Computer Networks and Internet

Basic Concepts and Terminology

Author: Iskra Popova

There are plenty of books and other literature on networking. However, I could not find anything that is short and explains the technology in a simple way so that it can be used as a literature for this course. Hence, I created this text to be used by students in the course.

1. *What is a network?*

Computer network is a collection of interconnected devices that can share information. They are built to provide sharing of equipment and communication among users. Devices that are usually shared are printers, scanners, data storage, other expensive equipment, Internet access, all kind of information etc. Internet is a worldwide network of networks. Latest trends in networking are towards convergence of networks used for sharing data, voice, and video, and in creating pervasive networks that provide access for the users to information they need everywhere and at any time. When possible, the infrastructure that is already built is used. From that prospective, a network can be viewed as collection of resources interconnected by an infrastructure that allows sharing of these resources among various users in particular locations.

Although connecting two computers may not look like a real network, we will use this example to define the main elements in a network. The following picture shows all the elements of this simple network.

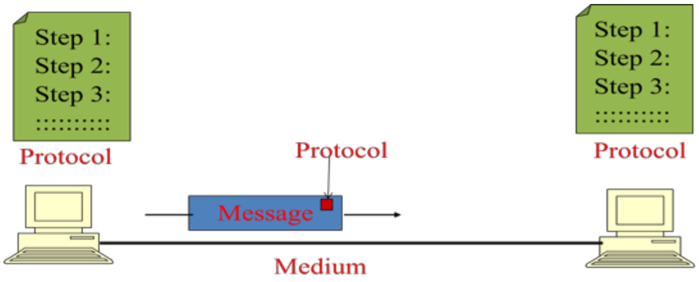

1. The sender and the receiver or the two devices that communicate. One of them sends information and the other receives it. Since the communication is usually both ways, the sender at times becomes a receiver and the receiver becomes a sender. Alternative words for them are source and destination. The common name for both is nodes in the network.
2. The medium through which information is transmitted. Data are transmitted as electric or light signals or electromagnetic waves. Hence, the medium has to be capable of transmitting the signals. It can be some cable (copper of fiber-optic) or the air in case of wireless connection.
3. The message or unit of information that travel from the sender to the receiver. The message is a string of 0s and 1s and it can represent a mail message, audio, video, executable or any kind of file, content of the whole web page, personal data conveyed to the social site or any other kind of data.

4. Rules or protocols that govern how the messages are sent, directed, received and interpreted. The control information sent within the message is a part of the protocol. This is a very important part of any communication. The protocols are a part of either the software if they are concerned with the application or firmware built in devices that are interface between the device and the medium.

2. What are protocols?

Different types of communication exist depending on the participants or the occasion for it. For example, chatting on the Internet with a friend is different than having a job interview.

Depending on the type of communication some rules are either established in advance, or both parties know how to behave for the particular occasion. These rules of behavior are called protocols. Among protocols that govern successful human communication are the following:

1. An identified sender and receiver (you need to know who communicates with whom)
2. The media for communication (telephone, e-mail, Skype, letter, face-to-face, etc.)
3. The common language and grammar so that the two parties can understand each other
4. The speed and timing of delivery so that they know when to expect the response
5. Confirmation or acknowledgement requirements which are very important for some types of communication and not necessary at all for others.

Machines (computers) communicate using protocols (rules) in the same way the humans do. However, there are certain differences. Let's consider the telephone network as an example. It can be used to connect two telephone sets and two computers. The human voice, as well as the messages sent by computers are transformed into electric signals and transferred via medium that is not perfect. This means that the electrical signals will be slightly modified when they travel through the medium. As a consequence, the received message will not be exactly the same than the original. In case of human communication, these changes do not necessarily affect the communication. If a human do not recognize a single word in a sentence, still the sentence can be understood. In case the receiver is the machine this is not the case. Therefore, the protocols used by computers need to be strict and take in consideration not only the imperfectness of the medium, but also all other things that might happen during the communication.

3. Network Types and Topology

Networks that are more complicated have many more than two devices. A link is considered to be the medium that connects two nodes. Devices involved in the network can be connected in different ways. The arrangement of links in a network is referred to as a topology of the network. There is difference between the physical and logical topology. The first defines the physical constellation of the nodes and links, the second determines how the communication is performed. The elements can be arranged in a star topology with a central device in the middle and all others connected to it. Some networks have bus or ring topology where there is no central device. Full mesh topology is rarely found because of the high cost for connecting devices in such constellation. Hence, partial mesh is an alternative. The picture below shows different types of topologies.

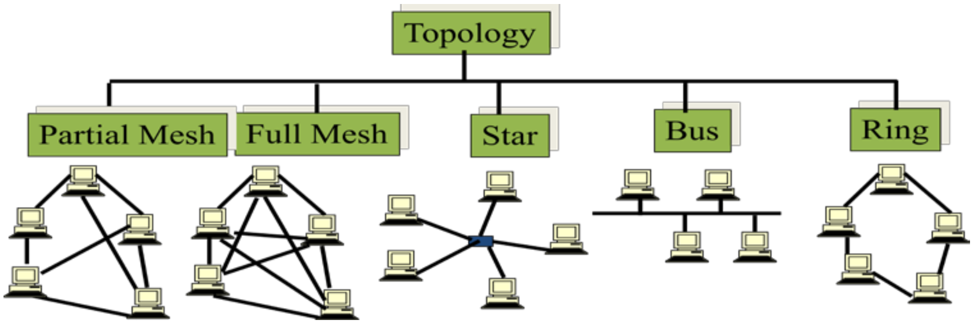

One way of defining types of computer networks is according to the geographical area they span. Local Area Network (LAN) is usually built into a single or several buildings; Metropolitan Area Network is usually inside a single city, while Wide Area Network connects different towns or countries. LANs are usually privately owned (home or company network), MANs are owned by city authorities or by private companies, ISPs (Internet Service Providers), WANs are in possession of telecom companies or large ISPs, and Internet is not owned by anyone. Everyone owns only a piece of it.

4. *Network Interface Card (NIC)*

Every device is connected to the network through a network interface card (NIC). It is usually plugged into the device and has an opening for the network cable or antenna for wireless connection. It is actually an interface and acts as a border between the device and the communication medium. NIC has a physical address that is unique. This address is also called media access (MAC) address and consists of 6 bytes. Each byte is presented with two hexadecimal digits. For example: 01-3B-85-6F-72-A0. Hexadecimal digit has 16 possible digits (0, 1, …., 9,A, B, …F). Each digit can be written as for bits. A bit is acronym for Binary DigIT. Binary digits are 0 and 1.

5. *What is Ethernet?*

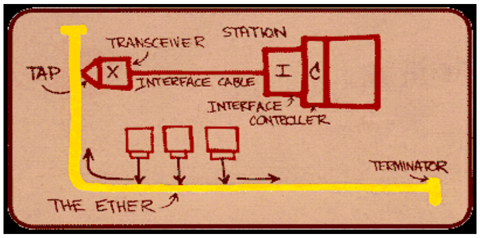

Ethernet is the oldest and today most commonly used technology for local and metropolitan networks. It defines the topology and the protocols used for communication. It was invented by Robert Metcalfe in 1973. The idea was to create a simple and cheap way for connecting computers. The wire which he called ”the ether” connects all the computers. The picture on the right side presents the original drawing of Metcalf's network. The first operational Ethernet based on this idea was with 5Mbps speed.

The main characteristic of the Ethernet is the shared channel or the ether and bus topology. It can be compared to the conversation of a group of people when only one talk and the others listen or with a satellite which transmits signals that can be received by many earth stations. This means that information sent by a single machine is conveyed through the channel. All computers can listen to it. In other words, bus is used as physical and logical topology. The picture below shows normal operation (the picture on the top) and operation under “collision” (the picture at the bottom).

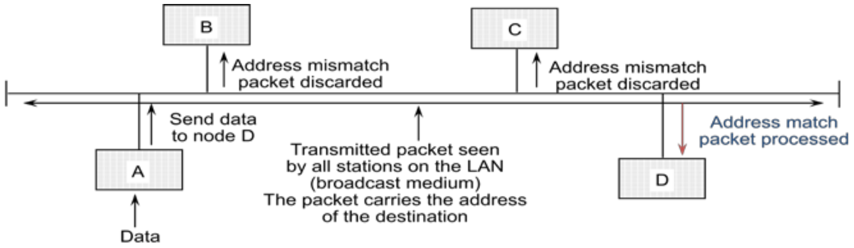

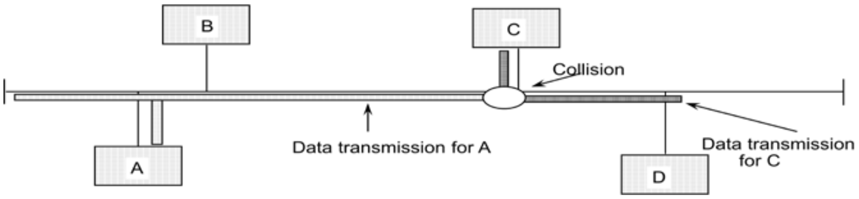

If machine A wants to send a packet to machine D in the control data it adds its own physical address as source address and the physical or MAC address of the node D as destination address. If none of the other machines on the wire is sending, then the packet will be transmitted through the wire. All other machines connected to the wire will be able to ”see” the packet (inspect what is in it). Actually the network interface card on every machine on the network will check the destination MAC address in the packet by comparing it with its own address. If it matches, the packet will be forwarded to the machine. If it does not match it will be ignored. In our example B and C will ignore the packet and D will accept it.

This protocol is called Carrier Sense Multiple Access (CSMA). Carrier is the electrical signal on the medium which can be access by all devices. Carrier sense refers to “listening” or sensing the medium to check if another device is sending a signal. This protocol does not work when two devices wants to start sending data at the same time. They both sense that there is no carrier present on the medium.

In case two packets from different machines are sent at the same time, there will be a collision and none of the packets will be able to reach its destination. This means that the two electrical signals will be merged and the original message cannot be retrieved from this signal. When collision occurs, all machines on the wire receive a special signal that indicates collision. Then, the machines that tried to send packets will wait for some random time before they try to send the same packets again. Hopefully these times will be different and collision will not happen. This part of the protocol is called Collision Detection (CD). The complete name for the protocol is CSMA CD (Carrier Sense Multiple Access with Collision Detection).

If the local network involves many devices, the collisions can happen often and the delays in the network can be large. Ethernet technology evolved with years. Several wireless Ethernet standards were created and the wired version was improved to reach the speed of 10GBps. Several network devices are used today to improve Ethernet operation. Repeaters are used to extend the length of the wire. The hubs act as repeaters, but provide a star physical topology which is easier to maintain. Bridges and switches are capable to reduce collisions. Bridges split a single collision domain into several and switches are even better, they do not allow collisions at all. They provide topology that is a star, physical and logical, too.

The picture below shows networks connected with a bridge and a switch. Using a switch requires more cabling, but provides much better performance.

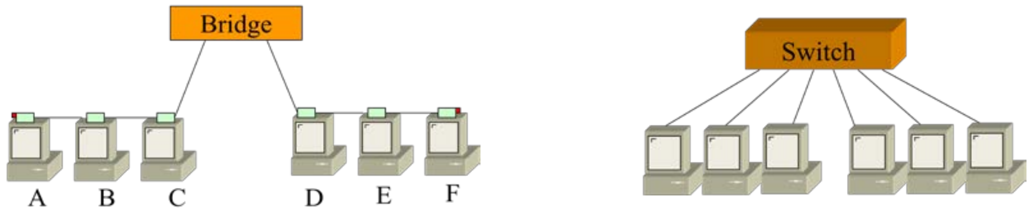

The bridges and switches have tables that help them forward messages from one node to another. The tables can be static or dynamic depending on whether they are filled by human or built automatically. Usually the plug and play algorithm is used to populate these tables. Each table has an entry for the MAC address of the NIC through which the device connects to the switch (bridge) and the port number of the switch (bridge). When the switch (bridge) is powered, the table is empty. Hence, when the first frame arrives it is sent to all other ports except the port on which it arrived. At the same time, the switch (bridge) fills in the row of the table with the port where the frame arrived. The source address in the frame is the NIC address for this particular port number. In this way, the next time a frame with this destination arrives it will be forwarded to this port only. Usually the whole table is populated within seconds and is updated every time a new frame is received. The switch (bridge) checks the destination address in each frame in order to forward the frame and the source address in order to keep the table updated.

Wireless Ethernet, mainly known as Wi-Fi technology or hot-spot is also a flavor of the Ethernet. The wireless Ethernet has a central point usually called Access Point (AP). The medium is the air. The electrical signals are in the form of electromagnetic waves that spread equally in each direction from the antenna of the device. Instead of using CSMA CD it uses CSMA CA (Carrier Sense Multiple Access with Collision Avoidance). This is because of the possibility two nodes to be in the radius of reach of AP and out of the radius of reach of each other. The phenomenon is known as hidden terminals. Because the devices can be out of reach of each other it is not possible to detect collision. Hence, each device communicates shortly the AP before transmitting in order to get information whether another device has a permission to communicate with AP. When confirmation is received the device can transmit the signal. During this transmission no other device can get a permission to transmit.

6. Types of Communication

Unicast or one-to-one communication is the simplest way to communicate. Conversation with a single person is one-to-one communication, regardless whether it takes place face-to-face or over the phone or over e-mail. On the Ethernet one-to-one communication is achieved by identifying the two machines that exchange data by their MAC addresses.

Multicast is one-to-many or many-to-many type of communication. For example a conversation among a number of people is many-to-many type of communication, sending e-mail to a group of people is one-to-many communication. In the Ethernet there is a special MAC address that is invented for machines to send information once if they want to reach several machines that belong to a group. For this purpose a special MAC group address is built in the NIC for machines in the group. Then the sending machine sends the information only once instead of separately to each machine in the group.

To broadcast information means that anyone who wants can get it. Good example is information broadcasted through the radio or TV. Every machine on Ethernet type of network has built in broadcast MAC address, so that information needs to be sent once for every machine to receive it.

Two types of communication can be distinguished depending on the time when the two communicating parties (the sender and the receiver) are active. One is called synchronous and in this case both the sender and the receiver are active at the same time. For example, the phone conversation is synchronous because when one person talks the other listens. Another example is chat over the Internet, because one person is writing and the other is reading at the same time. The other type of communication is asynchronous where the sender and the receiver are active at different times. Example for this type of communication is e-mail, sending letters, the answering machine on the phones. In all these examples when the sender is active, the receiver is not and vice versa.

7. *What are packets?*

Sometimes the information sent can be long. It is not convenient to send long chunks of data. Therefore users' data are usually split into smaller parts and control information is added to each of these parts. The data together with the control information is called a packet. The main two reasons for using packets are the following: First if there is a loss of a packet or it is corrupted during transmission, it is easier to resend only the corrupted packet instead of the whole message. Since every packet carries the control information, packets from different sources can share the same network resources. In order to avoid sending too much control information, it is recommended to have moderately large packets. The length of the packet is often limited by the properties of the media through which the packets are transmitted.

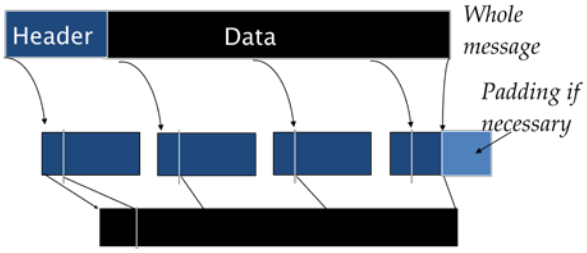

The picture on the right side illustrates how the message is divided into packets. The header is the control information carried by the message. All packets created from a single message need to have the same header and some additional control information that will help the receiver to combine the packets back into the message when they arrive at the destination. Sometimes, the length of all packets needs to be the same. If this is the case the shorter packets are stuffed with dummy content. This process is called padding.

8. *Packet Switching versus Circuit Switching*

Wide area network connect local area networks at different geographical locations. Most often the wide area networks consist of point-to-point connection or one node of a LAN in a single place is connected to a node of another LAN at a geographically remote place. These nodes are special devices used for connecting different networks and are called routers.

The long distance transmission is costly. Therefore it is important to do it in an efficient way. There are two paradigms with regards to the use of long distance circuits. One is called circuit switching and the other packet switching.

Circuit switching has been used for a long time by the telephone networks. A circuit between a sender and a receiver is established by request of the sender. Network resources (lines and switches) towards the receiver are reserved. Once the receiver answers the call the charges for these resources are made by the time the connection lasts. Even if there is no traffic through the connection, no one else can use it. Charges will be in effect until the connection is closed. This type of switching is not suited to data traffic which comes in bursts (a lot of packets from a communication between a single sender and single receiver are sent and then there are no packets from this communication at all).

Packet switching does not use dedicated connections. Packets are sent to the next router when they are generated by the sender whenever the link is idle. There is no request for establishing connection. Routers in the network move the packets closer to the destination until they reach the receiver. Packets belonging to a single message can travel different paths and may arrive at the receiver out of order. The ordering is done by the receiver based on the information sent by the sender. This type of switching is more efficient for data traffic that comes in bursts. The same resources (transmission lines and routers) can be used by any traffic whenever they are idle.

The picture bellow illustrates circuit and packet switching.

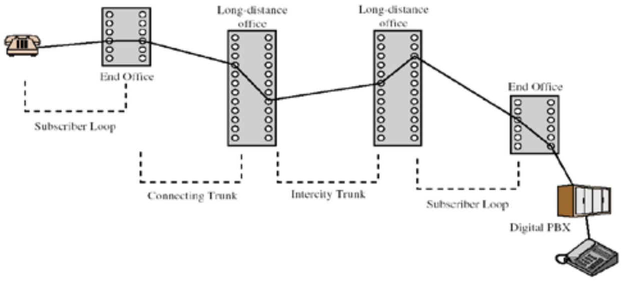

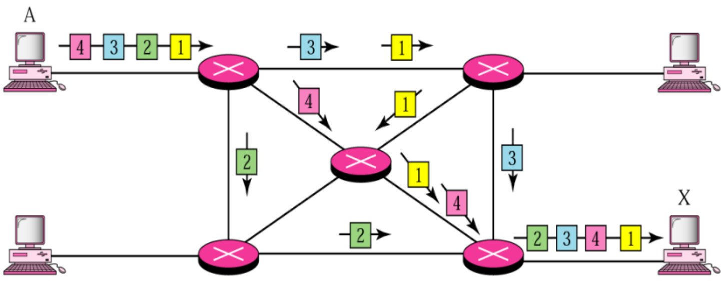

9. *Layering*

Networking is complex. Many issues need to be resolved when moving data from one node to another. Nodes and links are with different characteristics and there might be many networks in between the source and the destination. The solution to the complexity is to use the divide and conquer technique. This technique splits a big problem into smaller ones and deals with each of them separately. In context of computer networks the divide and conquer technique is implemented in such a way that each small communication task is incorporated into a single protocol. There are many protocols. Some of them perform tasks that are somehow similar to one another. Therefore they are grouped in a single layer. Depending on the nature of tasks they perform the protocols at a single layer are placed closer or further away to the physical medium for transmission. All the layers are one on the top of the other starting with the physical layer which is the closest to the physical medium for transmission and ending with the application layer which is the closest to the user running different applications. A set of protocols organized in layers is called a protocol suite or protocol stack or communication model.

As we defined them before, protocols are actually a combination of software that exists on the machines and the control information that is added to the packets so that the sender and receiver can negotiate how to handle the data sent. When the sender machine has a packet to send, the user's data are processed by the top layer (application layer). Here the particular control information is added to the user data as a header of the packet. This packet is passed to the lower layer. The lower layer is independent of the upper one and it adds its one control information to the packet. The same procedure is done with the other layers except the last one, the physical layer. The physical layer transforms the whole packet into electrical signal and sends it through the network.

At the receiver the procedure is opposite. The lowest layer looks at the control information added by the same layer at the sender, interprets this information, removes this part of the header and passes the packet to the higher layer. This is repeated with the other layers until the original user's data are obtained. The picture below illustrates the process that takes place at the sender (protocol construction) versus the process that takes place at the receiver (protocol reduction).

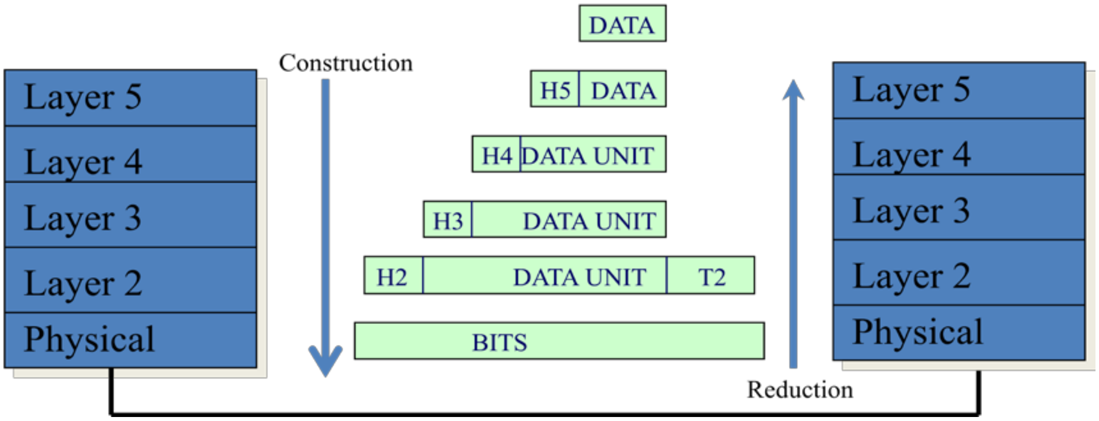

Not all devices need to have all the layers. For example routers operate only with the three lower layers and therefore these layers are present there. End devices need all the layers. A virtual connection is established between the same upper layer at the sender and the receiver. They communicate with each other although there is no direct physical connection between them. The information is sent through the lower layers at the sender, conveyed unchanged through the routers and received via lower layers at the receiver.

A theoretical model often used is the Open Standard Interconnection model (ISO model) created by the International Standard Organization. It consists of seven layers. The well known TCP/IP model used in the Internet has only five layers.

10. What is Internet?

A more complicated data network or internetwork consists of many connected networks. Some of them are LANs and some are WANs. Computers which belong to LANs are called end systems or hosts. Different networks are joined with devices called intermediate systems or routers. These types of networks use packet switching for moving information from one end device to the other. The picture is a simple illustration of interconnecting networks.

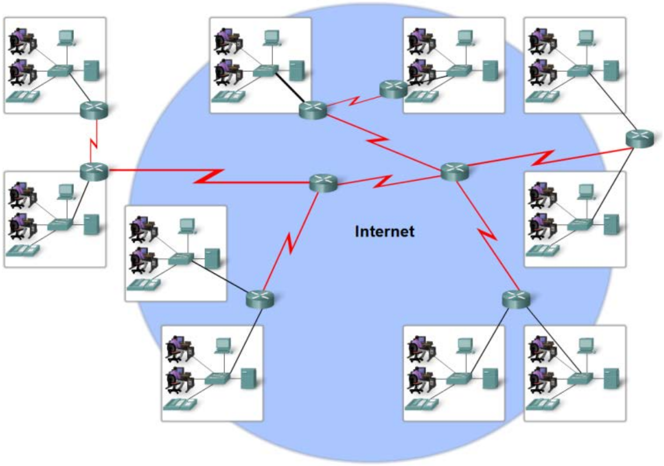

Internet is internetwork that uses a common protocol model. The name of the model is TCP/IP (Transmission Control Protocol/Internet Protocol). The network offers a myriad of services. The old ones are e-mail, WWW (World Wide Web), FTP (File Transfer Protocol), Telnet (a protocol for connecting to a remote computer). The new ones are protocols for using VoIP (Voice over IP), multimedia streaming, security issues, and many more.

11. The History

The predecessor of the Internet is ARPANET. It is a network created by the United States Department of Defense in 1969. The main goal in building the network was to achieve a high level of resilience, so that if nuclear attack destroys a part of it, the other part can still be operational. In the 1970s several US universities started developing and experimenting with these ideas. They had in mind to connect different types of computers and to provide access to expensive supercomputers located at some places to the academic community in the whole country. TCP/IP model was born and messages using these protocols were for the first time sent in 1984. Besides Telnet, the application for accessing a remote computer, other applications were developed. E-mail showed to be the most successful. Others followed, for example Gopher and Archie were predecessors of the searching engines. They were mainly with text based interface. The beginning of the 1990s is the period when the killer application emerged. The World Wide Web (Web 1.0) made the Internet a global network and gave boost to creation of many other applications. The graphical user interface attracted ordinary users to the network and paved the way to commercialization. Web 2.0 emerged in the last decade. It allows users not only to browse, but also to edit different types of content or to organize in many kinds of social networks.

12. Client/Server Paradigm

Most Internet applications, as well as many protocols in the TCP/IP model operate using client/server architecture. Server can be thought of as a process that is running on some machine, such that it waits to serve all the requests coming to it. Client is the process that sends requests. Usually these two processes exist on different machines; hence the machines are also referred as a server and a client. The machine where the server is running is often much more powerful and has big storage necessary to store data to be sent upon request. A client can request a web page (from a Web server) or a file (from FTP server). A single

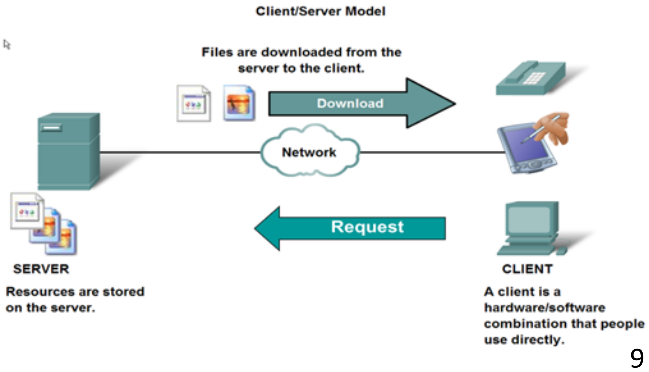
Client/Server Model
Files are downloaded from the server to the client.
Download
Network
Request
SERVER
Resources are stored on the server.
CLIENT
A client is a hardware/software combination that people use directly.

server serves many clients one after another. This is done quickly so that users have a feeling that the server is working simultaneously with all of them. On the other hand, clients can request almost simultaneously data from many servers. The picture on the left illustrates the client/server model.

13. TCP/IP Protocol Model

TCP and IP are not enough to address all the issues in a complex network such as the Internet. Internet is complicated not only because of its size, but also because of its structure that involves a large number of different devices (computers, mobile phones, PDAs, routers, switches, hubs) and thousands of applications. Moreover, it spans different technologies and is changing all the time. The picture below shows the layers in the TCP/IP protocol suite together with some of the protocols that belong to each layer.

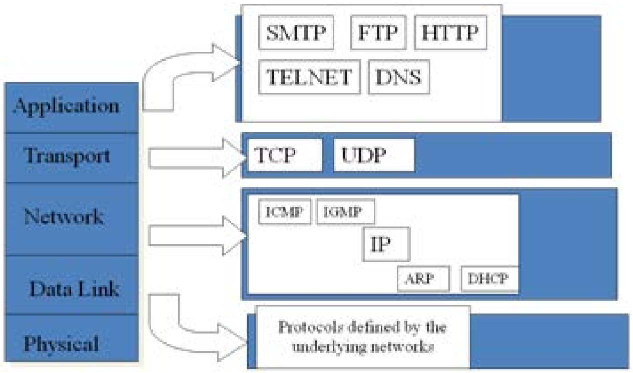

The protocols in the network layer are helping the IP. For example ARP (Address Resolution Protocol) resolves MAC address from the IP address, DHCP (Dynamic Host Configuration Protocol) dynamically configures hosts (assigns the IP address, informs about the IP address of the default gateway, DHCP server, closest DNS server), ICMP (Internet Control Message Protocol) is responsible for various controls. For example the programs *ping* and *traceroute* use ICMP. It is also used by the routers to inform the hosts that have send packets when their packets was lost and has many other control functions. IGMP (Internet Group Management Protocol) is used to gather information about groups who want to receive certain content. SNMP (Simple Mail Transport Protocol) is used for sending E-mails. It is very simple and efficient and is often misused by those who want either to send commercials in a cheap and easy way, or misuse people by sending them offers to disclose personal information like bank account numbers or passwords. These kinds of e-mail messages are called SPAM or pishing.

TCP is very complex and slow protocol and it is indispensable for certain applications. However, in certain occasions a simpler transport protocol is necessary. UDP (User Datagram Protocol) is used for that. For example DNS uses UDP because fast resolution of the IP address is crucial.

Internet protocol is in the center of the TCP/IP protocol stack and proved to be very resilient. The proponents of IP have a saying: You can run IP over anything and any application can be run using IP. We have witnessed many new data-link layer technology developed and a myriad of applications over the network. They all use IP.

14. IP addresses

Each and every device connected to the Internet has a unique IP address. There are two versions of IP addresses used today, version 4 which is still prevalent and version 6 which is emerging. IPv4 (IP version 4) addresses consist of 32 bits (0s and 1s). In order to make the addresses easily readable by humans the so called dotted decimal notation is used. The address is written as four numbers in the decimal number system each corresponding to the eight-bit binary number. The four numbers are separated by dots. For example, a valid IPv4 address is 152.23.45.6. Unlike the MAC or physical addresses the uniqueness of these addresses is assured through their allocation. They are logical addresses and are bind to the network card of the system. Internet authorities are responsible for allocating a range of addresses to each continent. Then parts of these addresses again in hierarchical manner are distributed to the ISPs and ISPs allocate them to the users. The uniqueness of the addresses is achieved in the way that each one in the hierarchy of distribution takes care not to allocate the same address to several entities.

The number of IPv4 addresses is very big. Since each bit in the address can be 0 or 1, there are 2^{32} possible different combinations. For various reasons all these addresses cannot be used. In addition, because of unwise allocation of the address space during the early days of the Internet and the unexpected exponential growth, IPv4 addresses are almost depleted today. Hence, many new ways were invented to save the address space. One way is to use a special server running DHCP (Dynamic Host Configuration Protocol) which allocates IP addresses to the hosts in the local network only when they are needed. This means that two machines can use the same address at different times. Another way is to use addresses defined by the Internet authorities as private IP address and NAT (Network Address Translation). These private addresses are not allowed to be used in the Internet. However, the Network Address Translation is software that makes possible to change the private address in the header of the packet with a public one in the first router on the way to Internet. In this way many private addresses can be used inside the local network and small number of public addresses when communication outside is needed. The IPv6 addresses are 128 bits and provide much bigger address space. The intention is to replace all IPv4 addresses. However, this process is very slow. Hence, in this text we will concentrate on IPv4 addresses only.

Each IP address has a network part and a host part. Hosts connected to a single local area network (for example Ethernet) have the same network part of the IP address and differ only in the host part. The subnet mask is used to define which part of the IP address identifies the network part. It consists of contiguous number of 1s and 0s and their total number is 32 (the same as in the IP address). The number of 1s defines the length of the network part of the address and the number of 0s the host part. Instead of writing the IP address and the subnet mask, another notation is used. The subnet mask can be written in dotted decimal notation or given as a number of 1s immediately after the IP address. If a decimal number is 255 it corresponds to 8 1s. For example, a subnet musk 255.255.0.0 has 16 1s. The picture below shows an IP address where the network part is 24 bits long. Hence, the subnet mask is 255.255.255.0 or the address of the host is 193.5.1.5/24. The network address is obtained if all bits in the host part are replaced with 0s. Then the address for the network is 193.5.1.0/24. This address is not used by any host, but is useful for the routers. When IP addresses are allocated these network addresses are used, too. Then the allocator determines the network part and the administrator of the network takes care that the host addresses are unique.

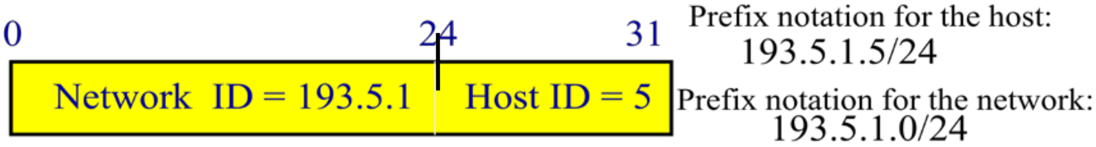
0
24
31
Network ID = 193.5.1
Host ID = 5
Prefix notation for the host:
193.5.1.5/24
Prefix notation for the network:
193.5.1.0/24

Once the host obtains answer from DNS about the IP address for the destination it can create the first packet to be sent with user data. A part of the control information sent is the IP addresses of the source and the destination. The network mask helps the source to decide whether the packet will be sent to the machine in the same network or not. In the first case the network address for the source and the destination will be the same. Hence, it is not necessary to send the packet to the closest router, usually called default gateway. Then, besides the IP addresses, the MAC address of the source and the destination are also placed as the control information and the packet can reach the destination using the local Ethernet. In the second case the source sends the packet to the default gateway which for the sender is also a machine on the same network. Hence, it can send the packet using the MAC address. The process of moving the packet from the router that is the closest to the source to the router that is the closest to the destination is known as routing.

15. Routing

Ordinary computers usually called hosts have one network interface card (NIC) to connect to the network. Since the routers connect several networks they need to have several NICs. The IP address of a host is actually associated with the NIC. The router NICs have IP addresses that belong to different networks. Each NIC is associated with the IP address that belongs to the network this NIC is connected to. Routers forward packets by switching them from one incoming port (NIC) to other outgoing port (NIC) so that the packet is forwarded closer to its destination.

The basic architectural components of a router are the input and output ports, the switching fabric and the forwarding table. Input and output ports are actually network interface cards for the router. Packets entering through the input ports are processed inside the router. Processing consists of finding the output port that moves the packet closer to its destination. Once the output port is determined the switching fabrics is used to switch the packet from the incoming to the appropriate outgoing port. The destination IP address carried as control information in the packet together with the forwarding table that needs to be present at the router is crucial for selecting the output port.

The forwarding table consists of at least two columns. The first column contains destination addresses and the second column contains the next hop. Destination address is an IP address of some network. The next hop is the next router to which the packet should be sent. Forwarding or routing tables are entered in the router by humans, usually network administrators, or they are automatically created by special purpose applications called routing protocols. These are applications run by the routers.

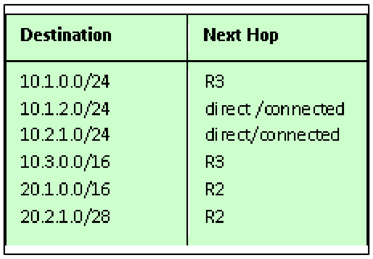

Destination	Next Hop
10.1.0.0/24	R3
10.1.2.0/24	direct /connected
10.2.1.0/24	direct/connected
10.3.0.0/16	R3
20.1.0.0/16	R2
20.2.1.0/28	R2

The table on the left is an example of a forwarding (routing) table for one of the routers on the picture shown below. The IP addresses of the networks are given in prefix notation. For simplicity, the names of the routers are written in the next hop column instead of the IP addresses of the ports on that router. The notation direct/connected means that the packet should be delivered directly to the destination (the destination is on the same network as the router).

The picture below presents an example of forwarding a single packet. Host H1 sends a packet to host H2 with IP address 20.2.1.2. The packet is first sent to the only router to which H1 is connected, router R3. This router compares the destination address 20.2.1.2 with all the entries in the first column of its forwarding table. It chooses the row with the closest match (the row where the network address matches the destination address in the largest number of bits). In this case it is the row with router R4 as the next hop. Therefore the packet is forwarded to router R4. In the same way R2 finds out that this packet should be forwarded to router R2 and R2 finds out that the destination is directly connected. Then, the packet can be delivered using the MAC or the physical address of the destination. A separate protocol ARP (Address Resolution Protocol) is used to resolve the MAC address from the IP address.

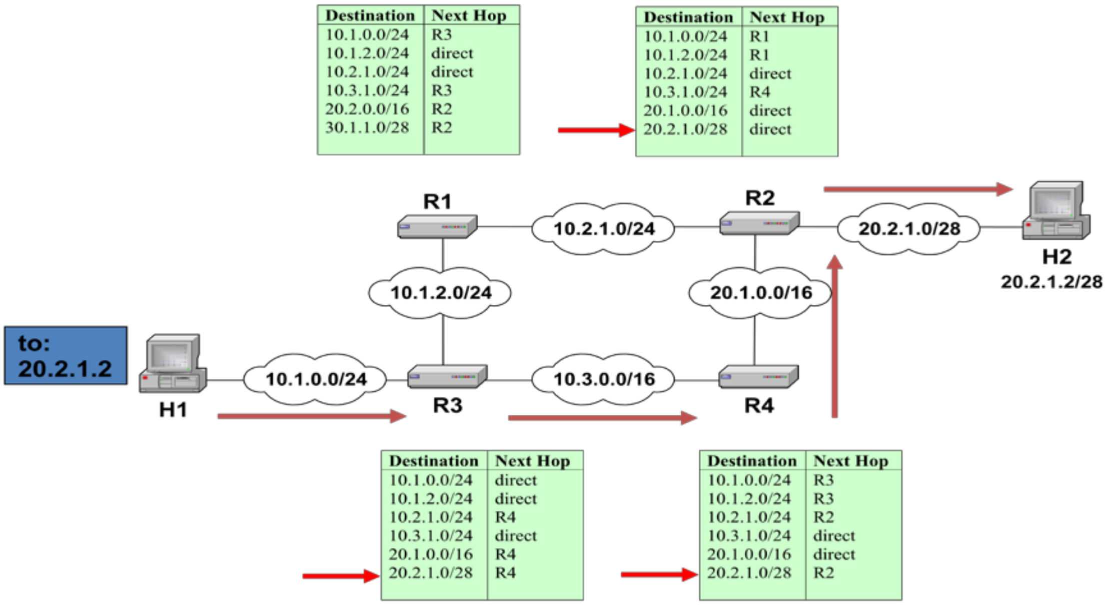

Destination	Next Hop
10.1.0.0/24	R3
10.1.2.0/24	direct
10.2.1.0/24	direct
10.3.1.0/24	R3
20.2.0.0/16	R2
30.1.1.0/28	R2

Destination	Next Hop
10.1.0.0/24	R1
10.1.2.0/24	R1
10.2.1.0/24	direct
10.3.1.0/24	R4
20.1.0.0/16	direct
20.2.1.0/28	direct

Destination	Next Hop
10.1.0.0/24	direct
10.1.2.0/24	direct
10.2.1.0/24	R4
10.3.1.0/24	direct
20.1.0.0/16	R4
20.2.1.0/28	R4

Destination	Next Hop
10.1.0.0/24	R3
10.1.2.0/24	R3
10.2.1.0/24	R2
10.3.1.0/24	direct
20.1.0.0/16	direct
20.2.1.0/28	R2

The correctness of the forwarding tables is crucial for the delivery of packets. For example, if the next hop in the last row of the forwarding table of the router R4 is R3 instead of R2, then the packets sent by host H1 to host H2 will go back and forth between routers R3 and R4. This means that a loop is created. The IP protocol introduces the so called TTL (Time-to-live) parameter to solve problems like this. The time to live is set to some value when the packet is sent and it is reduced by one at each router the packet traverses. When its value reaches zero the router discards the packet. In this way unnecessary traffic is removed from the network.

Packets that are removed by the routers because their time-to-live has value 0 are lost. Packets can be also lost if the router does not have enough buffers to store the incoming or outgoing packets in case there is too much traffic. Therefore IP is called an unreliable protocol and it is said that it offers the best effort service. The protocol that corrects this is called TCP (Transmission Control Protocol).

16. Delays in the Internet

When packet switching is used the delay that is experienced by each packet (time necessary to travel between the sender and the receiver) consists of four types of delays:

1. The transmission delay depends on the length of the packet and the throughput of the links (the speed of the links) through which the packet is transmitted. On a particular link it can be calculated as a ratio of these two. Total transmission time is the sum of the times for particular links.
2. The propagation delay is due to the time necessary for the electrical signal to propagate through the link. It depends on the distance or the length of the link and the propagation speed which depends on the type of the medium. For example for fiber and air it is equal to the speed of light while for copper it is 2/3 of this speed. Total propagation time is the sum of propagation times on particular links.
3. The nodal delay is the time necessary the control information to be processed at the nodes, before the packet is forwarded to the next node. It depends only on the processing power at the particular node. The total for this type of delay depends on how many nodes the packet traverses.
4. The queuing delay becomes important when there is a lot of traffic. Then, the packets need to wait in a queue to be processed at the nodes. This delay varies depending on the traffic at particular time.

The total delay is the sum of all four types of delay. You can check the delay of some probe packets by using the program “ping”. “Traceroute” program shows each node on the path together with the round trip time.

17. Transmission Control Protocol (TCP) and User Datagram Protocol (UDP)

IP is the protocol at the heart of the Internet and TCP (Transmission Control Protocol) is not less important. While IP provides host to host connection, TCP is responsible for supporting the application. It establishes virtual connection between the processes. For example, the protocol used for the Web is HTTP (Hyper Text Transport Protocol). It operates using TCP. In order to provide connection between more than one processes on the hosts TCP uses another types of addresses called ports. They are 16-bits addresses. The first thousand ports are reserved for different types of services and are called well-known ports. For example, the Web server always listens for requests on port 80 and HTTP requests are always sent to port 80. When a client asks for a service it usually knows what is the destination port for that particular service.

One of the main functions of the Transmission Control Protocols is to provide reliability and resolve the issue of lost packets. TCP achieves this by introducing acknowledgements for the packets. When the packets arrive at the destination, acknowledgement is sent back to the source. If the source does not receive the acknowledgement for a certain packet, it assumes the packet has been lost and sends it again. The ordering of the packets is taken care of by attaching sequence numbers to each of them. This way the receiver is capable of reordering the packets that come out of order and removing duplicate packets.

TCP also has methods to check whether a packet has been corrupted (some of its contents is not equal to the one that was sent). In that case the packet is discarded and a copy of the original is sent. Finally, TCP also takes care of how fast the sender generates packets. In case the network or the receiver is overloaded. It has mechanisms to slow down the speed at which packets are sent. Hence, it is often referred to “well-behaved” or “polite” protocol.

UDP (User Datagram Protocol) is another protocol at the transport layer designed at about the same time as TCP. Unlike TCP which is very complex and has many functions, UDP is very simple. The main control information it contains are the source and destination port addresses. It does not require acknowledgements. Therefore it is fast. If a packet is lost the application has to worry about it. UDP was designed this way to work with applications where it is important to get the answer very quickly. For example the domain name system (DNS), the application that maps names into addresses uses UDP as a transport protocol. Today, UDP is exploited for some applications that involve video and audio transmission. These types of data don’t care if some packets are lost, but for them it is very important that the packets are sent fast. Therefore they prefer UDP to TCP.

18. Doman Name System (DNS)

IP addresses consist of numbers and humans have hard time to memorize numbers. Names that humans associate with the content they want to get are much easier to remember. That is why we use names to reach a certain web page or e-mail address, or to send e-mail message. Whenever a web page is requested the host needs to find out which IP address corresponds to the name written in the address bar of the browser. The application that does this is called domain name system or shortly DNS.

In the early days of the Internet when the number of computers was rather small there was a single file called hosts.txt where the mapping between the names and addresses was stored. The file was updated manually whenever a new system was added to the network. Then it was broadcasted to all the computers. Each machine had updated information and could look for the address of the name whenever it was needed. As the size of the network increased this method proved to be impractical and in 1983 DNS was created. It splits this table into many pieces and distributes the authority for its management.

DNS consists of three elements: a name space, servers making that name space available, and clients which query the servers about the names. In order to deal with the complexity (large number of organizations and large number of hosts), the name space is hierarchically organized as the root of a tree. Under the root there are a number of top level domains, determining the type of the organization. For example, *com* is for commercial organizations, *edu* for educational institutions, *mil* for the military, *org* for non-governmental organizations etc. They are called generic top level domains. Later on, when all the countries joined the network, country top level domains were added. Each country has a top level domain consisting of two letters, like *se* for Sweden or *au* for Australia. The root-server keeps the information about the IP addresses of all servers for the top level domains. There are actually 13 main root servers and many copies to them. Ten are located in USA, one in UK, one in Sweden and one in Japan. Only Internet authorities can add new top level domains. The administrators of a particular top level domain have complete control over the names under that domain. They can add new names and sub-domains. There is no limit to the number of sub-domains or levels they can create. Moreover, the name space is not related to the physical interconnection. Instead it is of organizational nature. Any kind of hierarchy is allowed to be used under the top level domains (for example names of cities, regions or institutions). The picture below shows a part of the name space.

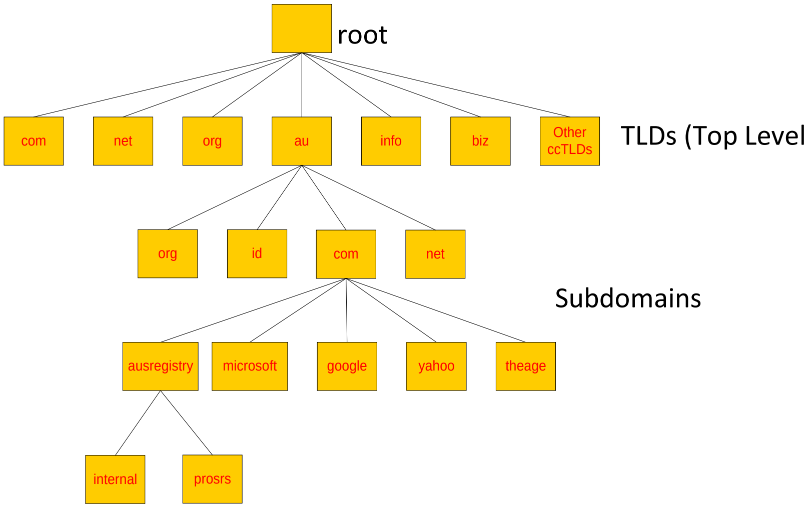

The domain name for a certain machine consists of several names from the name space. At the end is the top level domain, and at the beginning the name of the machine. The names in between are the names from the hierarchical levels on the path from the top level domain to the name of the machine. The parts of the name are separated by dots. For example, on the picture above the fully qualified name of the machine *internal* is *internal.ausregistry.com.au.* Sometimes machines can have several names which are called aliases. For example if a DNS server is running on this machine, the alias can be *dns.ausregistry.com.au.*

The responsibility for e certain TLD (Top Level Domain) can be further delegated to several organizations down through the tree hierarchy. Each organization that is responsible for a certain part of the name space maintains DNS server that stores the part of the table with names and corresponding IP addresses for that part of the name space.

DNS clients are located on the ordinary hosts and they are called resolvers. Often there is a local name server that has a task to resolve the addresses for all hosts on the LAN. It memorizes (cashes) the addresses that are resolved in the near past and sends immediately response to the hosts asking for these addresses. If there is a request for some other name, the root server is the first to be contacted. A request for the IP address of the DNS server for the top level domain in the fully qualified domain name is issued. Once the response is obtained, the top level domain server is requested to send the IP address of the DNS server for the next level. The procedure continues until the requested IP is obtained. The process of resolution for the fully qualified domain name www.ripe.net is presented on the picture below.

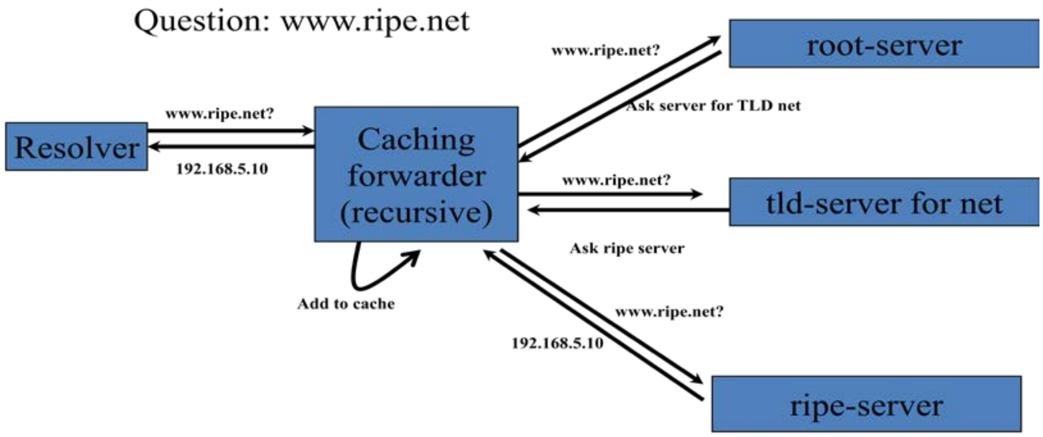

19. Name Space versus Directory Hierarchy

The arrangement of directories on the computer storage is similar to the DNS name space. The picture on the right shows an example of such arrangement on the computer with Windows operating system. The directories also have names and are arranged in a hierarchy that resembles a root of a tree. The path a certain directory is determined with the names of the directories on the path separated by slashes. For example, the path to the directory *drivers* will be *C:\Windows\System32\drivers*.

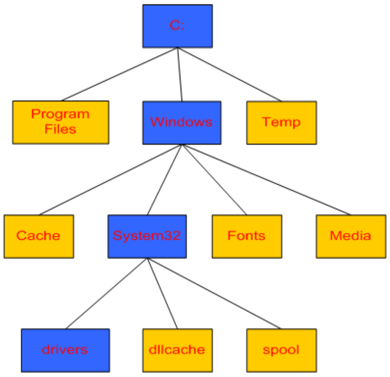

The name in the address bar of the browser sometimes includes the path to the page displayed. In that case the name is identified as URL (Uniform Resource Locator). It is actually complete identification of the path that leads to what is shown in the browser. It consists of the fully qualified domain name and the directory on the server where the files displayed are located.