UNIT IV

NETWORK LAYER

1. LOGICAL ADDRESSING

IPv4 ADDRESSES

An **IPv4** address is a 32-bit address that *uniquely* and *universally* defines the connection of a device (for example, a computer or a router) to the Internet.

An IPv4 address is 32 bits long.

IPv4 addresses are unique. They are unique in the sense that each address defines one, and only one, connection to the Internet. Two devices on the Internet can never have the same address at the same time. We will see later that, by using some strategies, an address may be assigned to a device for a time period and then taken away and assigned to another device.

On the other hand, if a device operating at the network layer has *m* connections to the Internet, it needs to have *m* addresses. We will see later that a router is such a device.

The IPv4 addresses are universal in the sense that the addressing system must be accepted by any host that wants to be connected to the Internet.

The IPv4 addresses are unique and universal.

Address Space

A protocol such as IPv4 that defines addresses has an address space. An address space is the total number of addresses used by the protocol. If a protocol uses *N* bits to define an address, the address space is *2N* because each bit can have two different values (0 or 1) and *N* bits can have *2N* values.

IPv4 uses 32-bit addresses, which means that the address space is 232 or 4,294,967,296 (more than 4 billion). This means that, theoretically, if there were no restrictions, more than 4 billion devices could be connected to the Internet. We will see shortly that the actual number is much less because of the restrictions imposed on the addresses.

The address space of IPv4 is 232 or 4,294,967,296.

Notations

There are two prevalent notations to show an IPv4 address: binary notation and dotted decimal notation.

Binary Notation

In binary notation, the IPv4 address is displayed as 32 bits. Each octet is often referred to as a byte. So it is common to hear an IPv4 address referred to as a 32-bit address or a 4-byte address. The following is an example of an IPv4 address in binary notation:

01110101 10010101 00011101 00000010

Dotted-Decimal Notation

To make the IPv4 address more compact and easier to read, Internet addresses are usually written in decimal form with a decimal point (dot) separating the bytes. The following is the dotted decimal notation of the above address:

117.149.29.2

Figure 19.1 shows an IPv4 address in both binary and dotted-decimal notation.

Note that because each byte (octet) is 8 bits, each number in dotted-decimal notation is a value ranging from 0 to 255.

Figure 19.1 *Dotted-decimal notation and binary notation for an IPv4 address*

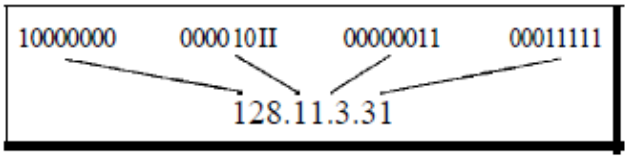

Classful Addressing

IPv4 addressing, at its inception, used the concept of classes. This architecture is called classful addressing. Although this scheme is becoming obsolete, we briefly discuss it here to show the rationale behind classless addressing.

In classful addressing, the address space is divided into five classes: A, B, C, D, and E. Each class occupies some part of the address space.

In classful addressing, the address space is divided into five classes:A, B, C, D, and E. We can find the class of an address when given the address in binary notation or dotted-decimal notation. If the address is given in binary notation, the first few bits can immediately tell us the class of the address. If the address is given in decimal-dotted notation, the first byte defines the class. Both methods are shown in Figure 19.2.

Figure 19.2 *Finding the classes in binary and dotted-decimal notation*

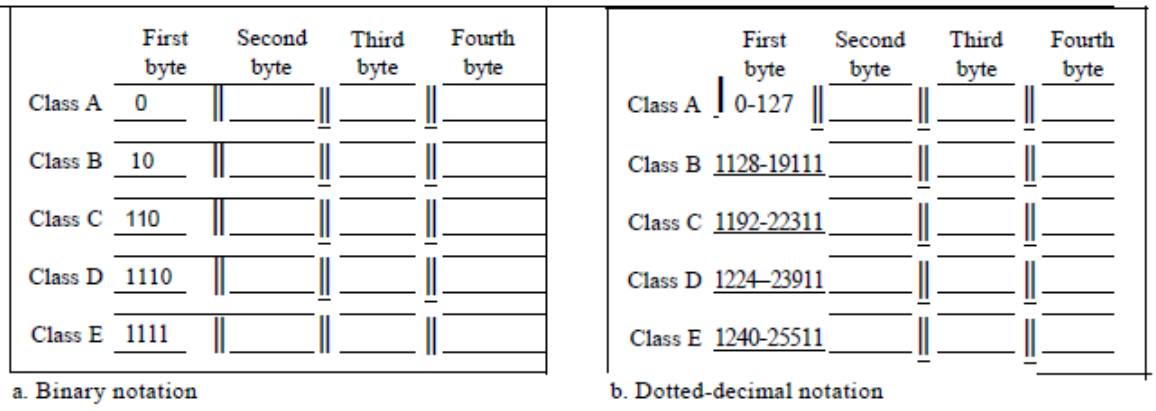

Limitations of classful addressing:

- A block in class A address is too large for almost any organization. This means most of the addresses in class A were wasted and were not used.
- A block in class B is also very large, probably too large for many of the organizations that received a class B block.
- A block in class C is probably too small for many organizations.
- Class D addresses were designed for multicasting which means each address in this class is used to define one group of hosts on the Internet.
- The Internet authorities wrongly predicted a need for 268,435,456 groups. This never happened and many addresses were wasted here too.

And lastly, the class E addresses were reserved for future use; only a few were used, resulting in another waste of addresses

Classes and Blocks

One problem with classful addressing is that each class is divided into a fixed number of blocks with each block having a fixed size as shown in Table 19.1.

Table 19.1 *Number of blocks and block size in classful IPv4 addressing*

Class	*Number of Blocks*	*Block Size*	*Application*
A	128	16,777,216	Unicast
B	16,384	65,536	Unicast
C	2,097,152	256	Unicast
D	1	268,435,456	Multicast
E	1	268,435,456	Reserved

Let us examine the table. Previously, when an organization requested a block of addresses, it was granted one in class A, B, or C. Class A addresses were designed for large organizations with a large number of attached hosts or routers. Class B addresses were designed for midsize organizations with tens of thousands of attached hosts or routers. Class C addresses were designed for small organizations with a small number of attached hosts or routers. We can see the flaw in this design. A block in class A address is too large for almost any organization. This means most of the addresses in class A were wasted and were not used. A block in class B is also very large, probably too large for many of the organizations that received a class B block. A block in class C is probably too small for many organizations. Class D addresses were designed for multicasting as we will see in a later chapter. Each address in this class is used to define one group of hosts on the Internet. The Internet authorities wrongly predicted a need for 268,435,456 groups.This never happened and many addresses were wasted here too. And lastly, the class E addresses were reserved for future use; only a few were used, resulting in another waste of addresses. « In c1assfnl addressing, a large part of the available addresses were wasted.

Netid and Hostid

In classful addressing, an IP address in class A, B, or C is divided into netid and hostid. These parts are of varying lengths, depending on the class of the address. Figure 19.2 shows some netid and hostid bytes. The netid is in color, the hostid is in white. Note that the concept does not apply to classes D and E. In class A, one byte defines the netid and three bytes define the hostid. In class B, two bytes define the netid and two bytes define the hostid. In class C, three bytes define the netid and one byte defines the hostid.

Mask

Although the length of the netid and hostid (in bits) is predetermined in classful addressing,

we can also use a mask (also called the default mask), a 32-bit number made of

Table 19.2 *Default masks for classful addressing*

Class	*Binary*	*Dotted-Decimal*	*CIDR*
A	11111111 00000000 00000000 00000000	255.0.0.0	*18*
B	11111111 11111111 00000000 00000000	255.255.0.0	*116*
C	11111111 11111111 11111111 00000000	255.255.255.0	*124*

Although the length of the netid and hostid is predetermined in we can also use a mask, which is a 32-bit number made of contiguous 1s followed by contiguous 0s.

- The masks for classes A, B, and C are shown in below table.
- The mask can help us to find the netid and the hostid.

SUBNETTING

- Subnetting was introduced for classful addressing.
- If an organization was granted a large block in class A or B,
- Then, it could divide the addresses into several contiguous groups and assign each group to smaller networks (called subnets) or, in rare cases, share part of the addresses with neighbors.
- Subnetting increases the number of Is in the
- mask, as we will see later when we discuss classless addressing.

SUPERNETTING

- In supernetting, an organization can combine several class C blocks to create a larger range of addresses.
- In other words, several networks are combined to create a supernetwork or a supemet.
- An organization can apply for a set of class C blocks instead of just one.
- Ex: An organization that needs 1000 addresses can be granted four contiguous class C blocks. The organization can then use these addresses to create one supernetwork.

- Supernetting decreases the number of 1s in the mask. For example, if an organization is given four class C addresses, the mask changes from /24 to /22.

ADDRESS DEPLETION:

- The flaws in classful addressing scheme combined with the fast growth of the Internet led to the near depletion of the available addresses.
- Yet the number of devices on the Internet is much less than the 232 address space.
- We have run out of class A and B addresses, and a class C block is too small for most midsize organizations.
- One solution that has alleviated the problem is the idea of classless addressing

CLASSLESS ADRESSING:

- To overcome address depletion and give more organizations access to the Internet, classless addressing was designed and implemented.
- In this scheme, there are no classes, but the addresses are still granted in blocks.

ADDRESS BLOCKS:

- In classless addressing, when an entity, small or large, needs to be connected to the Internet, it is granted a block (range) of addresses.
- The size of the block (the number of addresses) varies based on the nature and size of the entity.
- Example, a household may be given only two addresses; a large organization may be given thousands of addresses.

An ISP, as the Internet service provider, may be given thousands or hundreds of thousands based on the number of customers it may serve.

RESTRICTION:

- To simplify the handling of addresses, the Internet authorities impose three restrictions on classless address blocks:
    - The addresses in a block must be contiguous, one after another.
    - The number of addresses in a block must be a power of 2 (I, 2, 4, 8, ...).
    - The first address must be evenly divisible by the number of addresses.

MASK:

- A better way to define a block of addresses is to select any address in the block and the mask.
- As we discussed before, a mask is a 32-bit number in which the n leftmost bits are 1s and the 32 - n rightmost bits are Os.
- However, in classless addressing the mask for a block can take any value from 0 to 32.
- It is very convenient to give just the value of n preceded by a slash (CIDR notation).
- The address and the In notation completely define the whole block (the first address, the last address, and the number of addresses

In IPv4 addressing, a block of addresses can be defined as x.y.z.t /*n* in which x.y.z.t defines one of the addresses and the /*n* defines the mask.

NETWORK ADDRESSES:

- When an organization is given a block of addresses, the organization is free to allocate the addresses to the devices that need to be connected to the Internet.
- The first address in the class, however, is normally (not always) treated as a special address.
- The first address is called the network address and defines the organization network.
- It defines the organization itself to the rest of the world
- The organization network is connected to the Internet via a router.

- The router has two addresses. One belongs to the granted block; the other belongs to the network that is at the other side of the router.
- We call it as second address x.y.z.t/n because we do not know anything about the network it is connected to at the other side.
- All messages destined for addresses in the organization block (205.16.37.32 to 205.16.37.47) are sent, directly or indirectly, to x.y.z.t/n.
- We say directly or indirectly because we do not know the structure of the network to which the other side of the router is connected.

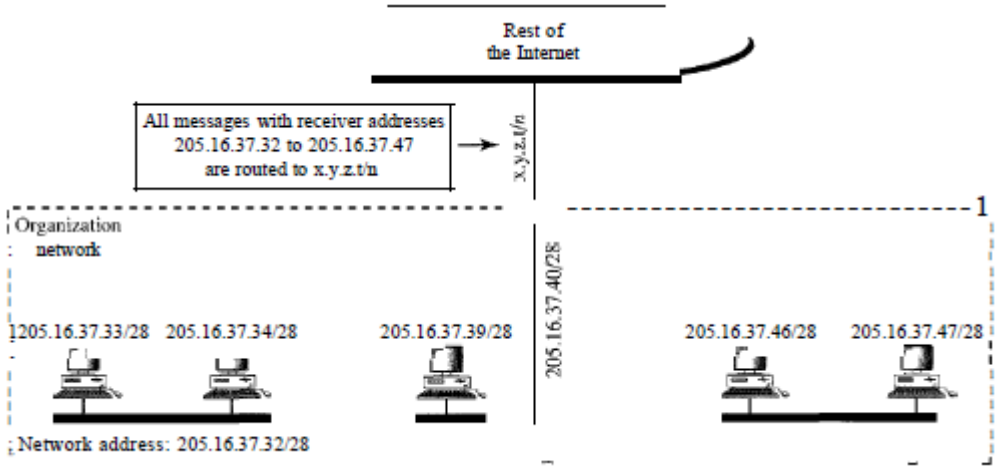

HIERARCHY

- IP addresses, like other addresses or identifiers can have levels of hierarchy.
- For example, a telephone network in North America has three levels of hierarchy.
- The leftmost three digits define the area code, the next three digits define the exchange, the last four digits define the connection of the local loop to the central office. Figure 19.5 shows the structure of a hierarchical telephone number

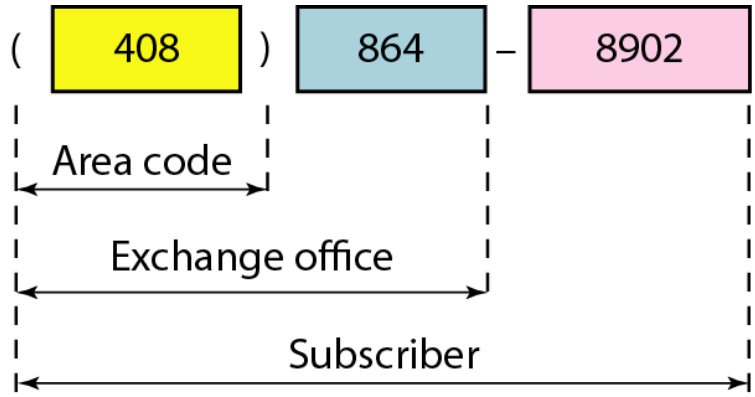

TWO-LEVEL HIERARCHY

- An IP address can define only two levels of hierarchy when not subnetted.
- The n leftmost bits of the address x.y.z.tJn define the network (organization network); the 32 – n rightmost bits define the particular host (computer or router) to the network.
- The two common terms are prefix and suffix.
- The part of the address that defines the network is called the prefix; the part that defines the host is called the suffix.

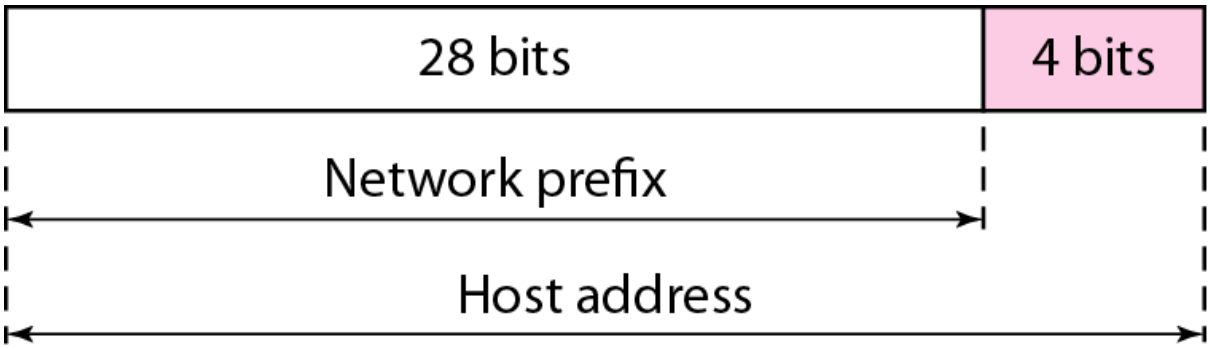

THREE-LEVEL HIERARCHY

- An organization that is granted a large block of addresses may want to create clusters of networks (called subnets) and divide the addresses between the different subnets.
- The rest of the world still sees the organization as one entity; however, internally there are several subnets.

- All messages are sent to the router address that connects the organization to the rest of the Internet; the router routes the message to the appropriate subnets.
- The organization, however, needs to create small subblocks of addresses, each assigned to specific subnets.
- The organization has its own mask; each subnet must also have its own.

Figure 19.7 ***Configuration and addresses in a subnetted network***

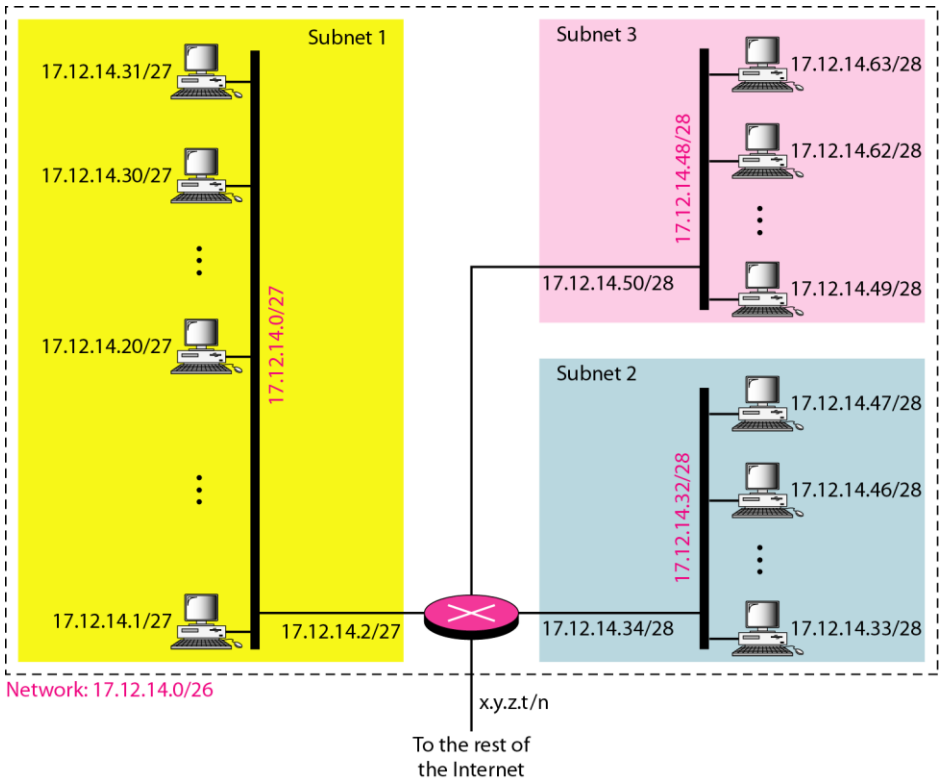

Figure 19.8 ***Three-level hierarchy in an IPv4 address***

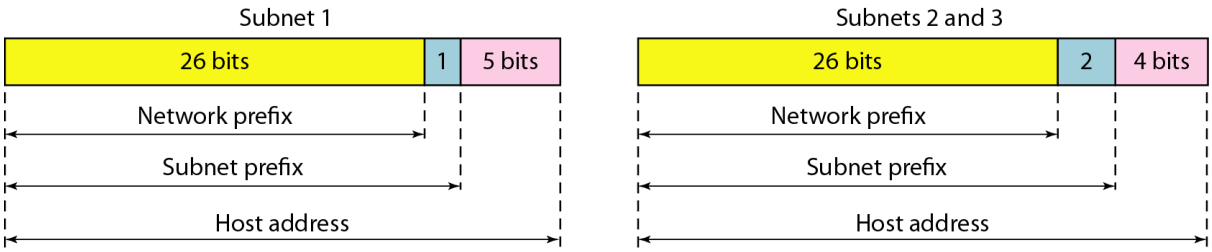

MORE LEVELS OF HIERARCHY:

The structure of classless addressing does not restrict the number of hierarchical levels.

An organization can divide the granted block of addresses into subblocks.

Each subblock can in turn be divided into smaller subblocks. And so on.

One example of this is seen in the ISPs. A national ISP can divide a granted large block into smaller blocks and assign each of them to a regional ISP. A regional ISP can divide the block received from the national ISP into smaller blocks and assign each one to a local ISP. A local ISP can divide the block received from the regional ISP into smaller blocks and assign each one to a different organization. Finally, an organization can divide the received blockand make several subnets out of it.

ADDRESS ALLOCATION:

- The next issue in classless addressing is address allocation. How are the blocks allocated?
- The ultimate responsibility of address allocation is given to a global authority called the Internet Corporation for Assigned Names and Addresses (**ICANN**).
- However, ICANN does not normally allocate addresses to individual organizations. It assigns a large block of addresses to an ISP.
- Each ISP, in turn, divides its assigned block into smaller subblocks and grants the subblocks to its customers.
- In other words, an ISP receives one large block to be distributed to its Internet users.
- This is called address aggregation: many blocks of addresses are aggregated in one block and granted to one ISP.

NETWORK ADDRESS TRANSLATION:

- The number of home users and small businesses that want to use the Internet is ever increasing.
- In the beginning, a user was connected to the Internet with a dial-up line, which means that she was connected for a specific period of time.
- An ISP with a block of addresses could dynamically assign an address to this user. An address was given to a user when it was needed. But the situation is different today.
- Home users and small businesses can be connected by an ADSL line or cable modem.
- In addition, many are not happy with one address; many have created small networks with several hosts and need an IP address for each host.
- With the shortage of addresses, this is a serious problem-solution to this problem is called network address translation (NAT).
- NAT enables a user to have a large set of addresses internally and one address, or a small set of addresses, externally.
- The traffic inside can use the large set; the traffic outside, the small set.
- To separate the addresses used inside the home or business and the ones used for the Internet, the Internet authorities have reserved three sets of addresses as private addresses, shown in below table.

Range			*Total*
10.0.0.0	to	10.255.255.255	2^{24}
172.16.0.0	to	172.31.255.255	2^{20}
192.168.0.0	to	192.168.255.255	2^{16}

- Any organization can use an address out of this set without permission from the Internet authorities.
- Everyone knows that these reserved addresses are for private networks.
- They are unique inside the organization, but they are not unique globally.
- No router will forward a packet that has one of these addresses as the destination address.
- The site must have only one single connection to the global Internet through a router that runs the NAT software. Below fig. shows a simple implementation of NAT.

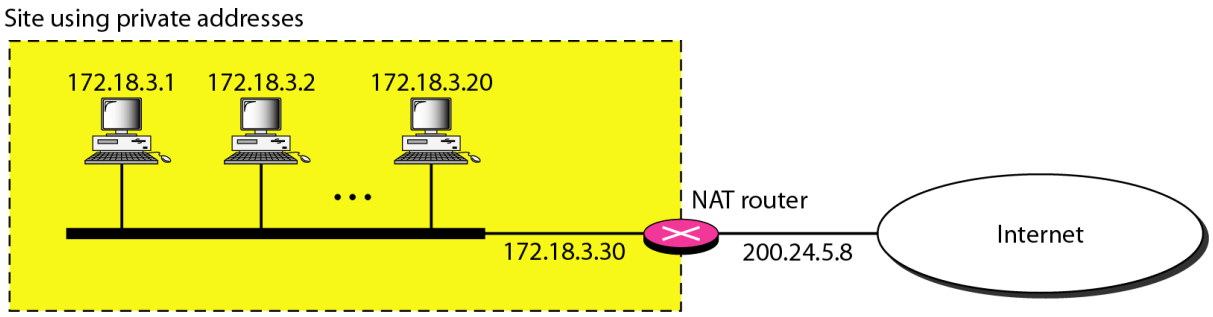

- The private network uses private addresses.
- The router that connects the network to the global address uses one private address and one global address.
- The private network is transparent to the rest of the Internet; the rest of the Internet sees only the NAT router with the address 200.24.5.8

ADDRESS TRANSLATION:

- All the outgoing packets go through the NAT router, which replaces the source address in the packet with the global NAT address.
- All incoming packets also pass through the NAT router, which replaces the destination address in the packet (the NAT router global address) with the appropriate private address.

Below fig. shows an example of address translation

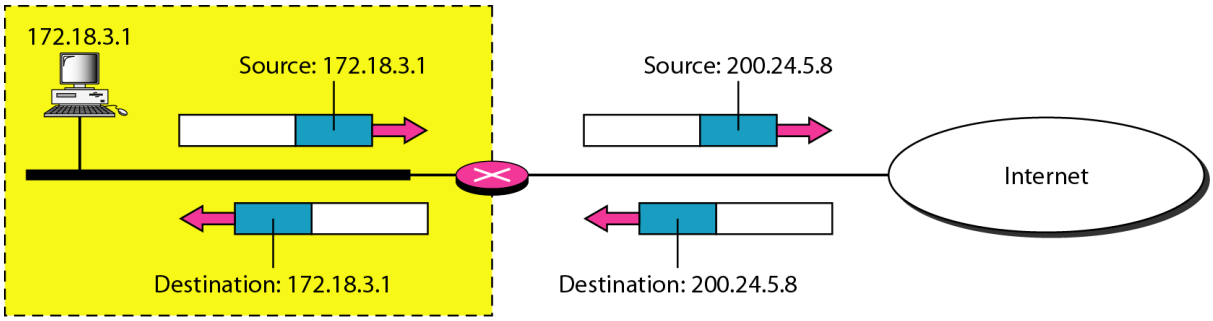

TRANSLATION TABLE:

- Translation table has only two columns: the private' address and the external address (destination address of the packet).
- When the router translates the source address of the outgoing packet, it also makes note of the destination address-where the packet is going.
- When the response comes back from the destination, the router uses the source address of the packet (as the external address) to find the private address of the packet.

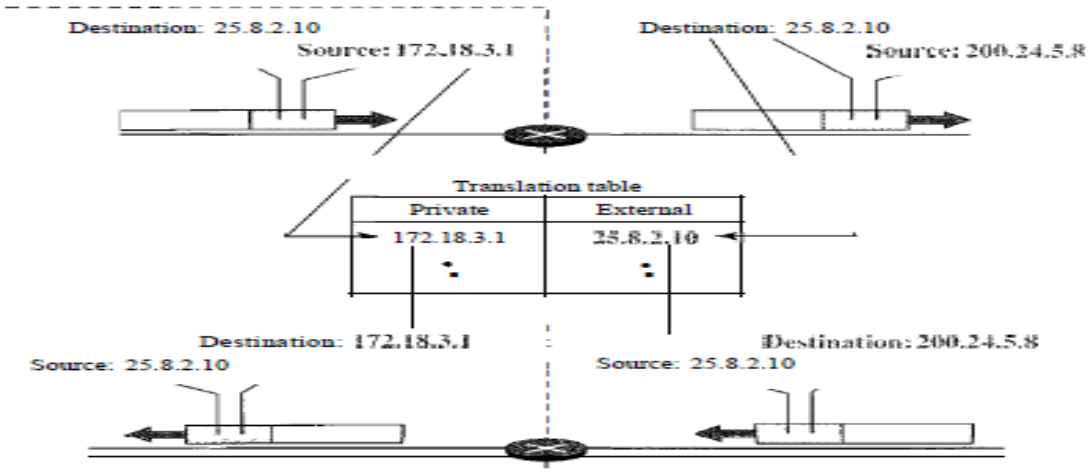

POOL OF IP ADDRESSES

- As NAT router has only one global address, only one private network host can access the same external host.

To remove this restriction, the NAT router uses a pool of global addresses

USING BOTH IP ADDRESSES AND PORT NUMBERS

- To allow a many-to-many relationship between private-network hosts and external server programs, we need more information in the translation table.
- For example, suppose two hosts with addresses 172.18.3.1 and 172.18.3.2 inside a private network need to access the HTTP server on external host 25.8.3.2.
- If the translation table has five columns, instead of two, that include the source and destination port numbers of the transport layer protocol, the ambiguity is eliminated.

Private Address	*Private Port*	*External Address*	*External Port*	*Transport Protocol*
172.18.3.1	1400	25.8.3.2	80	TCP
172.18.3.2	1401	25.8.3.2	80	TCP
. . .	. . .	. . .	. . .	. . .

NAT AND ISP

- An ISP that serves dial-up customers can use NAT technology to conserve addresses.
- For example, suppose an ISP is granted 1000 addresses, but has 100,000 customers. Each of the customers is assigned a private network address.
- The ISP translates each of the 100,000 source addresses in outgoing packets to one of the 1000 global addresses; it translates the global destination address in incoming packets to the corresponding private address. Below fig illustrates this concept.

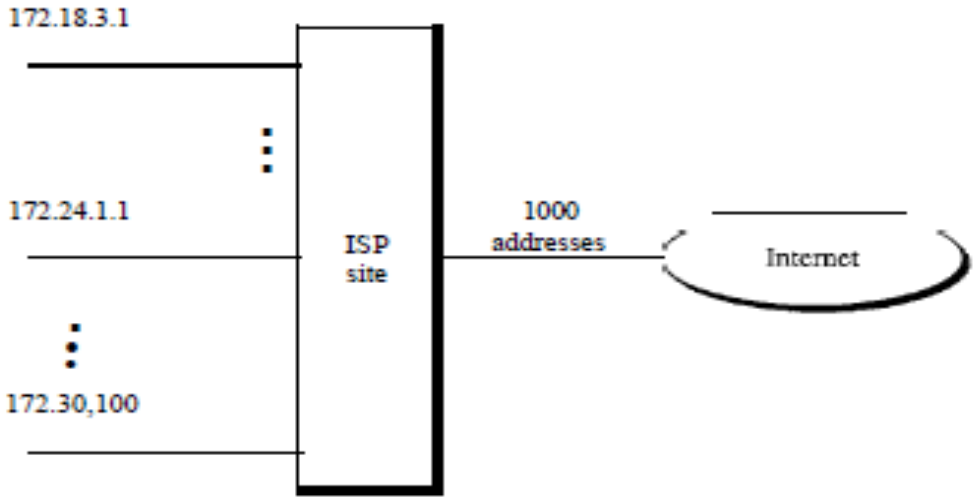

19-2 IPv6 ADDRESSES

Despite all short-term solutions, address depletion is still a long-term problem for the Internet. This and other problems in the IP protocol itself have been the motivation for IPv6.

Note: An IPv6 address is 128 bits long.

HEXADECIMAL COLON NOTATION

- To make addresses more readable, IPv6 specifies hexadecimal colon notation.
- In this notation, 128 bits is divided into eight sections, each 2 bytes in length.
- Two bytes in hexadecimal notation requires four hexadecimal digits.
- Therefore, the address consists of 32 hexadecimal digits, with every four digits separated by a colon, as shown in below fig.

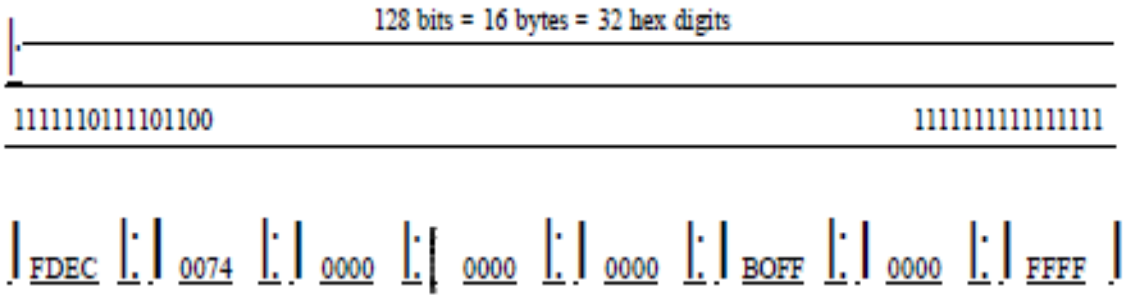

ABBREVIATION

- Although the IP address, even in hexadecimal format, is very long, many of the digits are zeros.
- In this case, we can abbreviate the address. The leading zeros of a section (four digits between two colons) can be omitted.
- Only the leading zeros can be dropped, not the trailing zeros (see below fig.)

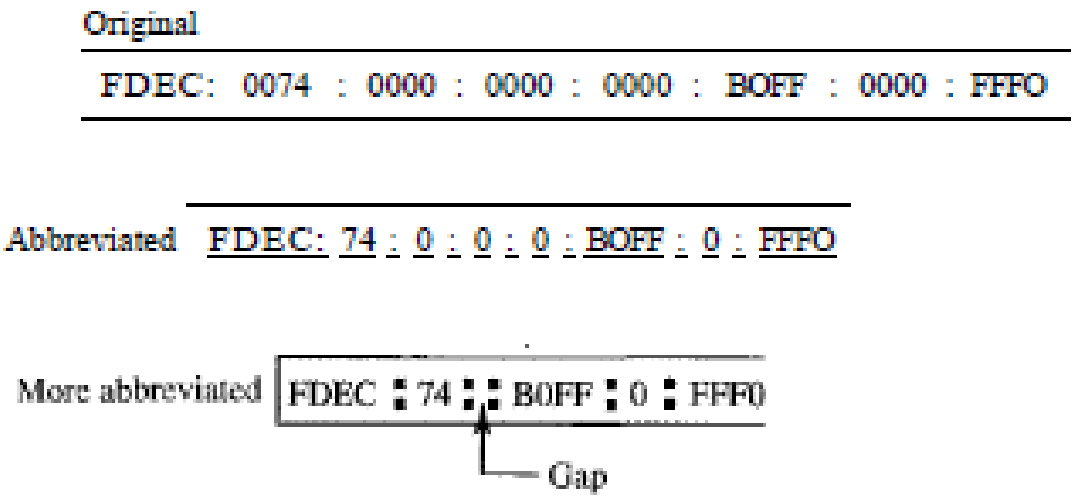

ADDRESS SPACE:

- IPv6 has a much larger address space; $\mathbf{2^{128}}$ addresses are available.
- The designers of IPv6 divided the address into several categories.
- A few leftmost bits, called the type prefix, in each address define its category.
- The type prefix is variable in length, but it is designed such that no code is identical to the first part of any other code.
- In this way, there is no ambiguity; when an address is given, the type prefix can easily be determined.
- Table below shows the prefix for each type of address.

- The third column shows the fraction of each type of address relative to the whole address space.

Type Prefix	*Type*	*Fraction*
00000000	Reserved	1/256
00000001	Unassigned	1/256
0000001	ISO network addresses	1/128
0000010	IPX (Novell) network addresses	1/128
0000011	Unassigned	1/128
00001	Unassigned	1/32
0001	Reserved	*1/16*
001	Reserved	1/8
010	Provider-based unicast addresses	*1/8*

Type Prefix	*Type*	*Fraction*
011	Unassigned	1/8
100	Geographic-based unicast addresses	1/8
101	Unassigned	1/8
110	Unassigned	1/8
1110	Unassigned	1116
11110	Unassigned	1132
1111 10	Unassigned	1/64
1111 110	Unassigned	1/128
11111110 a	Unassigned	1/512
1111 111010	Link local addresses	111024
1111 1110 11	Site local addresses	1/1024
11111111	Multicast addresses	1/256

2. INTER NETWORKING

- Physical and data link layers are jointly responsible for data delivery on the network from one node to the next node. Consider following figure, Assume a packet is being sent to D from A

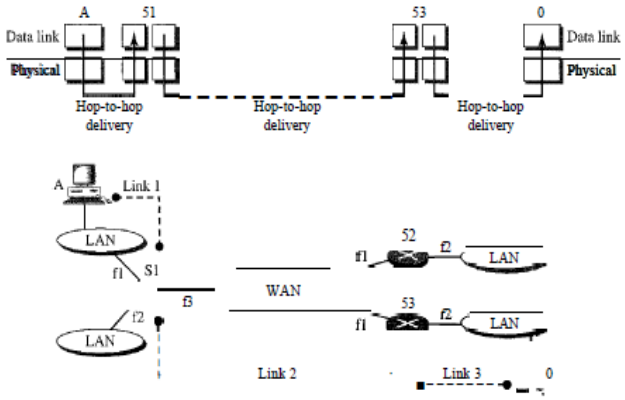

- How does interface of S3 know that the packet to be forwarded to f3
- This creates the necessity of network layer, which builds logical address in the packet, that gives routing information,

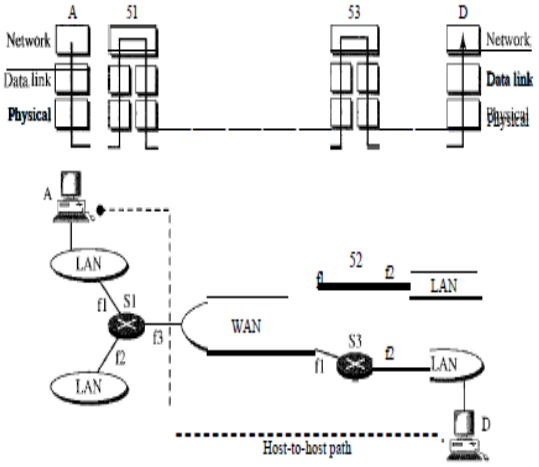

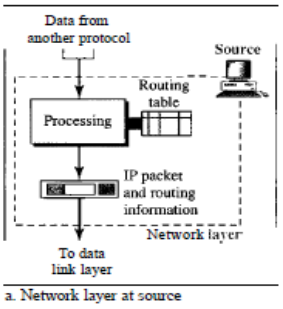

a. Network layer at source

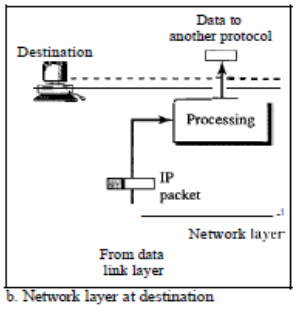

b. Network layer at destination

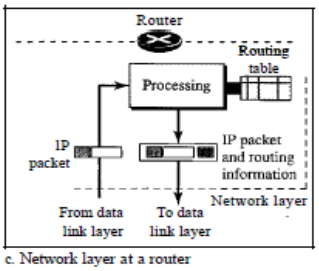

c. Network layer at a router

INTERNETWORKING AS DATAGRAM NETWORK

- The Internet, at the network layer, is a packet-switched network.
- In general, switching can be divided into three broad categories: circuit switching, packet switching, and message switching.
- Packet switching uses either the virtual circuit approach or the datagram approach.
- The Internet has chosen the datagram approach to switching in the network layer.
- It uses the universal addresses defined in the network layer to route packets from the source to the destination.

INTERNETWORKING: FRAGMENTATION

- ❑ Transparent fragmentation
    - o Strategy

- Gateway breaks large packet into fragments
- Each fragment addressed to same exit gateway
- Exit gateway does reassembly

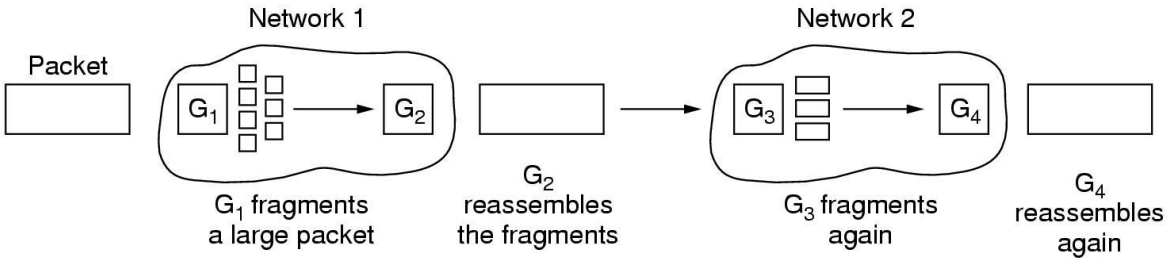

(a)

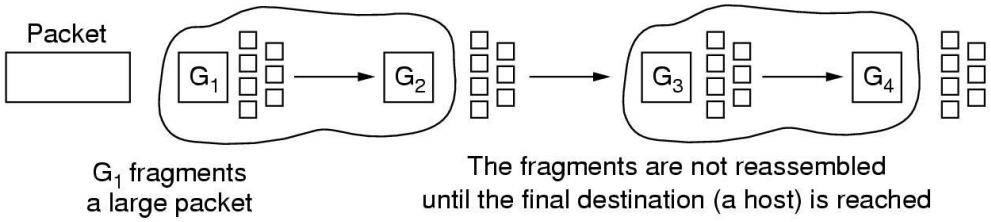

(b)

  - Simple, but some problems
    - Gateway must know when it has all pieces
    - Performance loss: all fragments through same gateway
    - Overhead: repeatedly reassemble and refragment
  - Example: ATM segmentation
- Non transparent fragmentation
  - Strategy
    - Gateway breaks large packet into fragments
    - Each fragment is forwarded to destination
  - problems

- Every host must be able to reassembly
- More headers

o Example: IP fragmentation

INTERNETWORKING AS CONNECTION LESS NETWORK

- In connection oriented circuit, there is a logical relationship between the packets, as they move in order across the network
- The order in which packets received is same as in order, how the source has emitted the packets.
- In conneetionless service, the network layer protocol treats each packet independently, with each packet having no relationship to any other packet.
- The packets in a message mayor may not travel the same path to their destination.
- This type of service is used in the datagram approach to packet switching. The Internet has chosen this type of service at the network layer.
- The reason for this decision is that the Internet is made of so many heterogeneous networks and it is not possible to create a connection from the source to the destination without knowing the nature of the networks.

3. TUNNELING: