorithm. Any active port that is not a root port or a designated port is a blocked port.

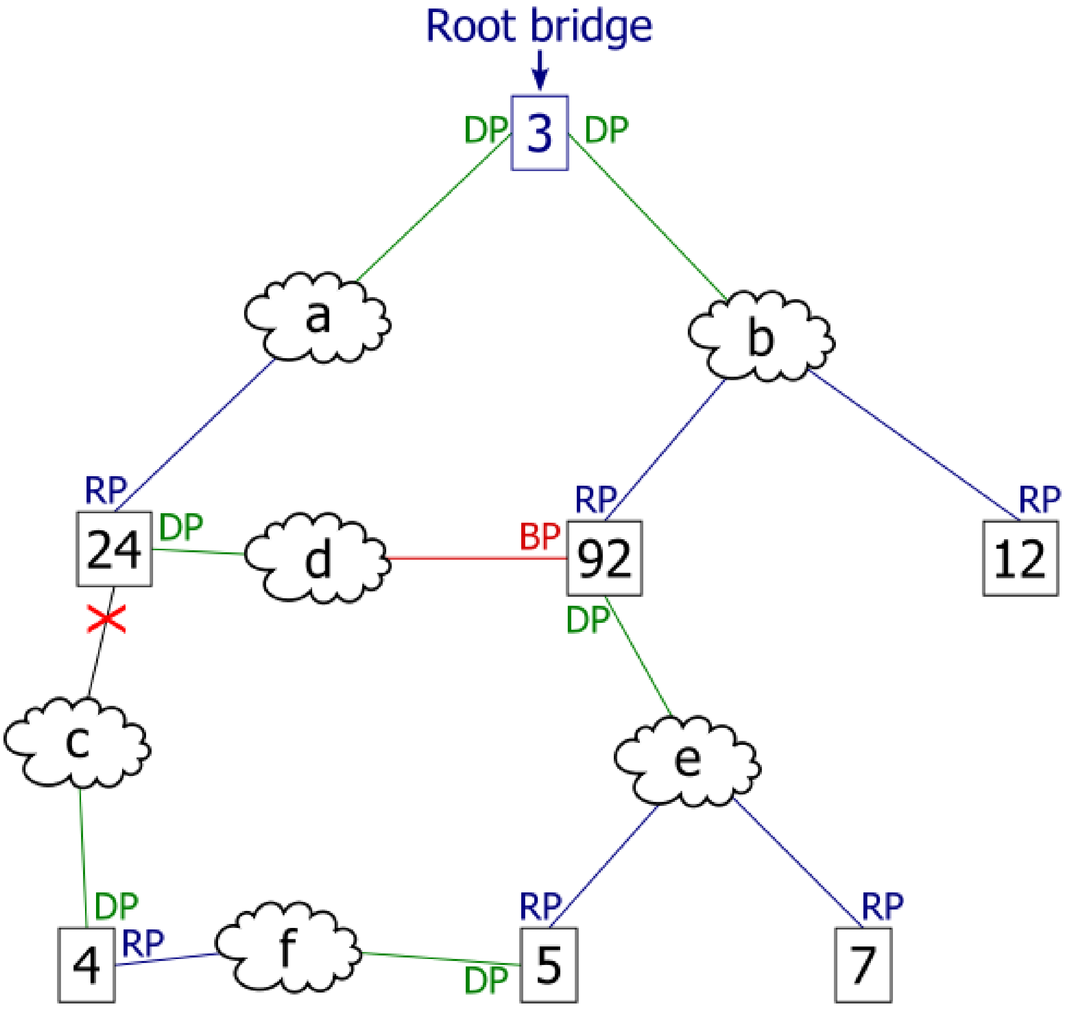

6. After link failure the spanning tree algorithm computes and spans new least-cost tree.

Select a root bridge. The *root bridge* of the spanning tree is the bridge with the smallest (lowest) bridge ID. Each bridge has a unique identifier (ID) and a configurable priority number; the bridge ID contains both numbers. To compare two bridge IDs, the priority is compared first. If two bridges have equal priority, then the MAC addresses are compared. For example, if switches A (MAC=0200.0000.1111) and B (MAC=0200.0000.2222) both have a priority of 10, then switch A will be selected as the root bridge. If the network administrators would like switch B to become the root bridge, they must set its priority to be less than 10.

Determine the least cost paths to the root bridge. The computed spanning tree has the property that messages from any connected device to the root bridge traverse a least cost path, i.e., a path from the device to the root that has minimum cost among all paths from the device to the root. The cost of traversing a path is the sum of the costs of the segments on the path. Different technologies have different default costs for network segments. An administrator can configure the cost of traversing a particular network segment. The property that messages always traverse least-cost paths to the root is guaranteed by the following two rules.

Least cost path from each bridge. After the root bridge has been chosen, each bridge determines the cost of each possible path from itself to the root. From these, it picks one with the smallest cost (a least-cost path). The port connecting to that path becomes the *root port* (RP) of the bridge.

Least cost path from each network segment. The bridges on a network segment collectively determine which bridge has the least-cost path from the network segment to the root. The port connecting this bridge to the network segment is then the *designated port* (DP) for the segment.

Disable all other root paths. Any active port that is not a root port or a designated port is a *blocked port* (BP).

Modifications in case of ties. The above rules over-simplify the situation slightly, because it is possible that there are ties, for example, two or more ports on a single bridge are attached to least-cost paths to the root or two or more bridges on the same network segment have equal least-cost paths to the root. To break such ties:

Breaking ties for root ports. When multiple paths from a bridge are least-cost paths, the chosen path uses the neighbor bridge with the lower bridge ID. The root port is thus the one connecting to the bridge with the lowest bridge ID. For example, in figure 3, if switch 4 were connected to network segment d, there would be two paths of length 2 to the root, one path going through bridge 24 and the other through bridge 92. Because there are two least cost paths, the lower bridge ID (24) would be used as the tie-breaker in choosing which path to use.

Breaking ties for designated ports. When more than one bridge on a segment leads to a least-cost path to the root, the bridge with the lower bridge ID is used to forward messages to the root. The port attaching that bridge to the network segment is the *designated port* for the segment. In figure 4, there are two least cost paths from network segment d to the root, one going through bridge 24 and the other through bridge 92. The lower bridge ID is 24, so the tie breaker dictates that the designated port is the port through which network segment d is connected to bridge 24. If bridge IDs were equal, then the bridge with the lowest MAC address would have the designated port. In either case, the loser sets the port as being blocked.

The final tie-breaker. In some cases, there may still be a tie, as when two bridges are connected by multiple cables. In this case, multiple ports on a single bridge are candidates for root port. In this case, the path which passes through the port on the neighbor bridge that has the lowest port priority is used.

Data rate and STP path cost

The table below shows the default cost of an interface for a given data rate.

Data rate	STP Cost (802.1D-1998)	STP Cost (802.1t-2001)
4 Mbit/s	250	5,000,000
10 Mbit/s	100	2,000,000
16 Mbit/s	62	1,250,000
100 Mbit/s	19	200,000
1 Gbit/s	4	20,000
2 Gbit/s	3	10,000
10 Gbit/s	2	2,000

Bridge Protocol Data Units (BPDUs)

The above rules describe one way of determining what spanning tree will be computed by the algorithm, but the rules as written require knowledge of the entire network. The bridges have to determine the root bridge and compute the port roles (root, designated, or blocked) with only the information that they have. To ensure that each bridge has enough information, the bridges use special data frames called **Bridge Protocol Data Units** (BPDUs) to exchange information about bridge IDs and root path costs.

A bridge sends a BPDU frame using the unique MAC address of the port itself as a source address, and a destination address of the STP multicast address 01:80:C2:00:00:00.

There are three types of BPDUs:

- Configuration BPDU (CBPDU), used for Spanning Tree computation
- Topology Change Notification (TCN) BPDU, used to announce changes in the network topology
- Topology Change Notification Acknowledgment (TCA)

BPDUs are exchanged regularly (every 2 seconds by default) and enable switches to keep track of network changes and to start and stop forwarding at ports as required.

When a device is first attached to a switch port, it will not immediately start to forward data. It will instead go through a number of states while it processes BPDUs and determines the topology of the network. When a host is attached such as a computer, printer or server the port will always go into the forwarding state, albeit after a delay of about 30 seconds while it goes through the listening and learning states. The time spent in the listening and learning states is determined by a value known as the forward delay (default 15 seconds and set by the root bridge). However, if instead another *switch* is connected, the port may remain in blocking mode if it is determined that it would cause a loop in the network. Topology Change Notification (TCN) BPDUs are used to inform other switches of port changes. TCNs are injected into the network by a non-root switch and propagated to the root. Upon receipt of the TCN, the root switch will set a Topology Change flag in its normal BPDUs. This flag is propagated to all other switches to instruct them to rapidly age out their forwarding table entries....

STP switch port states:

- **Blocking** - A port that would cause a switching loop, no user data is sent or received but it may go into forwarding mode if the other links in use were to fail and the spanning tree algorithm determines the port may transition to the forwarding state. BPDU data is still received in blocking state.
- **Listening** - The switch processes BPDUs and awaits possible new information that would cause it to return to the blocking state.
- **Learning** - While the port does not yet forward frames (packets) it does learn source addresses from frames received and adds them to the filtering database (switching database)
- **Forwarding** - A port receiving and sending data, normal operation. STP still monitors incoming BPDUs that would indicate it should return to the blocking state to prevent a loop.
- **Disabled** - Not strictly part of STP, a network administrator can manually disable a port

To prevent the delay when connecting hosts to a switch and during some topology changes, Rapid STP was developed and standardized by IEEE 802.1w, which allows a switch port to rapidly transition into the forwarding state during these situations.

BPDU fields

The bridge ID, or BID, is a field inside a BPDU packet. It is eight bytes in length. The first two bytes are the Bridge Priority, an unsigned integer of 0-65,535. The last six bytes are a MAC address supplied by the switch. In the event that MAC Address Reduction is used, the first two bytes are used differently. The first four bits are a configurable priority, and the last twelve bits carry either the VLAN ID or MSTP instance number.

Evolutions and extensions

The first spanning tree protocol was invented in 1985 at the Digital Equipment Corporation by Radia Perlman. In 1990, the IEEE published the first standard for the protocol as 802.1D, based on the algorithm designed by Perlman. Subsequent versions were published in 1998 and 2004, incorporating various extensions.

Although the purpose of a standard is to promote interworking of equipment from different vendors, different implementations of a standard are not guaranteed to work, due for example to differences in default timer settings. The IEEE encourages vendors to provide a "Protocol Implementation Conformance Statement", declaring which capabilities and options have been implemented, to help users determine whether different implementations will interwork correctly.

Also, the original Perlman-inspired Spanning Tree Protocol, called DEC STP, is not a standard and differs from the IEEE version in message format as well as timer settings. Some bridges implement both the IEEE and the DEC versions of the Spanning Tree Protocol, but their interworking can create issues for the network administrator, as illustrated by the problem discussed in an on-line Cisco document.

Rapid Spanning Tree Protocol (RSTP)

In 2001, the IEEE with document 802.1w introduced an evolution of the Spanning Tree Protocol: Rapid Spanning Tree Protocol (RSTP), which provides for faster spanning tree convergence after a topology change. Standard IEEE 802.1D-2004 now incorporates RSTP and obsoletes STP. While STP can take 30 to 50 seconds to respond to a topology change, RSTP is typically able to respond to changes within 3*Hello times (default: 6 seconds) or within a few milliseconds of a physical link failure. The so-called Hello time is an important and configurable time interval that is used by RSTP for several purposes; its default value is 2 seconds.

RSTP bridge port roles:

- **Root** - A forwarding port that is the best port from Nonroot-bridge to Rootbridge
- **Designated** - A forwarding port for every LAN segment
- **Alternate** - An alternate path to the root bridge. This path is different than using the root port.
- **Backup** - A backup/redundant path to a segment where another bridge port already connects.
- **Disabled** - Not strictly part of STP, a network administrator can manually disable a port

RSTP is a refinement of STP and therefore shares most of its basic operation characteristics. However there are some notable differences as summarized below:

- Detection of root switch failure is done in 3 hello times, which is 6 seconds if default hello times have not been changed.
- Ports may be configured as edge ports if they are attached to a LAN that has no other bridges attached. These edge ports transition directly to the forwarding state. RSTP still continues to monitor the port for BPDUs in case a bridge is connected. RSTP can also be configured to automatically detect edge

ports. As soon as the bridge detects a BPDU coming to an edge port, the port becomes a non-edge port.

- Unlike in STP, RSTP will respond to BPDUs sent from the direction of the root bridge. An RSTP bridge will "propose" its spanning tree information to its designated ports. If another RSTP bridge receives this information and determines this is the superior root information, it sets all its other ports to discarding. The bridge may send an "agreement" to the first bridge confirming its superior spanning tree information. The first bridge, upon receiving this agreement, knows it can rapidly transition that port to the forwarding state bypassing the traditional listening/learning state transition. This essentially creates a cascading effect away from the root bridge where each designated bridge proposes to its neighbors to determine if it can make a rapid transition. This is one of the major elements that allows RSTP to achieve faster convergence times than STP.
- As discussed in the port role details above, RSTP maintains backup details regarding the discarding status of ports. This avoids timeouts if the current forwarding ports were to fail or BPDUs were not received on the root port in a certain interval.

Per-VLAN Spanning Tree (PVST)

In Ethernet switched environments where multiple Virtual LANs exist, spanning tree can be deployed per Virtual LAN. Cisco's name for this is *per VLAN spanning tree* (PVST and PVST+, which is the default protocol used by Cisco switches). Both PVST and PVST+ protocols are Cisco proprietary protocols and they cannot be used on 3rd party switches, although Force10 Networks, Extreme Networks and Blade Network Technologies support PVST+, Extreme Networks does so with two limitations (lack of support on ports where the VLAN is untagged/native and also on the VLAN with ID 1). PVST works only with ISL (Cisco's proprietary protocol for VLAN encapsulation) due to its embedded Spanning tree ID. Due to high penetration of the IEEE 802.1Q VLAN trunking standard and PVST's dependence on ISL, Cisco defined a different PVST+ standard for 802.1Q encapsulation. PVST+ can tunnel across an MSTP Region.

Multiple Spanning Tree Protocol (MSTP)

The *Multiple Spanning Tree Protocol* (MSTP), originally defined in IEEE 802.1s and later merged into IEEE 802.1Q-2005, defines an extension to RSTP to further develop the usefulness of virtual LANs (VLANs). This "Per-VLAN" Multiple Spanning Tree Protocol configures a separate Spanning Tree for each VLAN group and blocks all but one of the possible alternate paths within each Spanning Tree.

If there is only one Virtual LAN (VLAN) in the network, single (traditional) STP works appropriately. If the network contains more than one VLAN, the logical network configured by single STP would work, but it is possible to make better use of the alternate paths available by using an alternate spanning tree for different VLANs or groups of VLANs.

MSTP allows formation of MST regions that can run multiple MST instances (MSTI). Multiple regions and other STP bridges are interconnected using one single common spanning tree (CST).

MSTP is similar to Cisco Systems' **Multiple Instances Spanning Tree Protocol** (MISTP), and is an evolution of the **Spanning Tree Protocol** and the Rapid Spanning Tree Protocol. It was introduced in IEEE 802.1s as an amendment to 802.1Q, 1998 edition. Standard IEEE 802.1Q-2005 now includes MSTP.

Unlike some proprietary per-VLAN spanning tree implementations, MSTP includes all of its spanning tree information in a single BPDU format. Not only does this reduce the number of BPDUs required on a LAN to communicate spanning tree information for each VLAN, but it also ensures backward compatibility with RSTP (and in effect, classic STP too). MSTP does this by encoding additional region information after the standard RSTP BPDU as well as a number of MSTI messages (from 0 to 64 instances, although in practice many bridges support less). Each of these MSTI configuration messages conveys the spanning tree information for each instance. Each instance can be assigned a number of configured VLANs and frames (packets) assigned to these VLANs operate in this spanning tree instance whenever they are inside the MST region. In order to avoid conveying their entire VLAN to spanning tree mapping in each BPDU, bridges encode an MD5 digest of their VLAN to instance table in the MSTP BPDU. This digest is then used by other MSTP bridges, along with other administratively configured values, to determine if the neighboring bridge is in the same MST region as itself.

MSTP is fully compatible with RSTP bridges, in that an MSTP BPDU can be interpreted by an RSTP bridge as an RSTP BPDU. This not only allows compatibility with RSTP bridges without configuration changes, but also causes any RSTP bridges outside of an MSTP region to see the region as a single RSTP bridge, regardless of the number of MSTP bridges inside the region itself. In order to further facilitate this view of an MST region as a single RSTP bridge, the MSTP protocol uses a variable known as remaining hops as a time to live counter instead of the message age timer used by RSTP. The message age time is only incremented once when spanning tree information enters an MST region, and therefore RSTP bridges will see a region as only one "hop" in the spanning tree. Ports at the edge of an MST region connected to either an RSTP or STP bridge or an endpoint are known as boundary ports. As in RSTP, these ports can be configured as edge ports to facilitate rapid changes to the forwarding state when connected to endpoints.

Rapid Per-VLAN Spanning Tree (R-PVST)

Cisco's proprietary protocol that combines the functionalities of RSTP and PVST. It is based on a per VLAN instance that creates a tree for each VLAN.

Chapter 9

Ethernet

A standard 8P8C (often called RJ45) connector used most commonly on cat5 cable, a type of cabling used primarily in Ethernet networks

Ethernet is a family of frame-based computer networking technologies for local area networks (LAN). It defines a number of wiring and signaling standards for the Physical Layer of the OSI networking model as well as a common addressing format and a variety of Medium Access Control procedures at the lower part of the Data Link Layer.

Ethernet is standardized as IEEE 802.3. The combination of the twisted pair versions of Ethernet for connecting end systems to the network, along with the fiber optic versions for site backbones, is the most widespread wired LAN technology. It has been used since around 1980 to the present, largely replacing competing LAN standards such as token ring, FDDI, and ARCNET.

History

Ethernet was developed at Xerox PARC between 1973 and 1975. It was inspired by ALOHAnet, which Robert Metcalfe had studied as part of his Ph.D. dissertation. In 1975, Xerox filed a patent application listing Metcalfe, David Boggs, Chuck Thacker and Butler Lampson as inventors. In 1976, after the system was deployed at PARC, Metcalfe and Boggs published a seminal paper.

Metcalfe left Xerox in 1979 to promote the use of personal computers and local area networks (LANs), forming 3Com. He convinced Digital Equipment Corporation (DEC), Intel, and Xerox to work together to promote Ethernet as a standard, the so-called "DIX" standard, for "Digital/Intel/Xerox"; it specified the 10 Mbit/s Ethernet, with 48-bit destination and source addresses and a global 16-bit Ethertype-type field and was first published on September 30, 1980 as "The Ethernet, A Local Area Network. Data Link Layer and Physical Layer Specifications". Version 2 of this document was published in November, 1982 and defines what has become known as **Ethernet II**. The Institute of Electrical and Electronics Engineers (IEEE) first published the 802.3 standard as a draft in 1983 and as a standard in 1985. Support of Ethernet's carrier sense multiple access with collision detection (CSMA/CD) in other standardization bodies (i.e., ECMA, IEC, and ISO) was instrumental in getting past delays of the finalization of the Ethernet standard due to the difficult decision processes in the IEEE, and due to the competitive Token Ring proposal strongly supported by IBM. Ethernet initially competed with two largely proprietary systems, Token Ring and Token Bus. These proprietary systems soon found themselves competing in a market inundated by Ethernet products. In the process, 3Com became a major company. 3Com shipped its first 10 Mbit/s Ethernet 3C100 transceiver in March 1981, and that year started selling adapters for DEC/PDP11 and VAXes, as well as Intel Multibus and Sun Microsystems machines. This was followed quickly by DEC's Unibus to Ethernet adapter, which DEC sold and used internally to build its own corporate network, which reached over 10,000 nodes by 1986; far and away the largest extant computer network in the world at that time.

Through the first half of the 1980s, DEC's Ethernet implementation, 10BASE5, used a coaxial cable 0.375 inches (9.5 mm) in diameter, later called "thick ethernet" or "thicknet" in contrast to its successor, 10BASE2, called "thin ethernet" or "thinnet". Thinnet uses a cable similar to cable television cable of the era. The emphasis was on making installation of the cable easier and less costly.

Shared cable Ethernet was always hard to install in offices because its bus topology was in conflict with the star topology cable plans designed into buildings for telephony. Modifying Ethernet to conform to twisted pair telephone wiring already installed in commercial buildings provided another opportunity to lower costs, expand the installed base, and leverage building design, and, thus, twisted-pair

Ethernet was the next logical development in the mid-1980s, beginning with StarLAN. UTP-based Ethernet became widely deployed with the 10BASE-T standard. This system replaced the coaxial cable systems with a system of full duplex switches linked via UTP.

With the advent of the 10BASE-T standard in 1990, Ethernet switches supplemented the half duplex CSMA/CD scheme with a full duplex system offering higher performance at a lower cost than routers. With the arrival of 100BASE-T, Ethernet switches capable of mixed speed and mixed duplex operation were built.

Standardization

Notwithstanding its technical merits, timely standardization was instrumental to the success of Ethernet. It required well-coordinated and partly competitive activities in several standardization bodies such as the IEEE, ECMA, IEC, and finally ISO.

In February 1980, IEEE started a project, IEEE 802, for the standardization of local area networks (LAN).

The "DIX-group" with Gary Robinson (DEC), Phil Arst (Intel), and Bob Printis (Xerox) submitted the so-called "Blue Book" CSMA/CD specification as a candidate for the LAN specification. Since IEEE membership is open to all professionals, including students, the group received countless comments on this brand-new technology.

In addition to CSMA/CD, Token Ring (supported by IBM) and Token Bus (selected and henceforward supported by General Motors) were also considered as candidates for a LAN standard. Due to the goal of IEEE 802 to forward only one standard and due to the strong company support for all three designs, the necessary agreement on a LAN standard was significantly delayed.

In the Ethernet camp, it put at risk the market introduction of the Xerox Star workstation and 3Com's Ethernet LAN products. With such business implications in mind, David Liddle (General Manager, Xerox Office Systems) and Metcalfe (3Com) strongly supported a proposal of Fritz Röscheisen (Siemens Private Networks) for an alliance in the emerging office communication market, including Siemens' support for the international standardization of Ethernet (April 10, 1981). Ingrid Fromm, Siemens representative to IEEE 802 quickly achieved broader support for Ethernet beyond IEEE by the establishment of a competing Task Group "Local Networks" within the European standards body ECMA TC24. As early as March 1982 ECMA TC24 with its corporate members reached agreement on a standard for CSMA/CD based on the IEEE 802 draft. The speedy action taken by ECMA decisively contributed to the conciliation of opinions within IEEE and approval of IEEE 802.3 CSMA/CD by the end of 1982.

Approval of Ethernet on the international level was achieved by a similar, cross-partisan action with Fromm as liaison officer working to integrate IEC TC83 and ISO TC97SC6, and the ISO/IEEE 802/3 standard was approved in 1984.

Evolution

Ethernet is an evolving technology. Evolutions have included higher bandwidth, improved media access control methods, and changes to the physical medium. Ethernet evolved into the complex networking technology that today underlies most LANs. The coaxial cable was replaced with point-to-point links connected by Ethernet repeaters or switches to reduce installation costs, increase reliability, and enable point-to-point management and troubleshooting. There are many variants of Ethernet in common use.

Ethernet stations communicate by sending each other data packets, blocks of data that are individually sent and delivered. As with other IEEE 802 LANs, each Ethernet station is given a 48-bit MAC address. The MAC addresses are used to specify both the destination and the source of each data packet. Network interface cards (NICs) or chips normally do not accept packets addressed to other Ethernet stations. Adapters come programmed with a globally unique address. Despite the significant changes in Ethernet from a thick coaxial cable bus running at 10 Mbit/s to point-to-point links running at 1 Gbit/s and beyond, all generations of Ethernet (excluding early experimental versions) use the same frame formats (and hence the same interface for higher layers), and can be readily interconnected through bridging.

Due to the ubiquity of Ethernet, the ever-decreasing cost of the hardware needed to support it, and the reduced panel space needed by twisted pair Ethernet, most manufacturers now build the functionality of an Ethernet card directly into PC motherboards, eliminating the need for installation of a separate network card.

Shared media

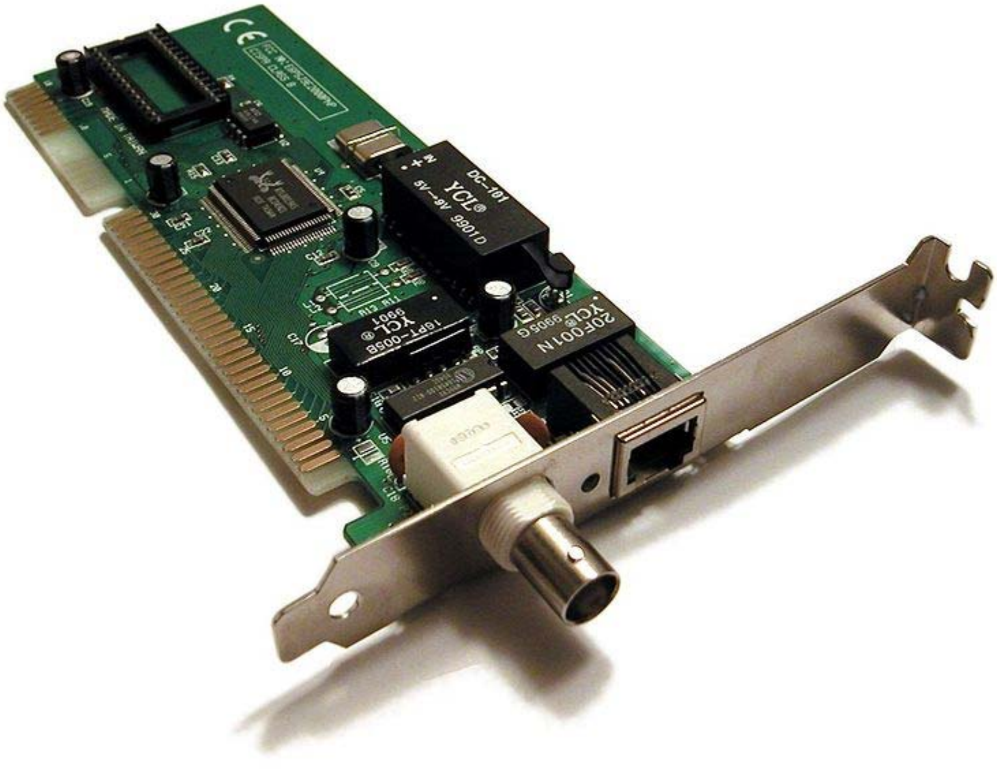

A 1990s network interface card supporting both coaxial cable-based 10BASE2 (BNC connector, left) and twisted pair-based 10BASE-T (8P8C connector, right).

Ethernet was originally based on the idea of computers communicating over a shared coaxial cable acting as a broadcast transmission medium. The methods used were similar to those used in radio systems, with the common cable providing the communication channel likened to the *Luminiferous Aether* in 19th century physics, and it was from this reference that the name "Ethernet" was derived.

Original Ethernet's shared coaxial cable (the shared medium) traversed a building or campus to every attached machine. A scheme known as carrier sense multiple access with collision detection (CSMA/CD) governed the way the computers shared the channel. This scheme was simpler than the competing token ring or token bus technologies. Computers were connected to an Attachment Unit Interface (AUI) transceiver, which was in turn connected to the cable (later with thin Ethernet the transceiver was integrated into the network adapter). While a simple passive wire was highly reliable for small networks, it was not reliable for large extended networks, where damage to the wire in a single place, or a single bad connector, could make the whole Ethernet segment unusable.

Since all communications happen on the same wire, any information sent by one computer is received by all, even if that information is intended for just one destination. The network interface card interrupts the CPU only when applicable packets are received: The card ignores information not addressed to it. Use of a single cable also means that the bandwidth is shared, so that network traffic can be very slow when many stations are simultaneously active.

Collisions reduce throughput by their very nature. In the worst case, when there are lots of hosts with long cables that attempt to transmit many short frames, excessive collisions can reduce throughput dramatically. However, a Xerox report in 1980 summarized the results of having 20 fast nodes attempting to transmit packets of various sizes as quickly as possible on the same Ethernet segment. The results showed that, even for the smallest Ethernet frames (64 Bytes), 90% throughput on the LAN was the norm. This is in comparison with token passing LANs (token ring, token bus), all of which suffer throughput degradation as each new node comes into the LAN, due to token waits. This report was controversial, as modeling showed that collision-based networks became unstable under loads as low as 40% of nominal capacity. Many early researchers failed to understand the subtleties of the CSMA/CD protocol and how important it was to get the details right, and were really modeling somewhat different networks (usually not as good as real Ethernet).

Repeaters and hubs

For signal degradation and timing reasons, coaxial Ethernet segments had a restricted size. Somewhat larger networks could be built by using an Ethernet repeater. Initial repeaters had only 2 ports, but they gave way to 4, 6, 8, and more ports. People recognized the advantages of cabling in a star topology, primarily that a fault in one of the legs affects operation of only the stations attached to that leg.

A twisted pair Cat-3 or Cat-5 cable is used to connect 10BASE-T Ethernet

Ethernet on unshielded twisted-pair cables (UTP), beginning with StarLAN and continuing with 10BASE-T, was designed for point-to-point links only, and all termination was built into the device. This changed repeaters from a specialist device used at the center of large networks to a device that every twisted pair-based network with more than two machines had to use. The tree structure that resulted from this made Ethernet networks more reliable by preventing faults with one peer or its associated cable from affecting other devices on the network.

Despite the physical star topology, repeater based Ethernet networks still use half-duplex and CSMA/CD, with only minimal activity by the repeater, primarily the Collision Enforcement signal, in dealing with packet collisions. Every packet is sent to every port on the repeater, so bandwidth and security problems are not addressed. The total throughput of the repeater is limited to that of a single link, and all links must operate at the same speed.

Bridging and switching

While repeaters could isolate some aspects of Ethernet segments, such as cable breakages, they still forwarded all traffic to all Ethernet devices. This created practical limits on how many machines could communicate on an Ethernet network. The entire network was one collision domain, and all hosts had to be able to detect collisions anywhere on the network. This limited the number of repeaters between the farthest nodes. Segments joined by repeaters had to all operate at the same speed, making phased-in upgrades impossible.

To alleviate these problems, bridging was created to communicate at the data link layer while isolating the physical layer. With bridging, only well-formed Ethernet packets are forwarded from one Ethernet segment to another; collisions and packet errors are isolated. Prior to discovery of network devices on the different segments, Ethernet bridges (and switches) work somewhat like Ethernet repeaters, passing all traffic between segments. However, as the bridge discovers the addresses associated with each port, it forwards network traffic only to the necessary segments, improving overall performance. Broadcast traffic is still forwarded to all network segments. Bridges also overcame the limits on total segments between two hosts and allowed the mixing of speeds, both of which became very important with the introduction of Fast Ethernet.

Early bridges examined each packet one by one using software on a CPU, and some of them were significantly slower than repeaters at forwarding traffic, especially when handling many ports at the same time. This was in part because the entire Ethernet packet would be read into a buffer, the destination address compared with an internal table of known MAC addresses, and a decision made as to whether to drop the packet or forward it to another or all segments.

In 1989, the networking company Kalpana introduced their EtherSwitch, the first Ethernet switch. This worked somewhat differently from an Ethernet bridge, in that only the header of the incoming packet would be examined before it was either dropped or forwarded to another segment. This greatly reduced the forwarding latency and the processing load on the network device. One drawback of this cut-through switching method was that packets that had been corrupted would still be propagated through the network, so a jabbering station could continue to disrupt the entire network. The eventual remedy for this was a return to the original store and forward approach of bridging, where the packet would be read into a buffer on the switch in its entirety, verified against its checksum and then forwarded, but using more powerful application-specific integrated circuits. Hence, the bridging is then done in hardware, allowing packets to be forwarded at full wire speed.

When a twisted pair or fiber link segment is used and neither end is connected to a repeater, full-duplex Ethernet becomes possible over that segment. In full-duplex mode, both devices can transmit and receive to and from each other at the same time, and there is no collision domain. This doubles the aggregate bandwidth of the link and is sometimes advertised as double the link speed (e.g., 200 Mbit/s). The elimination of the collision domain for these connections also means that all the link's bandwidth can be used by the two devices on that segment and that segment length is not limited by the need for correct collision detection.

Since packets are typically delivered only to the port they are intended for, traffic on a switched Ethernet is less public than on shared-medium Ethernet. Despite this, switched Ethernet should still be regarded as an insecure network technology, because it is easy to subvert switched Ethernet systems by means such as ARP spoofing and MAC flooding.

The bandwidth advantages, the slightly better isolation of devices from each other, the ability to easily mix different speeds of devices and the elimination of the chaining limits inherent in non-switched Ethernet have made switched Ethernet the dominant network technology.

Advanced networking

Simple switched Ethernet networks, while a great improvement over repeater-based Ethernet, suffer from single points of failure, attacks that trick switches or hosts into sending data to a machine even if it is not intended for it, scalability and security issues with regard to broadcast radiation and multicast traffic, and bandwidth choke points where a lot of traffic is forced down a single link.

Advanced networking features in switches and routers combat these issues through a number of means including spanning-tree protocol to maintain the active links of the network as a tree while allowing physical loops for redundancy, port security and protection features such as MAC lock down and broadcast radiation filtering, virtual LANs to keep different classes of users separate while using the same physical infrastructure, multilayer switching to route between different classes and link aggregation to add bandwidth to overloaded links and to provide some measure of redundancy.

Recent networking advances IEEE 802.1aq (SPB) include the use of the link-state routing protocol IS-IS to allow larger networks with shortest path routes between devices.

Varieties of Ethernet

The Ethernet physical layer evolved over a considerable time span and encompasses quite a few physical media interfaces and several magnitudes of speed. The most common forms used are 10BASE-T, 100BASE-TX, and 1000BASE-T. All three utilize twisted pair cables and 8P8C modular connectors. They run at 10 Mbit/s, 100 Mbit/s, and 1 Gbit/s, respectively. Fiber optic variants of Ethernet offer high performance, electrical isolation and distance (up to tens of kilometers with some versions). In general, network protocol stack software will work similarly on all varieties.

Ethernet frames

A data packet on the wire is called a frame. A frame begins with Preamble and Start Frame Delimiter, following which each Ethernet frame features an Ethernet header featuring source and destination MAC addresses. The middle section of the frame consists of payload data including any headers for other protocols (e.g., Internet Protocol) carried in the frame. The frame ends with a 32-bit cyclic redundancy check, which is used to detect any corruption of data in transit.

Autonegotiation

Autonegotiation is the procedure by which two connected devices choose common transmission parameters, such as speed and duplex mode. Autonegotiation was first introduced as an optional feature for Fast Ethernet, but it is also backward compatible with 10BASE-T. Autonegotiation is mandatory for Gigabit Ethernet.

Chapter 10

Link Aggregation

Link aggregation or **IEEE 802.1AX-2008** is a computer networking term which describes using multiple network cables/ports in parallel to increase the link speed beyond the limits of any one single cable or port, and to increase the redundancy for higher availability.

Most implementations now conform to what used to be clause 43 of IEEE 802.3-2005 Ethernet standard, usually still referred to by its working group name of "**IEEE 802.3ad**". The definition of link aggregation has since moved to a standalone IEEE 802.1AX standard.

Link aggregation is often abbreviated **LAG**. Other terms include **trunking**, **link bundling**, **Ethernet/network/NIC bonding**, **NIC teaming**, **port channel**, **EtherChannel**, **Multi-link trunking (MLT)**, **Network Fault Tolerance (NFT)**, **Smartgroup** (from ZTE), and **EtherTrunk** (from Huawei).

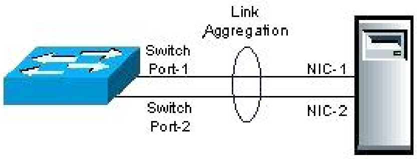

Link Aggregation between a switch and a server

Description

Link aggregation addresses two problems with Ethernet connections: bandwidth limitations and lack of resilience.

With regard to the first issue: bandwidth requirements do not scale linearly. Ethernet bandwidths historically have increased by an order of magnitude each generation: 10 Megabit/s, 100 Mbit/s, 1000 Mbit/s, 10,000 Mbit/s. If one started to bump into bandwidth ceilings, then the only option was to move to the next generation which could be cost prohibitive. An alternative solution, introduced by many of the network manufacturers in the early 1990s, is to combine two physical Ethernet links into one logical link via channel bonding. Most of these solutions required manual configuration and identical equipment on both sides of the aggregation.

The second problem involves the three single points of failure in a typical port-cable-port connection. In either the usual computer-to-switch or in a switch-to-switch configuration, the cable itself or either of the ports the cable is plugged into can fail. Multiple physical connections can be made, but many of the higher level protocols were not designed to failover completely seamlessly.

IEEE Link Aggregation

Standardization process

By the mid 1990s, most network switch manufacturers had included aggregation capability as a proprietary extension to increase bandwidth between their switches. But each manufacturer developed its own method, which led to compatibility problems. The IEEE 802.3 group took up a study group to create an inter-operable link layer standard in a November 1997 meeting. The group quickly agreed to include an automatic configuration feature which would add in redundancy as well. This became known as "Link Aggregation Control Protocol".

Initial release 802.3ad in 2000

As of 2000 most gigabit channel-bonding uses the IEEE standard of Link Aggregation which was formerly clause 43 of the IEEE 802.3 standard added in March 2000 by the IEEE 802.3ad task force. Nearly every network equipment manufacturer quickly adopted this joint standard over their proprietary standards.

Move to 802.1 layer in 2008

David Law noted in 2006 that certain 802.1 layers (such as 802.1X security) were positioned in the protocol stack above Link Aggregation which was defined as an 802.3 sublayer. This discrepancy was resolved with formal transfer of the protocol to the 802.1 group with the publication of IEEE 802.1AX-2008 on 3 November 2008.

Link Aggregation Control Protocol

Within the IEEE specification the **Link Aggregation Control Protocol (LACP)** provides a method to control the bundling of several physical ports together to form a single logical channel. LACP allows a network device to negotiate an automatic bundling of links by sending LACP packets to the peer (directly connected device that also implements LACP).

Advantages over static configuration

- Failover when a link fails and there is (for example) a Media Converter between the devices which means that the peer will not see the link down. With static link aggregation the peer would continue sending traffic down the link causing it to be lost.
- The device can confirm that the configuration at the other end can handle link aggregation. With Static link aggregation a cabling or configuration mistake could go undetected and cause undesirable network behavior.

Practical notes

LACP works by sending frames (LACPDUs) down all links that have the protocol enabled. If it finds a device on the other end of the link that also has LACP enabled, it will also independently send frames along the same links enabling the two units to detect multiple links between themselves and then combine them into a single logical link. LACP can be configured in one of two modes: active or passive. In active mode it will always send frames along the configured links. In passive mode however, it acts as "speak when spoken to", and therefore can be used as a way of controlling accidental loops (as long as the other device is in active mode).

Aggregation Modes in Linux (Bonding Modes)

Round-robin policy

> Transmit packets in sequential order from the first available slave through the last. This mode provides load balancing and fault tolerance.

Active-backup policy

> Only one slave in the bond is active. A different slave becomes active if, and only if, the active slave fails. The bond's MAC address is externally visible on only one port (network adapter) to avoid confusing the switch. This mode provides fault tolerance. The primary option affects the behavior of this mode.

XOR policy

> Transmit based on [(source MAC address XOR'd with destination MAC address) modulo slave count]. This selects the same slave for each destination MAC address. This mode provides load balancing and fault tolerance.

Broadcast policy

> transmits everything on all slave interfaces. This mode provides fault tolerance.

IEEE 802.3ad Dynamic link aggregation

> Creates aggregation groups that share the same speed and duplex settings. Utilizes all slaves in the active aggregator according to the 802.3ad specification.

Adaptive transmit load balancing

> channel bonding that does not require any special switch support. The outgoing traffic is distributed according to the current load (computed relative to the speed) on each slave. Incoming traffic is received by the current slave.

If the receiving slave fails, another slave takes over the MAC address of the failed receiving slave.

Adaptive load balancing

includes balance-tlb plus receive load balancing (rlb) for IPV4 traffic, and does not require any special switch support. The receive load balancing is achieved by ARP negotiation. The bonding driver intercepts the ARP Replies sent by the local system on their way out and overwrites the source hardware address with the unique hardware address of one of the slaves in the bond such that different peers use different hardware addresses for the server.

Usage

Network backbone

Link aggregation offers an inexpensive way to set up a high-speed backbone network that transfers much more data than any one single port or device can deliver. Although, in the past, various vendors used proprietary techniques, the preference today is to use the IEEE standard, which can either be set up statically or by using the Link Aggregation Control Protocol (LACP). This allows several devices to communicate simultaneously at their full single-port speed while not allowing any one single device to monopolize all available backbone capacity.

The actual benefits vary based on the load-balancing method used on each device (administrators can configure different balancing algorithms at each end and this is actually encouraged to avoid path polarization). Link aggregation also allows the network's backbone speed to grow incrementally as demand on the network increases, without having to replace everything and buy new hardware.

Most backbone installations install more cabling or fiber optic pairs than is initially necessary, even if they have no immediate need for the additional cabling. This is done because labor costs are higher than the cost of the cable, and running extra cable reduces future labor costs if networking needs change. Link aggregation can allow the use of these extra cables to increase backbone speeds for little or no extra cost if ports are available.

Order of frames

When balancing traffic, network administrators often wish to avoid reordering Ethernet frames. For example, TCP suffers additional overheads when dealing with out-of-order packets. This goal is approximated by sending all frames associated with a particular session across the same link. The most common implementations use L3 hashes (i.e. based on the IP address), ensuring that the same flow is always sent via the same physical link.

However, depending on the traffic, this may not provide even distribution across the links in the trunk. It effectively limits the client bandwidth in an aggregate to its single member's maximum bandwidth per session. Principally for this reason 50/50 load balancing is almost never reached in real-life implementations; around 70/30 is more usual. Advanced switches can employ an L4 hash (i.e. using TCP/UDP port numbers), which will bring the balance closer to 50/50 as different L4 flows between two hosts can make use of different physical links.

Efficiency of equipment

Aggregation becomes inefficient beyond a certain bandwidth — depending on the total number of ports on the switch equipment. A 24-port gigabit switch with two 8-gigabit trunks is using sixteen of its available ports just for the two interswitch connections, and leaves only eight of its 1-gigabit ports for other devices. This same configuration on a 48-port gigabit switch leaves 32 1-gigabit ports available, and so it is much more efficient (assuming of course that those ports are actually needed at the switch location).

When a switch utilizes 40-50% of its ports for backbone trunking, upgrading to a switch with either more ports or a higher base-operating speed may be a better option than simply adding more switches, especially if the old switch can be re-used elsewhere on a less performance-critical part of the network.

Use on network interface cards

Network interface cards (NICs) trunked together can also provide network links beyond the throughput of any one single NIC. For example, this allows a central file server to establish an aggregate 2-gigabit connection using two 1-gigabit NICs trunked together. Note the data signaling rate will still be 1Gb/s, which can be misleading depending on methodologies used to test throughput after link aggregation is employed.

Note that Microsoft Windows does not natively support link aggregation (at least up to Windows Server 2008). However, some manufacturers provide software for aggregation on their multiport NICs at the device-driver layer. Intel, for example, has released a package for Windows called Advanced Networking Services (ANS) to bind Intel Fast Ethernet and Gigabit cards. Nvidia also supports "teaming" with their Nvidia Network Access Manager/Firewall Tool. HP also has a very robust teaming tool for HP branded NICs which will allow for non-etherchanneled NIC teaming or which will also support several modes of etherchannel (port aggregation) including 802.3ad with LACP.

Linux, FreeBSD, NetBSD, OpenBSD, Mac OS X, OpenSolaris, Citrix XenServer, VMware ESX, and commercial Unix distributions such as AIX implement Ethernet bonding (trunking) at a higher level, and can hence deal with NICs from different manufacturers or drivers, as long as the NIC is supported by the kernel.

Limitations

Single switch

With modes balance-rr, balance-xor, broadcast and 802.3ad all physical ports in the link aggregation group must reside on the same logical switch, which in most scenarios will leave a single point of failure when the physical switch to which both links are connected goes offline. Modes active-backup, balance-tlb, and balance-alb can also be set up with two or more switches. But after failover (like all other modes), in some cases, active sessions may fail (due to arp problems) and have to be restarted.

However, almost all vendors have proprietary extensions that resolve some of this issue: they aggregate multiple physical switches into one logical switch. As of 2009, the IEEE has not yet committed resources to standardize this feature. The SMLT protocol allows multiple Ethernet links to be split across two devices, preventing any single point of failure, and additionally allowing the load to be balanced across the 2 aggregation switches from the single access system. These devices synchronize state across an Inter-Switch Trunk (IST) such that they appear to the connecting (access) device to be a single device (switch block) and prevent any packet duplication. SMLT's provide enhanced resiliency with sub-second failover and sub-second recovery for all speed trunks (10Mbps, 100Mbps, 1000Mbps, and 10Gbps) while operating transparently to end-devices.

Same link speed

In most implementations, all the ports used in an aggregation consist of the same physical type, such as all copper ports (CAT-5E/CAT-6), all multi-mode fiber ports (SX), or all single-mode fiber ports (LX). However, all the IEEE standard requires is that each link be full duplex and all of them have an identical speed (10, 100, 1000 or 10000 Mbps).

Many switches are PHY independent, meaning that a switch could have a mixture of copper, SX, LX, LX10 or other GBICs. While maintaining the same PHY is the usual approach, it is possible to aggregate a 1000BASE-SX fiber for one link and a 1000BASE-LX (longer, diverse path) for the second link, but the important thing is that the speed will be 1 Gbit/s full duplex for both links. One path may have a slightly longer transit time but the standard has been engineered so this will not cause an issue.

Ethernet aggregation mismatch

Aggregation mismatch refers to not matching the aggregation type on both ends of the link. Some switches do not implement the 802.1AX standard but support static configuration of link aggregation. Therefore link aggregation between similarly statically configured switches will work, but will fail between a statically configured switch and a device that is configured for LACP.

Chapter 11

Network Topology

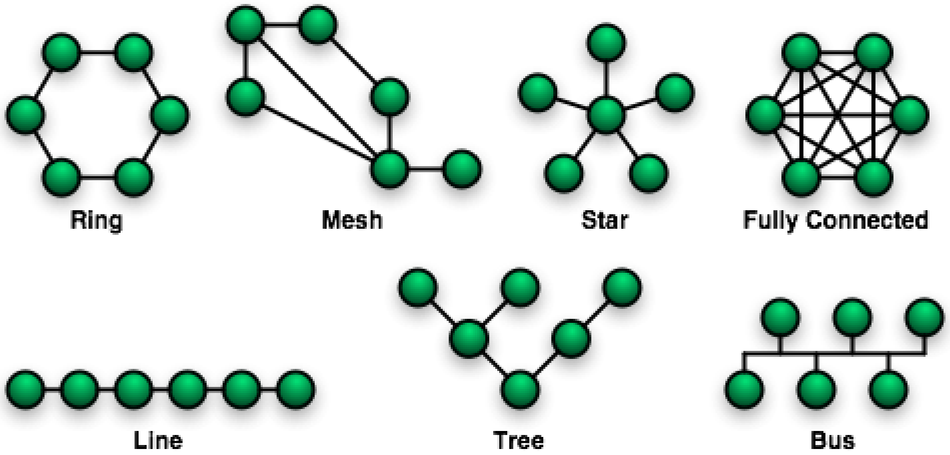

Diagram of different network topologies

Network topology is the layout pattern of interconnections of the various elements (links, nodes, etc.) of a computer network. Network topologies may be physical or logical. Physical topology means the physical design of a network including the devices, location and cable installation. Logical topology refers to how data is actually transferred in a network as opposed to its physical design.

Topology can be considered as a virtual shape or structure of a network. This shape does not correspond to the actual physical design of the devices on the computer network. The computers on a home network can be arranged in a circle but it does not necessarily mean that it represents a ring topology.

Any particular network topology is determined only by the graphical mapping of the configuration of physical and/or logical connections between nodes. The study of network topology uses graph theory. Distances between nodes, physical interconnections, transmission rates, and/or signal types may differ in two networks and yet their topologies may be identical.

A local area network (LAN) is one example of a network that exhibits both a physical topology and a logical topology. Any given node in the LAN has one or more links to one or more nodes in the network and the mapping of these links and nodes in a graph results in a geometric shape that may be used to describe the physical topology of the network. Likewise, the mapping of the data flow between the nodes in the network determines the logical topology of the network. The physical and logical topologies may or may not be identical in any particular network.

Basic topology types

The study of network topology recognizes seven basic topologies:

- Point-to-point topology
- Bus (point-to-multipoint) topology
- Star topology
- Ring topology
- Tree topology
- Mesh topology
- Hybrid topology

This classification is based on the interconnection between computers — be it physical or logical.

The physical topology of a network is determined by the capabilities of the network access devices and media, the level of control or fault tolerance desired, and the cost associated with cabling or telecommunications circuits.

Networks can be classified according to their physical span as follows:

- LANs (Local Area Networks)
- Building or campus internetworks
- Wide area internetworks

Classification of network topologies

There are also two basic categories of network topologies:

- Physical topologies
- Logical topologies

The shape of the cabling layout used to link devices is called the physical topology of the network. This refers to how the cables are laid out to connect many computers to one network. The physical topology you choose for your network influences and is influenced by several factors:

- Office Layout

- Troubleshooting Techniques
- Cost of Installation
- Type of cable used

Logical topology describes the way in which a network transmits information from network/computer to another and not the way the network looks or how it is laid out. The logical layout also describes the different speeds of the cables being used from one network to another.

Physical topologies

The mapping of the nodes of a network and the physical connections between them – the layout of wiring, cables, the locations of nodes, and the interconnections between the nodes and the cabling or wiring system.

Classification of physical topologies

Point-to-point

The simplest topology is a permanent link between two endpoints (the line in the illustration above). Switched point-to-point topologies are the basic model of conventional telephony. The value of a permanent point-to-point network is the value of guaranteed, or nearly so, communications between the two endpoints. The value of an on-demand point-to-point connection is proportional to the number of potential pairs of subscribers, and has been expressed as Metcalfe's Law.

Permanent (dedicated)
Easiest to understand, of the variations of point-to-point topology, is a point-to-point communications channel that appears, to the user, to be permanently associated with the two endpoints. Children's "tin-can telephone" is one example, with a microphone to a single public address speaker is another. These are examples of *physical dedicated* channels.
Within many switched telecommunications systems, it is possible to establish a permanent circuit. One example might be a telephone in the lobby of a public building, which is programmed to ring only the number of a telephone dispatcher. "Nailing down" a switched connection saves the cost of running a physical circuit between the two points. The resources in such a connection can be released when no longer needed, for example, a television circuit from a parade route back to the studio.
Switched:
Using circuit-switching or packet-switching technologies, a point-to-point circuit can be set up dynamically, and dropped when no longer needed. This is the basic mode of conventional telephony.

Bus

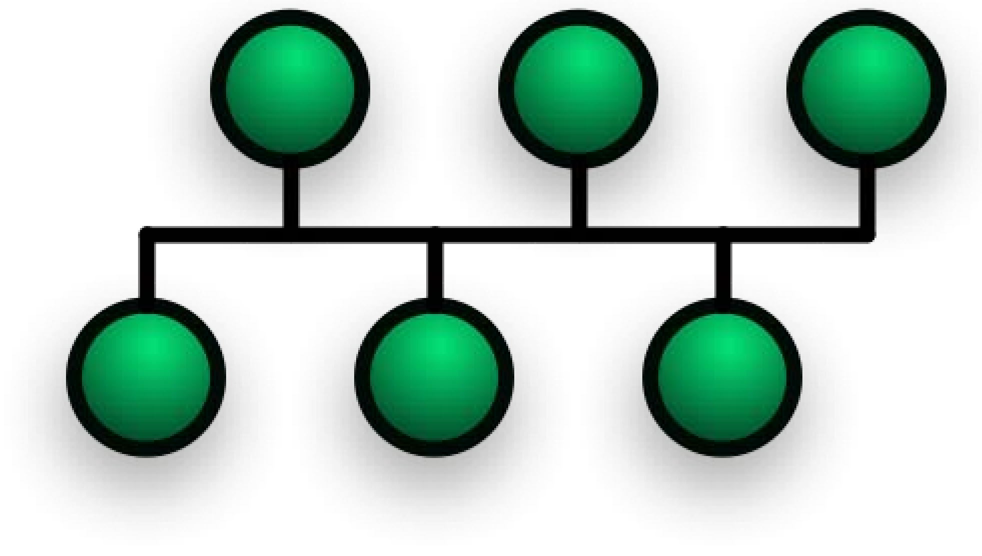

Bus network topology

In local area networks where bus topology is used, each machine is connected to a single cable. Each computer or server is connected to the single bus cable through some kind of connector. A terminator is required at each end of the bus cable to prevent the signal from bouncing back and forth on the bus cable. A signal from the source travels in both directions to all machines connected on the bus cable until it finds the MAC address or IP address on the network that is the intended recipient. If the machine address does not match the intended address for the data, the machine ignores the data. Alternatively, if the data does match the machine address, the data is accepted. Since the bus topology consists of only one wire, it is rather inexpensive to implement when compared to other topologies. However, the low cost of implementing the technology is offset by the high cost of managing the network. Additionally, since only one cable is utilized, it can be the single point of failure. If the network cable breaks, the entire network will be down.

Linear bus

The type of network topology in which all of the nodes of the network are connected to a common transmission medium which has exactly two endpoints (this is the 'bus', which is also commonly referred to as the backbone, or trunk) – all data that is transmitted between nodes in the network is transmitted over this common transmission medium and is able to be received by all nodes in the network virtually simultaneously (disregarding propagation delays).

Note: The two endpoints of the common transmission medium are normally terminated with a device called a terminator that exhibits the characteristic impedance of the transmission medium and which dissipates or absorbs the energy that remains in the signal to prevent the signal from being reflected or propagated back onto the transmission medium in the opposite direction, which would cause interference with and degradation of the signals on the transmission medium.

Distributed bus
The type of network topology in which all of the nodes of the network are connected to a common transmission medium which has more than two endpoints that are created by adding branches to the main section of the transmission medium – the physical distributed bus topology functions in exactly the same fashion as the physical linear bus topology (i.e., all nodes share a common transmission medium).

Notes:

1.) All of the endpoints of the common transmission medium are normally terminated with a device called a 'terminator'.

2.) The physical linear bus topology is sometimes considered to be a special case of the physical distributed bus topology – i.e., a distributed bus with no branching segments.

3.) The physical distributed bus topology is sometimes incorrectly referred to as a physical tree topology – however, although the physical distributed bus topology resembles the physical tree topology, it differs from the physical tree topology in that there is no central node to which any other nodes are connected, since this hierarchical functionality is replaced by the common bus.

Star

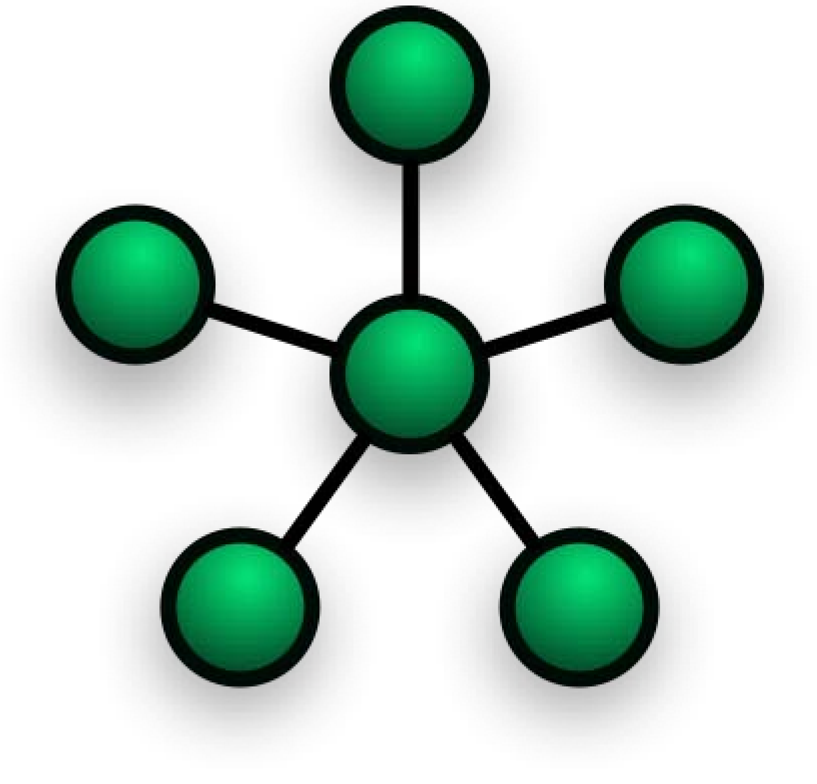

Star network topology

In local area networks with a star topology, each network host is connected to a central hub. In contrast to the bus topology, the star topology connects each node to the hub with a point-to-point connection. All traffic that traverses the network passes through the central hub. The hub acts as a signal booster or repeater. The star topology is considered the easiest topology to design and implement. An advantage of the star topology is the simplicity of adding additional nodes. The primary disadvantage of the star topology is that the hub represents a single point of failure.

Notes

- A point-to-point link (described above) is sometimes categorized as a special instance of the physical star topology – therefore, the simplest type of network that is based upon the physical star topology would consist of one node with a single point-to-point link to a second node, the choice of which node is the 'hub' and which node is the 'spoke' being arbitrary.

- After the special case of the point-to-point link, as in note 1.) above, the next simplest type of network that is based upon the physical star topology would consist of one central node – the 'hub' – with two separate point-to-point links to two peripheral nodes – the 'spokes'.

- Although most networks that are based upon the physical star topology are commonly implemented using a special device such as a hub or switch as the central node (i.e., the 'hub' of the star), it is also possible to implement a network that is based upon the physical star topology using a computer or even a simple common connection point as the 'hub' or central node – however, since many illustrations of the physical star network topology depict the central node as one of these special devices, some confusion is possible, since this practice may lead to the misconception that a physical star network requires the central node to be one of these special devices, which is not true because a simple network consisting of three computers connected as in note 2.) above also has the topology of the physical star.

- Star networks may also be described as either broadcast multi-access or nonbroadcast multi-access (NBMA), depending on whether the technology of the network either automatically propagates a signal at the hub to all spokes, or only addresses individual spokes with each communication.

Extended star

A type of network topology in which a network that is based upon the physical star topology has one or more repeaters between the central node (the 'hub' of the star) and the peripheral or 'spoke' nodes, the repeaters being used to extend the maximum transmission distance of the point-to-point links between the central node and the peripheral nodes beyond that which is supported by the transmitter power of the central node or beyond that which is supported by the standard upon which the physical layer of the physical star network is based.

If the repeaters in a network that is based upon the physical extended star topology are replaced with hubs or switches, then a hybrid network topology is created that is referred to as a physical hierarchical star topology, although some texts make no distinction between the two topologies.

Distributed Star

A type of network topology that is composed of individual networks that are based upon the physical star topology connected together in a linear fashion – i.e., 'daisy-chained' – with no central or top level connection point (e.g., two or more 'stacked' hubs, along with their associated star connected nodes or 'spokes').

Ring

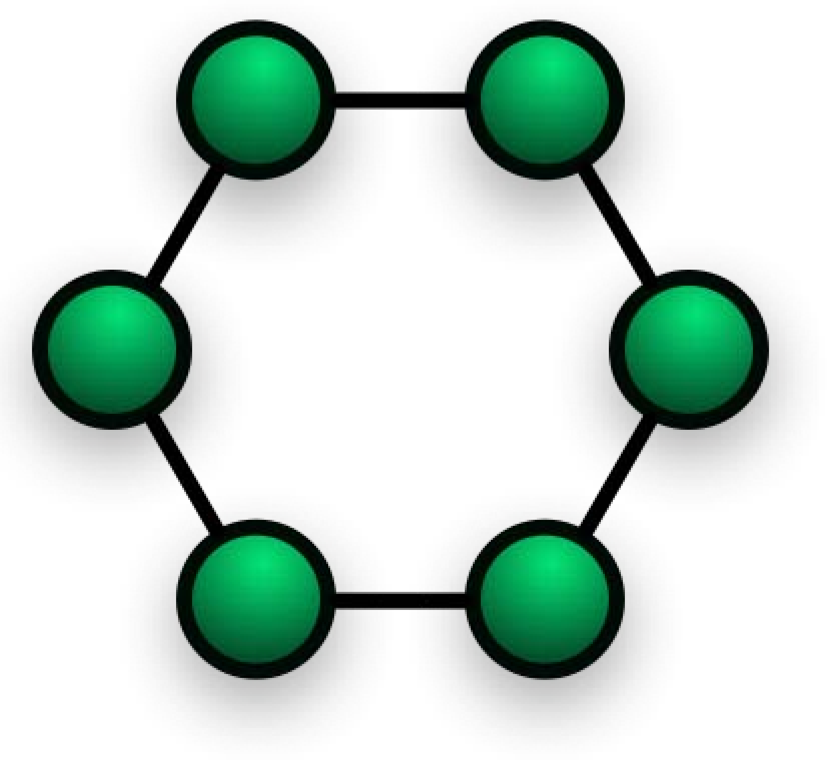

Ring network topology

A network topology that is set up in a circular fashion in which data travels around the ring in one direction and each device on the right acts as a repeater to keep the signal strong as it travels. Each device incorporates a receiver for the incoming signal and a transmitter to send the data on to the next device in the ring. The network is dependent on the ability of the signal to travel around the ring.

Mesh

The value of fully meshed networks is proportional to the exponent of the number of subscribers, assuming that communicating groups of any two endpoints, up to and including all the endpoints, is approximated by Reed's Law.

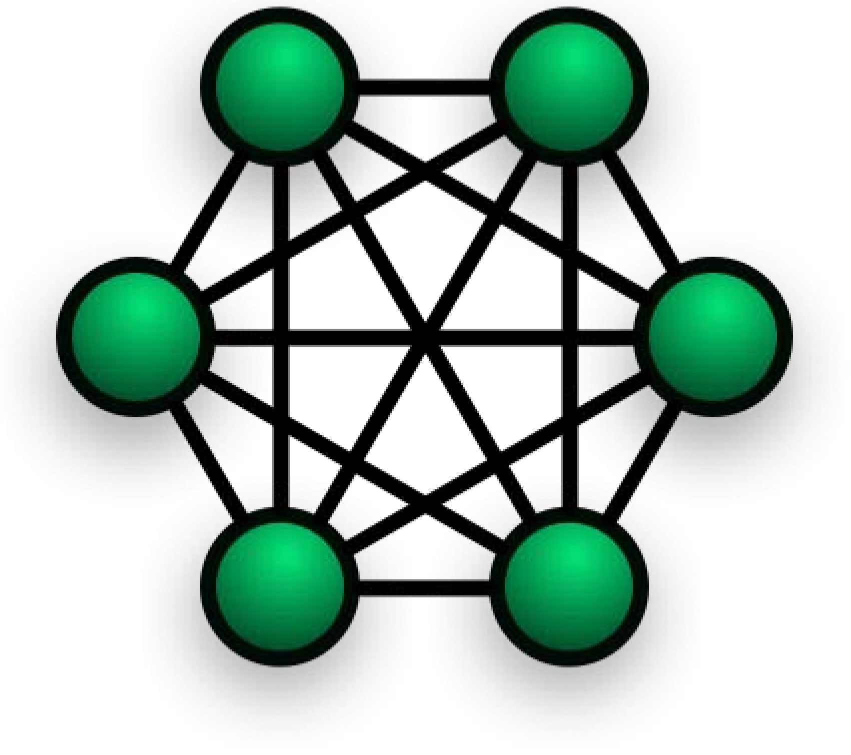

Fully connected mesh topology

The number of connections in a full mesh = n(n - 1) / 2

Fully connected

Note: The physical fully connected mesh topology is generally too costly and complex for practical networks, although the topology is used when there are only a small number of nodes to be interconnected.

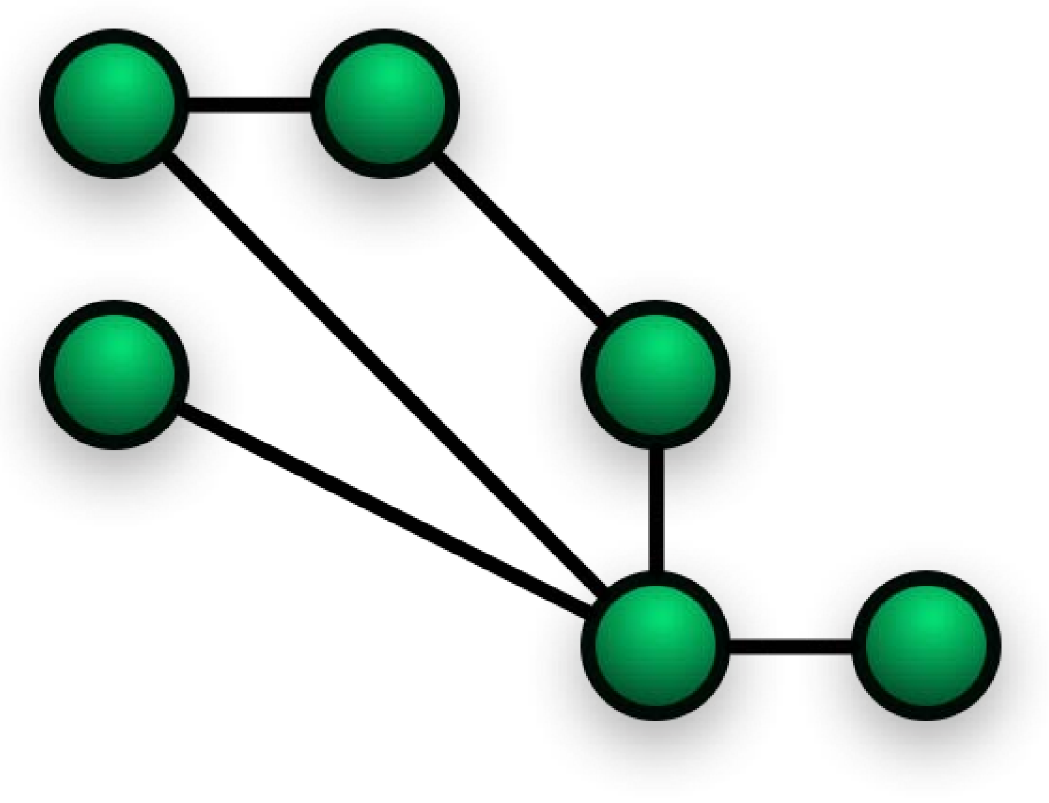

Partially connected mesh topology

Partially connected

The type of network topology in which some of the nodes of the network are connected to more than one other node in the network with a point-to-point link – this makes it possible to take advantage of some of the redundancy that is provided by a physical fully connected mesh topology without the expense and complexity required for a connection between every node in the network.

Note: In most practical networks that are based upon the physical partially connected mesh topology, all of the data that is transmitted between nodes in the network takes the shortest path (or an approximation of the shortest path) between nodes, except in the case of a failure or break in one of the links, in which case the data takes an alternative path to the destination. This requires that the nodes of the network possess some type of logical 'routing' algorithm to determine the correct path to use at any particular time.

Tree

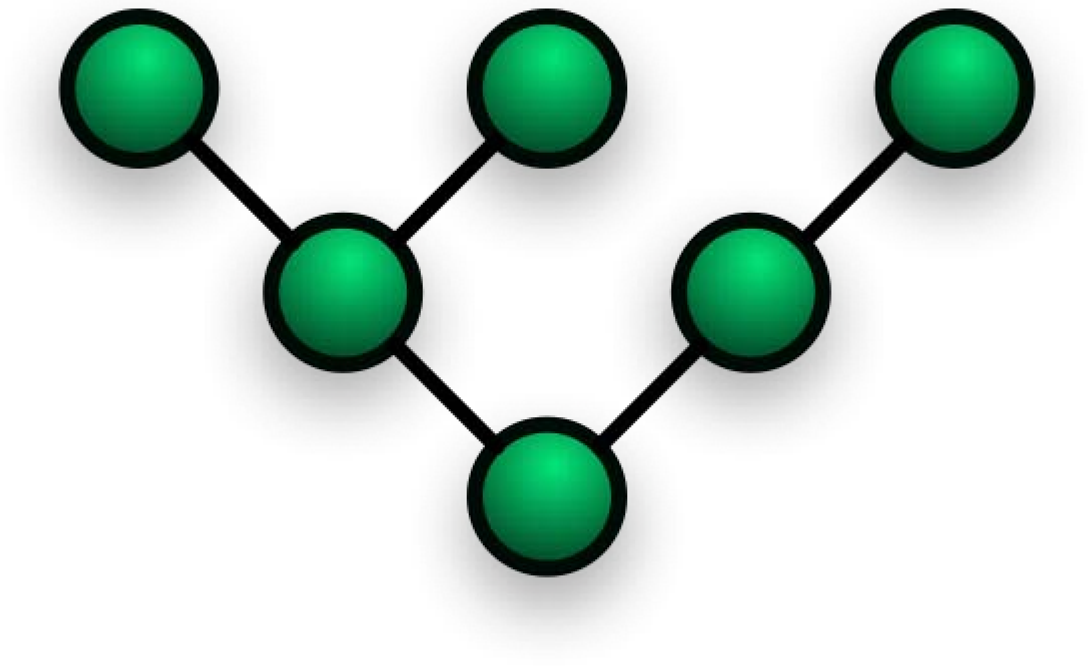

Tree network topology

Also known as a **hierarchical network**.

The type of network topology in which a central 'root' node (the top level of the hierarchy) is connected to one or more other nodes that are one level lower in the hierarchy (i.e., the second level) with a point-to-point link between each of the second level nodes and the top level central 'root' node, while each of the second level nodes that are connected to the top level central 'root' node will also have one or more other nodes that are one level lower in the hierarchy (i.e., the third level) connected to it, also with a point-to-point link, the top level central 'root' node being the only node that has no other node above it in the hierarchy (The hierarchy of the tree is symmetrical.) Each node in the network having a specific fixed number, of nodes connected to it at the next lower level in the hierarchy, the number, being referred to as the 'branching factor' of the hierarchical tree. This tree has individual peripheral nodes.

1.) A network that is based upon the physical hierarchical topology must have at least three levels in the hierarchy of the tree, since a network with a central 'root' node and only one hierarchical level below it would exhibit the physical topology of a star.
2.) A network that is based upon the physical hierarchical topology and with a branching factor of 1 would be classified as a physical linear topology.

3.) The branching factor, f, is independent of the total number of nodes in the network and, therefore, if the nodes in the network require ports for connection to other nodes the total number of ports per node may be kept low even though the total number of nodes is large – this makes the effect of the cost of adding ports to each node totally dependent upon the branching factor and may therefore be kept as low as required without any effect upon the total number of nodes that are possible.
4.) The total number of point-to-point links in a network that is based upon the physical hierarchical topology will be one less than the total number of nodes in the network.
5.) If the nodes in a network that is based upon the physical hierarchical topology are required to perform any processing upon the data that is transmitted between nodes in the network, the nodes that are at higher levels in the hierarchy will be required to perform more processing operations on behalf of other nodes than the nodes that are lower in the hierarchy. Such a type of network topology is very useful and highly recommended.

Logical topology

The logical topology, in contrast to the "physical", is the way that the signals act on the network media, or the way that the data passes through the network from one device to the next without regard to the physical interconnection of the devices. A network's logical topology is not necessarily the same as its physical topology. For example, twisted pair Ethernet is a logical bus topology in a physical star topology layout. While IBM's Token Ring is a logical ring topology, it is physically set up in a star topology.

Classification of logical topologies

The logical classification of network topologies generally follows the same classifications as those in the physical classifications of network topologies, the path that the *data* takes between nodes being used to determine the topology as opposed to the actual *physical* connections being used to determine the topology

Notes:
1.) Logical topologies are often closely associated with media access control (MAC) methods and protocols.
2.) The logical topologies are generally determined by network protocols as opposed to being determined by the physical layout of cables, wires, and network devices or by the flow of the electrical signals, although in many cases the paths that the electrical signals take between nodes may closely match the logical flow of data, hence the convention of using the terms 'logical topology' and 'signal topology' interchangeably.
3.) Logical topologies are able to be dynamically reconfigured by special types of equipment such as routers and switches.

Daisy chains

Except for star-based networks, the easiest way to add more computers into a network is by daisy-chaining, or connecting each computer in series to the next. If a message is intended for a computer partway down the line, each system bounces it along in sequence until it reaches the destination. A daisy-chained network can take two basic forms: linear and ring.

- A **linear topology** puts a two-way link between one computer and the next. However, this was expensive in the early days of computing, since each computer (except for the ones at each end) required two receivers and two transmitters.
- By connecting the computers at each end, a **ring topology** can be formed. An advantage of the ring is that the number of transmitters and receivers can be cut in half, since a message will eventually loop all of the way around. When a node sends a message, the message is processed by each computer in the ring. If a computer is not the destination node, it will pass the message to the next node, until the message arrives at its destination. If the message is not accepted by any node on the network, it will travel around the entire ring and return to the sender. This potentially results in a doubling of travel time for data.

Centralization

The **star topology** reduces the probability of a network failure by connecting all of the peripheral nodes (computers, etc.) to a central node. When the physical star topology is applied to a logical bus network such as Ethernet, this central node (traditionally a hub) rebroadcasts all transmissions received from any peripheral node to all peripheral nodes on the network, sometimes including the originating node. All peripheral nodes may thus communicate with all others by transmitting to, and receiving from, the central node only. The failure of a transmission line linking any peripheral node to the central node will result in the isolation of that peripheral node from all others, but the remaining peripheral nodes will be unaffected. However, the disadvantage is that the failure of the central node will cause the failure of all of the peripheral nodes also.

If the central node is *passive*, the originating node must be able to tolerate the reception of an echo of its own transmission, delayed by the two-way round trip transmission time (i.e. to and from the central node) plus any delay generated in the central node. An *active* star network has an active central node that usually has the means to prevent echo-related problems.

A **tree topology** (a.k.a. **hierarchical topology**) can be viewed as a collection of star networks arranged in a hierarchy. This tree has individual peripheral nodes (e.g. leaves) which are required to transmit to and receive from one other node only and are not required to act as repeaters or regenerators. Unlike the star network, the functionality of the central node may be distributed.

As in the conventional star network, individual nodes may thus still be isolated from the network by a single-point failure of a transmission path to the node. If a link connecting a leaf fails, that leaf is isolated; if a connection to a non-leaf node fails, an entire section of the network becomes isolated from the rest.

In order to alleviate the amount of network traffic that comes from broadcasting all signals to all nodes, more advanced central nodes were developed that are able to keep track of the identities of the nodes that are connected to the network. These network switches will "learn" the layout of the network by "listening" on each port during normal data transmission, examining the data packets and recording the address/identifier of each connected node and which port it's connected to in a lookup table held in memory. This lookup table then allows future transmissions to be forwarded to the intended destination only.

Decentralization

In a **mesh topology** (i.e., a partially connected mesh topology), there are at least two nodes with two or more paths between them to provide redundant paths to be used in case the link providing one of the paths fails. This decentralization is often used to advantage to compensate for the single-point-failure disadvantage that is present when using a single device as a central node (e.g., in star and tree networks). A special kind of mesh, limiting the number of hops between two nodes, is a hypercube. The number of arbitrary forks in mesh networks makes them more difficult to design and implement, but their decentralized nature makes them very useful. This is similar in some ways to a **grid network**, where a linear or ring topology is used to connect systems in multiple directions. A multi-dimensional ring has a toroidal topology, for instance.

A **fully connected network**, **complete topology** or **full mesh topology** is a network topology in which there is a direct link between all pairs of nodes. In a fully connected network with n nodes, there are n(n-1)/2 direct links. Networks designed with this topology are usually very expensive to set up, but provide a high degree of reliability due to the multiple paths for data that are provided by the large number of redundant links between nodes. This topology is mostly seen in military applications. However, it can also be seen in the file sharing protocol BitTorrent in which users connect to other users in the "swarm" by allowing each user sharing the file to connect to other users also involved. Often in actual usage of BitTorrent any given individual node is rarely connected to every single other node as in a true fully connected network but the protocol does allow for the possibility for any one node to connect to any other node when sharing files.

Hybrids

Hybrid networks use a combination of any two or more topologies in such a way that the resulting network does not exhibit one of the standard topologies (e.g., bus, star, ring, etc.). For example, a tree network connected to a tree network is still a tree network, but two star networks connected together exhibit a hybrid network topology. A hybrid topology is always produced when two different basic network topologies are connected. Two common examples for Hybrid network are: *star ring network* and *star bus network*

- A Star ring network consists of two or more star topologies connected using a multistation access unit (MAU) as a centralized hub.
- A Star Bus network consists of two or more star topologies connected using a bus trunk (the bus trunk serves as the network's backbone).

While grid networks have found popularity in high-performance computing applications, some systems have used genetic algorithms to design custom networks that have the fewest possible hops in between different nodes. Some of the resulting layouts are nearly incomprehensible, although they function quite well.

A Snowflake topology is really a "Star of Stars" network, so it exhibits characeristics of a hybrid network topology but is not composed of two different basic network topologies being connected together.

Chapter 12

Local Area Network and Metropolitan Area Network

Local area network

A **local area network (LAN)** is a computer network that connects computers and devices in a limited geographical area such as home, school, computer laboratory or office building. The defining characteristics of LANs, in contrast to wide area networks (WANs), include their usually higher data-transfer rates, smaller geographic area, and lack of a need for leased telecommunication lines.

ARCNET, Token Ring and other technologies have been used in the past, but Ethernet over twisted pair cabling, and Wi-Fi are the two most common technologies currently in use.

History

As larger universities and research labs obtained more computers during the late 1960s, there was an increasing pressure to provide high-speed interconnections. A report in 1970 from the Lawrence Radiation Laboratory detailing the growth of their "Octopus" network gives a good indication of the situation.

Cambridge Ring was developed at Cambridge University in 1974 but was never developed into a successful commercial product.

Ethernet was developed at Xerox PARC in 1973–1975, and filed as U.S. Patent 4,063,220. In 1976, after the system was deployed at PARC, Metcalfe and Boggs published their seminal paper, "Ethernet: Distributed Packet-Switching For Local Computer Networks."

ARCNET was developed by Datapoint Corporation in 1976 and announced in 1977. It had the first commercial installation in December 1977 at Chase Manhattan Bank in New York.

Standards evolution

The development and proliferation of CP/M-based personal computers from the late 1970s and then DOS-based personal computers from 1981 meant that a single site began to have dozens or even hundreds of computers. The initial attraction of networking these was generally to share disk space and laser printers, which were both very expensive at the time. There was much enthusiasm for the concept and for several years, from about 1983 onward, computer industry pundits would regularly declare the coming year to be "the year of the LAN".

In practice, the concept was marred by proliferation of incompatible physical Layer and network protocol implementations, and a plethora of methods of sharing resources. Typically, each vendor would have its own type of network card, cabling, protocol, and network operating system. A solution appeared with the advent of Novell NetWare which provided even-handed support for dozens of competing card/cable types, and a much more sophisticated operating system than most of its competitors. Netware dominated the personal computer LAN business from early after its introduction in 1983 until the mid 1990s when Microsoft introduced Windows NT Advanced Server and Windows for Workgroups.

Of the competitors to NetWare, only Banyan Vines had comparable technical strengths, but Banyan never gained a secure base. Microsoft and 3Com worked together to create a simple network operating system which formed the base of 3Com's 3+Share, Microsoft's LAN Manager and IBM's LAN Server - but none of these were particularly successful.

During the same period, Unix computer workstations from vendors such as Sun Microsystems, Hewlett-Packard, Silicon Graphics, Intergraph, NeXT and Apollo were using TCP/IP based networking. Although this market segment is now much reduced, the technologies developed in this area continue to be influential on the Internet and in both Linux and Apple Mac OS X networking—and the TCP/IP protocol has now almost completely replaced IPX, AppleTalk, NBF, and other protocols used by the early PC LANs.

Cabling

Early LAN cabling had always been based on various grades of coaxial cable. However shielded twisted pair was used in IBM's Token Ring implementation, and in 1984 StarLAN showed the potential of simple *unshielded* twisted pair by using Cat3—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. In addition, fiber-optic cabling is increasingly used in commercial applications. As cabling is not always possible, wireless Wi-Fi is now the most common technology in residential premises as the cabling required is minimal, and it is well suited to mobile laptops and smartphones.

Technical aspects

Switched Ethernet is the most common Data Link Layer and Physical Layer implementation for local area networks. At the higher layers, the Internet Protocol (TCP/IP) has become the standard. Smaller LANs generally consist of one or more switches linked to each other, often at least one is connected to a router, cable modem, or ADSL modem for Internet access.

Larger LANs are characterized by their use of redundant links with switches using the spanning tree protocol to prevent loops, their ability to manage differing traffic types via quality of service (QoS), and to segregate traffic with VLANs. Larger LANs also contain a wide variety of network devices such as switches, firewalls, routers, load balancers, and sensors.

LANs may have connections with other LANs via leased lines, leased services, or by tunneling across the Internet using virtual private network technologies. Depending on how the connections are established and secured in a LAN, and the distance involved, a LAN may also be classified as metropolitan area network (MAN) or wide area networks (WAN)

Metropolitan area network

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.

Chapter 13

Wide Area Network and Wireless Network

Wide area network

A **wide area network** (**WAN**) is a computer network that covers a broad area (i.e., any network whose communications links cross metropolitan, regional, or national boundaries). This is in contrast with personal area networks (PANs), local area networks (LANs), campus area networks (CANs), or metropolitan area networks (MANs) which are usually limited to a room, building, campus or specific metropolitan area (e.g., a city) respectively.

WAN design options

WANs are used to connect LANs and other types of networks together, so that users and computers in one location can communicate with users and computers in other locations. Many WANs are built for one particular organization and are private. Others, built by Internet service providers, provide connections from an organization's LAN to the Internet. WANs are often built using leased lines. At each end of the leased line, a router connects to the LAN on one side and a hub within the WAN on the other. Leased lines can be very expensive. Instead of using leased lines, WANs can also be built using less costly circuit switching or packet switching methods. Network protocols including TCP/IP deliver transport and addressing functions. Protocols including Packet over SONET/SDH, MPLS, ATM and Frame relay are often used by service providers to deliver the links that are used in WANs. X.25 was an important early WAN protocol, and is often considered to be the "grandfather" of Frame Relay as many of the underlying protocols and functions of X.25 are still in use today (with upgrades) by Frame Relay.

Academic research into wide area networks can be broken down into three areas: Mathematical models, network emulation and network simulation.

Performance improvements are sometimes delivered via WAFS or WAN optimization.

WAN connection technology options

There are also several ways to connect NonStop S-series servers to WANs, including via the ServerNet Wide Area Network (SWAN) or SWAN 2, 3, 4, 5, 6, 7, 8, 9, 10 concentrators, which provides WAN client connectivity to servers that have Ethernet ports and appropriate communications software. You can also use the Asynchronous Wide Area Network (AWAN) access server, which offers economical asynchronous-only WAN access. Several options are available for WAN connectivity:

Option:	Description	Advantages	Disadvantages	Bandwidth range	Sample protocols used
Leased line	Point-to-Point connection between two computers or Local Area Networks (LANs)	Most secure	Expensive		PPP, HDLC, SDLC, HNAS
Circuit switching	A dedicated circuit path is created between end points. Best example is dialup connections	Less Expensive	Call Setup	28 - 144 kbit/s	PPP, ISDN
Packet switching	Devices transport packets via a shared single point-to-point or point-to-multipoint link across a carrier internetwork. Variable length packets are transmitted over Permanent Virtual Circuits (PVC) or Switched Virtual Circuits (SVC)		Shared media across link		X.25 Frame-Relay
Cell relay	Similar to packet switching, but uses fixed length cells instead of variable length packets. Data is divided into fixed-length cells and then transported across	Best for simultaneous use of voice and data	Overhead can be considerable		ATM

virtual circuits

Transmission rates usually range from 1200 bit/s to 24 Mbit/s, although some connections such as ATM and Leased lines can reach speeds greater than 156 Mbit/s. Typical communication links used in WANs are telephone lines, microwave links & satellite channels.

Recently with the proliferation of low cost of Internet connectivity many companies and organizations have turned to VPN to interconnect their networks, creating a WAN in that way. Companies such as Cisco, New Edge Networks and Check Point offer solutions to create VPN networks.

Wireless network

Wireless network refers to any type of computer network that is wireless, and is commonly associated with a telecommunications network whose interconnections between nodes are implemented without the use of wires. Wireless telecommunications networks are generally implemented with some type of remote information transmission system that uses electromagnetic waves, such as radio waves, for the carrier and this implementation usually takes place at the physical level or "layer" of the network.

Types of wireless connections

Wireless PAN

Wireless Personal Area Networks (WPANs) interconnect devices within a relatively small area, generally within reach of a person. For example, Bluetooth provides a WPAN for interconnecting a headset to a laptop. ZigBee also supports WPAN applications. Wi-Fi PANs are also getting popular as vendors have started integrating Wi-Fi in variety of consumer electronic devices. Intel My WiFi and Windows 7 virtual Wi-Fi capabilities have made Wi-Fi PANs simpler and easier to set up and configure.

Wireless LAN

A wireless local area network (WLAN) links two or more devices using a wireless distribution method (typically spread-spectrum or OFDM radio), and usually providing a connection through an access point to the wider internet. This gives users the mobility to move around within a local coverage area and still be connected to the network.

- Wi-Fi: Wi-Fi is increasingly used as a synonym for 802.11 WLANs, although it is technically a certification of interoperability between 802.11 devices.
- Fixed Wireless Data: This implements point to point links between computers or networks at two locations, often using dedicated microwave or laser beams over line of sight paths. It is often used in cities to connect networks in two or more buildings without physically wiring the buildings together.

Wireless MAN

Wireless Metropolitan area networks are a type of wireless network that connects several Wireless LANs.

- WiMAX is the term used to refer to wireless MANs and is covered in IEEE 802.16d/802.16e.

Wireless WAN

wireless wide area networks are wireless networks that typically cover large outdoor areas. These networks can be used to connect branch offices of business or as a public internet access system. They are usually deployed on the 2.4 GHz band. A typical system contains base station gateways, access points and wireless bridging relays. Other configurations are mesh systems where each access point acts as a relay also. When combined with renewable energy systems such as photo-voltaic solar panels or wind systems they can be stand alone systems.

Mobile devices networks

With the development of smart phones, cellular telephone networks routinely carry data in addition to telephone conversations:

- Global System for Mobile Communications (GSM): The GSM network is divided into three major systems: the switching system, the base station system, and the operation and support system. The cell phone connects to the base system station which then connects to the operation and support station; it then connects to the switching station where the call is transferred to where it needs to go. GSM is the most common standard and is used for a majority of cell phones.
- Personal Communications Service (PCS): PCS is a radio band that can be used by mobile phones in North America and South Asia. Sprint happened to be the first service to set up a PCS.
- D-AMPS: Digital Advanced Mobile Phone Service, an upgraded version of AMPS, is being phased out due to advancement in technology. The newer GSM networks are replacing the older system.

Uses

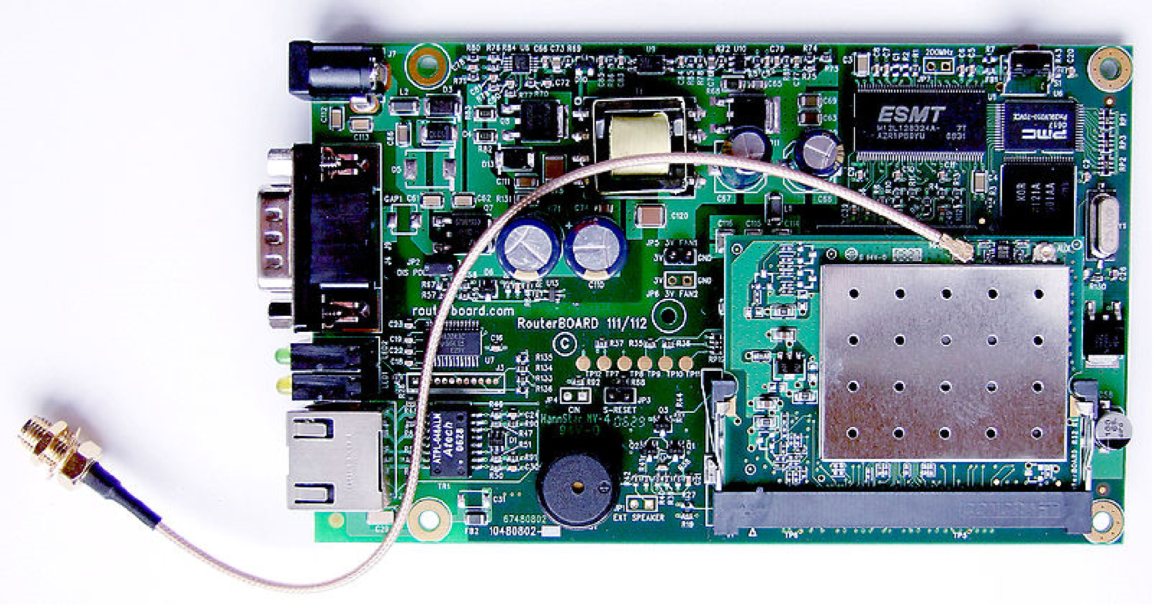

An embedded RouterBoard 112 with U.FL-RSMA pigtail and R52 mini PCI Wi-Fi card widely used by wireless Internet service providers (WISPs) in the Czech Republic.

Wireless networks have continued to develop and their uses have grown significantly. Cellular phones are part of huge wireless network systems. People use these phones daily to communicate with one another. Sending information overseas is possible through wireless network systems using satellites and other signals to communicate across the world. Emergency services such as the police department utilize wireless networks to communicate important information quickly. People and businesses use wireless networks to send and share data quickly whether it be in a small office building or across the world.

Another important use for wireless networks is as an inexpensive and rapid way to be connected to the Internet in countries and regions where the telecom infrastructure is poor or there is a lack of resources, as in most developing countries.

Compatibility issues also arise when dealing with wireless networks. Different components not made by the same company may not work together, or might require extra work to fix these issues. Wireless networks are typically slower than those that are directly connected through an Ethernet cable.

A wireless network is more vulnerable, because anyone can try to break into a network broadcasting a signal. Many networks offer WEP - Wired Equivalent Privacy - security systems which have been found to be vulnerable to intrusion. Though WEP does block some intruders, the security problems have caused some businesses to stick with wired networks until security can be improved. Another type of security for wireless networks is WPA - Wi-Fi Protected Access. WPA provides more security to wireless networks than a WEP security set up. The use of firewalls will help with security breaches which can help to fix security problems in some wireless networks that are more vulnerable.

Environmental concerns and health hazard

In recent times, there have been increased concerns about the safety of wireless communications, despite little evidence of health risks so far. The president of Lakehead University refused to agree to installation of a wireless network citing a California Public Utilities Commission study which said that the possible risk of tumors and other diseases due to exposure to electromagnetic fields (EMFs) needs to be further investigated.

Chapter 14

Bootstrap Protocol and Preboot Execution Environment

Bootstrap Protocol

In computer networking, the **Bootstrap Protocol**, or **BOOTP**, is a network protocol used by a network client to obtain an IP address from a configuration server. The BOOTP protocol was originally defined in RFC 951.

BOOTP is usually used during the bootstrap process when a computer is starting up. A BOOTP configuration server assigns an IP address to each client from a pool of addresses. BOOTP uses the User Datagram Protocol (UDP) as a transport on IPv4 networks only.

Historically, BOOTP has also been used for Unix-like diskless workstations to obtain the network location of their boot image in addition to an IP address, and also by enterprises to roll out a pre-configured client (e.g., Windows) installation to newly installed PCs.

Originally requiring the use of a boot floppy disk to establish the initial network connection, manufacturers of network cards later embedded the protocol in the BIOS of the interface cards as well as system boards with on-board network adapters, thus allowing direct network booting.

Recently, users with an interest in diskless stand-alone media center PCs have shown new interest in this method of booting a Windows operating system.

The Dynamic Host Configuration Protocol (DHCP) is a more advanced protocol for the same purpose and has superseded the use of BOOTP. Most DHCP servers also function as BOOTP servers.

History

The BOOTP protocol was first defined in RFC 951 as a replacement for the Reverse Address Resolution Protocol RARP, published in RFC 903 in June 1984. The primary motivation for replacing RARP with BOOTP is that RARP was a data link layer protocol. This made implementation difficult on many server platforms, and required that a server be present on each individual IP subnet. BOOTP introduced the innovation of a **relay agent**, which allowed BOOTP packets to be forwarded from the local network using standard IP routing, so that one central BOOTP server could serve hosts on many subnets.

Related RFCs

BOOTP Related RFC's

Note that grayed out RFCs are obsolete

RFC #	Title	Date	Obsolete and Update Information
RFC 3942	Reclassifying Dynamic Host Configuration Protocol version 4 (DHCPv4) Options	Nov-04	Updates RFC 2132
RFC 2132	DHCP Options and BOOTP Vendor Extensions	Mar-97	Obsoletes RFC 1533, Updated by RFC 3442, RFC 3942, RFC 4361, RFC 4833, RFC 5494
RFC 1542	Clarifications and Extensions for the Bootstrap Protocol	Oct-93	Obsoletes RFC 1532, Updates RFC 951
RFC 1534	Interoperation Between DHCP and BOOTP	Oct-93	
RFC 1533	DHCP Options and BOOTP Vendor Extensions	Oct-93	Obsoletes RFC 1497, RFC 1395, RFC 1084, RFC 1048, Obsoleted by RFC 2132
RFC 1532	Clarifications and Extensions for the Bootstrap Protocol	Oct-93	Obsoleted by RFC 1542, Updates RFC 951
RFC 1497	BOOTP Vendor Information Extensions	Aug-93	Obsoletes RFC 1395, RFC 1084, RFC 1048, Obsoleted by RFC 1533, Updates RFC 951
RFC 1395	BOOTP Vendor Information Extensions	Jan-93	Obsoletes RFC 1084, RFC 1048, Obsoleted by RFC 1497, RFC 1533, Updates RFC 951
RFC 1084	BOOTP vendor information extensions	Dec-88	Obsoletes RFC 1048, Obsoleted by RFC 1395, RFC 1497, RFC 1533
RFC 1048	BOOTP vendor information extensions	Feb-88	Obsoleted by RFC 1084, RFC 1395, RFC 1497, RFC 1533
RFC 0951	Bootstrap Protocol	Sep-85	Updated by RFC 1395, RFC 1497, RFC 1532, RFC 1542, RFC 5494

Preboot Execution Environment

The **Preboot eXecution Environment** (**PXE**, and also known as Pre-Execution Environment) is an environment to boot computers using a network interface independently of data storage devices (like hard disks) or installed operating systems.

PXE was introduced as part of the Wired for Management framework by Intel and is described in the specification (version 2.1) published by Intel and Systemsoft on September 20, 1999. It makes use of several network protocols like Internet Protocol (IP), User Datagram Protocol (UDP), Dynamic Host Configuration Protocol (DHCP) and Trivial File Transfer Protocol (TFTP) and of concepts like Globally Unique Identifier (GUID), Universally Unique Identifier (UUID) and Universal Network Device Interface and extends the firmware of the PXE client (the computer to be bootstrapped via PXE) with a set of predefined Application Programming Interfaces (APIs).

The term *PXE client* only refers to the role that the machine takes in the PXE boot process. A *PXE client* can be a server, desktop, laptop or any other machine that is equipped with PXE boot code.

Chain

The firmware on the client tries to locate a PXE redirection service on the network (Proxy DHCP) in order to receive information about available PXE boot servers. After parsing the answer, the firmware will ask an appropriate boot server for the file path of a network bootstrap program (NBP), download it into the computer's random-access memory (RAM) using TFTP, possibly verify it, and finally execute it. If only one NBP is used among all PXE clients it could be specified using BOOTP without any need of a proxy DHCP, but a TFTP boot server is still required.

Availability

PXE was designed to be applicable to many system architectures. The 2.1 version of the specification assigns architecture identifiers to six system types, including IA-64 and DEC Alpha. However, the specification only completely covers IA-32. Intel included PXE in the EFI for IA-64, creating a de-facto standard with the implementation.

Protocol

The PXE protocol is approximately a combination of DHCP and TFTP, albeit with subtle modifications to both. DHCP is used to locate the appropriate boot server or servers, with TFTP used to download the initial bootstrap program and additional files.

To initiate a PXE bootstrap session the PXE firmware broadcasts a DHCPDISCOVER packet extended with PXE-specific options (*extended DHCPDISCOVER*) to port 67/UDP (DHCP server port). The PXE options identify the firmware as capable of PXE, but they will be ignored by standard DHCP servers. If the firmware receives DHCPOFFERs from such servers, it may configure itself by requesting one of the offered configurations.

Proxy DHCP

If a PXE redirection service (Proxy DHCP) receives an *extended DHCPDISCOVER*, it replies by sending a DHCPOFFER packet extended with PXE-specific options (*extended DHCPOFFER*) to the client to port 68/UDP (DHCP client port).

An *extended DHCPOFFER* contains mainly:

- a PXE Discovery Control field to decide whether Multicasting, Broadcasting, or Unicasting is to be used for contacting PXE boot servers
- a list of IP addresses of each available PXE Boot Server Type
- a PXE Boot Menu with each entry representing a PXE Boot Server Type
- a PXE Boot Prompt telling the user to press a certain key to see the boot menu
- a timeout to launch the first boot menu entry if it expires.

The Proxy DHCP service may also be run on the same host as the standard DHCP service. Since both services cannot share port 67/UDP, the Proxy DHCP runs on port 4011/UDP and expects the *extended DHCPDISCOVER packets* from PXE Clients to be DHCPREQUESTs. The standard DHCP service has to send a special combination of PXE options in its DHCPOFFER, so the PXE client knows to look for a Proxy DHCP on the same host, port 4011/UDP.

Boot server contact

To contact a PXE Boot Server the booting system must have an IP address (perhaps from a DHCP server).

It multicasts or unicasts a *DHCPREQUEST* packet extended with PXE-specific options (*extended DHCPREQUEST*) to port 4011/UDP or broadcasts it to port 67/UDP. This packet contains the PXE Boot Server **type** and the PXE Boot Layer, allowing multiple boot server types to run from one daemon. The *extended DHCPREQUEST* may be a *DHCPINFORM*.

A PXE Boot Server receiving an *extended DHCPREQUEST* configured for the requested **type** and client architecture responds with an *extended DHCPACK* including:

- the complete file path to download the NBP via TFTP.
- PXE Boot Server **type** and PXE Boot Layer it answered

- the multicast TFTP configuration, if MTFTP as described in the PXE specification should be used.

The booting system accepts information from only one *extended DHCPOFFER*.

A 2.1 version PXE Boot Server supports "Boot Integrity Services" () allowing the Client to verify downloaded NBPs using a checksum file which is downloaded from the same boot server as the NBP.

To get the file path of this *credentials* file another exchange of *extended DHCPREQUEST* and *extended DHCPACK* is required.

Network bootstrap program

After receiving the requested *extended DHCPACK*, the *Network Bootstrap Program* is uploaded into RAM and after it is verified or if verification is not required, the NBP will be executed. It has access to the APIs of the PXE firmware extension (Pre-boot, UDP, TFTP, Universal Network Device Interface (UNDI)). Its functions or tasks are not described in the PXE specification.

Integration

The *PXE Client/Server Protocol* was designed so:

- it can be used in the same network as an existing DHCP environment without interference
- it can be integrated completely into standard DHCP services
- it can be easily extended at the most important points without a call for papers
- every service (DHCP, Proxy DHCP, Boot Server) can be implemented standalone or in any combination of them.

Additionally the PXE firmware extension was designed as an Option ROM for the IA-32 BIOS so you can get a personal computer (PC) PXE-capable by installing a NIC that provides a PXE Option ROM. Note, this procedure also applies to the newer AMD64 processor standard for PC.

The design goal of utilizing existing DHCP and TFTP servers cannot be achieved in a strictly conforming implementation. Some aspects of the PXE protocol require that the DHCP and TFTP servers be modified and communicate. One specific example is using multicast, where DHCP packets provide the multicast group information rather than an opening RFC-2090 multicast TFTP exchange. The impact of this is minimal as the most common PXE client implementation (written by Intel and provided at no cost as a linkable IA32 binary module) interoperates with a combination of isolated DHCP and unicast TFTP servers.

Chapter 15

Computer Networks and Internet Technology

Computer networks

When two or more computers are interconnected to each other via some kind of medium to share resources then the state is called computer network.

Connection method

Computer networks can also be classified according to the hardware and software technology that is used to interconnect the individual devices in the network, such as Optical fiber, Ethernet, Wireless LAN, HomePNA, Power line communication or G.hn. Ethernet uses physical wiring to connect devices. Frequently deployed devices include hubs, network switches, network bridges and/or routers.

Wireless LAN technology is designed to connect devices without wiring. These devices use radio waves or infrared signals as a transmission medium.

Ethernet over coax technology uses existing home wiring (coaxial cable, phone lines and power lines) to create a high-speed (up to 1 Gigabit/s) local area network.

Wired technologies

Twisted pair - This is the most widely used medium for telecommunication. Twisted-pair wires are ordinary telephone wires which consist of two insulated copper wires twisted into pairs and are used for both voice and data transmission. The use of two wires twisted together helps to reduce crosstalk and electromagnetic induction. The transmission speed range from 2 million bits per second to 100 million bits per second.

Coaxial cable – These cables are widely used for cable television systems, office buildings, and other worksites for local area networks. The cables consist of copper or aluminum wire wrapped with insulating layer typically of a flexible material with a high dielectric constant, all of which are surrounded by a conductive layer. The layers of insulation help minimize interference and distortion. Transmission speed range from 200 million to more than 500 million bits per second.

Fiber optics – These cables consist of one or more thin filaments of glass fiber wrapped in a protective layer. It transmits light which can travel over long distance and higher bandwidths. Fiber-optic cables are not affected by electromagnetic radiation. Transmission speed could go up to as high as trillions of bits per second. The speed of fiber optics is hundreds of times faster than coaxial cables and thousands of times faster than twisted-pair wire.

Wireless technologies

Terrestrial microwave – Terrestrial microwaves use Earth-based transmitter and receiver. The equipment look similar to satellite dishes. Terrestrial microwaves use low-gigahertz range, which limits all communications to line-of-sight. Path between relay stations spaced approx. 30 miles apart. Microwave antennas are usually placed on top of buildings, towers, hills, and mountain peaks.

Communications satellites – The satellites use microwave radio as their telecommunications medium which are not deflected by the Earth's atmosphere. The satellites are stationed in space, typically 22,000 miles above the equator. These Earth-orbiting systems are capable of receiving and relaying voice, data, and TV signals.

Cellular and PCS Systems – Use several radio communications technologies. The systems are divided to different geographic area. Each area has low-power transmitter or radio relay antenna device to relay calls from one area to the next area.

Wireless LANs – Wireless local area network use a high-frequency radio technology similar to digital cellular and a low-frequency radio technology. Wireless LANS use spread spectrum technology to enable communication between multiple devices in a limited area. Example of open-standard wireless radio-wave technology is IEEE 802.11b.

Bluetooth – A short range wireless technology. Operate at approx. 1Mbps with range from 10 to 100 meters. Bluetooth is an open wireless protocol for data exchange over short distances.

The wireless Web – The wireless web refers to the use of the World Wide Web through equipments like cellular phones, pagers, PDAs, and other portable communication devices. The wireless web service offers anytime/anywhere connection.

Scale

Networks are often classified as LAN, WAN, MAN, PAN, VPN, CAN, SAN, etc. depending on their scale, scope and purpose. Usage, trust levels and access rights often differ between these types of network - for example, LANs tend to be designed for internal use by an organization's internal systems and employees in individual physical locations (such as a building), while WANs may connect physically separate parts of an organization to each other and may include connections to third parties.

Functional relationship (network architecture)

Computer networks may be classified according to the functional relationships which exist among the elements of the network, e.g., Active Networking, Client-server and Peer-to-peer (workgroup) architecture.

Network topology

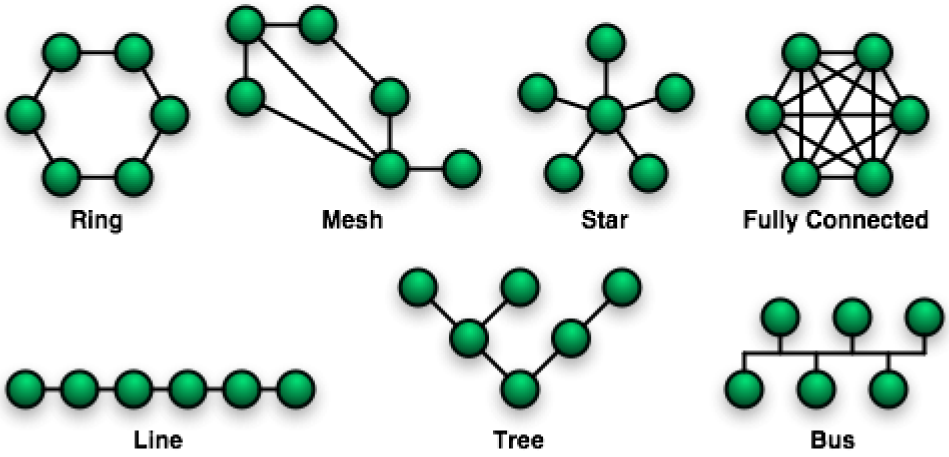

Network Topologies examples

Computer networks may be classified according to the network topology upon which the network is based, such as bus network, star network, ring network, mesh network, star-bus network, tree or hierarchical topology network. Network topology signifies the way in which devices in the network see their logical relations to one another. The use of the term "logical" here is significant. That is, network topology is independent of the "physical" layout of the network. Even if networked computers are physically placed in a linear arrangement, if they are connected via a hub, the network has a Star topology, rather than a bus topology. In this regard the visual and operational characteristics of a network are distinct; the logical network topology is not necessarily the same as the physical layout. Networks may be classified based on the method of data used to convey the data, these include digital and analog networks.

OSI-ISO Model

The Open System Interconnection Reference Model (OSI Reference Model or OSI Model) is an abstract description for layered communications and computer network protocoldesign. It was developed as part of the Open Systems Interconnection (OSI) initiative.[1 In its most basic form, it divides network architecture into seven layers which, from top to bottom, are the Application, Presentation, Session, Transport, Network, Data-Link, and Physical Layers. It is therefore often referred to as the **OSI Seven Layer Model**.

A layer is a collection of conceptually similar functions that provide services to the layer above it and receives service from the layer below it. On each layer an instance provides services to the instances at the layer above and requests service from the layer below. For example, a layer that provides error-free communications across a network provides the path needed by applications above it, while it calls the next lower layer to send and receive packets that make up the contents of the path. Conceptually two instances at one layer are connected by a horizontal protocol connection on that layer.

General communication model

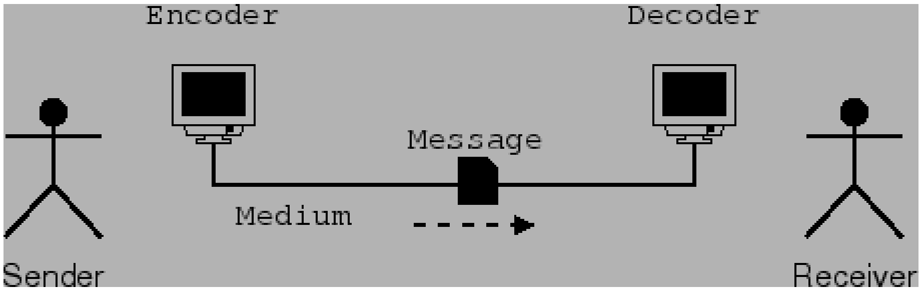

General Communication Model

Communication model

The parts of this model are as follows:

Sender: The sender is what or who is trying to send a message to the receiver.

Encoder: In the general case, it is not possible to directly insert the message onto the communications medium. For instance, when you speak on the telephone, it is not possible to actually transmit sound (vibrations in matter) across the wire for any distance. In your phone is a microphone, which converts the sound into electrical impulses, which can be transmitted by wires. Those electrical impulses are then manipulated by the electronics in the phone so they match up with what the telephone system expects.

Message: Since this is a communication engineer's model, the message is the actual encoded message that is transmitted by the medium.

Medium: The medium is what the message is transmitted on. The phone system, Internet, and many other electronic systems use wires. Television and radio can use electromagnetic radiation. Even bongo drums can be used as a medium.

Decoder: The decoder takes the encoded message and converts it to a form the receiver understands, since for example a human user of the phone system does not understand electrical impulses directly.

Chapter 16

Computer Network Diagram and Electronic Data Interchange

Computer network diagram

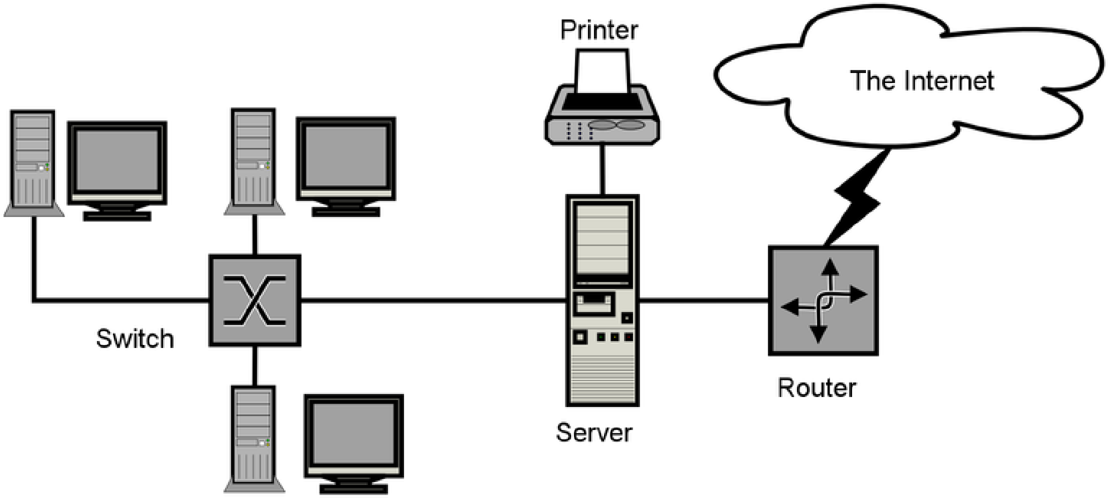

A sample network diagram

A **computer network diagram** is a schematic depicting the nodes and connections amongst nodes in a computer network or, more generally, any telecommunications network.

Symbolization

Readily identifiable icons are used to depict common network appliances e.g. Router, and the style of lines between them indicate the type of connection. Clouds are used to represent networks external to the one pictured for the purposes of depicting connections between internal and external devices, without indicating the specifics of the outside network. For example, in the hypothetical local area network pictured to the right, three personal computers and a server are connected to a switch; the server is further connected to a printer and a gateway router, which is connected via a WAN link to the Internet.

Depending on whether the diagram is intended for formal or informal use, certain details may be lacking and must be determined from context. For example, the sample diagram does not indicate the physical type of connection between the PCs and the switch, but since a modern LAN is depicted, Ethernet may be assumed. If the same style of line was used in a WAN (wide area network) diagram, however, it may indicate a different physical connection.

At different scales, or sizes, diagrams may represent various levels of network granularity. At the LAN level, individual nodes may represent individual physical devices, such as hubs or file servers, while at the WAN level, individual nodes may represent entire cities. In addition, when the scope of a diagram crosses the common LAN/MAN/WAN boundaries, representative hypothetical devices may be depicted instead of showing all actually existing nodes. For example, if a network appliance is intended to be connected through the Internet to many end-user mobile devices, only a single such device may be depicted for the purposes of showing the general relationship between the appliance and any such device.

Cisco Symbolization

Cisco uses its own brand of networking symbols. Since Cisco has a large Internet presence and designs a broad variety of network devices, its list of symbols ("Network Topology Icons") is exhaustive.

Topology

The *physical* network topology can be directly represented in a network diagram, as it is simply the physical graph (mathematics) represented by the diagrams, with network nodes as vertices and connections as undirected or direct edges (depending on the type of connection). The *logical* network topology can be inferred from the network diagram if details of the network protocols in use are also given.

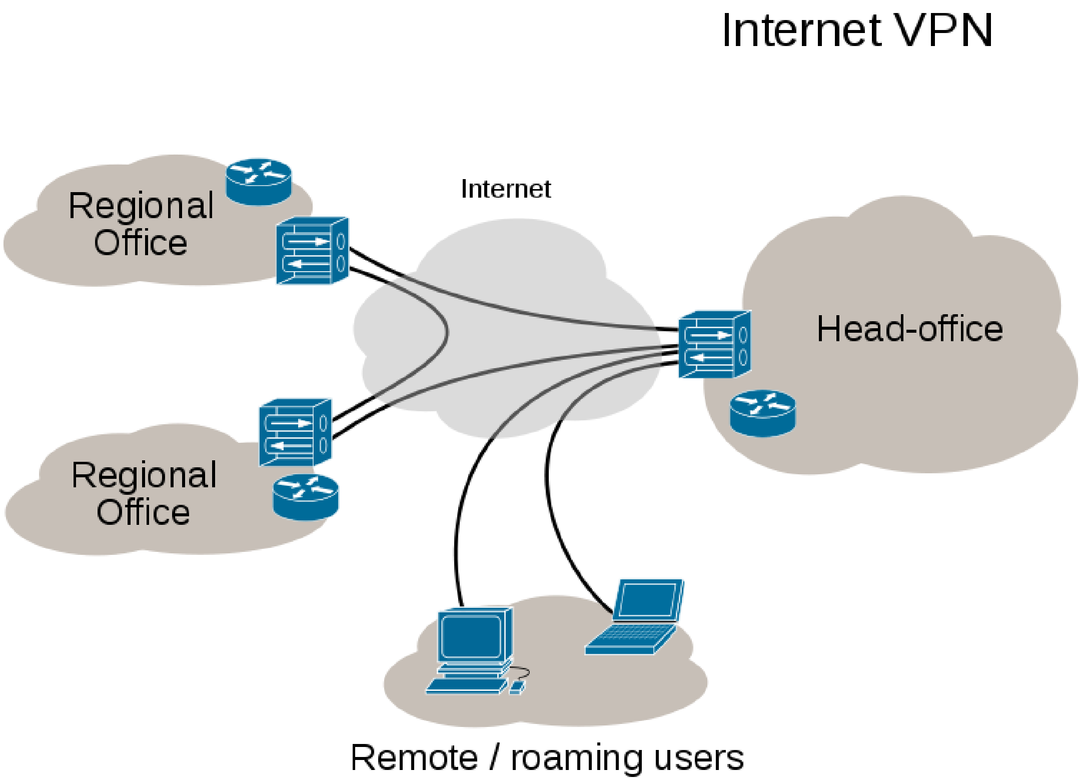

Virtual Private Network

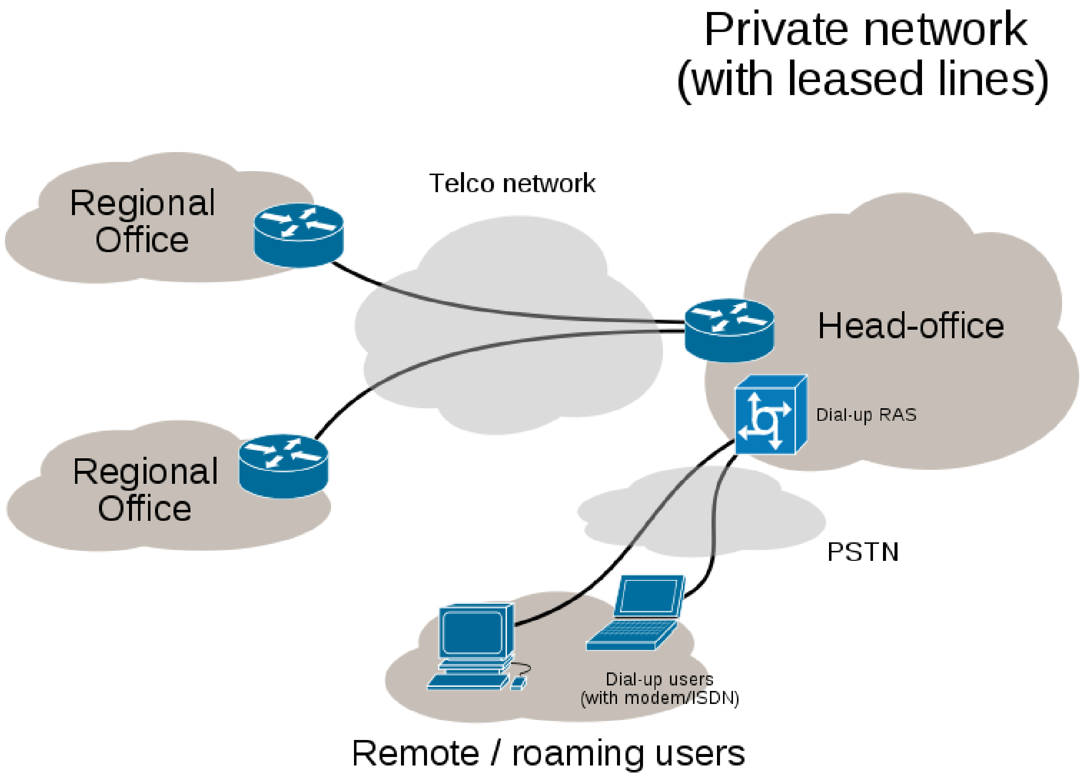

Enterprise private network

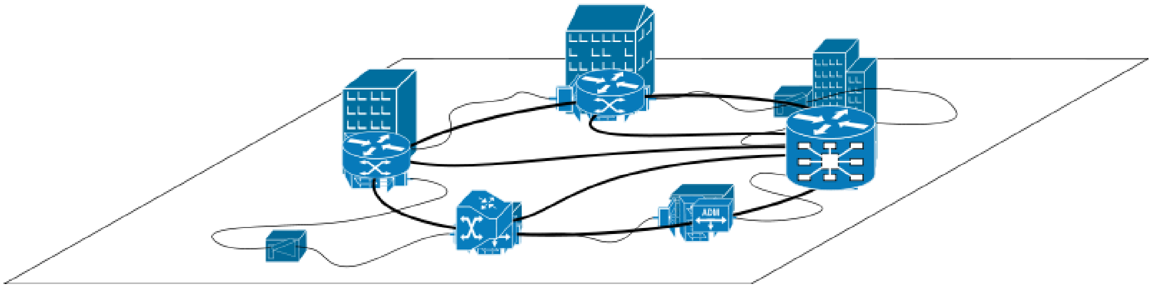

Overlay network collapsed

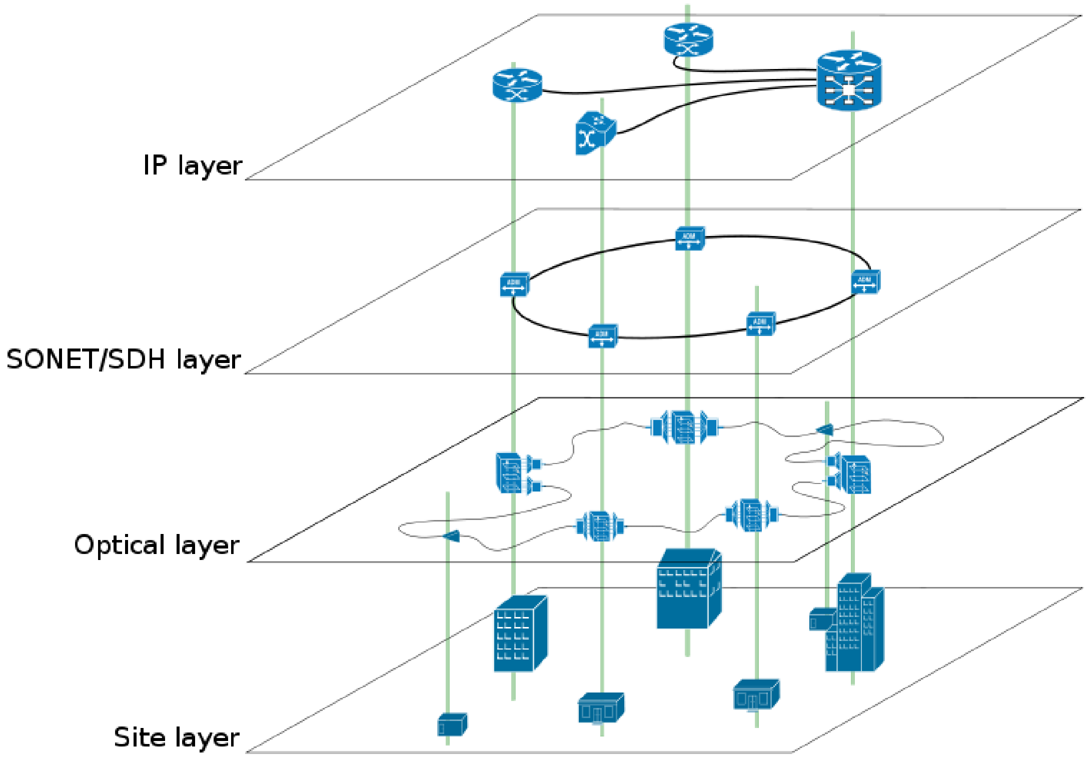

Overlay network broken-up

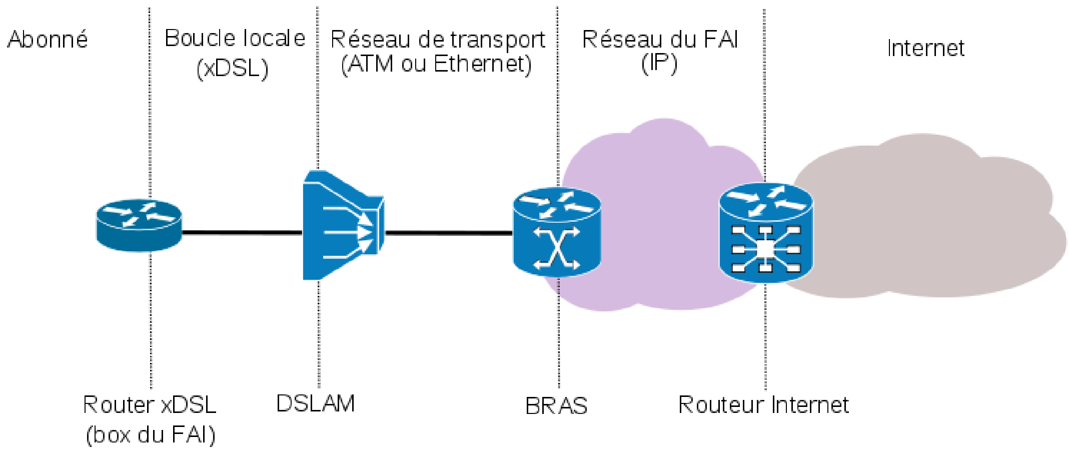

xDSL to Internet connectivity

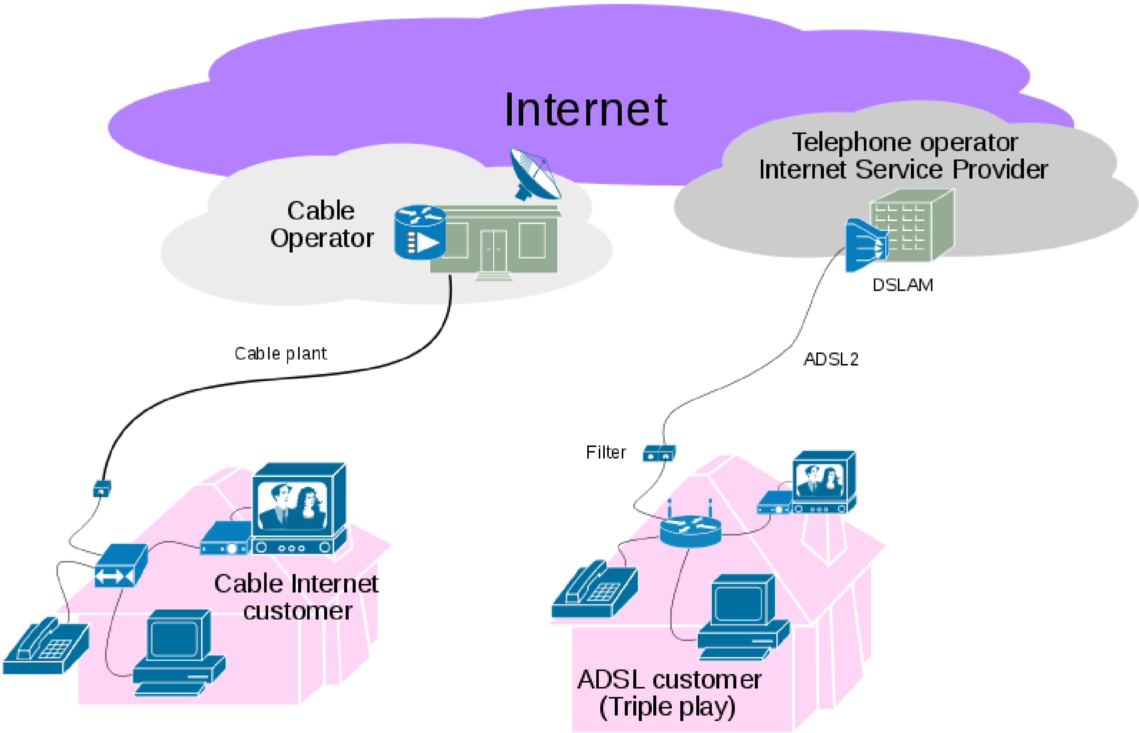

Triple play illustration

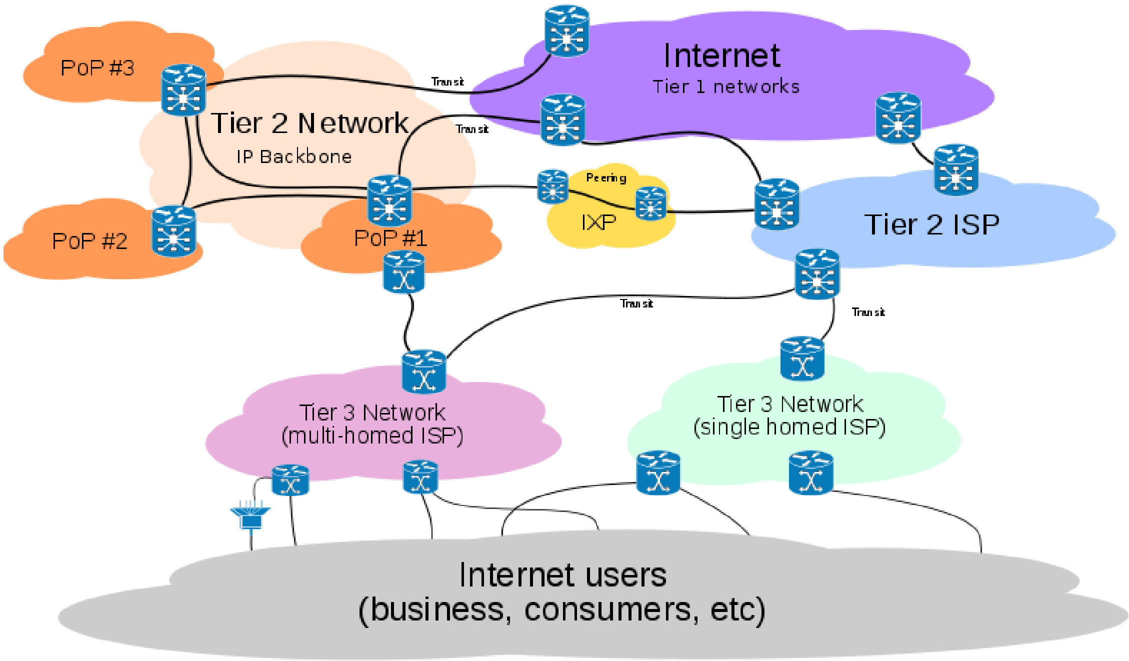

Internet Distribution and Core

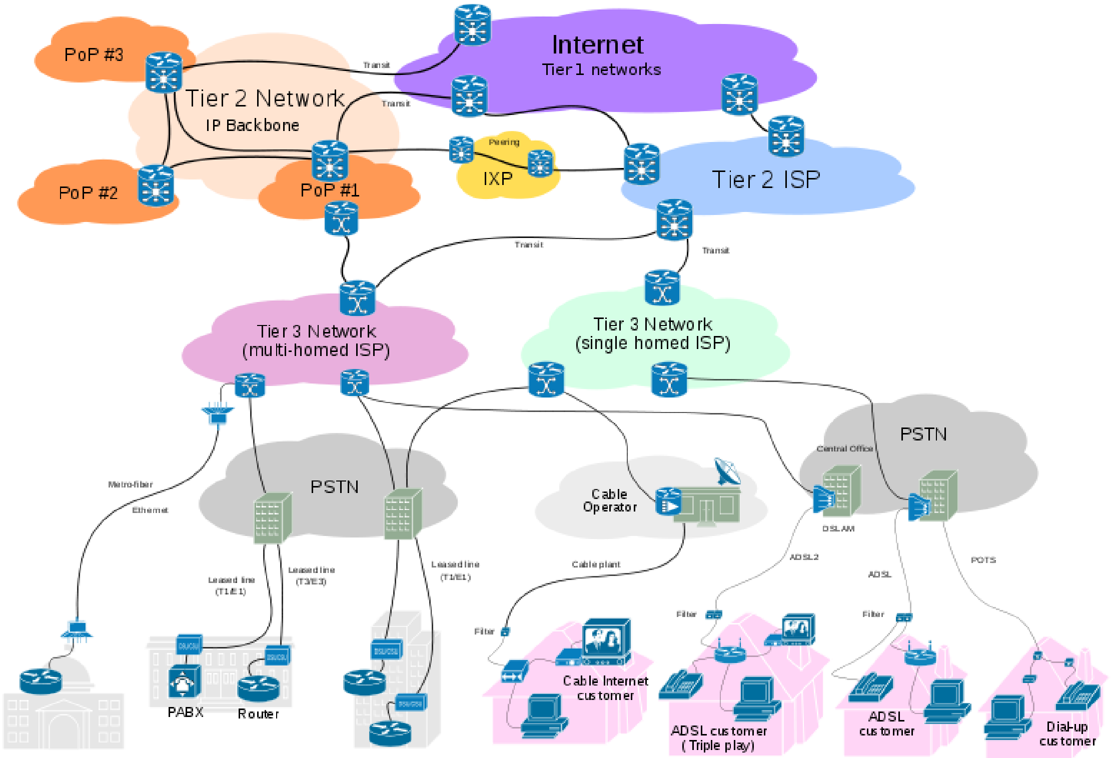

Internet, Access to core

Electronic Data Interchange

Electronic data interchange (EDI) is the structured transmission of data between organizations by electronic means. It is used to transfer electronic documents or business data from one computer system to another computer system, i.e. from one trading partner to another trading partner without human intervention.

It is more than mere e-mail; for instance, organizations might replace bills of lading and even cheques with appropriate EDI messages. It also refers specifically to a family of standards, e.g. UN/EDIFACT, ANSI X12.

The National Institute of Standards and Technology in a 1996 publication defines electronic data interchange as "the computer-to-computer interchange of strictly formatted messages that represent documents other than monetary instruments. EDI implies a sequence of messages between two parties, either of whom may serve as originator or recipient. The formatted data representing the documents may be transmitted from originator to recipient via telecommunications or physically transported on electronic storage media.". It goes on further to say that "In EDI, the usual processing of received messages is by computer only. Human intervention in the processing of a received message is typically intended only for error conditions, for quality review, and for special situations. **For example, the transmission of binary or textual data is not EDI as defined here unless the data are treated as one or more data elements of an EDI message and are not normally intended for human interpretation as part of online data processing."**

EDI can be formally defined as 'The transfer of structured data, by agreed message standards, from one computer system to another without human intervention'. Most other definitions used are variations on this theme. Even in this era of technologies such as XML web services, the Internet and the World Wide Web, EDI may be the data format used by the vast majority of electronic commerce transactions in the world.

Standards

EDI is considered to be a technical representation of a business conversation between two entities, either internal or external. Note that there is a perception that "EDI" constitutes the entire electronic data interchange paradigm, including the transmission, message flow, document format, and software used to interpret the documents. EDI is considered to describe the rigorously standardized format of electronic documents. EDI is very useful in supply chain.

The EDI standards were designed to be independent of communication and software technologies. EDI can be transmitted using any methodology agreed to by the sender and recipient. This includes a variety of technologies, including modem (asynchronous, and bisynchronous), FTP, E-mail, HTTP, AS1, AS2, etc. It is important to differentiate between the EDI documents and the methods for transmitting them. When they compared the bisynchronous protocol 2400 bit/s modems, CLEO devices, and value-added networks used to transmit EDI documents to transmitting via the Internet, some people equated the non-Internet technologies with EDI and predicted erroneously that EDI itself would be replaced along with the non-Internet technologies. These non-internet transmission methods are being replaced by Internet Protocols such as FTP, telnet, and E-mail, but the EDI documents themselves still remain.

As more trading partners use the Internet for transmission, standards have emerged. In 2002, the IETF published RFC 3335, offering a standardized, secure method of transferring EDI data via e-mail. On July 12, 2005, an IETF working group ratified RFC4130 for MIME-based HTTP EDIINT (aka. AS2) transfers, and is preparing a similar RFC for FTP transfers (aka. AS3). While some EDI transmission has moved to these newer protocols, the providers of the value-added networks remain active.

EDI documents generally contain the same information that would normally be found in a paper document used for the same organizational function. For example an EDI 940 ship-from-warehouse order is used by a manufacturer to tell a warehouse to ship product to a retailer. It typically has a ship to address, bill to address, a list of product numbers (usually a UPC code) and quantities. Another example is the set of messages between sellers and buyers, such as request for quotation (RFQ), bid in response to RFQ, purchase order, purchase order acknowledgment, shipping notice, receiving advice, invoice, and payment advice . However, EDI is not confined to just business data related to trade but encompasses all fields such as medicine (e.g., patient records and laboratory results), transport (e.g., container and modal information), engineering and construction, etc. In some cases, EDI will be used to create a new business information flow (that was not a paper flow before). This is the case in the Advanced Shipment Notification (856) which was designed to inform the receiver of a shipment, the goods to be received and how the goods are packaged.

There are four major sets of **EDI standards**:

- The UN-recommended UN/EDIFACT is the only international standard and is predominant outside of North America.
- The US standard ANSI ASC X12 (X12) is predominant in North America.
- The TRADACOMS standard developed by the ANA (Article Numbering Association) is predominant in the UK retail industry.
- The ODETTE standard used within the European automotive industry

All of these standards first appeared in the early to mid 1980s. The standards prescribe the formats, character sets, and data elements used in the exchange of business documents and forms. The complete X12 Document List includes all major business documents, including purchase orders (called "ORDERS" in UN/EDIFACT and an "850" in X12) and invoices (called "INVOIC" in UN/EDIFACT and an "810" in X12).

The EDI standard says which pieces of information are mandatory for a particular document, which pieces are optional and give the rules for the structure of the document. The standards are like building codes. Just as two kitchens can be built "to code" but look completely different, two EDI documents can follow the same standard and contain different sets of information. For example a food company may indicate a product's expiration date while a clothing manufacturer would choose to send color and size information.

Specifications

Organizations that send or receive documents between each other are referred to as "trading partners" in EDI terminology. The trading partners agree on the specific information to be transmitted and how it should be used. This is done in human readable specifications (also called Message Implementation Guidelines). While the standards are analogous to building codes, the specifications are analogous to blue prints. (The specification may also be called a mapping but the term mapping is typically reserved for specific machine readable instructions given to the translation software.) Larger trading "hubs" have existing Message Implementation Guidelines which mirror their business processes for processing EDI and they are usually unwilling to modify their EDI business practices to meet the needs of their trading partners. Often in a large company these EDI guidelines will be written to be generic enough to be used by different branches or divisions and therefore will contain information not needed for a particular business document exchange. For other large companies, they may create separate EDI guidelines for each branch/division.

Transmission

Trading partners are free to use any method for the transmission of documents. In the past one of the more popular methods was the usage of a bisync modem to communicate through a value added network (VAN). Some organizations have used direct modem to modem connections and bulletin board systems (BBS), and recently there has been a move towards using some of the many Internet protocols for transmission, but most EDI is still transmitted using a VAN. In the healthcare industry, a VAN is referred to as a "clearinghouse".

Value-added networks

In the most basic form, a VAN (value-added network) acts as a regional post office. They receive transactions, examine the 'from' and the 'to' information, and route the transaction to the final recipient. VANs provide a number of additional services, e.g. retransmitting documents, providing third party audit information, acting as a gateway for different transmission methods, and handling telecommunications support. Because of these and other services VANs provide, businesses frequently use a VAN even when both trading partners are using Internet-based protocols. Healthcare clearinghouses perform many of the same functions as a VAN, but have additional legal restrictions that govern protected healthcare information.

VANs also provide an advantage with certificate replacement in AS2 transmissions. Because each node in a traditionally business-related AS2 transmission usually involves a security certificate, routing a large number of partners through a VAN can make certificate replacement much easier. Value Added Networks

- Value Added Networks are the go-between in EDI communications.
- The VAN is responsible for routing, storing and delivering EDI messages. They also provide delivery reports
- Depending on the VAN type, messages may need extra envelopes or may be routed using intelligent *VANs which are able to read the EDI message itself.
- VANs may be operated by various entities:
    - telecom companies;
    - industry group consortia;
    - a large company interacting with its suppliers/vendors.

Internet/AS2

Until recently, the Internet transmission was handled by nonstandard methods between trading partners usually involving FTP or email attachments. There are also standards for embedding EDI documents into XML. Many organizations are migrating to this protocol to reduce costs. For example, Wal-Mart is now requiring its trading partners to switch to the AS2 protocol (Wal-Mart EDI Requirement).

AS2 (Applicability Statement 2) is the draft specification standard by which vendor applications communicate EDI or other business-to-business data (such as XML) over the Internet using HTTP, a standard used by the World Wide Web. AS2 provides security for the transport payload through digital signatures and data encryption, and ensures reliable, non-repudiable delivery through the use of receipts.

EDI via the Internet (Web EDI)

The Internet, as with VAN providers, uses its own communications protocols to ensure that EDI documents are transmitted securely. The most popular protocols are File Transfer Protocol Secure (FTPS), Hyper Text Transfer Protocol Secure (HTTPS), and AS2.

The Internet has provided a means for any company, no matter how small or where they are located in the world, to become part of a major supply chain initiative hosted by a global retailer or manufacturing company. Many companies around the world have shifted production of labour intensive parts to low-cost, emerging regions such as Brazil, Russia, India, China, and Eastern Europe. Web-based EDI, or webEDI, allows a company to interact with its suppliers in these regions without the worry of implementing a complex EDI infrastructure.

In its simplest form, webEDI enables small to medium-sized businesses to receive, turn around, create and manage electronic documents using just a web browser. This service seamlessly transforms your data into EDI format and transmits it to your trading partner. Simple pre-populated forms enable businesses to communicate and comply with their trading partners' requirements using built-in business rules. Using a friendly web-based interface, EDI transactions can be received, edited and sent as easily as an email. You will also be able to receive EDI documents and send EDI invoices and shipping documents with no software to install. All you require is an Internet connection. WebEDI has the added advantages that it is accessible anywhere in the world and you do not need a dedicated IT person to manage any software installation.

Even though VANs offer a very secure and reliable service to companies wishing to trade electronically, the Internet is making EDI more available to all. This is especially important in the emerging markets where IT awareness and infrastructure are very limited. WebEDI is traditionally based on the "hub and spoke'"model, with major trading partners or Application Service Providers (ASPs) being the hubs and smaller partners being the spokes.

- Hubs or ASPs implement EDI using email or virtual mailboxes
- Trading partners can send EDI messages directly to a web-enabled EDI messaging site, via the hub. EDI messages are simply sent using a web browser
- Systems that are currently being developed will enable EDI messages to be displayed in a web browser and directed via open standard XML, directly into the user's accounts system
- WebEDI-based users can interact with VANs without incurring the costs of setting up a dedicated VAN connection

Interpreting data

Often missing from the EDI specifications (referred to as EDI Implementation Guidelines) are real world descriptions of how the information should be interpreted by the business receiving it. For example, suppose candy is packaged in a large box that contains 5 display boxes and each display box contains 24 boxes of candy packaged for the consumer. If an EDI document says to ship 10 boxes of candy it may not be clear whether to ship 10 consumer packaged boxes, 240 consumer packaged boxes or 1200 consumer packaged boxes. It is not enough for two parties to agree to use a particular qualifier indicating case, pack, box or each; they must also agree on what that particular qualifier means.

EDI translation software provides the interface between internal systems and the EDI format sent/received. For an "inbound" document the EDI solution will receive the file (either via a Value Added Network or directly using protocols such as FTP or AS2), take the received EDI file (commonly referred to as a "mailbag"), validate that the trading partner who is sending the file is a valid trading partner, that the structure of the file meets the EDI standards and that the individual fields of information conforms to the agreed upon standards. Typically the translator will either create a file of either fixed length, variable length or XML tagged format or "print" the received EDI document (for non-integrated EDI environments). The next step is to convert/transform the file that the translator creates into a format that can be imported into a company's back-end business systems or ERP. This can be accomplished by using a custom program, an integrated proprietary "mapper" or to use an integrated standards based graphical "mapper" using a standard data transformation language such as XSLT. The final step is to import the transformed file (or database) into the company's back-end enterprise resource planning (ERP) system.

For an "outbound" document the process for integrated EDI is to export a file (or read a database) from a company's back-end ERP, transform the file to the appropriate format for the translator. The translation software will then "validate" the EDI file sent to ensure that it meets the standard agreed upon by the trading partners, convert the file into "EDI" format (adding in the appropriate identifiers and control structures) and send the file to the trading partner (using the appropriate communications protocol).

Another critical component of any EDI translation software is a complete "audit" of all the steps to move business documents between trading partners. The audit ensures that any transaction (which in reality is a business document) can be tracked to ensure that they are not lost. In case of a retailer sending a Purchase Order to a supplier, if the Purchase Order is "lost" anywhere in the business process, the effect is devastating to both businesses. To the supplier, they do not fulfill the order as they have not received it thereby losing business and damaging the business relationship with their retail client. For the retailer, they have a stock outage and the effect is lost sales, reduced customer service and ultimately lower profits.

In EDI terminology "inbound" and "outbound" refer to the direction of transmission of an EDI document in relation to a particular system, not the direction of merchandise, money or other things represented by the document. For example, an EDI document that tells a warehouse to perform an outbound shipment is an inbound document in relation to the warehouse computer system. It is an outbound document in relation to the manufacturer or dealer that transmitted the document.

Advantages of using EDI over paper systems

EDI and other similar technologies save a company money by providing an alternative to, or replacing information flows that require a great deal of human interaction and materials such as paper documents, meetings, faxes, etc. Even when paper documents are maintained in parallel with EDI exchange, e.g. printed shipping manifests, electronic exchange and the use of data from that exchange reduces the handling costs of sorting, distributing, organizing, and searching paper documents. EDI and similar technologies allow a company to take advantage of the benefits of storing and manipulating data electronically without the cost of manual entry. Another advantage of EDI is reduced errors, such as shipping and billing errors, because EDI eliminates the need to rekey documents on the destination side. One very important advantage of EDI over paper documents is the speed in which the trading partner receives and incorporates the information into their system thus greatly reducing cycle times. For this reason, EDI can be an important component of just-in-time production systems.

According to the 2008 Aberdeen report "A Comparison of Supplier Enablement around the World", only 34% of purchase orders are transmitted electronically in North America. In EMEA, 36% of orders are transmitted electronically and in APAC, 41% of orders are transmitted electronically. They also report that the average paper requisition to order costs a company $37.45 in North America, $42.90 in EMEA and

$23.90 in APAC. With an EDI requisition to order costs are reduced to $23.83 in North America, $34.05 in EMEA and $14.78 in APAC.

Barriers to implementation

There are a few barriers to adopting electronic data interchange. One of the most significant barriers is the accompanying business process change. Existing business processes built around slow paper handling may not be suited for EDI and would require changes to accommodate automated processing of business documents. For example, a business may receive the bulk of their goods by 1 or 2 day shipping and all of their invoices by mail. The existing process may therefore assume that goods are typically received before the invoice. With EDI, the invoice will typically be sent when the goods ship and will therefore require a process that handles large numbers of invoices whose corresponding goods have not yet been received.

Another significant barrier is the cost in time and money in the initial set-up. The preliminary expenses and time that arise from the implementation, customization and training can be costly and therefore may discourage some businesses. The key is to determine what method of integration is right for the company which will determine the cost of implementation. For a business that only receives one P.O. per year from a client, fully integrated EDI may not make economic sense. In this case, businesses may implement inexpensive "rip and read" solutions or use outsourced EDI solutions provided by EDI "Service Bureaus". For other businesses, the implementation of an integrated EDI solution may be necessary as increases in trading volumes brought on by EDI force them to re-implement their order processing business processes.

The key hindrance to a successful implementation of EDI is the perception many businesses have of the nature of EDI. Many view EDI from the technical perspective that EDI is a data format; it would be more accurate to take the business view that EDI is a system for exchanging business documents with external entities, and integrating the data from those documents into the company's internal systems. Successful implementations of EDI take into account the effect externally generated information will have on their internal systems and validate the business information received. For example, allowing a supplier to update a retailer's Accounts Payables system without appropriate checks and balances would be a recipe for disaster. Businesses new to the implementation of EDI should take pains to avoid such pitfalls.

Chapter 17

Load Balancing (Computing) and Home Network

Load balancing (computing)

In networking, **load balancing** is a technique to distribute workload evenly across two or more computers, network links, CPUs, hard drives, or other resources, in order to get optimal resource utilization, maximize throughput, minimize response time, and avoid overload. Using multiple components with load balancing, instead of a single component, may increase reliability through redundancy. The load balancing service is usually provided by a dedicated program or hardware device (such as a multilayer switch or a DNS server).

It is commonly used to mediate internal communications in computer clusters, especially high-availability clusters. If the load is more on a server, then the secondary server takes some load while the other is still processing requests.

Internet-based services

One of the most common applications of load balancing is to provide a single Internet service from multiple servers, sometimes known as a server farm. Commonly, load-balanced systems include popular web sites, large Internet Relay Chat networks, high-bandwidth File Transfer Protocol sites, Network News Transfer Protocol (NNTP) servers and Domain Name System (DNS) servers.

For Internet services, the load balancer is usually a software program that is listening on the port where external clients connect to access services. The load balancer forwards requests to one of the "backend" servers, which usually replies to the load balancer. This allows the load balancer to reply to the client without the client ever knowing about the internal separation of functions. It also prevents clients from contacting backend servers directly, which may have security benefits by hiding the structure of the internal network and preventing attacks on the kernel's network stack or unrelated services running on other ports.

Some load balancers provide a mechanism for doing something special in the event that all backend servers are unavailable. This might include forwarding to a backup load balancer, or displaying a message regarding the outage.

An alternate method of load balancing, which does not necessarily require a dedicated software or hardware node, is called **round robin DNS**. In this technique, multiple IP addresses are associated with a single domain name clients themselves are expected to choose which server to connect to. Unlike the use of a dedicated load balancer, this technique exposes to clients the existence of multiple backend servers. The technique has other advantages and disadvantages, depending on the degree of control over the DNS server and the granularity of load balancing desired.

Another technique for load-balancing using DNS, which is far more intelligent than the simple "round robin", is to delegate `www.example.org` as a sub-domain whose zone will be served out by each of the same servers that are serving the web site. This technique works particularly well where individual servers are spread around the Internet. For example,

```
one.example.org A 1.1.1.1
two.example.org A 2.2.2.2
www.example.org NS one.example.org
www.example.org NS two.example.org
```

However, the zone file for `www.example.org` on each server will be different such that each server will give out its own IP Address as the A-record. On "one" the zone file for `www.example.org` will say:

```
@ in a 1.1.1.1
```

On "two" the same zone file will say:

```
@ in a 2.2.2.2
```

This way, if a server is down, its DNS will not respond and so the web service will not receive any traffic. Also if the line to one server becomes congested the unreliability of DNS will ensure less HTTP traffic will reach that server, further the DNS response that gets back to the resolver the quickest will nearly always be from the network closest server, ensuring geo-sensitive load-balancing. A short TTL on the A-record will also help to ensure traffic is quickly diverted if a server goes down. Consideration must be given the possibility that this technique may cause individual clients to switch between individual servers mid-session.

A variety of scheduling algorithms are used by load balancers to determine which backend server to send a request to. Simple algorithms include random choice or round robin. More sophisticated load balancers may take into account additional factors, such as a server's reported load, recent response times, up/down status (determined by a monitoring poll of some kind), number of active connections, geographic location, capabilities, or how much traffic it has recently been assigned. High-performance systems may use multiple layers of load balancing.

In addition to using dedicated hardware load balancers, software-only solutions are available, including open source options. Examples of the latter include the Apache web server's mod_proxy_balancer extension and the Pound reverse proxy and load balancer.

In a Multitier architecture, terminology for designs behind a load balancer or network dispatcher may include **Bowties** and **Stovepipes**. A stovepipe presents a situation such that a transaction that is dispatched at a top tier follows a static path through the stack of devices and software behind the load balancer to its final destination. Alternatively, if Bowties are used, at each tier the transaction could take one of many paths after being serviced by the applications at a particular tier. Network diagrams with transaction flows resemble Stovepipes or Bowties, or hybrid architectures based on need at each tier.

Persistence

An important issue when operating a load-balanced service is how to handle information that must be kept across the multiple requests in a user's session. If this information is stored locally on one backend server, then subsequent requests going to different backend servers would not be able to find it. This might be cached information that can be recomputed, in which case load-balancing a request to a different backend server just introduces a performance issue...

One solution to the session data issue is to send all requests in a user session consistently to the same backend server. This is known as "persistence" or "stickiness". A significant downside to this technique is its lack of automatic failover: if a backend server goes down, its per-session information becomes inaccessible, and any sessions depending on it are lost. The same problem is usually relevant to central database servers; even if web servers are "stateless" and not "sticky", the central database is (see below).

Assignment to a particular server might be based on a username, client IP address, or random assignment. Owing to DHCP, Network Address Translation, and web proxies, the client's IP address may change across requests, and so this method can be somewhat unreliable. Random assignments must be remembered by the load balancer, which creates a storage burden. If the load balancer is replaced or fails, this information can be lost, and assignments may need to be deleted after a timeout period or during periods of high load to avoid exceeding the space available for the assignment table. The random assignment method also requires that clients maintain some state, which can be a problem, for example when a web browser has disabled storage of cookies. Sophisticated load balancers use multiple persistence techniques to avoid some of the shortcomings of any one method.

Another solution is to keep the per-session data in a database. Generally this is bad for performance since it increases the load on the database: the database is best used to store information less transient than per-session data. To prevent a database from becoming a single point of failure, and to improve scalability, the database is often replicated across multiple machines, and load balancing is used to spread the query load across those replicas. Microsoft's ASP.net State Server technology is an example of a session database. All servers in a web farm store their session data on State Server and any server in the farm can retrieve the data.

Fortunately there are more efficient approaches. In the very common case where the client is a web browser, per-session data can be stored in the browser itself. One technique is to use a browser cookie, suitably time-stamped and encrypted. Another is URL rewriting. Storing session data on the client is generally the preferred solution: then the load balancer is free to pick any backend server to handle a request. However, this method of state-data handling is not really suitable for some complex business logic scenarios, where session state payload is very big or recomputing it with every request on a server is not feasible, and URL rewriting has major security issues, since the end-user can easily alter the submitted URL and thus change session streams.

Load balancer features

Hardware and software load balancers can come with a variety of special features.

- **Asymmetric load:** A ratio can be manually assigned to cause some backend servers to get a greater share of the workload than others. This is sometimes used as a crude way to account for some servers being faster than others.
- **Priority activation:** When the number of available servers drops below a certain number, or load gets too high, standby servers can be brought online
- **SSL Offload and Acceleration:** SSL applications can be a heavy burden on the resources of a Web Server, especially on the CPU and the end users may see a slow response (or at the very least the servers are spending a lot of cycles doing things they weren't designed to do). To resolve these kinds of issues, a Load Balancer capable of handling SSL Offloading in specialized hardware may be used. When Load Balancers are taking the SSL connections, the burden on the Web Servers is reduced and performance will not degrade for the end users.
- **Distributed Denial of Service (DDoS) attack protection:** load balancers can provide features such as SYN cookies and delayed-binding (the back-end servers don't see the client until it finishes its TCP handshake) to mitigate SYN flood attacks and generally offload work from the servers to a more efficient platform.
- **HTTP compression:** reduces amount of data to be transferred for HTTP objects by utilizing gzip compression available in all modern web browsers
- **TCP offload:** different vendors use different terms for this, but the idea is that normally each HTTP request from each client is a different TCP connection.

This feature utilizes HTTP/1.1 to consolidate multiple HTTP requests from multiple clients into a single TCP socket to the back-end servers.

- **TCP buffering:** the load balancer can buffer responses from the server and spoon-feed the data out to slow clients, allowing the server to move on to other tasks.
- **Direct Server Return:** an option for asymmetrical load distribution, where request and reply have different network paths.
- **Health checking:** the balancer will poll servers for application layer health and remove failed servers from the pool.
- **HTTP caching:** the load balancer can store static content so that some requests can be handled without contacting the web servers.
- **Content Filtering:** some load balancers can arbitrarily modify traffic on the way through.
- **HTTP security:** some load balancers can hide HTTP error pages, remove server identification headers from HTTP responses, and encrypt cookies so end users can't manipulate them.
- **Priority queuing:** also known as rate shaping, the ability to give different priority to different traffic.
- **Content aware switching:** most load balancers can send requests to different servers based on the URL being requested.
- **Client authentication:** authenticate users against a variety of authentication sources before allowing them access to a website.
- **Programmatic traffic manipulation:** at least one load balancer allows the use of a scripting language to allow custom load balancing methods, arbitrary traffic manipulations, and more.
- **Firewall:** direct connections to backend servers are prevented, for network security reasons
- **Intrusion Prevention System:** offer application layer security in addition to network/transport layer offered by firewall security.

In telecommunications

Load balancing can be useful when dealing with redundant communications links. For example, a company may have multiple Internet connections ensuring network access even if one of the connections should fail.

A failover arrangement would mean that one link is designated for normal use, while the second link is used only if the first one fails.

With load balancing, both links can be in use all the time. A device or program decides which of the available links to send packets along, being careful not to send packets along any link if it has failed. The ability to use multiple links simultaneously increases the available bandwidth.

Major telecommunications companies have multiple routes through their networks or to external networks. They use more sophisticated load balancing to shift traffic from one path to another to avoid network congestion on any particular link, and sometimes to minimize the cost of transit across external networks or improve network reliability.

Relationship with failover

Load balancing is often used to implement failover — the continuation of a service after the failure of one or more of its components. The components are monitored continually (e.g., web servers may be monitored by fetching known pages), and when one becomes non-responsive, the load balancer is informed and no longer sends traffic to it. And when a component comes back on line, the load balancer begins to route traffic to it again. For this to work, there must be at least one component in excess of the service's capacity. This is much less expensive and more flexible than failover approaches where a single "live" component is paired with a single "backup" component that takes over in the event of a failure. Some types of RAID systems can also utilize hot spare for a similar effect.

Home network

IEEE Home networking Recommendations	
Recommendations	
HomePlug AV	
HomePlug AV2	
HomePlug Green PHY	
Recommendations	**GHz**
Wi-Fi 802.11a	5 GHz
Wi-Fi 802.11b	2.4 GHz
Wi-Fi 802.11g	2.4 GHz
Wi-Fi 802.11n	2.4GHz and/or 5GHz
ITU-T Home networking Recommendations	
Common Name	**Recommendations**
HomePNA 2.0	G.9951, G.9952, G.9953
HomePNA 3.0	G.9954 (02/05)
HomePNA 3.1	G.9954 (01/07)
G.hn/HomeGrid	G.9960, G.9961
G.cx	G.9972
G.hnta	G.9970

A **home network** or **home area network** (**HAN**) is a residential local area network (LAN). It is used for communication between digital devices typically deployed in the home, usually a small number of personal computers and accessories, such as printers and mobile computing devices. An important function is the sharing of Internet access, often a broadband service through a cable tv or Digital Subscriber Line (DSL) provider. Additionally, a home server may be added for increased functionality.

Home networks may use wired or wireless technologies. Wired systems typically use shielded or unshielded twisted pair cabling, such as any of the Category 3 (CAT3) through Category 6 (CAT6) classes, but may also be implemented with coaxial cable, or over the existing electrical power wiring within homes.

One of the most common ways of creating a home network is by using wireless radio signal technology; the 802.11 network as certified by the IEEE. Most products that are wireless-capable operate at a frequency of 2.4 GHz under 802.11b and 802.11g or 5 GHz under 802.11a. Some home networking devices operate in both radio-band signals and fall within the standard 802.11n.

A wireless network can be used for communication between many electronic devices, to connect to the Internet or to wired networks that use Ethernet technology. Wi-Fi is a marketing and compliance certification for IEEE 802.11 technologies.. The WiFi Alliance has tested compliant products certifies them for interoperability.

As an alternative to wireless networking, the existing home wiring (coax in North America, telephone wiring in multi dwelling units (MDU) and power-line in Europe and USA) can be used as a network medium. With the installation of a home networking device, the network can be accessed by simply plugging the Computer into a wall socket. The ITU-T G.hn and IEEE Powerline standard, which provide high-speed (up to 1 Gbit/s) local area networking over existing home wiring, are examples of home networking technology designed specifically for IPTV delivery. Recently, the IEEE passed proposal P1901 which grounded a standard within the Market for wireline products produced and sold by companies that are part of the HomePlug Alliance. The IEEE is continuously working to push for P1901 to be completely recognized worldwide as the sole standard for all future products that are produced for Home Networking.

Network devices

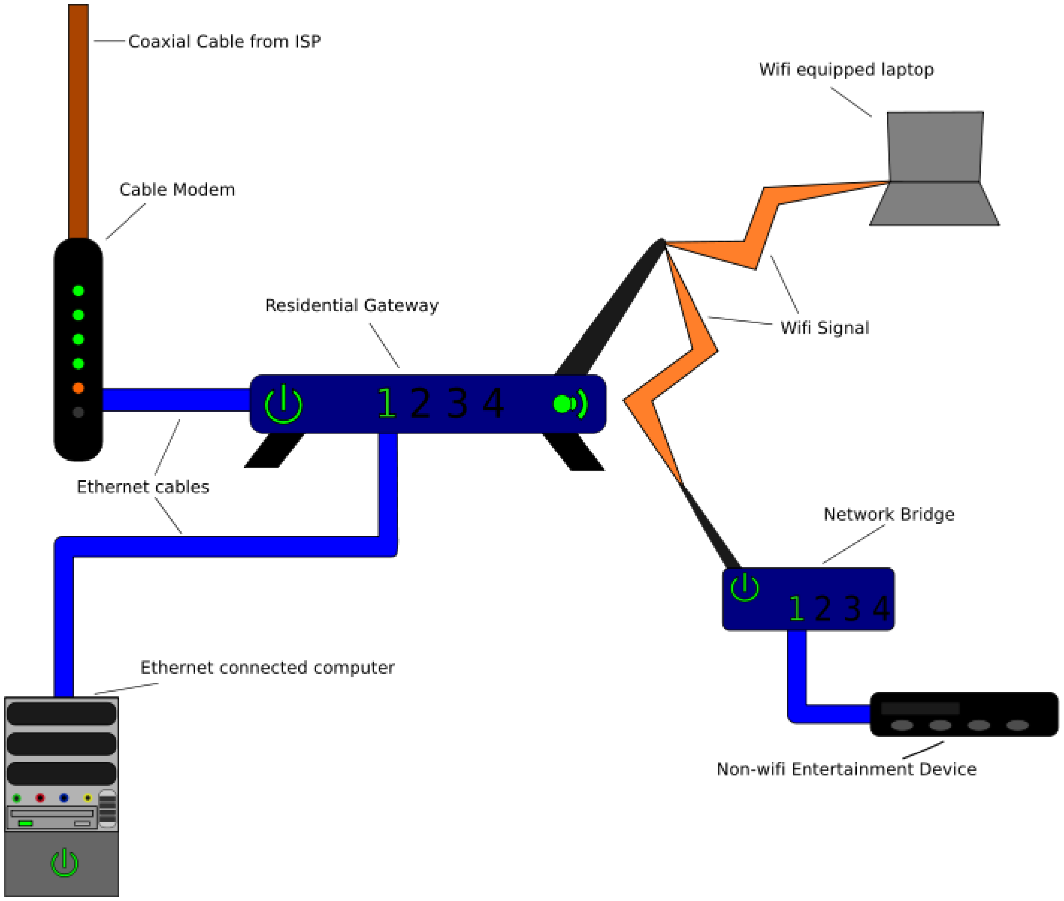

An example of a simple home network

A home network may consist of the following components:

Infrastructure Devices

- A broadband modem for connection to the internet (either a DSL modem using the phone line, or cable modem using the cable internet connection).
- A residential gateway (sometimes called a broadband router) connected between the broadband modem and the rest of the network. This enables multiple devices to connect to the internet simultaneously. Residential gateways, hubs/switches, DSL modems, and wireless access points are often combined.
- A wireless access point, usually implemented as a feature rather than a separate box, for connecting wireless devices

Client Devices

- A PC, or multiple PCs including laptops, Netbooks and Tablet PC’s

- Entertainment peripherals - an increasing number of devices can be connected to the home network, including DVRs like TiVo, digital audio players, games machines, stereo system, and IP set-top box as well as TVs themselves.
- Internet Phones (VoIP)
- Smart Phones connected via Wifi.
- A network bridge connects two networks together, often giving a wired device, e.g. Xbox, access to a wireless network.
- A network hub/switch - a central networking hub containing a number of Ethernet ports for connecting multiple networked devices
- A network attached storage (NAS) device can be used for storage on the network.
- A print server can be used to share printers among computers on the network.

Older devices may not have the appropriate connector to the network. USB and PCI network controllers can be installed in some devices to allow them to connect to networks.

Network devices may also be configured from a computer. For example, broadband modems are often configured through a web client on a networked PC. As networking technology evolves, more electronic devices and home appliances are becoming Internet ready and accessible through the home network. Set-top boxes from cable TV providers already have USB and Ethernet ports "for future use".

Network media

Ethernet cables are the standard medium for networks. However, homes are often more difficult to wire than office environments, and other technologies are being developed which don't require new wires.

Home networking may use

- Ethernet Category 5 cable, Category 6 cable - for speeds of 10 Mbit/s, 100 Mbit/s, or 1 Gbit/s.
- Wi-Fi Wireless LAN connections - for speeds up to 450 Mbit/s, dependent on signal strength and wireless standard.
- Coaxial cables (TV antennas) - for speeds of 270 Mbit/s
- Electrical wiring - for speeds of 14 Mbit/s to 200 Mbit/s
- Phone wiring - for speeds of 160 Mbit/s
- Fiber optics - although rare, new homes are beginning to include fiber optics for future use. Optical networks generally use Ethernet.
- All home wiring (coax, powerline and phone wires) - future standard for speeds up to 1 Gbit/s being developed by the ITU-T

Ethernet and Wireless are the most common standards. As the demand for home networks has increased, the other alliances have formed to produce standards for networking alternatives.

Home Coverage

Challenges

Wireless Signal Loss

- The Wireless signal strength may not be powerful enough to cover the entire house or may not be able to get through to all floors of multiple floor residences.

Wired Background "Noise"

- One of the largest challenges posed for those that wanted to utilize the home electrical system for networking is how to combat other electrical “noise” that would be around due to the use of a power outlet to transfer information. Whenever any appliance is turned on or turned off it creates noise that could possibly dissrupt data transfer through the wiring. IEEE products past the HomePlug 1.0 stage have combated this problem and no longer interfere with, or receive interference from, other devices plugged into a power outlet.

”Leaky” WiFi

- As can be construed, WiFi often extends beyond the boundaries of a home and can create coverage and holes where it is least wanted and it can also allow a way for people to compromise a system and retrieve personal data. The most forward and most widespread way to combat this is the investment of an authentication, encryption, or VPN that requires a password to access the WiFi.

Chapter 18

Mylogon and Router

Mylogon

MyLogon is a network-authentication applet for Microsoft Windows.

Mainly aimed at small-business networks, it offers a simple, useraccount-based method of connecting a Microsoft Windows computer to a fileserver.

The current release (2.02) is available under the GPL licence, and is open-source. Previous versions were free-to-use but closed source, under a proprietary licence.

Background

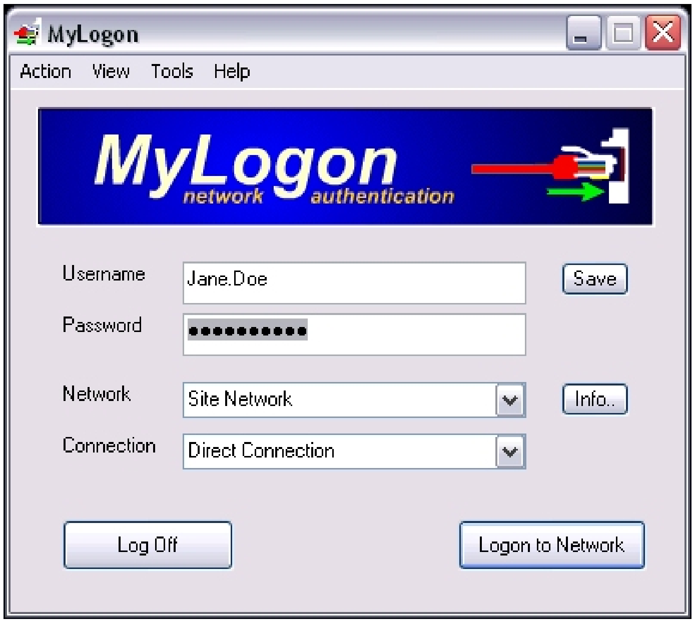

Username: Jane.Doe

Password: ●●●●●●●●●●●

Network: Site Network

Connection: Direct Connection

Typical Dialog-box

As the MyLogon website describes, the standard options for Microsoft Windows networks are either Workgroup or Active Directory membership.

Workgroups, also known as peer-groups, while easy to establish and use, offer very little in the way of security, and have a tendency to become disorganised, owing to the lack of centralisation or administrative control. While acceptable for very small networks, peer file-sharing begins to show its limitations with as few as five computers.

The Active Directory Domain, on the other hand, is primarily aimed at the large corporate client, with a wealth of features intended to make the management of very large networks easier. Active Directory membership provides very tight control over both the users and computers in a network, and also confers many powerful remote-management options onto the site administrator, including the ability to automatically install software packages without physically visiting workstations.

While providing a near-ideal solution for large networks, the complex nature of the Active Directory, and the need for an in-depth knowledge of DNS, LDAP, Group Policy, etc. in order to manage it effectively weigh against its use in small networks, where the steep learning curve involved in understanding its use may result in most of its features remaining unused.

MyLogon's development came about in response to a specific requirement for an intermediate solution -one with greater security and organisation than workgroup arrangements, but one having a more manageable level of complexity than the Active Directory.

How it Works

MyLogon sits as software shim in between the standard Windows logon-process (winlogon.exe) and the launching of the Desktop environment by Explorer. The standard Windows Logon, which would otherwise assign a specific user profile for the current session, is set to automatically select one standard profile regardless of actual user. MyLogon then authenticates the person at the computer against an account on the fileserver, and if the credentials match, connects to network resources as defined in a logon script, then permits access to the Windows Desktop.

An alternative mode of working allows use of the computer itself without the need to log-in, and connection to one of several configured networks on an as-required basis. This may suit laptop-users who wish to work at multiple sites.

A side-effect of MyLogon's method of working is that the settings and behaviour of the user's own computer are not altered by the process of logging-on. The logon purely determines their right to access any company computer, and their right to use network resources. The site explains that in many small offices this is preferable to the Windows default behaviour.

The author goes on to describe an undesirable situation found in many small offices, password-less working. The computers in small firms being typically allocated to a particular task or department, it is a requirement that the computer shall perform its alloted task correctly rather than being a general-purpose resource, as it might be in a corporate cubicle-farm. In this task-oriented environment the compulsory user-profiling of the standard Windows Logon creates a problem, in that a change of user will default the settings of any specialist software, often rendering it useless. This, as he has observed in the process of site service-visits, leads to many small sites -even those which use the full Active Directory topology- working without passwords, so as to avoid the need to ever change username. Working passwordless is universally recognised as a poor security practice. MyLogon overcomes this security issue by making user-controlled access possible without the associated reprofiling, or loss of settings.

A MyLogon workstation requires no specific DNS settings to connect to a server within the same subnet, and the fileserver need not in fact be running a DNS process. This eliminates one of the most complex and troublesome aspects of Active Directory setup.

Connections

MyLogon's approach to creating network connections is essentially script-based, and in this respect it is similar to traditional products such as Windows NT, or Novell Netware. The logon script may use the standard NT command-line syntax, or instead may use a syntax akin to that of .ini files. The preference for this traditional approach is based on the observation that most users are comfortable with the idea of network-resources being denoted by additional drive letters, but do not understand UNC shortcuts. A second argument for this approach, perusal of the posts on the Microsoft helpdesk forums confirms that My Network Places – the Windows tool provided for browsing non-drivemapped resources – has a very poor reliability record, whereas mapped driveletters are seldom problematic.

Limitations

The present version gives the user no way to change his/her password.

Logging-on to a server does not necessarily grant automatic rights to access other peer-computers, as does a Domain logon.

MyLogon is not suitable for use with roaming profiles.

Compatible Clients

Designed for Windows XP Home Or Professional, or Windows 2000. Also suitable for Windows Vista and Windows 7.

Compatible Fileservers

Any which supports NetBIOS (SMB) networking, including all versions of Windows Server from NT4 to Server 2008, and Linux/Samba (software).

Router

A **router** is an electronic device that intercepts signals on a computer network. The router determines where the signals have to go. Each signal it receives is called a data packet. The packet contains address information that the router uses to divert signals appropriately.

Application

When multiple routers are used in a large collection of interconnected networks, the routers exchange information, so that each router can build up a reference table showing the preferred paths between any two systems on the interconnected networks.

A router can have many interface connections, for different physical types of network (such as copper cables, fiber optic, or wireless transmission). It may contain firmware for different networking protocol standards. Each network interface device is specialized to convert computer signals from one protocol standard to another.

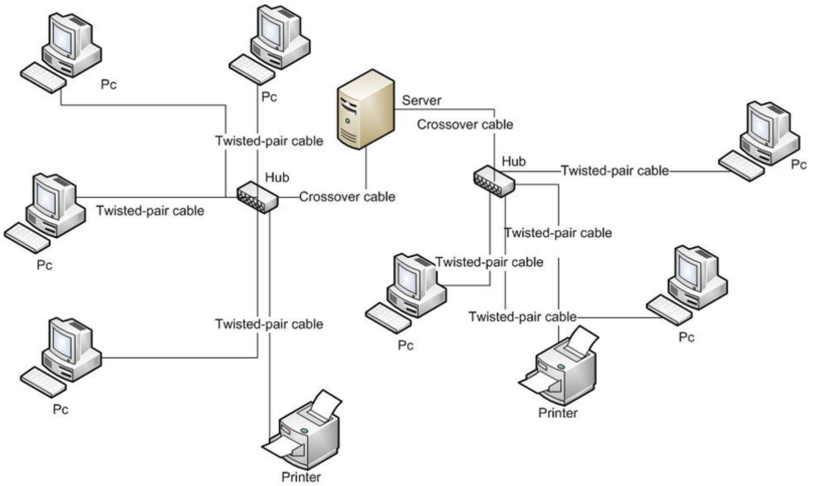

Two small computer networks connected with HUBS, these are not ROUTERS, but simply connectors between computers. SWITCHES may be used to connect HUBS together to help transfer signals more efficently between groups of users.

Routers can be used to connect two or more logical subnets, each having a different network address. The subnets addresses in the router do not necessarily map directly to the physical interfaces of the router. The term "layer 3 switching" is often used interchangeably with the term "routing". The term switching is generally used to refer to data forwarding between two network devices with the same network address. This is also called layer 2 switching or LAN switching.

Conceptually, a router operates in two operational planes (or sub-systems):

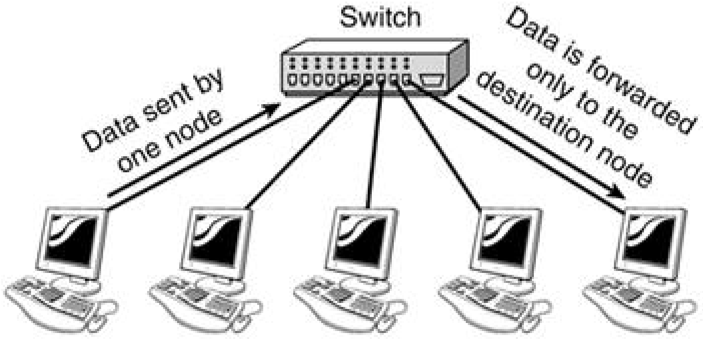

How a switch makes a direct signal exchange connection between only the two required computers.

- Control plane: where a router builds an address table (called routing table) that records where a packet should be forwarded, and through which physical interface.It does this by using either statically configured statements (called static routes), or alternatively, by exchanging information with other routers in the network through a dynamical routing protocol.

- Forwarding plane: The router actually forwards traffic, (called data packets in Internet Protocol language) from incoming interfaces to outgoing interfaces destination addresses that the packet header contains. It performs this function by following rules derived from the routing table that has been recorded in the control plane.

Types of routers

A typical home router showing the ADSL telephone line and ETHERNET network cable connections.

Routers may provide connectivity inside enterprises, between enterprises and the Internet, and inside internet service providers (ISPs). The largest routers (for example the Cisco CRS-1 or Juniper T1600) interconnect ISPs, are used inside ISPs, or may be used in very large enterprise networks. The smallest routers provide connectivity for small and home offices.

Routers for Internet connectivity and internal use

Routers intended for ISP and major enterprise connectivity almost invariably exchange routing information using the Border Gateway Protocol (BGP). RFC 4098 defines several types of BGP-speaking routers according to the routers' functions:

- *Edge router* (ER): An ER is placed at the edge of an ISP network. The router speaks external BGP (EBGP) to a BGP speaker in another provider or large enterprise Autonomous System(AS). This type of router is also called PE (Provider Edge) routers.
- *Subscriber edge router* (SER): An SER is located at the edge of the subscriber's network, it speaks EBGP to its provider's AS(s). It belongs to an end user (enterprise) organization. This type of router is also called CE (Customer Edge) routers.

- *Inter-provider border router*: Interconnecting ISPs, this is a BGP-speaking router that maintains BGP sessions with other BGP speaking routers in other providers' ASes.
- Core router: A *core router* is one that resides within an AS as back bone to carry traffic between edge routers.

  Within an ISP: Internal to the provider's AS, such a router speaks internal BGP (IBGP) to that provider's edge routers, other intra-provider core routers, or the provider's inter-provider border routers.
  "Internet backbone:" The Internet does not have a clearly identifiable backbone, as did its predecessors. Nevertheless, the major ISPs' routers make up what many would consider the core. These ISPs operate all four types of the BGP-speaking routers described here. In ISP usage, a "core" router is internal to an ISP, and used to interconnect its edge and border routers. Core routers may also have specialized functions in virtual private networks based on a combination of BGP and Multi-Protocol Label Switching (MPLS).

Routers are also used for port forwarding for private servers.

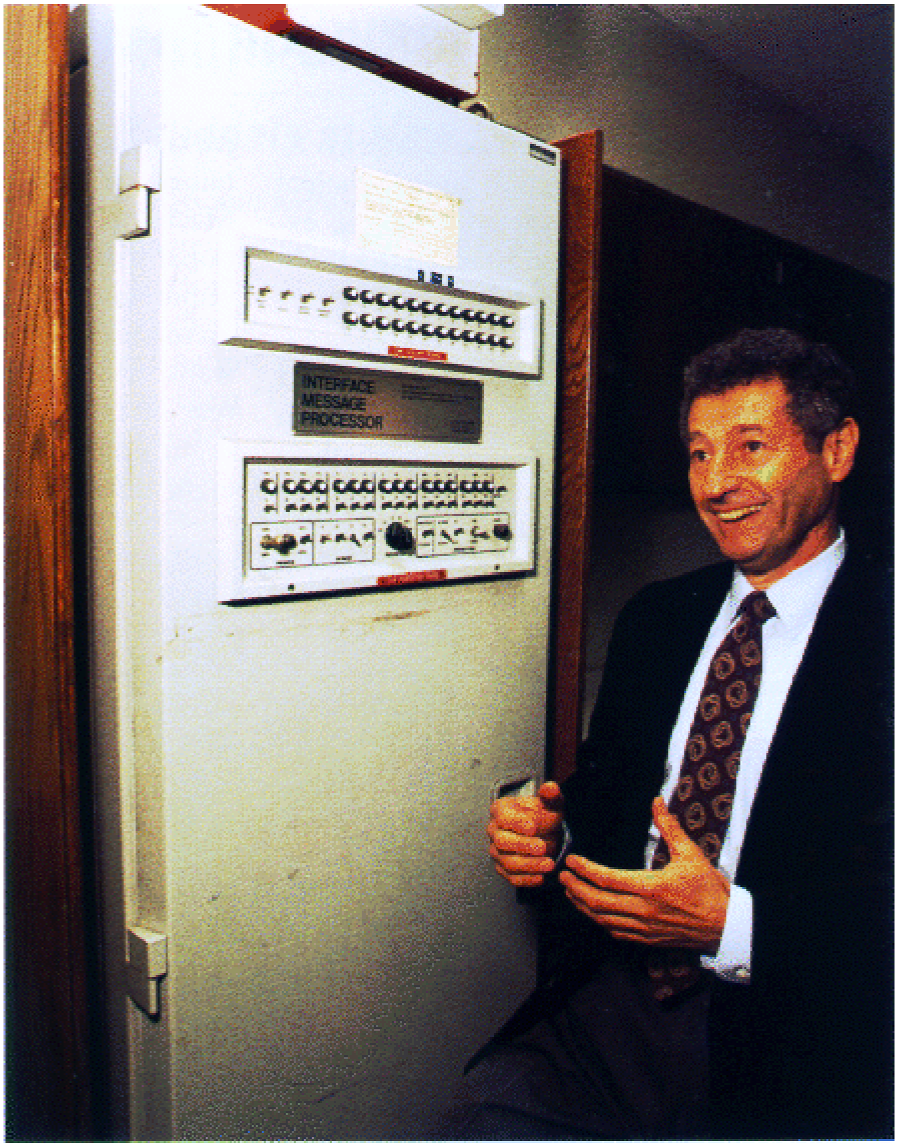

Leonard Kleinrock and the first IMP.

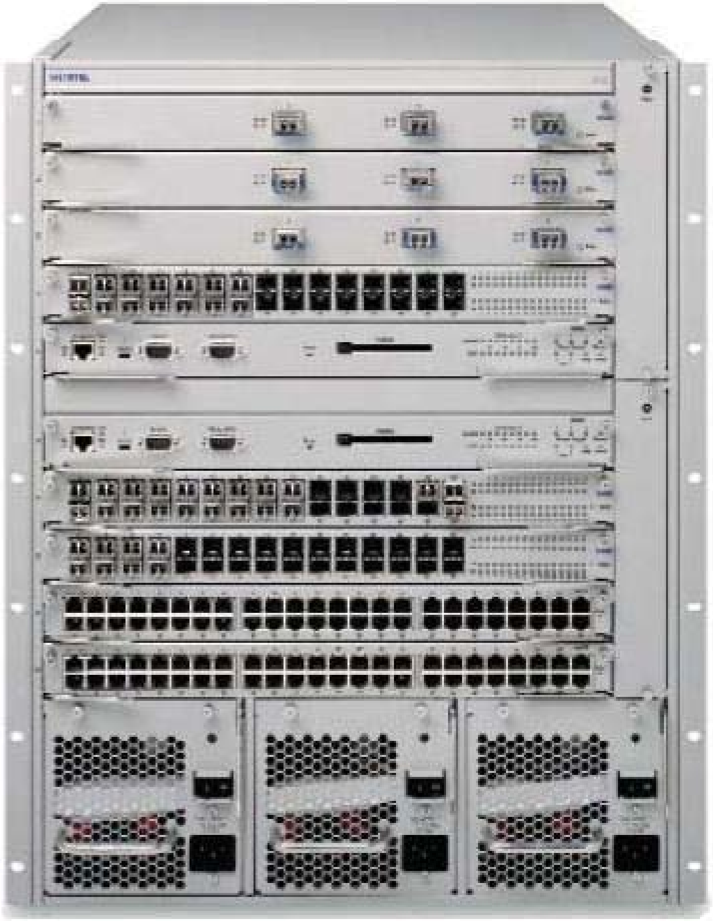

Avaya ERS 8600 (2010)

The very first device that had fundamentally the same functionality as a router does today, i.e a packet switch, was the Interface Message Processor (IMP); IMPs were the devices that made up the ARPANET, the first packet switching network. The idea for a router (although they were called "gateways" at the time) initially came about through an international group of computer networking researchers called the International Network Working Group (INWG). Set up in 1972 as an informal group to consider the technical issues involved in connecting different networks, later that year it became a subcommittee of the International Federation for Information Processing.

These devices were different from most previous packet switches in two ways. First, they connected dissimilar kinds of networks, such as serial lines and local area networks. Second, they were connectionless devices, which had no role in assuring that traffic was delivered reliably, leaving that entirely to the hosts (although this particular idea had been previously pioneered in the CYCLADES network).

The idea was explored in more detail, with the intention to produce a real prototype system, as part of two contemporaneous programs. One was the initial DARPA-initiated program, which created the TCP/IP architecture of today. The other was a program at Xerox PARC to explore new networking technologies, which produced the PARC Universal Packet system, although due to corporate intellectual property concerns it received little attention outside Xerox until years later.

The earliest Xerox routers came into operation sometime after early 1974. The first true IP router was developed by Virginia Strazisar at BBN, as part of that DARPA-initiated effort, during 1975-1976. By the end of 1976, three PDP-11-based routers were in service in the experimental prototype Internet.

The first multiprotocol routers were independently created by staff researchers at MIT and Stanford in 1981; the Stanford router was done by William Yeager, and the MIT one by Noel Chiappa; both were also based on PDP-11s.

As virtually all networking now uses IP at the network layer, multiprotocol routers are largely obsolete, although they were important in the early stages of the growth of computer networking, when several protocols other than TCP/IP were in widespread use. Routers that handle both IPv4 and IPv6 arguably are multiprotocol, but in a far less variable sense than a router that processed AppleTalk, DECnet, IP, and Xerox protocols.

In the original era of routing (from the mid-1970s through the 1980s), general-purpose mini-computers served as routers. Although general-purpose computers can perform routing, modern high-speed routers are highly specialized computers, generally with extra hardware added to accelerate both common routing functions, such as packet forwarding and specialised functions such as IPsec encryption.

Still, there is substantial use of Linux and Unix machines, running open source routing code, for routing research and other applications. While Cisco's operating system was independently designed, other major router operating systems, such as those from Juniper Networks and Extreme Networks, are extensively modified but still have Unix ancestry.

Enterprise routers

All sizes of routers may be found inside enterprises. The most powerful routers tend to be found in ISPs and academic & research facilities. Large businesses may also need powerful routers.

A three-layer model is in common use, not all of which need be present in smaller networks.

Access

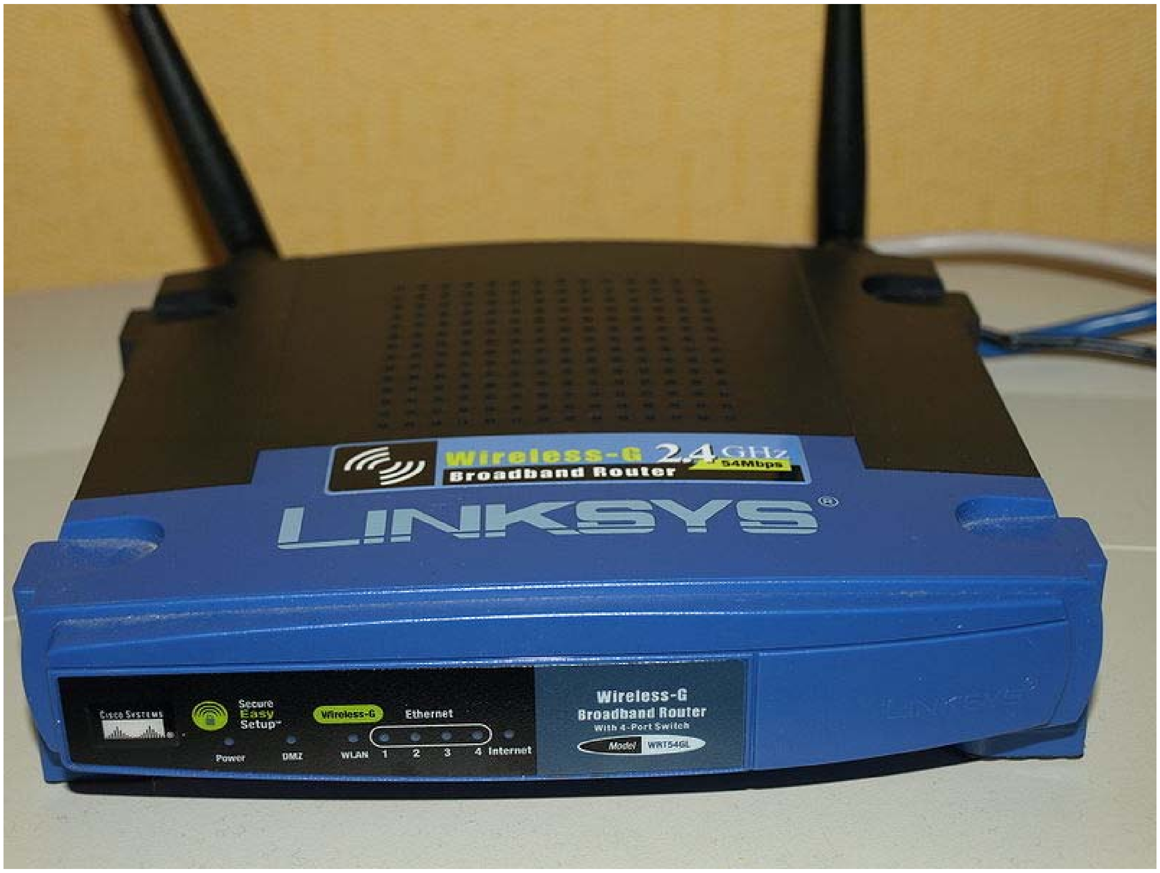

Linksys by Cisco WRT54GL SoHo Router

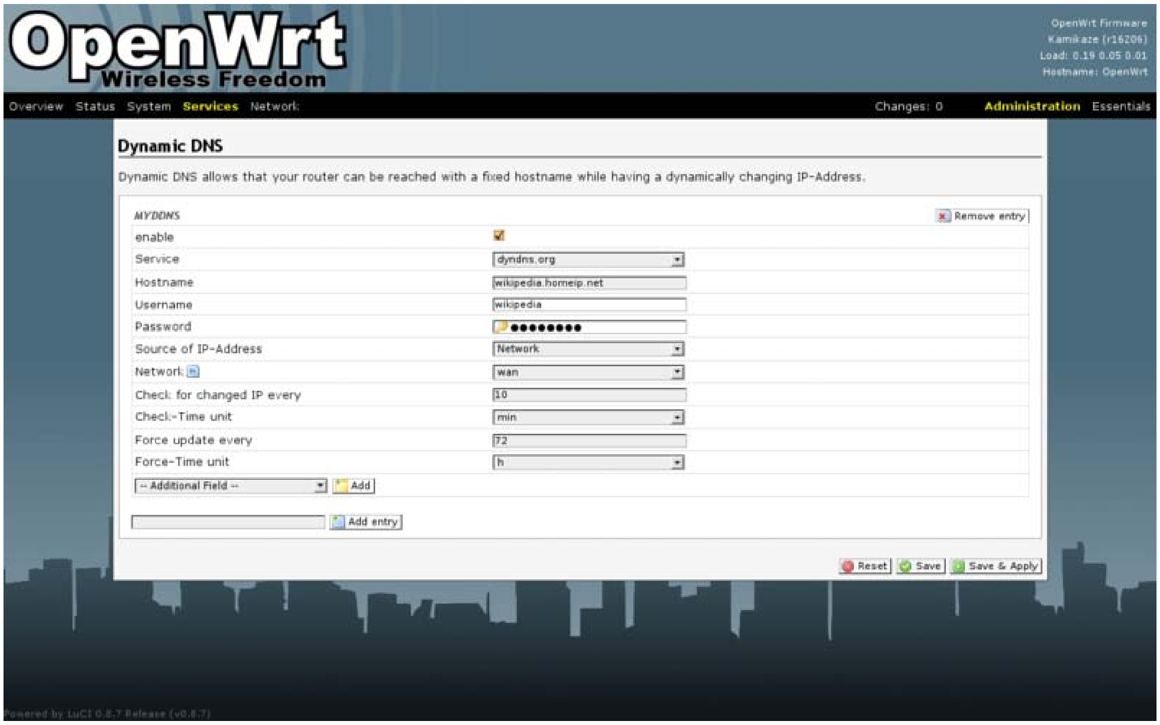

Dynamic DNS

Dynamic DNS allows that your router can be reached with a fixed hostname while having a dynamically changing IP-Address.

MYDDNS

enable	☑
Service	dyndns.org
Hostname	wikipedia.homeip.net
Username	wikipedia
Password	●●●●●●●●●
Source of IP-Address	Network
Network	wan
Check for changed IP every	10
Check-Time unit	min
Force update every	72
Force-Time unit	h

A screenshot of the LuCI web interface used by OpenWrt. Here it is being used to configure Dynamic DNS.

Access routers, including 'small office/home office' (SOHO) models, are located at customer sites such as branch offices that do not need hierarchical routing of their own. Typically, they are optimized for low cost. Some SOHO routers are capable of running alternative free Linux-based firmwares like OpenWrt.

Distribution

Distribution routers aggregate traffic from multiple access routers, either at the same site, or to collect the data streams from multiple sites to a major enterprise location. Distribution routers often are responsible for enforcing quality of service across a WAN, so they may have considerable memory, multiple WAN interfaces, and substantial processing intelligence.

They may also provide connectivity to groups of servers or to external networks. In the latter application, the router's functionality must be carefully considered as part of the overall security architecture. Separate from the router may be a firewall or VPN concentrator, or the router may include these and other security functions.

Core

In enterprises, a core router may provide a "collapsed backbone" interconnecting the distribution tier routers from multiple buildings of a campus, or large enterprise locations. They tend to be optimized for high bandwidth.

Forwarding plane (a.k.a. data plane)

For pure Internet Protocol (IP) forwarding function, a router is designed to minimize the state information on individual packets. The main purpose of a router is to connect multiple networks and forward packets destined either for its own networks or other networks. A router is considered a Layer 3 device because its primary forwarding decision is based on the information in the Layer 3 IP packet, specifically the destination IP address. This process is known as routing. When each router receives a packet, it searches its routing table to find the best match between the destination IP address of the packet and one of the network addresses in the routing table. Once a match is found, the packet is encapsulated in the layer 2 data link frame for that outgoing interface. A router does not look into the actual data contents that the packet carries, but only at the layer 3 addresses to make a forwarding decision, plus optionally other information in the header for hint on, for example, QoS. Once a packet is forwarded, the router does not retain any historical information about the packet, but the forwarding action can be collected into the statistical data, if so configured.

Forwarding decisions can involve decisions at layers other than the IP internetwork layer or OSI layer 3. A function that forwards based on data link layer, or OSI layer 2, information, is properly called a bridge or switch. This function is referred to as layer 2 switching, as the addresses it uses to forward the traffic are layer 2 addresses in the OSI layer model.

Besides making decision as which interface a packet is forwarded to, which is handled primarily via the routing table, a router also has to manage congestion, when packets arrive at a rate higher than the router can process. Three policies commonly used in the Internet are tail drop, random early detection, and weighted random early detection. Tail drop is the simplest and most easily implemented; the router simply drops packets once the length of the queue exceeds the size of the buffers in the router. Random early detection (RED) probabilistically drops datagrams early when the queue is exceeds a pre-configured size of the queue until a pre-configured max when it becomes tail drop. Weighted random early detection requires a weight on the average queue size to act upon when the traffic is about to exceed the pre-configured size, so that short bursts will not trigger random drops.

Another function a router performs is to decide which packet should be processed first when multiple queues exist. This is managed through Quality of service (QoS), which is critical when VoIP (Voice over IP) is deployed, so that delays between packets do not exceed 150ms to maintain the quality of voice conversations.

Yet another function a router performs is called "policy based routing" where special rules are constructed to override the rules derived from the routing table when a packet forwarding decision is made.

These functions may be performed through the same internal paths that the packets travel inside the router. Some of the functions may be performed through an application-specific integrated circuit (ASIC) to avoid overhead caused by multiple CPU cycles, and others may have to be performed through the CPU as these packets need special attention that cannot be handled by an ASIC.