Chapter 6
Computer Networking

Introduction

There are two computer network architectures: the client server and the peer-to-peer architecture. In the client/server architecture, all communication happens between the server and the client only. A client does not communicate directly with another client. If a client wishes to share a file with another client, it must place the file in a commonly shared area at the server. In case of an application, the client makes the request and the server processes the request and returns the result to the client. In the peer-to-peer architecture, any node can act as a server as well as a client. Only very limited security is available on a peer-to-peer network.

A computer network generally has six components: computers, communication ports, cables, hubs or switches, routers, and the network software. The computers in a network, nodes, can range from a handheld computer to a supercomputer. While several communication ports could be used, for this chapter the description will be limited to the Ethernet card. Depending on the network card or ports used different types of cables would be required. For the Ethernet card coaxial or twisted pair or fiber cables can be used. Only twisted pair cables will be discussed in this chapter. Connecting more than two computers together will require a device to connect all the cables coming from each machine. This device may be a hub or a switch. In order to connect two computers that belong two different networks there needs to be a pair of routers that connect one network to the other. A Network Operating System (NOS) software will be required to make the communication to occur.

Computers in a network

There are no special requirements for the computers in a peer-to-peer network. However, in a client/server architecture, the server will have to be robust—must be able to run for months or years without being shut down. It must be fast enough to handle all the requests from the clients, and must have large fast hard drives to store information for all users. There must be enough memory to cache all the directory information as well as most used pages. The computers must be reliable in case of power failures and hardware failures—sufficient redundancy must exist. There are different types of servers, all having different requirements. Some examples of servers are, fileserver, print server, database server, domain name server, and communication server.

Network Interface Card

Each computer to be connected to a network must have a network interface card (NIC) installed. While Ethernet and Token Ring cards are used today, the Ethernet NIC is by far the most popular. Most computers come with a built-in Ethernet card; otherwise, an Ethernet card must be purchased and installed. With the plug and play (PNP) feature of the hardware and operating systems today, installing a card is as easy as opening the case, plugging in the card, and inserting the driver disk that comes with the card when asked by the operating system. If for some reason the plug and play feature does not detect the NIC, the interrupt request line (IRQ), base memory address, and the base I/O port must be manually set up on the card by setting the jumpers or DIP switches, or through the use of the manufacturer provided setup software. The connector on most Ethernet cards today is the RJ-45 having 8 pins. Older Ethernet cards may have a Bayonet Nut Connector (BNC), an Attachment Unit Interface (AUI), and/or a RJ-45 connector. A jumper or switch will specify the type of connector used.

The Ethernet uses a shared medium and all connected computers can hear all transmission, and all computers can transmit on the medium. This could introduce a lot of confusion and wasted bandwidth. To ameliorate this situation the Carrier Sense Multiple Access with Collision Detection (CSMA/CD) protocol is implemented in the Ethernet. This protocol dictates that a sender must first listen to the media and may send only if there is no activity on the media. It is possible for two or more listening computers to start communication at the same time and cause collisions. This collision can be detected as a spike in the voltage on the line. The collision detection mandates a backing off period. The Ethernet cards are capable of transmitting 10 Megabits per second (Mbps), 100 Mbps (Fast Ethernet) or 1000 Mbps (Gigabit Ethernet) depending on the card purchased. The faster cards are downward compatible. In order to take advantage of the full speed, every component in the entire network path should be able to communicate at the full speed. If the components are mixed, the network will fall back to the lowest speed of the slowest component in the communication path.

Every computer connected to a network should have two unique identifying numbers, a physical address and an Internet Protocol (IP) address. The physical address is used to send messages within a local area network and the IP address is used to send messages to the outside world. The physical address, a 6 Byte number, is encoded on a ROM chip on the Ethernet card at the time of manufacturing. The IP address is typed in on each machine at the time of network configuration. The IP address scheme will be covered at length later in this chapter.

Hub/Switch

The next step is to obtain a hub or switch with sufficient ports to connect the wires from all the computers on the network. All hubs sold today are active hubs. An active hub amplifies all signals received and sends the amplified signals on all its ports. Passive hubs are signal splitters and are used only in wiring panels. Even though a little more expensive, a switch will provide faster communication and great deal of diagnostic features. Switches provide link management through physical address identification. As the network grows, multiple switches or hubs can be connected with each other using an uplink port or a direct connect cable. Most modern switches have the auto-sensing feature, which will allow any port to work as an uplink. The best location for a switch or hub is a centrally located closet to which cables from all computers are brought.

Network Cabling

Cabling is the most time consuming part of networking, particularly when cables need to be hidden inside the walls. For that reason it is advisable to have the cables installed or conduits placed at the time of the building construction. To hide wires inside existing buildings, a tape fish will be needed. It may be well worth the time to investigate the wireless technology, particularly when a small area is being networked. Three types of cables are used today: coaxial, twisted pair, and fiber-optic. This discussion only deals with the twisted pair cable. Twisted pair cables either can be unshielded or shielded with a foil. Unshielded twisted pair (UTP) cable can transmit 10 Megabits per second and the shielded can transmit more than 100 Mbps, perhaps all the way up to 1000 Mbps. Wires in each pair are twisted to reduce cross-talk and minimize the effect of external electromagnetic interference. Twisted pair cables are categorized into five categories, from Cat 1 to Cat 5, based on the bandwidth capabilities. The maximum length of a segment is 100 meters.

The UTP cable has 4 pairs of color-coded (orange, green, blue and brown) wires. The color of one of the wires in a pair will be solid and the other will be striped with white. Pins in the RJ-45 jack are numbered from 1 to 8. An acceptable wiring color scheme to place the wires into the jack is given in Table __ . Upon placing the wires in the jack a crimp-tool is used to crimp wires into the jack. After crimping both ends of a cable, it must be checked for continuity of each wire using a cable tester. When both ends are connected according to the pin order given in Table ___, that cable is called a straight-through cable. When two computers are connected directly without the use of a hub or a switch some wires need to be crossed (Transmit to Receive). Such a cable is called a crossover cable. Crossover cables need to be used when two hubs or switches that lack uplink ports are connected together.

Pin nu mb	**Color of the wire**	**Activity**

er		
1	Orange/White	Transmit
2	Orange	Receive
3	Green/White	Transmit
4	Blue	Receive
5	Blue/White	Transmit
6	Green	Receive
7	Brown/White	Transmit
8	Brown	Receive

Table ___
TIA/EIA T 568B standard

Connector 1	Connector 2	Diagram
1	3	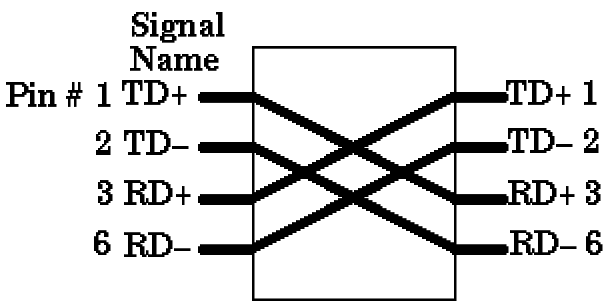 Signal Name Pin # 1 TD+ 2 TD– 3 RD+ 6 RD– TD+ 1 TD– 2 RD+ 3 RD– 6
2	6	
3	1	
4	Not used	
5	Not used	
6	2	
7	Not used	
8	Not used	

Table ____
A crossover cable

Network Operating System

The computer networking had its beginning in the1960s when the Advance Research Project Agency (ARPA), the research arm of the Department of Defense, became very interested in computer networking. The formation of this agency was part of the U.S. reaction to the then Soviet Union's launch of Sputnik in 1957. As a result, the ARPAnet was born in 1967, which eventually became what is known today as the Internet. By 1972, ARPAnet had grown to approximately 15 nodes. The protocols derived out of ARPAnet is known as Transmission Control Protocol/Internet Protocol (TCP/IP).

IBM developed its own network protocols known as the SNA, and the Digital Equipment Corporation had its proprietary network called DECnet. The International Standards Organization (ISO) began developing standards for networking, and published the Open Systems Interconnection model in 1978. Meanwhile, with the explosive growth of PCs in the 1980s, many private companies developed network technologies such as the BITnet, CSNET, NSFNET, Novell, Banyon Vines, Invisible Net, and Lantastic, to name a few. Most of these were sold as a software and hardware combination, and the companies were unwilling to standardize, and many went out of business. Novell, the IPX/SPX protocol suite, the most popular networking software for the PCs in the 1980s, finally adopted the

TCP/IP in 1990s. The OSI model essentially became academic and all predominant network software uses the TCP/IP protocol suite. The IP part of the suite transmits packets with no reliability measures, and hopes that they will get to the destination. A reliable transmission protocol was the most logical addition, which is the TCP. It did not take long to realize that transmission such as a voice package required timely arrival rather than reliable arrival of every packet since human being is capable of filling in the missing parts, which gave way to the user datagram protocol (UDP). The IP now can operate on top of either TCP or UDP.

There exists three TCP/IP network environment today, Novell Networking, Microsoft Windows Networking, and the UNIX/LINUX networking. The Novell Networking incorporates its original IPX/SPX protocol suite as well as the TCP/IP protocol suite. The other two operating systems include TCP/IP networking as part of their operating system. All versions of Windows provide the peer-to-peer networking capabilities, while the NT, 2000 professional, and the XP professional provide for Client/Server architecture. The UNIX/LINUX has consistently adhered to the Client/Server architecture. It is beyond the scope of this book to compare and contrast these different types of networks.

Windows peer-to-peer networking

The peer-to-peer network may use the TCP/IP protocol suite or some other simpler protocols such as the NetBIOS and the NetBEUI. The NetBIOS, an application program interface (API) extends the BIOS, discussed in chapter 3, to include the support for I/O calls over a network. The NetBIOS Extnded User Interface (NetBEUI), developed by IBM and Microsoft, is the protocol used by Windows Workgroup networking. This protocol is non-routable, therefore could not reach beyond the local physical network. It does not require an IP address, rather works with the name registration within a workgroup. In the newer Windows operating systems, the NetBEUI is encapsulated inside the TCP/IP, referred to as NetBIOS over TCP/IP (NBT), and uses the IP address.

Once all the hardware and software have been installed, the network portion of the software need to be configured. In windows, launch the network property window and configure each item given there. Different versions of Windows have different ways to get to the network properties window. In Windows 98, right click the **Network Neighborhood** icon. In windows XP right click the **Network Connections** icon. . If the installed NIC does not appear there, it needs to be manually configured by installing the driver that came with the NIC. The next step is to install the protocols used. In case of windows workgroups and an IP address is not available, choose the NetBEUI. Multiple protocols can be added. But, adding unnecessary protocols increases traffic on the media. If one needs to connect to a TCP/IP network or the Internet, the TCP/IP protocol should be added with the appropriate frame type. After adding all protocols needed, configure each protocol in the properties. For TCP/IP, for example, a static IP address, subnet mask, DNS address, and default gateway can be entered. Or configure it to obtain these parameters automatically at boot time from a DHCP server. Once the NIC and the protocols have been configured, click the add button and add a client. There are several choices such as Novell, Microsoft, Banyon, etc. Choose Microsoft. Allow file sharing and/or printer sharing. Click on **My Computer** and click on each drive, allow sharing of the drive, directory, or printer based on necessity. The guest may be required to login with a name and password and files could have full privileges or read only privileges. Sharing data on another computer could be accomplished by mapping that drive to a drive letter. In order to do that, in Windows XP, click on the **view workgroup computers** from the network connections screen. In Windows 98, click on **network neighborhood.** Highlight the drive and directory that needs sharing, and from the **File** pull down menu, **select map network drive,** and give a drive letter. Click the option buttons appropriately, for example, if the drive needs to shared upon boot, check that box.

Novel Network

Novell Netware no longer provides for peer-to-peer networking. The server should be setup first and then the clients. In case of Novell network, the IPX/SPX protocol should be added. At the time of Novell server installation these configurations are entered interactively. More recent Novell installations require TCP/IP configuration to get full functionality of the operating system, such as clustering, and tape backup. In case of a Novell client installation all configuration is done automatically. However, configuring for Novell using Windows configuration will require NIC, protocol, and client configurations. After server and client installations, login as the admin into the desired server under a tree and add users, groups, directories, and rights to files. Failsafe servers can be setup using redundant drives (RAID) and redundant servers (Novell clustering).

Printing on a network

Network ready printers come already installed with a NIC. Configure an IP address for the printer either using the printer console buttons or using software provided by the manufacturer. Under Windows there check the list of printer ports for the TCP/IP port. If it is not there add it giving the appropriate IP address of the network printer. Install the printer driver to complete the printer installation. Alternatively, print queues can be setup on a print server. A client can then print to a print queue.

Linux Network

The Linux will automatically detect the network card during installation and the device name “eth0” will be assigned to it. Using the ifconfig, assign the IP address and subnet mask to the interface. Next edit the /etc/hosts file and add the IP addresses and the corresponding symbolic name such as cs.panam.edu. To convert a symbolic name to the IP address try running the nslookup (example nslookup panam.edu). The nslookup command uses the /etc/resolv.conf file to find the host runs the name server software. The resolv.conf file needs to be edited to include the domain name server (DNS). For example, the file should include these lines:

```
domain cs.panam.edu
nameserver 129.113.132.237
search cs.panam.edu
```

Some examples of Internet services that can be added are: Telnet, FTP, Secure Shell, Name Server, Web Server, Mail Server, and Samba. The /etc/services file will have a list of installed services. The Internet services can be invoked by issuing the inetd command if it is not already running.

The Internet can be probed by using the ping command. For example to probe panam.edu issue the command: ping panam.edu. Another valuable diagnostic tool is the traceroute, which will list all computers along the path to the destination.

TCP/IP Protocol Suite

The TCP/IP protocol suite is divided into five layers, physical, network, Internet, transport and application. The peer layers on the communicating computers interact with each other. For example, the network layer on one node interacts with the corresponding network layer on the other node. All communication must take place through the physical layer at the bit level by generating electrical signals. Each layer adds its own required headers and trailers to communicate with the peer layer.

The physical layer is concerned with the cable, connector and electrical signal specifications. The network layer takes care of details about making frames, obtaining access to a shared medium of a local area network (LAN), and the physical addressing. The Internet layer specifies details about forming, and delivering packets from a source network to a destination network through routers. The transport layer can be either TCP or UDP. It is responsible for reliable transfer of messages from one process to another process. The application layer is concerned with providing environment for running the application programs such as email, FTP, etc.

The IP has gone through several revisions. The most recent version is the IPv6. This version rectifies all of the addressing problems that existed in IPv4 and provides for audio and video applications. It uses a 16 byte address written in 8 hexadecimal notations separated by colons. The version 5 was to implement the OSI model, which never materialized. It is important to note that the version still most commonly used is the IPv4.

IP Address, Subnetting and Supernetting

Computers and other devices connected to the Internet are identified by unique IP addresses, and these addresses are included in the source address and destination address fields of all IP packets. An IP address provides sufficient information to route a packet from the source to the destination network and deliver it to the appropriate node. The IP address uses a 32 bit binary number allowing for a total of 4,294,967,295 possible nodes to be divided among the anticipated number of local networks.

An IP address is 32 bits long and for ease of human understanding, it is written as four octets connected by dots, each octet ranging from 0 to 11111111 binary or 0 to 255 decimal. The IP address has two distinct parts: the network address (prefix) and the host address (suffix). The prefix portion of the address identifies the physical network to which a host is attached, while the suffix portion identifies an individual computer on that network Number of bits used for the network portion and host portion depends upon the class of the IP address. All zeros and all ones have special meanings and cannot be used for normal communication. Therefore, the number of networks and hosts each class (as described next) can have is total possible binary alternatives minus two ($2^{n\text{-bits}}$ -2).

There are five classes of IP addresses**,** named A through E. Class D is used for multicast addressing and class E is reserved for future use. Class A addresses begin with 1 and end with 127 in the first octet portion of the IP address. Defined another way, the most significant bit of the first octet of the IP address is always zero for all class A addresses. Class B addresses begin with 128 and end with 191; the two most significant bits of the first octet are always 10. Class C addresses begin with 192 and end with 223; the three most significant bits of the first octet are always 110. Classes D and E begin with 1110 and 1111 respectively. Class A uses 8 bits for network address and 24 bits for host address; class B uses 16 bits for network address and 16 bits of host address; and class C uses 24 bits of network address and 8 bits for host address. Figure__ summarizes the various classes and network and host portions of the total address space.

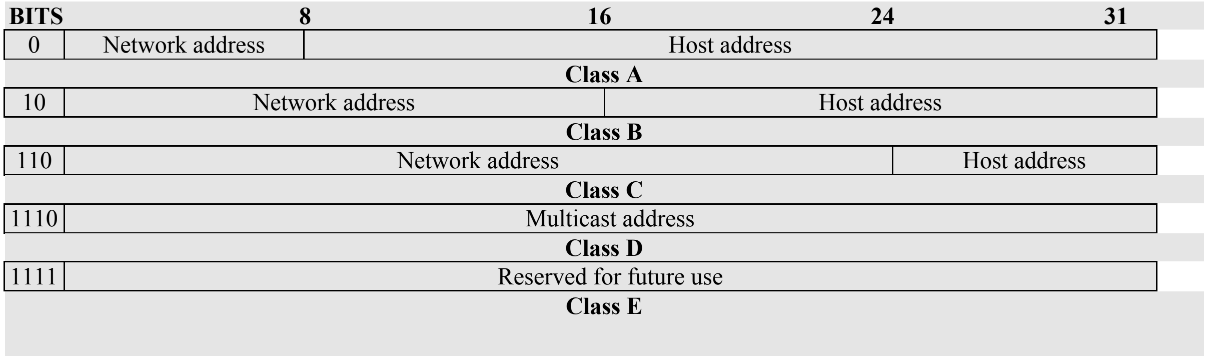

Figure
Classes of IP addresses and number of bits used for network and host address portions.

Class	**Available bits in network portion**	**Number of networks**	**Available bits in host portion**	**Number of hosts**
A	7	128	24	16,777,214
B	14	16,384	16	65534
C	21	2,097,152	8	254

Figure
Number of Networks and Hosts for each class of IP address

Out of the 8 bits allocated for the network address portion of the Class A IP address, the first bit should remain 0, and the remaining 7 bits can be used for assigning networks yielding a maximum of 127 Class A networks. Each network of Class A can have a maximum of 16,777,216 minus 2 hosts. All zeros and all ones have special meanings and may not be used for host addressing. All zeros in the host portion has the special meaning, "this computer or this network", and all ones are used for broadcasting a message to every host on a network. Figure __ reveals the number of networks and hosts each of the three Classes of IP address can have. Not all networks are used; for example, networks 0 and 127 of Class A are reserved giving a total of 126 usable Class A networks.

The network addresses are distributed by the Internet Assigned Numbers Authority (IANA) and the American Registry for Internet Numbers (ARIN). One Class A network with almost 17 million hosts would be extremely unmanageable. Even the 254 hosts available in a Class C would be difficult to manage. Furthermore, there are restrictions on the number of nodes a particular cable can have. A network can be divided into smaller more manageable networks using subnet masks.

Subnet masks determine if a destination address can be found within the local physical network or outside the local physical network. Routers are specialized computers that find paths to destination addresses. Routers connect to multiple physical networks and are called multi-homed hosts. When a router receives a packet from outside for one of the physical networks connected to it, it matches the IP address with the physical address of the host and sends the packet to that host. Each host has a physical address on its Ethernet card. When the router receives a packet from one of the hosts connected to its physical network, the router must determine if the destination can be found on one of its physical networks or it must send it out. Subnet mask provides necessary information to make this decision. Once a network address is assigned to an organization by IANA, that organization must decide how many of the host portion of the bits would be used for subnet masking. The remaining bits can be used for host addressing. For example, if the Class B IP address 129.113.0.0 is assigned to an organization, it has 16 bits of the host portion to work with. These 16 bits can be divided into two 8 bit portions, the first half for subnet addressing and the second for host addressing. Based on this decision, that organization can have 254 different networks each having 254 hosts. For this example, the subnet mask would be all binary ones for the first three octets and 0 for the last octet or decimal 255.255.255.0. Continuing with this example, let us suppose that host 129.113.200.111 sends a packet to destination 129.113.200.120. Does the destination host reside on the local physical network or outside? This determination is made by ANDing the destination address with the subnet mask. It is important to keep in mind that for this example, the source subnet address is **129.113.200.0.**

	10000001.01110001.11001000.01111000	destination address
AND	11111111.11111111.11111111.00000000	subnet mask
	10000001.01110001.11001000.00000000	destination on the same subnet
	129.113.200.0 – same as the source subnet address.	

Figure
ANDing IP Address and Subnet Mask on the same Subnet

	10000001.01110001.10111111.01101111	destination address
AND	11111111.11111111.11111111.00000000	subnet mask
	10000001.01110001.10111111.00000000	destination not on the same subnet
	129.113.191.0 – not same as the source subnet address.	

Figure
ANDing IP Address and Subnet Mask on Different Networks.

The ANDing in Figure 3 reveals that the destination address is on the same physical subnet as the source address. Changing the destination address to 129.113.191.121 in Figure 4 reveals that the destination address is not found on the same subnet and the packet must be routed to the appropriate network.

With the depletion of Class A and Class B addresses, it has become necessary for grouping Class C addresses together to make a supernet. In order to supernet the number of Class C address blocks used should be multiples of 2, the network address be contiguous, and the third byte of the first address be evenly divisible by the number of blocks. Once these conditions are satisfied a supernet mask can be calculated.

Delivery of a Packet

In order to send packets between two physical networks, a router, a bridge, or a switch must be used, the most common one being the router. A router keeps a routing table to look up which direction the packet must be sent or to determine if the packet belongs to self. The routing table is constantly updated based on if the target is reachable or unreachable, shortest path to the destination, traffic congestion, etc. The subnets are only visible within a network and not to the outside world. In the above example, ANDing the subnet mask with destination address gets rid of the host portion of the address. The remaining portion, the network address portion, can be looked up in the routing table. A router may have several cable segments connected to it (multi-homed), each having its own network address and subnet mask. Based on the result of the above described ANDing, the router makes the decision whether the destination is on the same physical cable as the source, or on another segment connected to it, or outside its immediate reach. If the destination is outside, based on the routing table, it must send the packet to the appropriate segment to which it is connected. When a router is going to forward a packet, it must determine whether it can send it directly to its destination, or whether it needs to pass it through another router. If the latter, it needs to determine which router to use. If the next hop is not known, a request may be sent to the next hop resolution protocol (NHRP) server for the next hop resolution.

Once a packet arrives at the router to which the destination host is connected, it must be delivered to that host. Within the local area network all communication takes place using the physical address (also known as the MAC address). However, the packet to be delivered contains the destination IP address, not the physical address. A host communicating to another within a LAN will build a table in the memory containing IP addresses and corresponding MAC addresses. These entries are gradually built using the Address Resolution Protocol (ARP) broadcasts, and the table itself is called the ARP cache. If an entry is not found in the ARP cache, it will issue an ARP broadcast, essentially requesting the one host that holds the particular IP address to respond with its MAC address. All listening hosts now can update their cache for future communication. Once the MAC address is obtained a frame is constructed using that MAC address and sent to the destination host. With time some entries in the ARP cache will become stale and would be deleted.

DNS, DHCP and Proxy Servers

If an IP address cannot be resolved from the hosts file, a Domain Name Server (DNS) will be requested to resolve the address. There can be a primary and a secondary name server. The DNS provides for mapping names to addresses and mapping addresses to names. The large amount of information about the symbolic names and IP addresses is divided into smaller parts and stored in different name servers throughout the world. A full domain name is a sequence of labels separated by dots, starting with the node up to the root as in cs.panam.edu. The period after the edu is the root and can be omitted. After the root level, com, edu, gov, int, mil, net, org, aero, biz, coop, info, pro, etc. generic domains are implemented. Alternatively, a two letter country domain name may be used.

The Dynamic Host Configuration Protocol (DHCP) server is very useful when numerous hosts need to be configured with IP, subnet mask, and default gateway. The DHCP can be set up to issue these addresses when a device is turned on. The range of addresses given out and lease time can be configured on the DHCP server. The DHCP provides a temporary IP address for a limited period of time. This also allows for a node to be moved from one network to another without reconfiguration.

A Proxy server allows clients to access the Internet from behind a firewall. A firewall is a security system implemented to protect a local area network from intruders. Furthermore, Proxy servers can cache web pages retrieved and can serve these pages when requested by the same user or other users using the Proxy server. Caching documents reduces external traffic.

Network Security

A network needs to be secure enough to block intruders while allowing legitimate users to perform required tasks without the hassle of security. A user attempting to gain access to a network should be authenticated to be a legitimate user rather than someone masquerading. Once a user is allowed access to the network, that individual should only be able to perform tasks for which he/she has rights. The data transmitted between the network and the user should be kept confidential so that eavesdroppers will not be able to make sense of what is being transmitted. Precautions must be taken so that data transmitted will not be tampered with or altered in transit. Finally, a receiver of information should not be able to deny that the information was received.