5 Network Connectivity

5.1. Introduction

All but the most basic of networks require devices to provide connectivity and functionality. Understanding how these networking devices operate and identifying the functions they perform are essential skills for any network administrator and requirements for a Network + candidate.

This chapter introduces commonly used networking devices and, although it is true that you are not likely to encounter all of the devices mentioned in this chapter on the exam, you can be assured of working with at least some of them.

5.2. Network Connectivity Devices

Network devices are components used to connect computers or other electronic devices together so that they can share files or resources like printers or fax machines. Devices used to setup a Local Area Network (LAN) are the most common types of network devices used by the public. A LAN requires a hub, router, cabling or radio technology, network cards, and if online access is desired, a high-speed modem. This is much less complicated than it might sound to someone new to networking. In the smooth functioning of networks, many devices play important roles. Here, in this section, we are going to discuss a few of them.

Let us begin with NICs.

5.2.1. Network Interface Unit

The network interface unit is a microprocessor based device containing hardware and software which supply the intelligence to control access to and communications across the network and to perform all communications processing. It is the means by which the workstations are connected functionally and physically to the network.

On most microcomputer networks, the network interface is a printed circuit board installed in the microcomputer. Depending on the vendor, it may be called a network card, network adapter, or network interface unit.

On some networks, the network interface may be implemented as a stand alone box, termed a wiring centre, or hub, attached between the main network cable and the workstation.

An assumption is made that networking ability is not already present in the microcomputer, but has to be added. Microcomputers with built-in communications and networking will not need added hardware; the functions of the network interface already will be present in the machine. The crucial factor is not where or how the interface is located, but what functions it serves.

Network interface functions are realized through chips on the interface unit : Network bus drivers, communications controller chips, specialized microprocessors, RAM buffers and ROM code that is executed by the workstation itself. For most LANs, the network interface unit for all user workstations is identical. Server interface units may include additional ROM code to implement additional functions. Physical connection to the network is provided through a standard communications or input/output interface.

Through the network interface, data on the medium is available to all attached workstations and peripherals. System users never need to know what it takes to get from one point to another; they simply indicate the desired destination. The network interface unit provides transmission and data control, formats the data into manageable units, translates the data rate and protocols of the attached workstation to that of the network communication medium and vice-versa and supplies address recognition capabilities. Details of network operation are hidden from users of the attached workstations.

Technically, two parts of the network interface can be identified : the communication interface, containing network oriented functions, and the host interface, containing computer specific functions. Both parts of the interface are shown in the following figure.

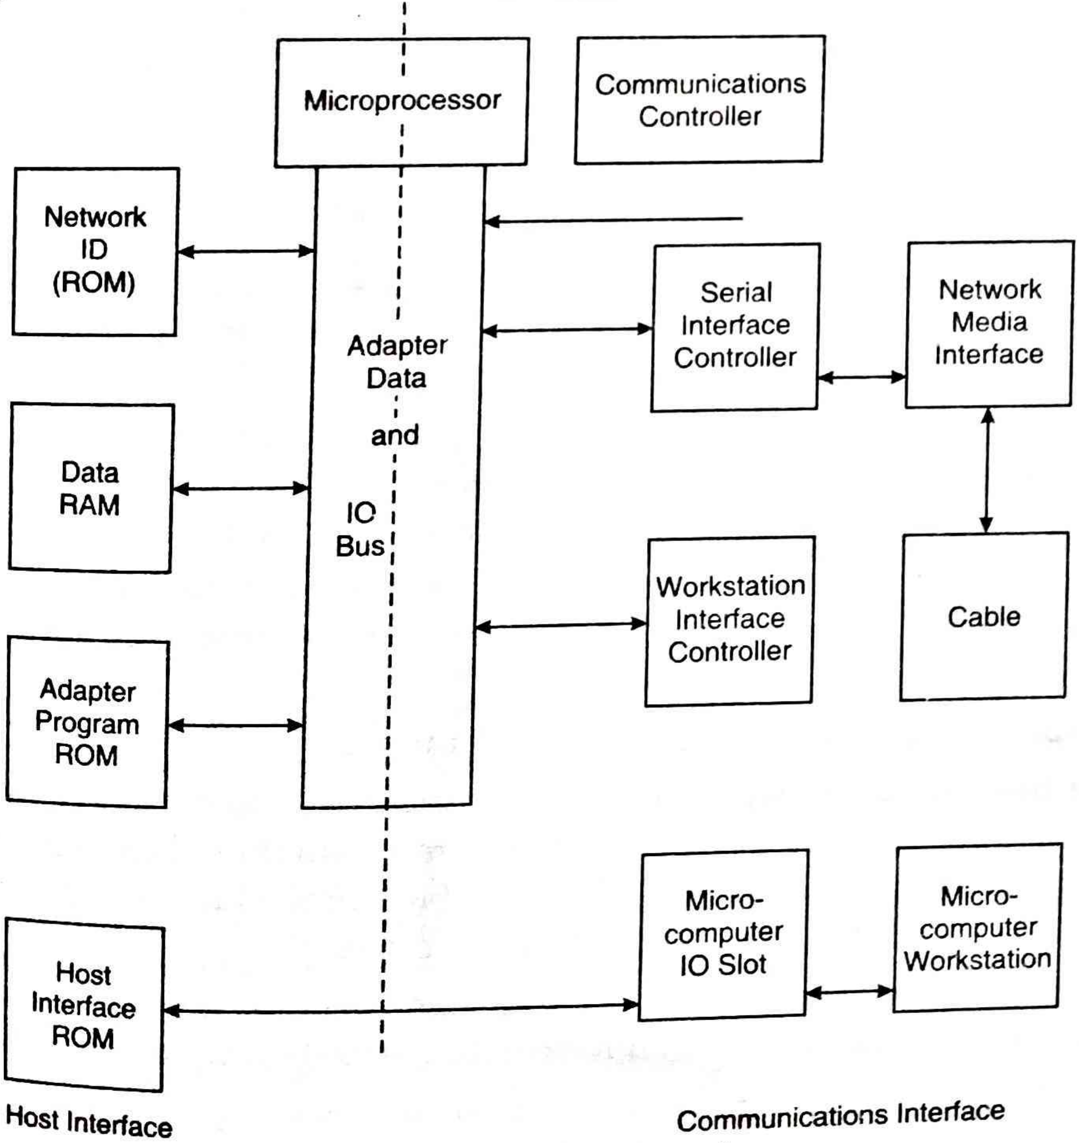

Fig. 5.1. Parts of a Network Interface

The communications interface is the unit which logically interfaces to the network. It performs all transmission related functions. It accepts data from the attached workstation, buffers the data until the communication channel is avilable, and then transmits the data. The communications interface also monitors the channel for messages addressed to its workstation, stores the data, and transfers the data to the device.

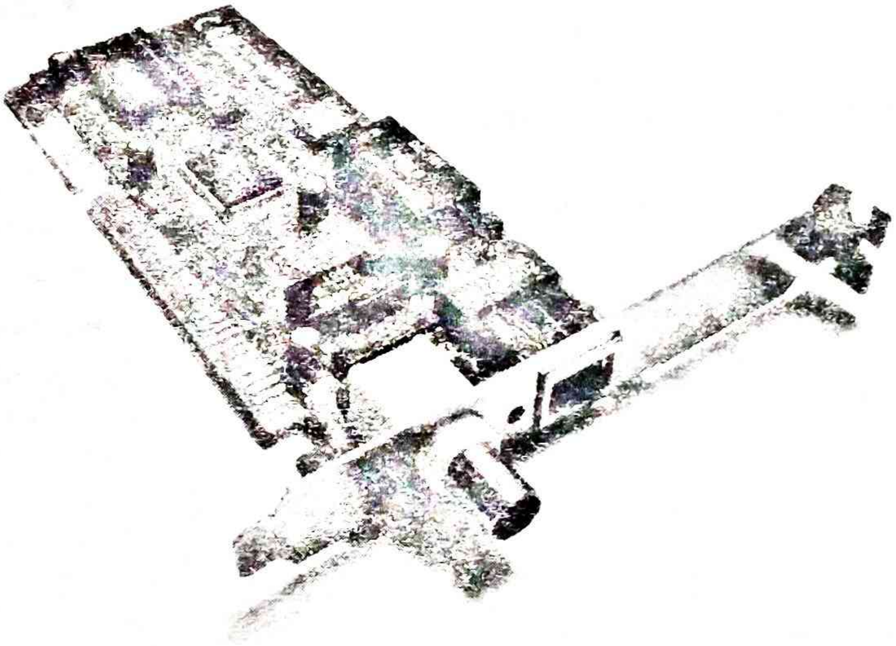

Fig. 5.2. Network Interface Card

The actual physical connection between the workstation and the network is achieved by running a secondary cable between the communications interface and the main network cable. The two cables are joined by a tap.

The host interface supplies the connection between a specific workstations's internal circuitry and the communications interface unit. It fits into the input/output structure of a particular computer, and governs all data exchange between the workstation and the communications-oriented portion of the network interface. Because of the many methods of implementing network and workstation functions, the host interface is workstation and vendor-specific.

When specifying or installing a NIC, you must consider the following issues :

System bus compatibility : If the network interface you are installing is an internal device, bus compatibility must be verified. The most common bus system in use is the Peripheral Component Interconnect (PCI) bus, but some older systems might still use Industry Standard Architecture (ISA) expansion cards.

System resources : Network cards, like other devices, need IRQ and memory I/O addresses. If the network card does not operate correctly after installation, there might be a device conflict.

Media compatibility : Today, the assumption is that networks use twisted-pair cabling, so if you need a card for coaxial or fiber-optic connections, you must specify this. Wireless network cards are also available.

Even more than the assumption you are using twisted-pair cabling is that the networking system being used is Ethernet. If you require a card for another networking system such as Token Ring, this must be specified when you order.

To install or configure a network interface, you will need drivers of the device, and might need to configure it, although many devices are now plug and play. Most network cards are now software configured. Many of these software configuration utilities also include testing capabilities. The drivers and software configuration utilities supplied with the cards are often not the latest available, so it is best practice to log on to the Internet and download the latest drivers and associated software.

5.2.2. Hubs

At the bottom of the networking food chain, so to speak, are hubs. Hubs are used in networks that use twisted-pair cabling to connect devices. Hubs can also be joined together to create larger networks. Hubs are simple devices that direct data packets to all devices connected to the hub, regardless of whether the data package is destined for the device. This makes them inefficient devices can create a performance bottleneck on busy networks.

In its most basic form, a hub does nothing except provide a pathway for the electrical signals to travel along. Such a device is called a passive hub. For more common now-a-days is an active hub, which, as well as providing a path for the data signals, regenerates the signal before it forwards it to all of the connected devices. A hub does not perform any processing on the data that it forwards, nor does it perform any error checking.

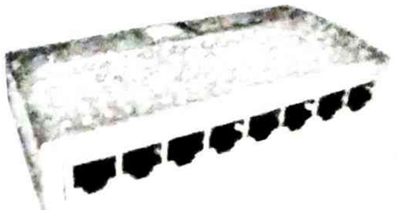

Fig. 5.3. Hubs

Hubs come in a variety of shapes and sizes. Small hubs with five or eight connection ports are commonly referred to as workgroup hubs. Others can accommodate larger numbers of devices (normally up to 32). These are referred to as high-density devices. Because hubs don't perform any processing, they do little except enable communication between connected devices. For today's high-demand network applications, something with a little more intelligence is required. That's where switches come in.

A hub is a hardware device used to connect several computers togather. A hub that contains multiple independent but connected modules of network and inter-networked equipment. A similar term is **concentrator.** A concentrator is a device that provides a central connection point for cables from workstations, servers, and peripherals. In a star topology, twisted-pair wire is run from each workstation to a central concentrator.

Basically, hubs are multi-slot concentrators into which a number of multi-port cards can be plugged to provide additional access as the network grows in size. Hubs can be either passive or active.

Active hubs electrically amplify the signal as it moves from one connected device to another. Active concentrators are used like repeaters to extend the length of a network.

Passive hubs allow the signal to pass from one computer to another without any change. Hubs usually can support 8, 12, or 24 RJ-45 ports. These are often used in a star or star-wired ring topology and require specialized software for port management.

How a Hub Functions : Hubs are simple devices that interconnect groups of users. Hubs forward any data packets including e-mail, word processing documents, spreadsheets, graphics, print requests-they receive over one port from one workstation to all of their remaining ports. All users connected to a single hub or stack of connected hubs are in the same **segment**, sharing the hub's bandwidth or data-carrying capacity. As more users are added to a segment, they complete for a finite amount of bandwidth devoted to that segment.

To understand how a hub serves a network, imagine a hotel with just one phone line available to all guests. Let's say one guest wants to call another. She picks up her phone and the phone rings in all rooms. All the other guests have to answer the phone and determine whether or not the call is intended for them. Then, as long as the conversation lasts, no one else can use the line. With a few guests, this system is marginally acceptable. However, at peak times of the day –say, when everyone returns to his or her room in the evening–it becomes difficult to communicate. The phone line is always busy.

5.2.3. Switches

Like hubs, switches are the connectivity points of an Ethernet network. Devices connect to switches via twisted-pair cabling, one cable for each device. The difference between hubs and switches is in how the devices deal with the data that they receive. Whereas a hub forwards the data it receives to all of the ports on the device, a switch forwards it only to the port that connects to the destination device. It does this by learning the MAC address of the devices attached to it, and then by matching the destination MAC address in the data it receives. Figure 5.4 shows how a switch works

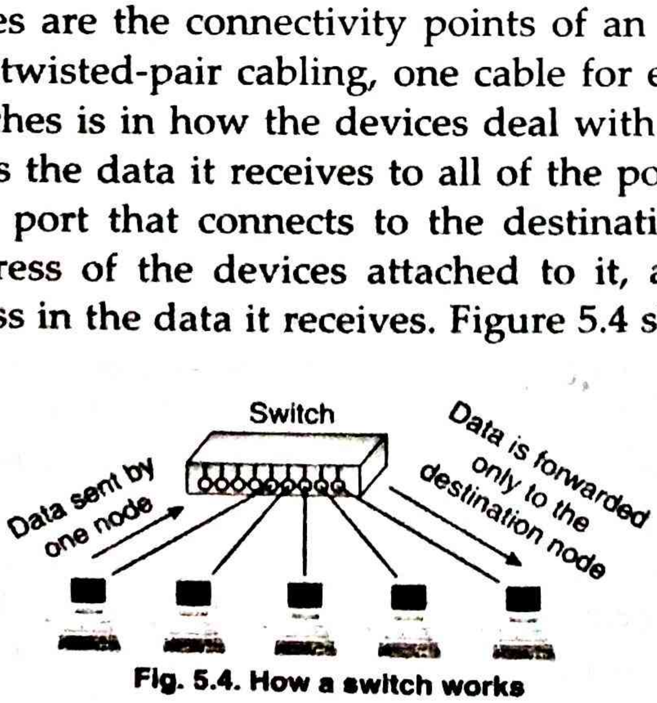

Fig. 5.4. How a switch works

By forwarding data only to the connection that should receive it, the switch can improve network performance in two ways. First, by creating a direct path between two devices and controlling their communication, it can greatly reduce the number of collisions on the network. As you might recall, collisions occur on Ethernet networks when two devices attempt to transmit at exactly the same time. In addition, the lack of collisions enables switches to communicate with devices in full-duplex mode. In a full-duplex configuration, devices can send and receive data from the switch at the same time. Contrast this with half-duplex communication, in which communication can occur in only one direction at a time. Full-duplex transmission speeds are double that of a standard, half-duplex, connection So, a 10 Mbps connection becomes 20 Mbps, and a 100 Mbps connection becomes 200 Mbps.

The net result of these measures is that switches can offer significant performance improvements over hub-based networks, particularly when network use is high.

Irrespective of whether a connection is at full or half duplex, the method of switching dictates how the switch deals with the data it receives. The following is a brief explanation of each method :

Cut-through : In a cut-through switching environment, the packet begins to be forwarded as soon as it is received. This method is very fast, but creates the possibility of errors being propagated through the network, as there is no error checking.

Store-and-forward : Unlike cut-through, in a store-and-forward switching environment, the entire packet is received and error checked before being forwarded. The upside of this method is that errors are not propagated through the network. The downside is that the error checking process takes a relatively long time, and store-and-forward switching is considerably slower as a result.

Fragment free : To take advantage of the error checking of store-and-forward switching, but still offer performance levels nearing that of cut through switching. Fragment Free switching can be used in a Fragment Free-switching environment, enough of the packet is read so that the switch can determine whether the packet has been involved in a collision. As soon as the collision status has been determined, the packet is forwarded.

Hub and Switch Cabling

In addition to acting as a connection point for network devices, hubs and switches can also be connected to create larger networks. This connection can be achieved through standard ports with a special cable or by using special ports with a standard cable.

The ports on a hub to which computer systems are attached are called Medium Dependent Interface-Crossed (MDI-X). The crossed designation is derived from the fact that two of the wires within the connection are crossed so that the send signal wire on one device becomes the receive signal of the other. Because the ports are crossed internally, a standard or straight-through cable can be used to connect devices.

Another type of port, called a Medium Dependent Interface (MDI) port, is often included on a hub or switch to facilitate the connection of two switches or hubs. Because the hubs or switches are designed to see each other as simply an extension of the network, there is no need for the signal to be crossed. If a hub or switch does not have an MDI port, hubs or switches can be connected by using a crossover cable between two MDI-X ports. The crossover cable serves to uncross the internal crossing. You can see diagrams of the cable pinouts for both a straight-through and crossover cable in Figures 5.5 and 5.6 respectively.

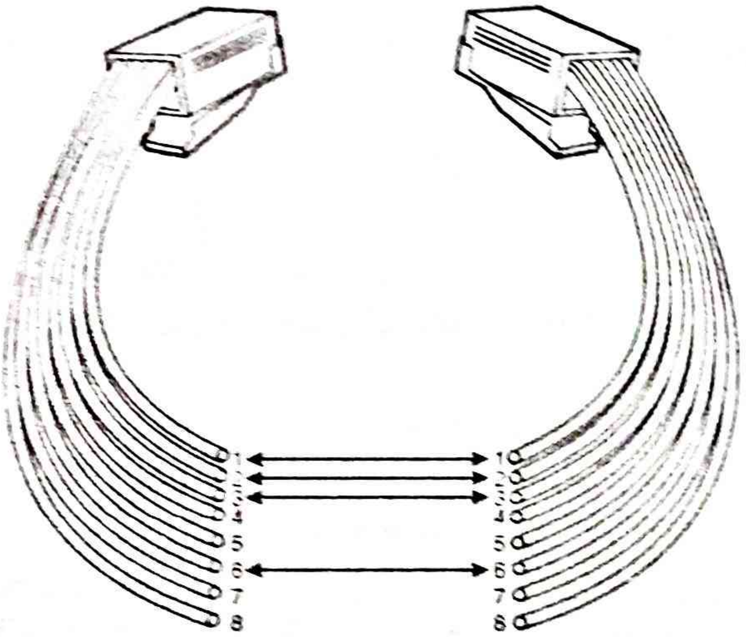

Fig. 5.5. The pinouts for a straight-through cable

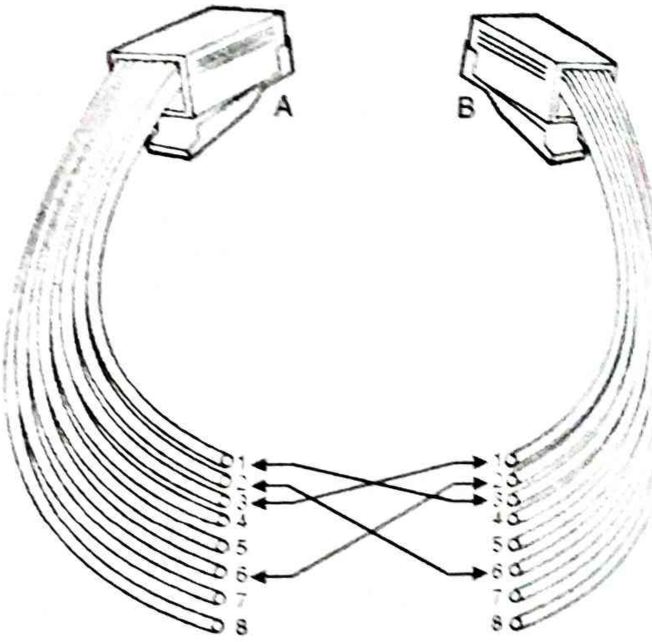

Fig. 5.6. The pinouts for a crossover cable

How a Switch Functions

To insulate the transmission from the other ports, the switch establishes a temporary connection between the source and destination, and then terminates the connection once the conversation is done.

A switch would be like a phone system with private lines in place of the hub's party line. For instance, Meira Sen at the Maurya Hotel calls Ibrahim Soz in another room, and the operator or phone switch connects the two of them on a dedicated line. This allows more conversations at any one time thereby allowing more guests to communicate.

5.2.4. Repeater

A repeater is a device that amplifies a signal being transmitted on the network. It is used in long network lines, which exceed the maximum rated distance for a single run.

Over distance, the cables connecting a network lose the signal transmitted. If the signal degrades too much, it fails to reach the destination. Or if it does arrive, the degradation of the message makes it useless. Repeaters can be installed along the way to ensure that data packets reach their destination. Repeaters are of two kinds; amplifie and signal repeater. The first merely amplifies all incoming signals over the network. However, it amplifies both the signal and any concurrent noise. The second type collects the inbound packet and then retransmits the packet as if it were starting from the source station.

5.2.5. Gateway

Any device that translates one data format to another is called a gateway. Some examples of gateways include a router the translates data from one network protocol to another, a bridge that converts between two networking systems, and a software application that converts between two dissimilar formats. The key point about a gateway is that only the data format is translated, not the data itself. In many cases, the gateway functionality is incorporated into another device.

A gateway is a device that connects dissimilar networks. A gateway operates at the highest layer of network abstraction. It expands the functionality of routers by performing data translation and protocol conversion. It is needed to convert Ethernet traffic from the LAN, to SNA (Systems Network Architecture) traffic on a legacy system. It then routes the SNA traffic to the mainframe. When the mainframe answers, the reverse process occurs.

A gateway is actually a node on a network that serves as an entrance to another network. In enterprises, the gateway is the computer that routes the traffic from a workstation to the outside network that is serving the Web pages. In homes, the gateway is the ISP that connects the user to the Internet.

In enterprises, the gateway node often acts as a proxy server (a machine that is not actually a server but appears as a server) and a firewall (a system designed to prevent unauthorized access to or form a private network). The gateway is also associated with both a router, which use headers and forwarding tables to determine where packets are sent, and a switch, which provides the actual path for the packet in and out of the gateway.

5.2.6. Bridges

Bridges are used to divide larger networks into smaller section. They do this by sitting between two physical network segments and managing the flow of data between the two. By looking at the MAC address of the devices oonnected to each segment, bridges can elect to forward the data (if they believe that the destination address is on another interface), or block it from crossing (if they can verify that it is on the interface from which it came). Figure 5.7 shows how a bridge can be used to segregate a network.

When bridges were introduced, the MAC addresses of the devices on the connected networks had to be entered manually, a time-consuming process that had plenty of opportunity for error. Today, almost all bridges can build a list of the MAC addresses on an interface by watching the traffic on the network. Such devices are called learning bridges because of this functionality.

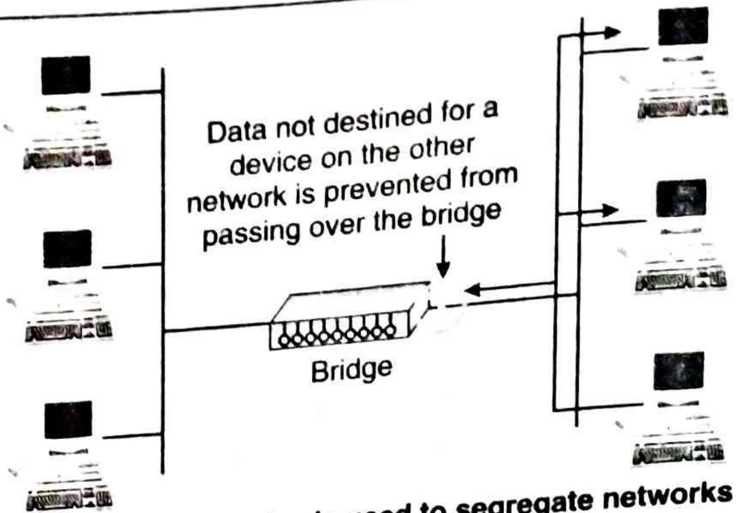
Data not destined for a device on the other network is prevented from passing over the bridge
Bridge

Fig. 5.7. How a bridge is used to segregate networks

Bridge Placement and Bridging Loops

There are two issues that you must consider when using bridges. The first is the bridge placement, and the other is the elimination of bridging loops :

Placement : Bridges should be positional in the network using the 80/20 rule. This rule dictates that 80% of the data should be local and that the other 20% should be destined for devices on the other side of the bridge.

Bridging loops : Bridging loops can occur when more than one bridge is implemented on the network. In this scenario, the bridges can confuse each other by leading one another to believe that a device is located on a certain segment when it is not. To combat the bridging loop problem, the IEEE 802.1d Spanning Tree protocol enables bridge interfaces to be assigned a value that is then used to control the bridge-learning process.

Types of Bridges

Three types of bridges are used in network :

Transparent bridge : Derives its name from the fact that the devices on the network are unaware of its existence. A transparent bridge does nothing except block or forward data based on the MAC address.

Source route bridge : Used in Token Ring networks. The source route bridge derives its name from the fact that the entire path that the packet is to take through the network is embedded within the packet.

Translational bridge : Used to convert one networking data format to another; for example, from Token Ring to Ethernet and vice-versa.

Today, bridges are slowly but surely falling out of favor. Ethernet switches offer similar functionality ; they can provide logical divisions, or segments, in the network. In fact, switches are sometimes referred to as multiport bridges because of the way they operate.

5.2.7. Modems

A modem, short for modulator/demodulator, is a device that converts the digital signals generated by a computer into analog signals that can travel over conventional phone lines. The modem at the receiving end converts the signal back into a format the computer can understand. Modems can be used as a means to connect to an ISP or as a mechanism for dialing up to a LAN.

Modems can be internal add-in expansion cards, external devices that connect to the serial or USB port of a system, PCMCIA cards designed for use in laptops, or proprietary devices designed for use on other devices such as portables and handhelds.

The configuration of a modem depends on whether it is an internal or external device. For internal devices, the modem must be configured with an interrupt request (IRQ) and a memory I/O address. It is common practice, when installing and internal modem, to disable the built-in serial interfaces and assign the modem the resources of one of those (typically COM2). Table 5.1 shows the resources associated with serial (COM) port assignments.

Table 5.1 Common Serial (COM) Port Resource Assignments

Port ID	IRQ	I/0 Address	Associated Serial I/F Number
COM1	4	03F8	1
COM2	3	02F8	2
COM3	4	03E8	1
COM4	3	02E8	2

For external modems, you need not concern yourself directly with these port assignments, as the modem connects to the serial port and uses the resources assigned to it. This is a much more straight forward approach and one favored by those who work with modems on a regular basis. for PCMCIA and USB modems, the plug-and-play nature of these devices makes them simple to configure, and no manual resource assignment is required. Once the modem is installed and recognized by the system, drivers must be configured to enable use of the device.

Two factors directly affect the speed of the modem connection—the speed of the modem itself and the speed of the Universal Asynchronous Receiver/Transmitter (UART) chip in the computer that is connected to the modem. The UART chip controls the serial communication of a computer, and although modern systems have UART chips that can accommodate far greater speeds than the modem is capable of, older systems should be checked to make sure that the UART chip is of sufficient speed to support the modem speed. The UART chip installed in the system can normally be determined by looking at the documentation that comes with the system.

Table 5.2 Shows the maximum speed of the commonly used UART chip types.

Table 5.2 UART Chip Speeds

UART Chip	Speed (Kbps)
8250	9600
16450	9600
16550	115,200
16650	430,800
16750	921,600
16950	921,600

5.2.8. Router

A Router is a network device that connects multiple networks irrespective of their protocol. This is because a router can handle different protocols, otherwise it works similar to that of a bridge.

A Router forwards data packets from one connected network to another depending upon their IP addresses and not their MAC addresses.

So, now you know two differences of bridge and router :

(i) A bridge connot handle multiple protocols whereas a router can.

(ii) A bridge works with MAC address whereas a router works with IP addresses.

Now that you have fair idea about what constitutes a network, let us now discuss how communication takes place over a network. This is being given in following information box.

In a common configuration, routers are used to create larger networks by joining two network segments. Such as a SOHO router used to connect a user to the Internet. A router can be a dedicated hardware device or a computer system with more than one network interface and the appropriate routing software. All modern network operating systems include the functionality to act as a router.

A router derives its name from the fact that it can route data it receives from one network onto another. When a router receives a packet of data, it reads the header of the packet to determine the destination address. Once it has determined the address, it looks in its routing table to determine whether it knows how to reach the destination and, if it does, it forwards the packet to the next hop on the route. The next hop might be the final destination, or it might be another router. Figure 5.8 shows, in basic terms, how a router works.

As you can see from this example, routing tables play a very important role in the routing process. They are the means by which the router makes its decisions. For this reason, a routing table needs to be two things. It must be up-to-date, and it must be complete. There are two ways that the router can get the information for the routing table—through static routing or dynamic routing.

Static Routing

In environments that use static routing, routes and route information are entered into the routing tables manually. Not only can this be a time-consuming task, but also errors are more common. Additionally, when there is a change in the layout, or topology, of the network, statically configured routers must be manually updated with the changes. Again, this is a time consuming and potentially error-laden task. For these reasons, static routing is suited to only the smallest environments with perhaps just one or two routers. A far more practical solution, particularly in larger environments, is to use dynamic routing.

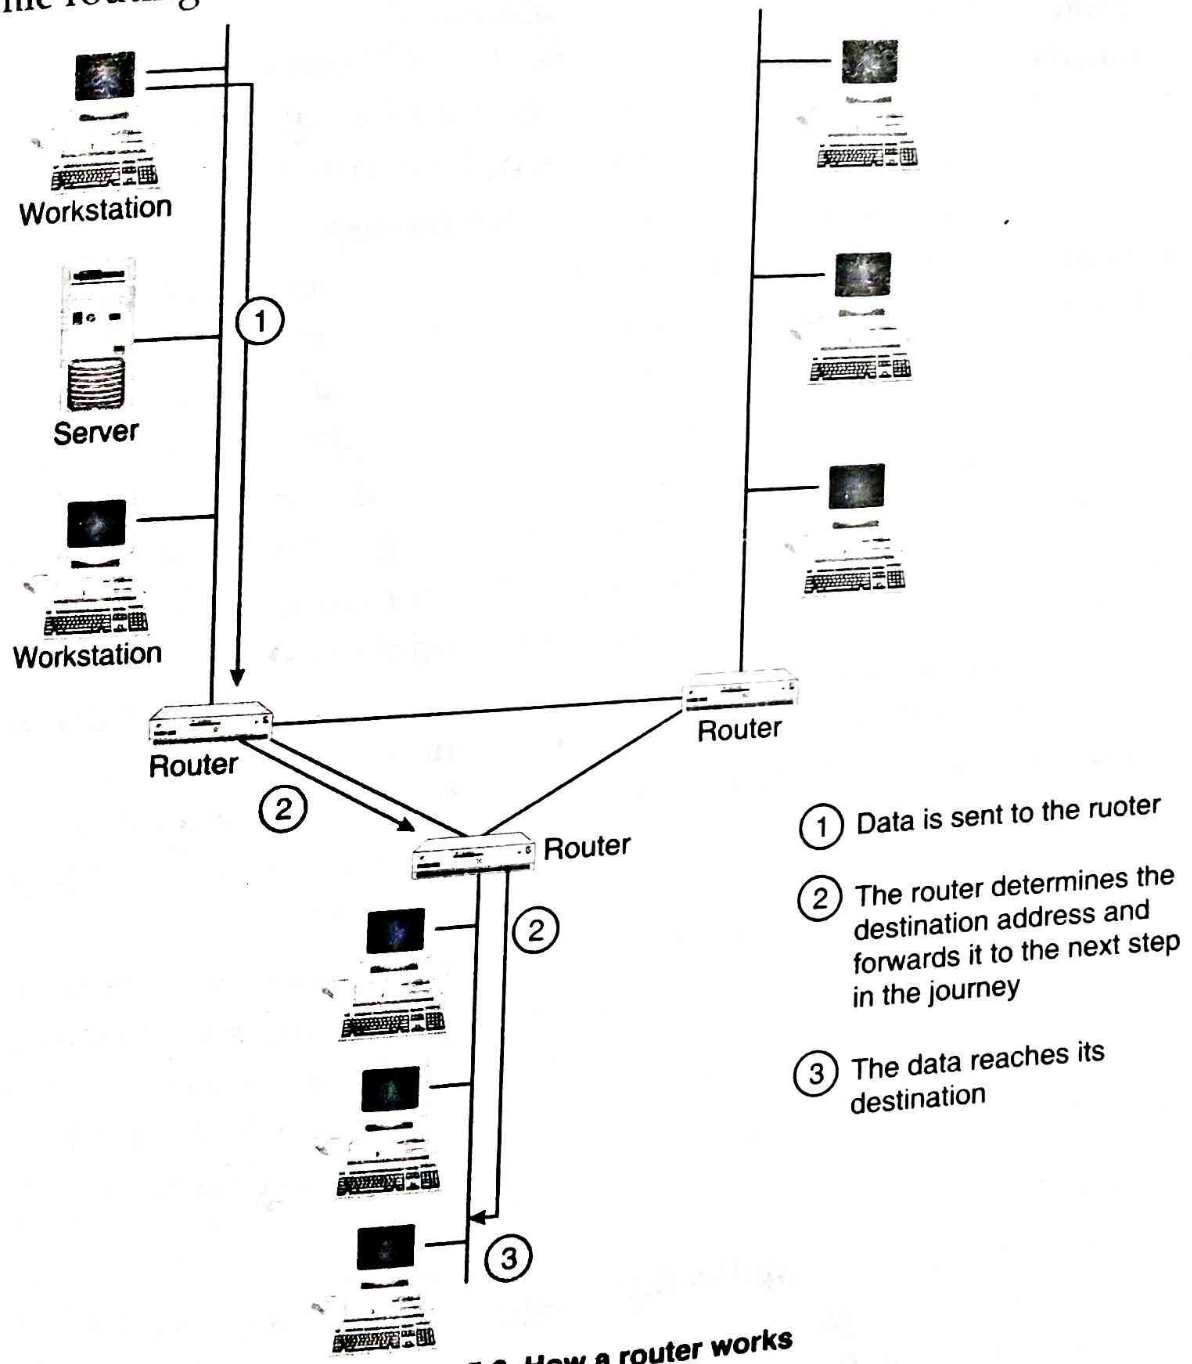
Workstation
Server
Workstation
1
Router
Router
2
Router
2
3
1 Data is sent to the ruoter
2 The router determines the destination address and forwards it to the next step in the journey
3 The data reaches its destination

Fig. 5.8. How a router works

Dynamic Routing

In a dynamic routing environment, routers use special routing protocols to communicate. The purpose of these protocols is simple; they enable routers to pass on information about themselves to other routers so that other routers can build routing tables. There are two types of routing protocols used—the older distance vector protocols and the newer link state protocols.

Distance Vector Routing

The two most commonly used distance vector routing protocols are both called Routing Information Protocol (RIP). One version is used on networks running TCP/IP. The other, sometimes referred to as IPX RIP, is designed for use on networks running the IPX/SPX protocol.

RIP works on the basis of hop counts. A hop is defined as one step on the journey to the data's destination. Each router that the data has to cross to reach its destination constitutes a hop. The maximum number of hops that RIP can accommodate is 15. That is to say that in a network that uses RIP, all routers must be within 15 hops of each other to communicate. Any hop count that is in excess of 15 is considered unreachable.

Distance vector routing protocols operate by having each router send updates about all the other routrers it knows about to the routers directly connected to it. These updates are used by the routers to compile their routing tables. The updates are sent out automatically every 30 or 60 seconds. The actual interval depends on the routing protocol being used. Apart from the periodic updates, routers can also be configured to send a triggered update if a change in the network topology is detected. The process by which routers learn of a change in the network topology is knwon as convergence.

Although distance vector protocols are capable of maintaining routing tables, they have three problesm. The first is that the periodic update system can make the update process very slow. The second problem is that the periodic updates can create large amounts of network traffic—much of the time unnecessarily as the topology of the network should rarely change. The last, and perhaps more significant, problem is that because the routers only know about the next hop in the journey, incorrect information can be propagated between routers, creating routing loops.

Two strategies are used to combat this last problem. One, split horizon, works by preventing the router from advertising a route back to the other router from which it was learned. The other, poison reverse (also called split horizon with poison reverse), dictates that the route is advertised back on the interface from which it was learned, but that it has a metric of 16. Recall that a metric of 16 is considered an unreachable destination.

Link Stae Routing

Link state routing works quite differently from distance vector-based routing. Rather than each router telling each other connected router about the routes it is aware of, routers in a link state environment send out special pakcets, called link state advertisements (LSA), which contain information only about that router. These LSAs are forwarded to all the routers on the network, which enables them to build a map of the entire network. The advertisements are sent when the router is first brought onto the network and when a change in the topology is detected.

Of the two (distance vector and link state), distance vector routing is better suited to small networks and link state routing to larger ones. Link state protocols do not suffer from the constant updates and limited hop count, and they are also quicker to correct themselves (to converge) when the network topology changes.

On TCP/IP networks, the most commonly used link state routing protocol is the Open Shortest Path First (OSPF). On IPX networks, the NetWare Link State Protocol (NLSP) is used. Table 5.3 summarizes the distance vector and link state protocols used with each network protocol.

Table 5.3 Routing Protocols

Network Protocol	Distance Vector	Link State
TCP/IP	RIP	OSPF
IPX/SPX	RIP*	NLSP

Routing Protocol

A routing protocol is the language a router speaks with other routers in order to share information about the reach ability and status of networks. We can say that a routing protocol specifies how routers communicate with each other, disseminating information that enables them to select routes between any two nodes on a computer network. Routing algorithms determine the specific choice of route. Each router has a prior knowledge only of networks attached to it directly. A routing protocol shares this information first among immediate neighbors, and then throughout the network. This way, routers gain knowledge of the topology of the network.

Although there are many types of routing protocols. three major classes are in widespread use on IP networks :

- Interior gateway protocols type 1, link-state routing protocols, such as OSPF and IS-IS
- Interior gateway protocols type 2, distance-vector routing protocols, such as Routing Information Protocol, RIPv2, IGRP
- Exterior gateway protocols are routing protocols used on the internet for exchanging routing information between Autonomous systems, such as Border Gateway Protocol (BGP), Path Vector Routing Protocol.

Note that the term "Exterior gateway protocol" has two meanings. It could mean a category of protocols used to exchange routing information between autonomous systems (see: exterior gateway protocol). It could also mean a specific RFC-described Protocol (See: Exterior Gateway Protocol).

Many routing protocols are defined in documents called RFCs.

Some versions of the Open System Interconnection (OSI) networking model distinguish routing protocols in a special sublayer of the Network Layer (Layer 3).

The specific characteristics of routing protocols include the manner in which they avoid routing loops, the manner in which they select preferred routes, using information about hop costs, the time they require to reach routing convergence, their scalability, and other factors.

Link State Routing Protocol

A link-state routing protocol is one of the two main classes of routing protocols used in packet switching networks for computer communications (the other is the distance-vector routing protocol). Examples of link-state routing protocols include open shortest path first (OSPF) and intermediate system to intermediate system (IS-IS).

The link-state protocol is performed by every switching node in the network (i.e. nodes that are prepared to forward packets; in the Internet, these are called routers). The basic concept of link-state routing is that every node constructs a map of the connectivity to the network, in the form of a graph, showing which nodes are connected to which other nodes. Each node then independently calculates the next best logical path from it to every possible destination in the network. The collection of best paths will then form the node's routing table.

This contrasts with distance-vector routing protocols, which work by having each node share its routing table with its neighbours. In a link-state protocol the only information passed between nodes is connectivity related.

Link-state algorithms are sometimes characterized informally as each router "telling the word about its neighbours".

Distance Vector Routing Protocol

In computer communication theory relating to packet-switched networks, a distance-vector routing protocol is one of the two major classes of routing protocols, the other major class being the link state protocol. Distance-vector routing protocols use the Bellman—Ford algorithm, Ford—Fulkerson algorithm, or DUAL FSM (in the case of Cisco Systems's protocols) to calculate paths.

A distance-vector routing protocol requires that a router informs its neighbors of topology changes periodically. Compared to link-state protocols, which require a router to inform all the nodes in a network of topology changes, distance-vector routing protocols have less computational complexity and message overhead.

The term distance vector refers to the fact that the protocol manipulates vectors (arrays) of distances to other nodes in the network. The vector distance algorithm was the original ARPANET routing algorithm and was also used in the internet under the name of RIP (Routing Information Protocol).

Examples of distance-vector routing protocols include RIPv1 and RIPv2 and IGRP.

Interior Gateway Protocols

Interior gateway protocols (IGPs) exchange routing information within a single routing domain.

Examples of IGPs include :

- Open Shortest Path First (OSPF)
- Routing Information Protocol (RIP)
- Intermediate System to Intermediate System (IS-IS)
- Enhanced Interior Gateway Routing Protocol (EIGRP)
- Exterior gateway protocols
- Exterior gateway protocols exchange routing information between autonomous systems. Examples include:
- Exterior Gateway Protocols (EGP)
- Border Gateway Protocol (BGP)

5.3. Firewall

A firewall is a networking device, either hardware or software based, that controls access to your organization's network. This controlled access is designed to protect data and resources from an outside threat. To do this, firewalls are typically placed at entry/exit points of a network—for example, placing a firewall between an internal network and the Internet. Once there, it can control access in and out of that point.

Although firewalls typically protect internal networks from public networks, they are also used to control access between specific network segments within a network—for example, placing a firewall between the Accounts and the Sales departments.

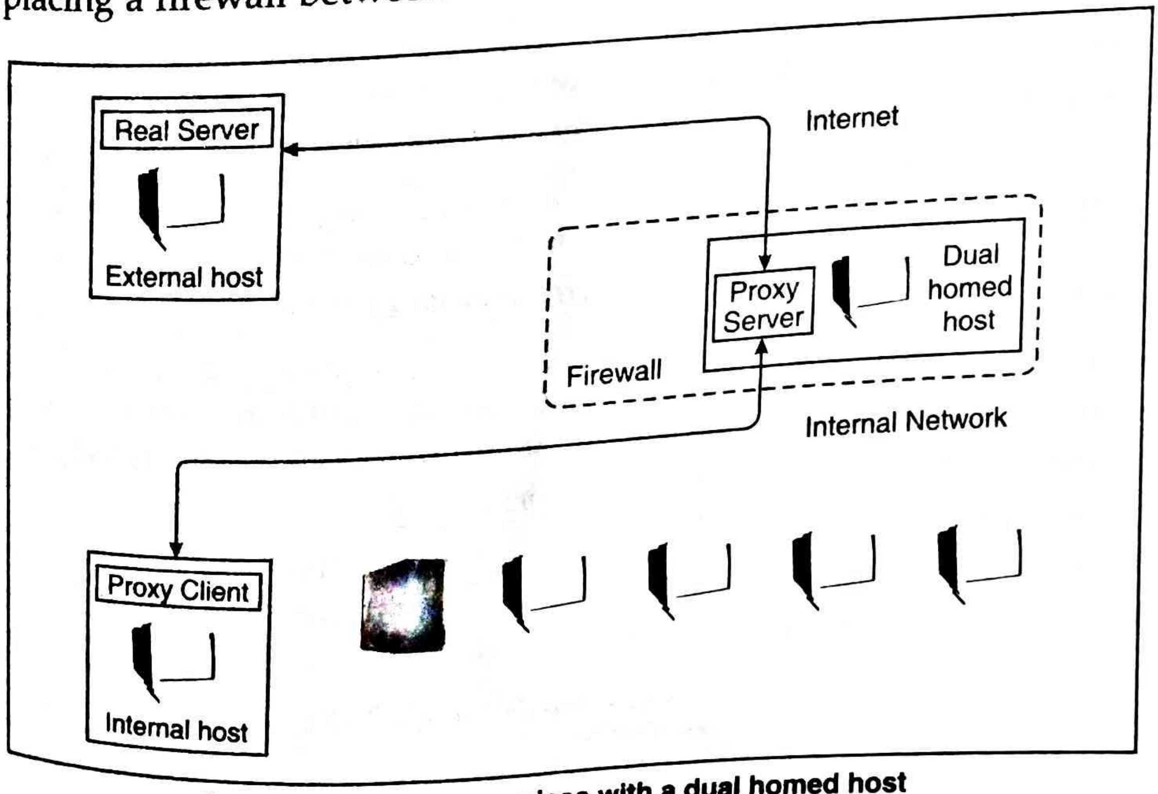

Fig. 5.9. Using proxy services with a dual homed host

As mentioned, firewalls can be implemented through software or through a dedicated hardware device. Organizations implement software firewalls through network operating systems (NOS) such as Linux/UNIX, Windows servers, and Mac OS servers. The firewall is configured on the server to allow or permit certain types of network traffic. In small offices and for regular home use, a firewall is commonly installed on the local system and configured to control traffic. Many third-party firewalls are available.

Hardware firewalls are used in networks of all sizes today. Hardware firewalls are often dedicated network devices that can be implemented with very little configuration and protect all systems behind the firewall from outside sources. Hardware firewalls are readily available and often combined with other devices today. For example, many broadband routers and wireless access points have firewall functionality built in. In such case, the router or WAP might have a number of ports available to plug systems in to.

5.3.1. Firewalls and Internet Security

Interest and knowledge about computer and network security is growing along with the need for it. This interest is, no doubt, due to the continued expansion of the Internet and the increase in the number of businesses that are migrating their sales and information channels to the Internet. The growth in the use of networked computers in business, especially for e-mail, has also fueled this interest. Many people are also presented with the postmortems of security breaches in high-profile companies in the nightly news and are given the impression that some bastion of defense had failed to prevent some intrusion. One result of these influences is that many people feel that Internet security and Internet firewalls are synonymous. Although we should know that no single mechanism or method will provide for the entire computer and network security needs of an enterprise, many still put all their network security eggs in one firewall basket.

Computer networks may be vulnerable to many threats along many avenues of attack, including :

- Social engineering, wherein someone tries to gain access through social means (pretending to be a legitimate system user or administrator, tricking people into revealing secrets, etc).
- War dialing, wherein someone uses computer software and a modem of search for desktop computers equipped with modems that answer, providing a potential path into a corporate network.
- Denial-of-service attacks, including all types of attacks intended to overwhelm a computer or a network in such a way that legitimate users of the computer or network can not use it.
- Protocol-based attacks, which take advantage of known (or unknown) weaknesses in network services.

- Host attacks, which attack vulnerabilities in particular computer operating systems or in how the system is set up and administered.
- Password guessing.
- Eavesdropping of all sorts, including stealing e-mail messages, files, passwords and other information over a network connection by listening on the connection.

Internet firewalls have been around for a hundred years-in Internet time. Firewalls can help protect against some of these attacks, but certainly not all. Firewalls can be very effective at what they do. The people who set up and use them must have the knowledge of how they work, and also be aware of what they can and cannot protect. In this article, we examine the Internet firewall, touch on its history, see how firewalls are used today, and discuss changes that are in place for the next hundred years.

5.3.2. Three Myths of Firewalls

1. Firewalls make the assumption that the only way in or out of a corporate network is through the firewalls; that there are no "back doors" to your network. In practice, this is rarely the case, especially for a network which spans a large enterprise. Users may setup their own backdoors, using modems, terminal servers, or use such programs as "PC Anywhere" so that they can work from home. The more inconvenient a firewall is to your user community, the more likely someone will set up their own "back door" channel to their machine, thus bypassing your firewall.
2. Firewalls make the assumption that all of the bad guys are on the outside of the firewall, and everyone on the inside of the can be considered trustworthy. This neglects the large number of computer crimes which are committed by insiders.
3. Newly evolving systems are blurring the lines between data and executables more and more. With macros, JavaScript, Java and other forms executable fragments which can be embedded inside data, a security model which neglects this will leave you wide open to a wide range of attacks.

5.3.3. Benefits of Using a Firewall

The main benefits of using a firewall are :

(a) Protection from services which are inherently more prone to attacks.

(b) Access to host in the network can be strictly controlled.

(c) Security is concentrated on a single firewall system. This leads to better implementation of authentication procedures.

(d) Logging and statistics of network use and misuse.

(e) Policy enforcement.

5.3.4. What Constitutes a Good Firewall System?

Firewalls can protect network environment. But what constitutes a good firewall? The answer actually depends on site security requirements. However, one should always check for the following attributes in a firewall :

(a) The firewall should be able to support a "deny all services except those specifically permitted" design policy, even if that is not the policy used.

(b) The firewall should be flexible. It should be able to accommodate new services and needs if the security policy of the organization warrants so.

(c) The firewall should contain advanced authentication measures.

(d) The firewall should employ filtering techniques to permit or to deny services to specified host systems as and when needed.

(e) The firewall should use proxy services for File Transfer Protocol (FTP) and TELNET (TELecommunication NETwork), so that advanced authentication measures can be employed and centralized. If services such as gopher or HTTP are required, the firewall should contain the corresponding proxy services.

(f) The firewall should accommodate public access to the site, such that public information servers can be protected by the firewall but can be segregated from site systems that do not require the public access.

(g) The firewall should contain mechanisms for logging traffic and suspicious activity, and should contain mechanisms for log reduction so that logs are readable and understandable.

(h) If the firewall requires an operating system such as Unix, a secured version of the operating system should be part of the firewall.

5.3.5. Firewall Types

Firewall uses a variety of architectures to manage access control. These are :

(a) Packet-Filtering Firewalls

(b) Proxy Firewalls

(c) Stageful Inspection

Packet-Filtering Firewalls

This type of firewalls examines all the packets it comes across. It forwards them or drops them based on pre-defined rules. This rudimentary firewall provides only basic protection. Packet-filtering firewalls are restrictive since network managers can only define a few parameters.

Many routers and proxy servers use some form of packet filtering that provides firewall capabilities for protecting the network from unauthorized traffic. Administrators can create rules for filtering out unwanted packets and can arrange these rules in the most efficient order. A packet that passes all the rules is only allowed through, while a packet that violates any rule is dropped.

Packet filtering can be implemented on routers and other devices in two ways :

(a) Static filtering

(b) Dynamic filtering

Static Filtering : Static packet filtering provides limited security by configuring selected ports as either permanently open or permanently closed. For example, to deny outside packets access to a company intranet server on port 80 (the standard port number for the Hypertext Transfer Protocol, or HTTP) one could configure the router or firewall to block all incoming packets directed toward port 80.

Dynamic Filtering : Dynamic packet filtering provides enhanced security. It acts by allowing selected ports to be opened at the start of a legitimate session and then closes them at the end of the session. This is particularly useful for protocols that allocate ports dynamically-for example, with the File Transfer Protocol (FTP).

Application-level Firewall (Application Gate Way)

Application-level firewall (or application gateway) is part of a proxy server. Application gateways do not allow any packet to pass directly between the two networks they connect. Instead, proxy applications running on the firewall computer forward requests to services on the private network. Then forward responses to the originators on the unsecured public network. Application gateways authenticate the credentials of a user before allowing access to the network. They use auditing and logging mechanisms as part of the security policy. Application gateways require some configuration on the part of users to enable their client machines to function properly. For example, if a File Transfer Protocol (FTP) proxy is configured on an application gateways, it can be configured to allow some FTP commands but deny others. One could configure an SMTP (Simple Mail Transfer Protocol) proxy on an application gateway that would accept mail from the outside (without revealing internal e-mail addresses), and then forward the mail to the internal mail server. However, because of the additional processing overhead, application gateways have advanced hardware requirements and are slower than network-level firewalls.

Proxy Firewalls

This type of firewall acts as an intermediary of user requests, setting up a second connection to the desired resource either at the application layer (an application proxy) or at the session or transport layer (a circuit relay). Proxy firewalls tend to take a performance hit since it relies on application layer. They are restrictive when it comes to allowing or denying evolving or new types of applications.

Stageful Inspection Firewalls

These are the new generation of firewall technology patented by Check Point Software Technologies. Stageful Inspection provides full application-layer awareness without requiring a separate proxy for every service to be secured. This results in multiple benefits to customers including excellent performance, reliability, and the ability to support new and custom applications and services quickly and easily. Stageful inspection architecture is unique in that it understands the state of any communication through the firewall machine, including packet, connection and application information.

Packet filters do not track application or connection state. Application proxies track only application state, not packet or connection state which may introduce some vulnerabilities.

5.4. ATM

ATM is a switching and multiplexing technology that employs small, fixed-length cells to very quickly and efficiently move all types of traffic. ATM is fast and efficient because its cells fit into spaces too small for larger packets or frames, traffic routes are preplanned, switching is done without the need for time-consuming software, and payload error checking and correction is performed only at the destination node, not at every hop along the way.

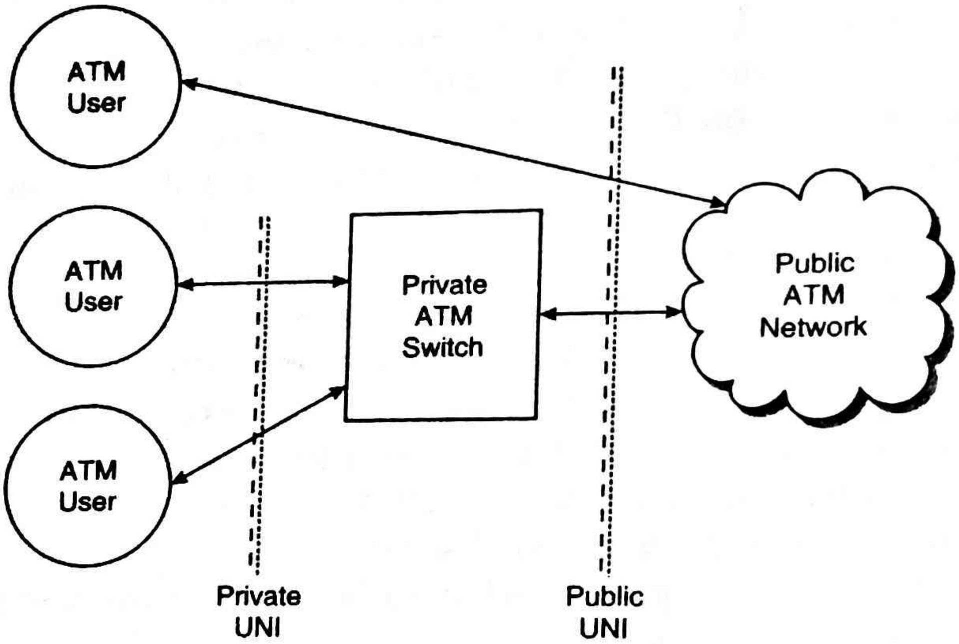

Fig. 5.10. ATM Network

ATM was designed to be the protocol of choice for future Broadband-Integrated Services Digital Network (B-ISDN) services. Because ATM is asynchronous, it provides true bandwidth-on-demand. Additionally, its small cell size makes ATM adaptable to any form of information—data, voice, video, audio, e-mail, faxes—and capable of moving this information amazingly fast across a network that can provide millions of virtual paths and channels between end user equipment.

Characteristically, ATM has two dimensions; transport and switching. In the transport dimension, ATM can move no faster or slower than any other digital communication technology. It is in the switching dimension that ATM shines. Packets and frames of various sizes need smart switches controlled by slow-moving software to move them through a network. Small, uniformly sized cells on an ATM network move through switches without needing software assistance. The cells already know the route to take and do not need to slow down to look for road signs or stop to get directions.

ATM allows the user to select the level of service it needs, provides guaranteed service quality and makes reservations and preplans routes so those signals needing the most attention are given the best service. Whether the signal travels first class or standby, ATM can accommodate the user.

5.4.1. ATM Cells

An ATM cell is a 53-octet packet of information consisting of two main parts (see Figure 5.11) :

- **Header** : 5 octets reserved for :
  - Routing (GFC)
  - Addressing (VPI, VCI, PTI)
  - Flow control (CLP, HEC)

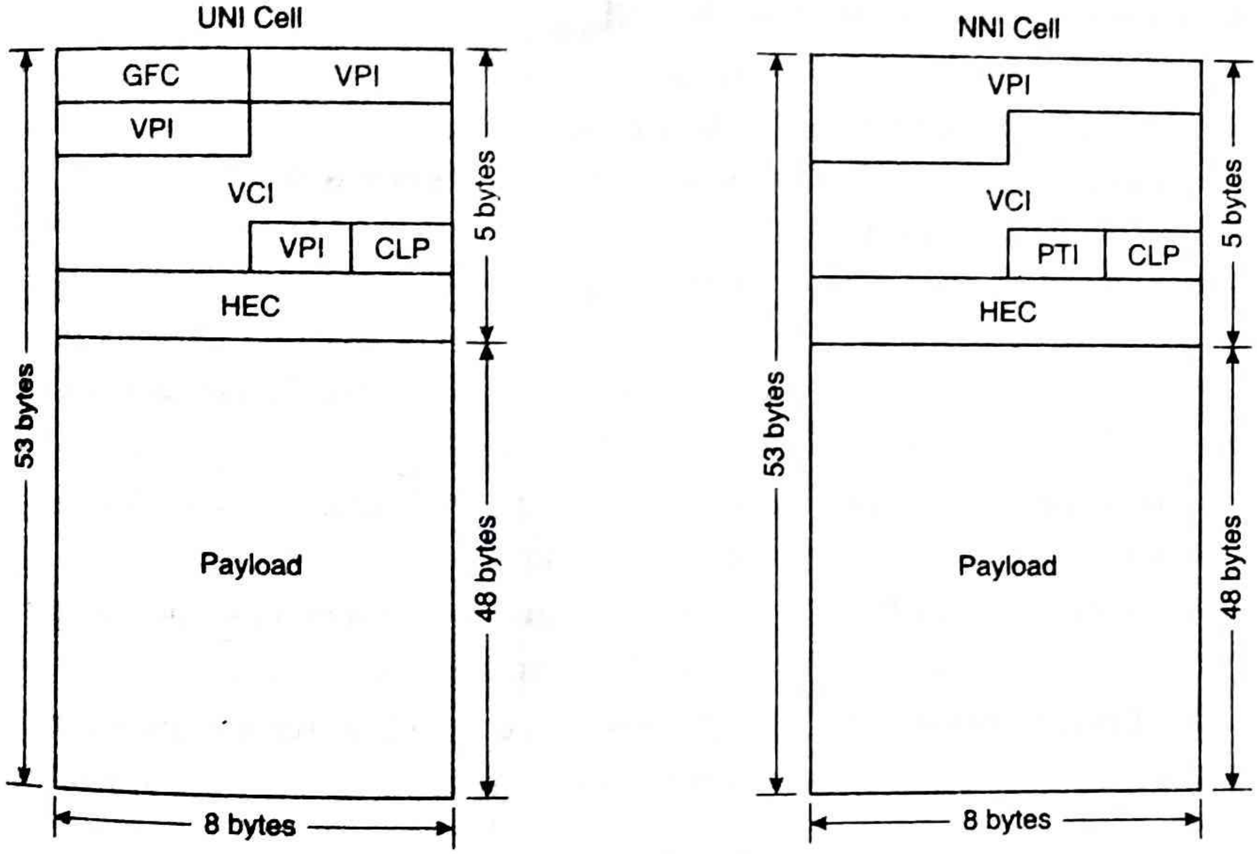

Fig. 5.11. ATM cells

- **Payload** : 48 octets reserved for voice, video, autio and data (user or service)

The cell size is a compromise between what Europe and the United States wanted. Europe liked a 32-octet payload to reduce delay (smaller cells get through switches more quickly). The US preferred a 64-octet payload to increase bandwidth efficiency (because of a better header-to-payload ratio). The ITU settled the issue with a 48-octet compromise, giving both sides a portion of what they wanted. After adding five header octets, the ATM cell size is 53 octets (this gives an approximate 1 : 10 ratio between header and payload cells).

Header Cell Structure : There are 40 bits in each ATM header. These bits are subdivided into various-sized groupings designed to move payload through the ATM network to its destination. These subgroups are :

Generic Flow Control (GFC) : Four bits that control traffic flow between the ATM network and terminal equipment. These are gatekeeper bits that do not travel with the cells across the ATM network but are used to establish connections with end user equipment.

Additionally, GFC bits :

- Manage access conflicts, giving each user fair access to the ATM network.
- Ensure that proper quality of service is allotted to each user (see ATM Traffic Contract).
- Support up to 100 users on each UNI.

Virtual Path Identifier (VPI) : The address for up to 256 UNI virtual paths (VPs) (8 VPI bits) or up to 4096 NNI VPs (12 VPI bits). The path is fixed at connection but is shared with multiple other calls. Because NNI VPIs overwrite UNI GFC bits, more than 4000 virtual paths can be used within the ATM network.

Virtual Channel Identifier (VCI) : The rest of the VPI address that identifies virtual channels with in each virtual path. Sixteen bits make possible 65,536 virtual channels. The combination of VPI and VCI fields allow for 16,777,216 simultaneous UNI calls and up to 268,435,456 simultaneous NNI calls.

Payload Type Identifier (PTI) : Three bits that identify the cell as carrying information for the user or as carrying service information.

Cell Loss Priority (CLP) : One bit that determines if a cell can be discarded if the network becomes too congested (0 = keep, 1 = discard).

Header Error Control (HEC) : Eight bits that do cyclical redundancy checks on the first four header octets. The HEC ensures multiple bit error detection and single bit error correction.

5.4.2. ATM Protocol

The ATM protocol layer model consists of four layers and three planes (see Figure 5.12). The layers are closely interrelated, but each layer addresses a specific set of functions. The physical layer and ATM layer can be compared with the physical layer of the OSI reference model. As with the OSI model, the various layers function independently, but continuous interaction among the layers is highly coordinated.

Physical Layer

The physical layer has four functions :

- Converts cells to a bit stream.
- Controls transmission and receipt of bits on the physical medium.

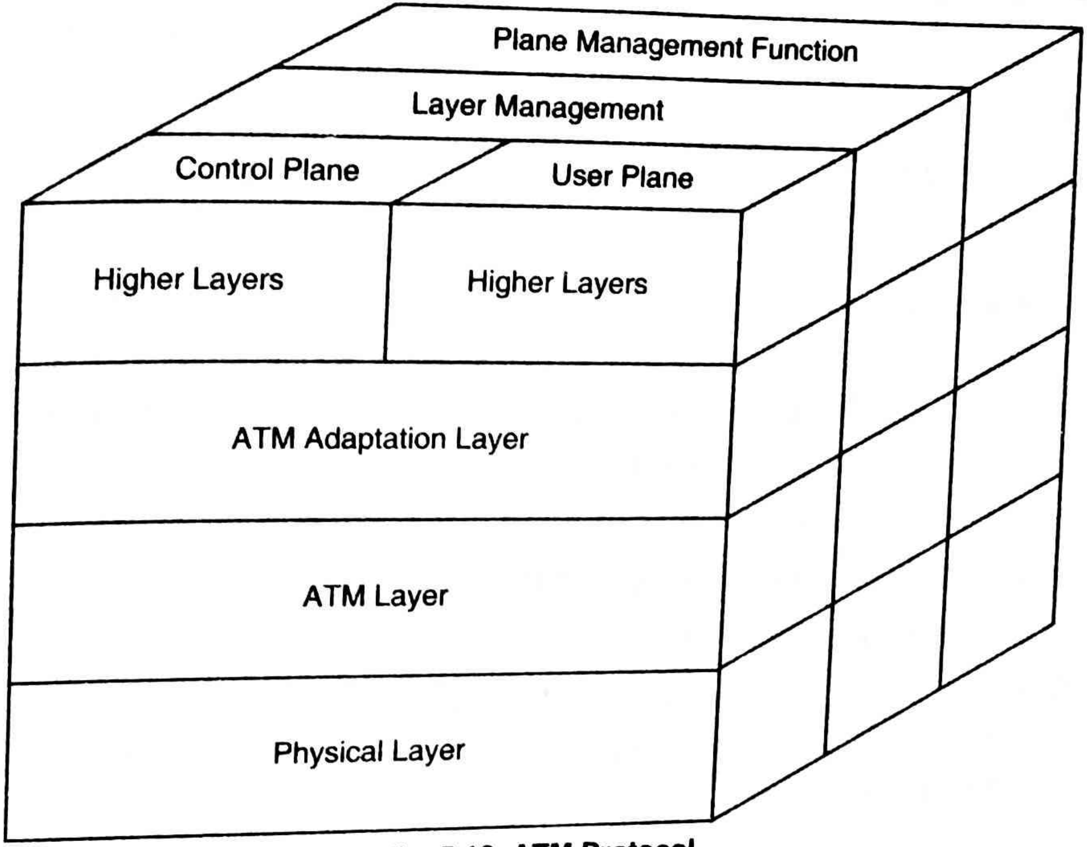

Fig. 5.12. ATM Protocol

- Tracks ATM cell boundaries.
- Packages cells into frame types that fit the physical medium (SONET cells are packaged differently than cells going to or coming from a DS3 line).

The physical layer is divided into two sublayers that perform these four functions :

- **Physical Medium-Dependent (PMD)** : Syncs the bits and gets them to or from the correct medium, down to the correct cable and connector types.
- **Transmission Convergence (TC)** : Performs error checking, maintains cell boundaries and synchronization and packages the cells for the particular physical medium.

ATM Layer

The ATM layer performs many very critical functions essential to the exchange of end-to-end communications :

- First the ATM layer takes the 48-byte payload from the ATM adaptation layer and adds the 5-byte addressing header.
- Then it multiplexes all the cells from various connections, prepares a single-cell stream for the physical layer and puts in idle cells, if needed, as fillers for synchronous transmission systems (for example, SDH or SONET).
- Next the ATM layer provides translation (directional coding) for every cell to get the cells switched through multiple virtual connections. The ATM layer can do this because it knows the capabilities of virtual connections carrying the cells. These capabilities vary according to :

- Bandwidth
- Delay
- Delay variation
- Cell loss

5.4.3. ATM Adaptation Layer

The ATM adaptation layer (AAL) is the top ATM layer in the protocol stack. This layer interacts with higher layers to get such customer information as voice, video and data into and out of the payload portion of a 53-byte ATM cell.

AAL functions are divided within two sublayers (see Figure 5.13) :

- **Segmentation and Reassembly (SAR)**
- **Convergence Sublayer (CS)**

The SAR sublayer takes a continuous bit stream of customer data, slices it up and puts it into small ATM cells. At the other end of the network, the SAR sublayer unwraps the ATM cells and exactly reconstructs the bit stream.

The CS provides different classes of service (A, B, C, D or X) and performs a variety of tasks that are dependent on the AAL type in which the CS resides. AALS types are described in the following paragraphs.

AAL Type 1 (AAL1) : AAL Type 1 is class A service that is connection-oriented and capable of handling constant bit rate (CBR) traffic such as voice and video conferencing. AAL1 requires exact timing between source and destination, so it needs to travel over a synchronous network (for example, SONET or SDH).

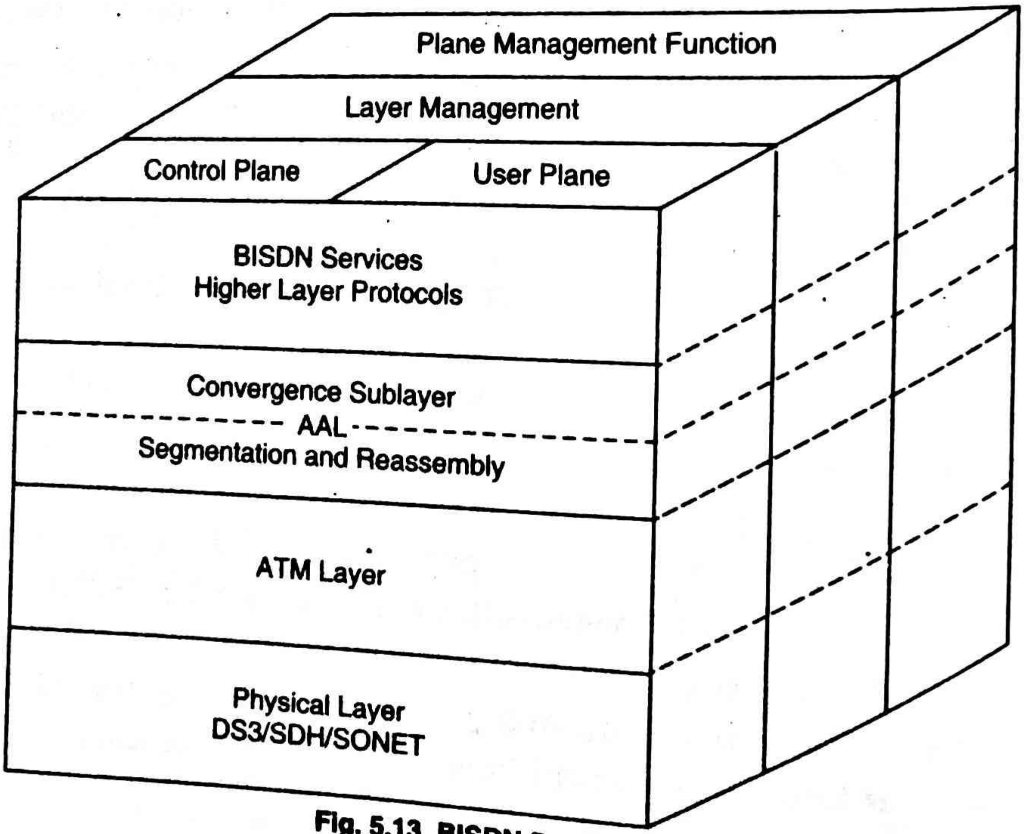

Fig. 5.13. BISDN Protocol

Synchronization samples are inserted into the payload field and one payload byte is used for sequencing (4 bits each for a sequence number [SN] and sequence number protection [SNP]).

AAL Type 2 (AAL2) : AAL2 supports class B traffic. This too is connection-oriented; however, AAL2 can have a variable bit rate in real time and does not require end-to-end timing. AAL2 is great for compressed audio and video and travels at a high priority through the ATM network. Forty-four bytes make up the AAL2 payload. Four bytes are reserved to support AAL2 processes.

AAL Type 3/4 (AAL3/4) : After AAL types 3 and 4 were developed, it was determined they so closely resembled each other that they were combined to form AAL3/4. This combined type supports class C or class D non-real time, variable bit rate traffic that requires no timing.

Class C traffic (such as Frame Relay and X.25 data packets) require connections. Connectionless traffic (such as LAN and SMDS data) is considered Class D. Both classes are sensitive to loss but not especially sensitive to delay.

AAL3/4 provides two types of mode services; message and streaming. Message mode has one interface data unit (IDU), a single framed cell of up to 65,535 octets. Streaming mode contains multiple IDUs that are transported asynchronously. Both modes are tagged with 10-bit SAR CRC trailers for error checking and a CS protocol data unit (PDU) that is prepended with begin/end tags and a length field.

AAL Type 5 (AAL5) : AAL5 is the primary AAL for data, both connection-oriented and connectionless. AAL5 is known as the simple and efficient layer (SEAL) because nothing extra is appended to the CS-PDU that goes into the 48-octet payload.

AAL5 supports class C and class X traffic, including LAN Emulation (LANE) and IP, with unspecified or available bit rates (UBR or ABR). As no header or trailer is added, AAL5 traffic cannot be interleaved.

5.4.4. ATM Service Categories

ATM offers five classes of service. Each class is designed to accommodate data bursts according to customer needs and provide the appropriate quality of service (QoS) for each service class. The five service categories are :

- **Constant Bit Rate (CBR)**
  - Provides a continuous rate of flow
  - Supports traffic sensitive to delay and loss
  - Emulates circuit switching
  - Carries uncompressed voice and video
- **Real-time Variable Bit Rate (rt-VBR)**
  - Supports traffic dependent on timing and control information
  - Carries compressed voice, video and audio

- **Non-real-time Variable Bit Rate (nrt-VBR)**
  - ➢ Supports traffic at rates that vary with time
  - ➢ Unaffected by loss or delay because of time to recover
  - ➢ Carries data and buffered voice and video
- **Unspecified Bit Rate (UBR)**
  - ➢ Provides no assurance the data will be delivered (best effort only)
  - ➢ Carries file transfers and E-mail

Table 5.4 : ATM Service Classes

Service Class	Class A	Class B	Class C	Class D	Class X
AAL Type	1	2	3/4; 5 in Message Mode	3/4	5
ATM Forum	CBR	rt-VBR	nrt-VBR UBR	UBR ABR	ABR
Bit Rate	CBR	VBR			
Timing	Required	Required			
Connection Mode	Connection oriented	Connectionless	Connection or connectionless		
Traffic Contract Parameters	PCR, CDVT	PCR, CDVT, SCR, MBS, BT	PCR, CDVT	PCR, CDVT, MCR	
QoS Parameters	CDV, CTD, CLR	CLR	Not specified		
Some Applications	Uncompres-sed voice, video and audio	Compressed voice, video and audio	X.25, Frame Relay, Transaction Processing	SMDS, LAN, nrt buffered video	Network management, E-mail, FTP, WAN, LAN, IP

Available Bit Rate (ABR)

- Provides no assurance the data will be delivered (best effort only).
- Supports nrt-VBR traffic with flow control.

QoS standards have been established for each service category. Table 5.4 lists ATM service classes and Table 5.5 lists QoS parameters for each service category.

Table 5.5 : QoS Parameters

Service Category	Quality of Service
CBR	CDV, CTD, CLR
rt-VBR	CLR
nrt-VBR	CLR
UBR	none
ABR	none

5.4.5. ATM Access

A rate of DS1 or greater is needed to access an ATM backbone. The access can be made through customer premises equipment (CPE) or by way of a switched multibit data service (SMDS), Frame Relay, or X.25 network switch (or more commonly, a DS3 or SONET/SDH connection) through a User Network Interface (UNI).

A UNI is where the access network stops and the ATM network begins. The UNI is the point between the user and the public network service provider, and it is here that specifications for procedures and protocols between these two networks are established. Also, the ATM network takes responsibility for converting user protocol data units (PDUs) to ATM PDUs and cells.

A private UNI that allows end user access to an ATM network is a data exchange interface (DXI). DXI access is through a router, bridge or ATM data service unit. DXI allows protocol sharing between the user and network provider and reduces ATM protocol responsibility, such as breaking out and reconstructing ATM cells.

Network-to-network interface (NNI) is a public UNI. The NNI provides access to the public network along permanent virtual circuits established at the UNI.

UNIs and NNIs can be connected in the following arrangements :

- UNI to UNI (subscriber-to-subscriber)
- UNI to NNI (subscriber-to-network)
- NNI to UNI (network-to-subscriber)
- NNI to NNI (network-to-network)

ATM Interface Connections : Each interface permits a maximum number of virtual path (VP) and virtual channel (VC) connections for user data. These are the physical limits due to hardware, not software. Table 3 lists the total number of VPs and VCs available at each interface connection.

ATM Connections : A permanent virtual circuit (PVC) is a logical connection between two end users established by administative procedures. A PVC is usually created long before it is used and remains in place until the connection is deprovisioned. PVCs can be virtual path or virtual channel connections (VPC or VCC). Bandwidth is allocated for a PVC whether it is used or not.

A switched virtual circuit (SVC) is a logical connection between two subscribers that is established and deconstructed with access and network signaling procedures. Cell transfer instructions for user traffic are established as each SVC is created. Bandwidth is allocated dynamically as it is needed.

ATM Connection Topologies : ATM topologies can be point-to-point, point-to-multipoint, unidirectional and bidirectional. In a point-to-multipoint configuration, the primary source is only from a UNI, while the multiple end points can be UNI or NNI. This configuration also makes use of multicasting to replicate and distribute data to multiple subscribers. Multicasting types are logical and spatial.

In logical multicasting, the multiple endpoints are on the same physical interface. In spatial multicasting, the end points are on separate physical interfaces.

5.4.6. Virtual Paths and Virtual Channels

In ATM, cells move on the transmission protocol using known end-to-end routes called virtual connections. These virtual connections link one communicating entity with another through what is sometimes a very complex system of physical medium links. Every transmission link can support thousands of virtual connections, depending on when and how much traffic needs to be sent. The pathways between virtual connections are virtual paths and virtual channels.

Traffic from various sources is bundled in a pipe and directed as a whole until it reaches its destination. While the path in the network can be redirected, all bundled channels arrive at their destination. The ATM pipe is a virtual path (VP) and is a concatenation of circuits in a link see (Figure 5.14). The action of redirecting the pipe is called VP switching (VPS).

Various sources or channels request connection to specific destinations. These channels are not fixed throughout the network and can jump from one VP to another to reach their destination. These are called virtual channels (VCs), and the redirecting of VCs is called VC switching (VCS).

ATM cells contain header VPI and VCI that keep track of where the cells are from one node or switch to another. These identifiers do not remain constant from endpoint to endpoint but change from hop to hop.

Virtual Path and Channel Switches : Switches join transmission resources together and make sure a cell gets from one transmission link to another. If needed, switches use buffers that hold cells until time for them to move onto the next link.

The switch checks the incoming cell's VPI and VCI, determines the next outgoing VP and VC, and switches accordingly. The switch has the intelligence to move the cells to the appropriate next node by using switching tables that provide information for permanent or switched virtual connections.

Switched virtual connections are the preferred mode of operation as they are dynamically established. They are much like the connections established each time a

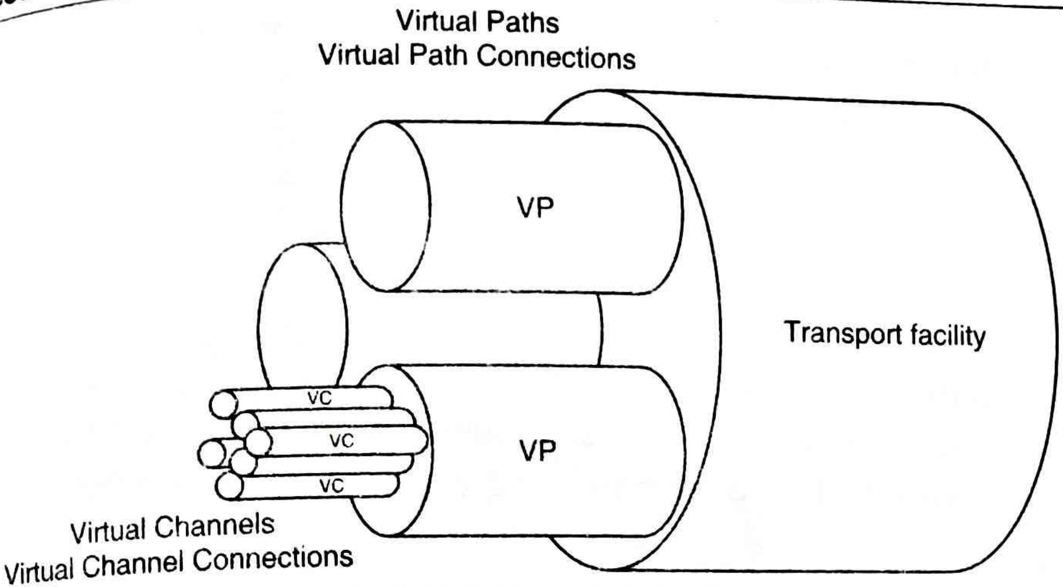

Fig. 5.14. Paths and channels

phone call is made. The destination of the call is determined by dialing the number, and the network determines the links to that destination and establishes the connection. The connection is permanent until the call is finished and all call traffic is delivered in the same order it was sent.

5.4.7. ATM Advantages

ATM provides several advantages :

- ATM fixed-length cells require lower processing overhead and produce higher transmission speeds than traditional packet switching methods.
- ATM transmits asynchronous data in a synchronous network while prioritizing time-sensitive traffic ahead of delay-tolerance traffic to ensure that quality of service is maintained.
- ATM delivers true bandwidth-on-demand (a big plus for high-speed voice, data and video service) and uses statistical multiplexing techniques to efficiently use resources.
- ATM is application-independent, meaning it can be used as a common infrastructure for many network types, including public, private, LAN and campus backbones.
- ATM is designed for high-performance, multimedia networking on a broad range of devices :
  - PC, workstation, server network interface cards
  - Switched Ethernet and token ring workgroup hubs
  - Workgroup and campus ATM switches
  - Enterprise network switches, multiplexers, edge and backbone switches

International standards compliance in central office and customer premises environments allow multi-vendor operation and interoperability.

5.4.8. ATM Disadvantages

While ATM has several advantages, ATM disadvantages can be seen in three areas :

- Cost
- Complexity
- Availability

Compared with voice switches, ATM switches are still much more expensive per line. ATM was implemented before the designers intended and ATM standards are still trying to catch up, which should help to bring down the cost of ATM equipment.

ATM equipment is very complex and intelligent. It takes a very intelligent management team and system to operate ATM successfully (that is, efficiently and cost-effectively).

ATM is not as widely available as SONET and SDH. These are older protocols that are widely accepted and very well standardized. Additionally, with the proliferation of Fast Ethernet and Gigabit Ethernet, ATM may not be necessary. This is especially true for those who transport Ethernet over already ubiquitous SONET/SDH networks.

5.5. Voice over Internet Protocol

Voice over Internet Protocol (VoIP), is a technology that allows you to make voice calls using a broadband Internet connection instead of a regular (or analog) phone line. Some VoIP services may only allow you to call other people using the same service, but others may allow you to call anyone who has a telephone number-including local, long distance, mobile and international numbers. Also while some VoIP services only work over your computer or a special VoIP phone, other services allow you to use a traditional phone connected to a VoIP adapter.

What is VoIP?

Voice over Internet Protocol (VoIP) is a form of communication that allows you to make phone calls over a broadband internet connection instead of typical analog telephone lines. Basic VoIP access usually allows you to call others who are also receiving calls over the internet.

Interconnected VoIP services also allow you to make and receive calls to and from traditional landline numbers, usually for a service fee. Some VoIP services require a computer or a dedicated VoIP phone, while others allow you to use your landline phone to place VoIP calls through a special adapter.

VoIP is becoming an attractive communications option for consumers. Given the trend towards lower fees for basic broadband service and the brisk adoption of even faster internet offerings, VoIP usage should only gain popularity with time. However, as VoIP usage increases, so will the potential threats to the typical user. While VoIP vulnerabilities are typically similar to the ones users face on the internet, new threats, scams and attacks unique to IP telephony are now emerging.

How VoIP/Internet Voice Works

VoIP services convert your voice into a digital signal that travels over the internet. If you are calling a regular phone number, the signal is converted to a regular telephone signal before it reaches the destination. VoIP can allow you to make a call directly from a computer, a special VoIP phone, or a traditional phone connected to special adapter. In addition, wireless "hot spots" in locations such as airports, parks and cafes allow you to connect to the Internet and may enable you to use VoIP service wirelessly.

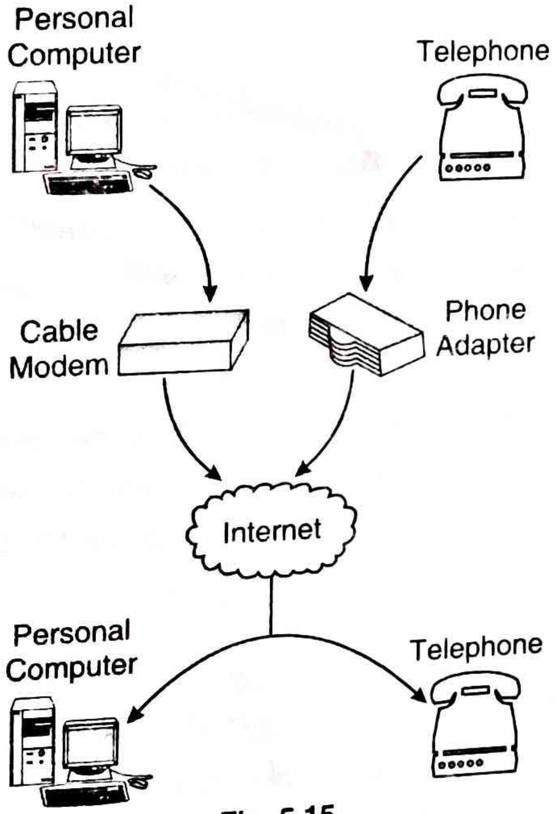

Fig. 5.15

What Kind of Equipment Do I Need?

A broadband (high speed internet) connection is required. This can be through a cable modem, or high speed services such as DSL or a local area network. **A computer, adaptor, or specialized phone is required.** Some VoIP services only work over your computer or a special VoIP phone, while other services allow you to use a traditional phone connected to a VoIP adapter. If you use your computer, you will need some software and an inexpensive microphone. Special VoIP phones plug directly into your broadband connection and operate largely like a traditional telephone. If you use a telephone with a VoIP adapter, you'll be able to dial just as you always have, and the service provider may also provide a dial tone.

VoIP Configurations

Dedicated routers : These devices allow you to use your traditional phone to place VoIP calls. They are connected to cable/DSL modems (or any high-speed internet source) and allow you to attach an ordinary telephone. Once configured, and with an appropriate VoIP provider and service plan, these devices require no special software or interaction with a computer. In fact, you only need to pick up your phone and dial a number at the dial tone. you also may bring your adapter with you when you travel and make calls wherever broadband internet access is available.

Adapters (USB)

These devices also allow you to use a traditional phone to place VoIP calls. They usually come in the form of USB adapters that are slightly larger than the typical thumb drive. They feature a standard modular phone jack to which you can attach an ordinary phone line. Once connected, your phone behaves as if it were connected to standard phone service. Behind the scenes, however, the included software is actually setting up a VoIP call.

Software-controlled VoIP applications : "softphones"

There are many software applications ("softphones") that allow you to place VoIP phone calls directly from an ordinary computer with a headset, microphone, and sound card. Internet telephony service providers usually give away their softphones but require that you use their service. Together, these applications and services enable users to talk to other people using the same service at no cost, and to the rest of the world for a fee. Software based VoIP applications are quite attractive to consumers because they often already have most of the components necessary to get started at little to no cost.

Dedicated VoIP phones

A VoIP phone looks like an ordinary corded or cordless telephone, but it connects directly to a computer network rather than a traditional phone line. A dedicated VoIP phone many consist of a phone and base station that connects to the internet or it may also operate on a local wireless network. Like the VoIP adapters mentioned above, dedicated VoIP phones also require a provider and service plan.

Requirements, Availability, and Service Limitations

When considering VoIP service, you should not assume that its features, functionality and options will equal those of traditional landlines; you should be familiar with the requirements, availability, and possible service limitations of VoIP service before switching to VoIP as either a primary means of communication or an enhancement to your current services.

Requirements

VoIP requires a connection to the Internet through an ISP, a VoIP service to extend the reach to traditional landlines, and VoIP software to actually place calls. Plain Old Telephone Service (POTS) requires none of these prerequisites. It is important to note that Digital Subscriber Line (DSL) internet service uses traditional phone lines for your internet connection; in this case, you already have telephone service to begin with. You may wish to weigh the expected benefits of VoIP against these costs given your current operating environment.

Availability due to Power Outages

During a typical power outage, VoIP becomes unavailable because VoIP devices (computers, routers, adapters) usually rely on a power source to function. Traditional phone lines are usually still available during such as outage, which is a major advantage in an emergency. Ultimately, it may be necessary to use an uninterruptible power supply (UPS) with a VoIP installation if connectivity is desired during a power outage or some other kind of emergency.

Availability due to Bandwidth

VoIP communication nearly always requires a high-speed (broadband) internet connection for reliable functionality. Even given typical broadband connection speeds, though, service interruptions or degradation of quality is possible due to high internet traffic. For example, if you are trying to place a VoIP call while other people are using a lot of bandwidth on the same internet connection, the sound quality of your VoIP call or general VoIP availability may be affected.

Threats/Risks

Many of the threats associated with VoIP are similar to the threats inherent to any internet application. Internet users are already familiar with the nuisance of email abuse in the form of spam and phishing attempts. VoIP opens yet another pathway for these annoyances, which can lead to spam over internet telephony (SPIT), spoofing and identity theft. Additionally, the confidentiality of VoIP conversations themselves has come into question, depending on service type or VoIP configuration.

Spam over Internet Telephony (SPIT)

As VoIP usage increases, so will the pesky marketing strategies associated with it. Perennial annoyances like telemarketing and spam have been plaguing consumers and internet users for years. A new sort of hybrid of these two concepts is SPIT, or spam over internet telephony. Like email spamming, sending commercial messages via VoIP is fast and cheap. Unlike traditional telemarketing, though, VoIP offers the potential for large volumes of unsolicited calls, due to the wide array of tools already available to attackers on the internet. Telemarketers could easily send large amounts of messages to VoIP customers. Unlike traditional spam email messages, which average only 10-20 kilobytes in file size, unwanted VoIP voicemails can require megabytes of storage.

Spoofing

It is technically possible for an attacker to masquerade as another VoIP caller. For example, an attacker could possibly inject a bogus caller ID into an ordinary VoIP call so that the receiver believes the call to be coming from a known and trusted source (a bank, for example). The receiver, fooled by the electronic identification of the caller, may place unwarranted trust in the person at the other end. In such an exchange, the receiver may be tricked into disclosing personal information like account numbers, social security numbers, or secondary authentication factor : a mother's maiden name, for example. This scheme is essentially the VoIP version of traditional phishing, where a user follows links in an unsolicited email and is tricked into providing personal information on a bogus web site. Attackers may use these bits and pieces of personal information to complete partial identity records of victims of identity theft.

Confidentiality Concerns

Many critics of VoIP question its confidentiality. The concern is that VoIP data sometimes travels unencrypted over the internet. Therefore, it is technically possible for someone to collect VoIP data and attempt to reconstruct a conversation. Although it is extremely difficult to achieve, some software programs are designed to piece together bits and pieces of VoIP data in an effort to reconstruct conversations. While such activity is currently rare, you should be aware of this possibility as it many increase as VoIP becomes more widespread.

How to Protect Against Risks

Many of the principles and practices for safe VoIP usage are the same as those you may already be practicing with other internet applications. Ignoring these general principles could allow attackers to gain control of your computer operating system by means of an existing software flaw or a misconfiguration unrelated to your VoIP application. It may then be possible for them to exploit flaws in your VoIP configuration thereby possibly gaining access to personal information you share when using VoIP. Here are some of the key practices of good personal computing.

- Use and maintain anti-virus and anti-spyware programs.
- Be cautious about opening files attached to email messages or instant messages.
- Verify the authenticity and security of downloaded files and new software.
- Configure your web browser(s) securely.
- Use a firewall.
- Identify, back-up and secure your personal or financial data.
- Create and use strong passwards.
- Patch and update your application software
- Do not divulge personal information to people you don't know.
- If you are using a software VoIP application, consider using encryption software for both your installation and for those you wish to talk to.

5.6. Internet Telephony

A category of hardware and software that enables people to use the internet as the transmission medium for telephone calls. For users who have free, or fixed-price Internet access, Internet telephony software essentially provides free telephone calls anywhere in the world. To date, however, Internet telephony does not offer the same quality of telephone service as direct telephone connections.

There are many Intenet telephony applications available. Some like Cool Talk and Net Meeting, come bundled with popular Web browsers. Others are stand-alone products. Internet telephony products are sometimes called IP telephony, Voice over the Internet (VoI) or Voice over IP (VoIP) products.

Net telephony means making a call via pc/device through VoIP technology. The advantage of Net telephony is the ability to place long-distance phone calls at very cheaper charges than regular ISD call.

EXERCISE

1. Define network device.
2. Define the following terms : Hub, Switch, Repeater, Gateway, Bridge.
3. Explain the concept of hub and switch cabling.
4. Name three types of bridges.
5. Explain in brief the working of router and its different types of routing techniques.
6. What is routing protocol? Name some routing protocols.
7. Define the concept of firewall and internet security.
8. What are the benefits of using firewall?
9. Explain in brief types of firewalls.
10. Write a short note on ATM.
11. What are the ATM services? Explain in brief.
12. Mention some advantages of ATM.
13. What is VoIP? How does it work?
14. Define the term spoofing.
15. Write short note on Internet Telephony.

●