

## Syllabus Content:

### 2.1 - IP addressing



Explain the use of IP addresses in the transmission of data over the internet  
Including:

- format of an IP address including IPv4 and IPv6
- use of sub-netting in a network
- how an IP address is associated with a device on a network
- difference between a public IP address and a private IP address and the implications for security
- difference between a static IP address and a dynamic IP address



Explain how a Uniform Resource Locator (URL) is used to locate a resource on the World Wide Web (WWW) and the role of the Domain Name Service (DNS)

## Internet Protocol (IP) Address

Each device on the internet is given a unique address known as the INTERNET PROTOCOL (IP) ADDRESS. This is a 32-bit number which is usually written in the form:

**109.108.158.1**

A home computer is given an IP address when it connects to the internet. This is assigned by the ISP and is unique for that particular internet session. The only IP addresses that remain fairly unchanged are web servers. An IP address can be used instead of typing in the full URL. For example: `http://109.108.158.1` would take you straight to the device corresponding to this address. IP addresses and MAC addresses


You will recall the term **MEDIA ACCESS CONTROL (MAC) ADDRESS** from earlier chapters. This is a unique number that identifies a device connected to the internet. So what is the difference between an IP address and a MAC address?

The IP address gives the location of a device on the internet and IP address is given to software, whereas the MAC address identifies the device connected to the internet and MAC address is given to hardware (Network Interface Card NIC).

An **Internet Protocol address (IP address)** is a numerical label assigned to each device (e.g., computer, printer) participating in a Computer network.

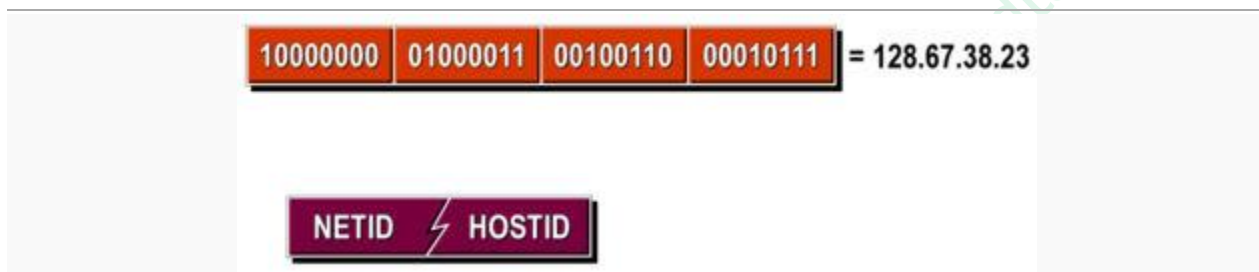
## IPv4 addresses

An IPv4 address (dotted-decimal notation)



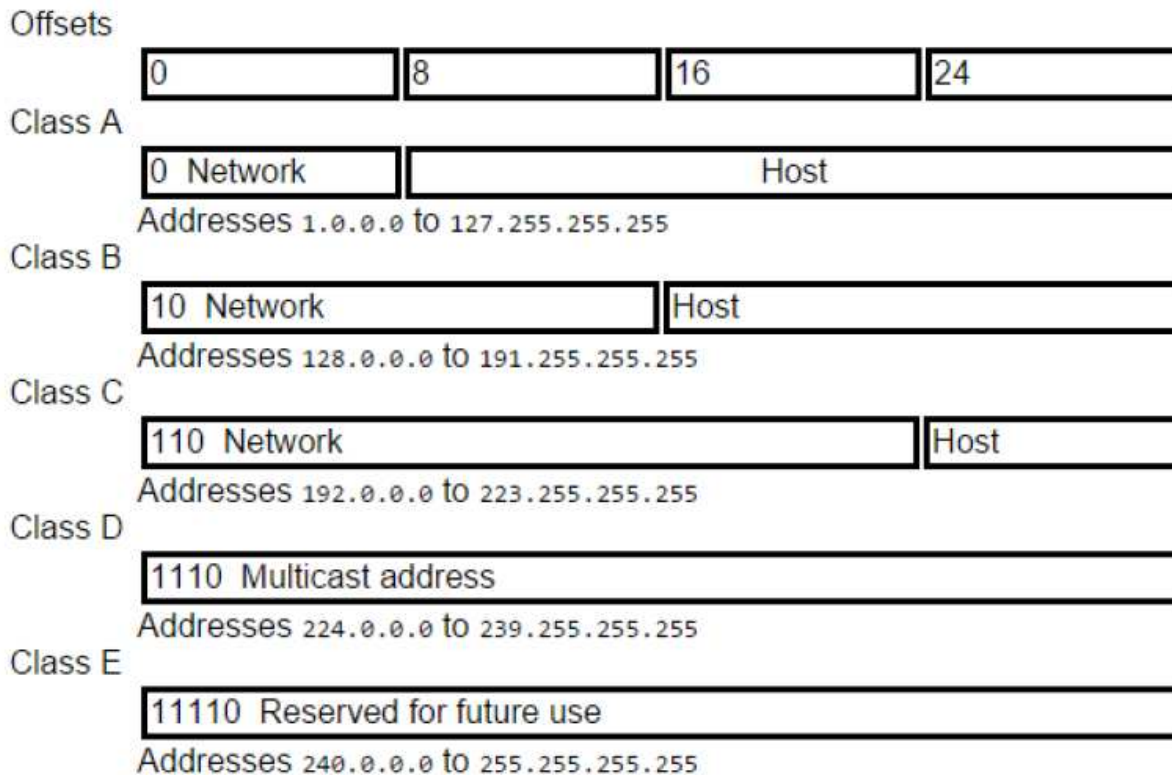
- In IPv4 an address consists of 32 bits which limits the address space to 4294967296 ( $2^{32}$ ) possible unique addresses.
- IPv4 reserves some addresses for special purposes such as private networks (~18 million addresses) or multicast addresses (~270 million addresses).
- IPv4 addresses are canonically represented in dot-decimal notation, which consists of four decimal numbers, each ranging from 0 to 255, separated by dots, e.g., 172.16.254.1. Each part represents a group of 8 bits (octet) of the address.
- In some cases of technical writing, IPv4 addresses may be presented in various hexadecimal, octal, or binary representations.

## IP address



## IPv4 Addresses:

- Each device attached to a TCP/IP-based network must be given a unique address. These addresses are carried in the IP packet to identify the source and destination hosts.
- IP defines 32 bit (four octet) addresses, which permit over four billion hosts to attach to a single internetwork.
- IP addresses are represented in a format known as dotted decimal, where each octet is represented by its decimal equivalent and the four numbers are separated by a dot (.). With eight bits, you can represent numbers between 0 and 255. Therefore, the IP address represented in binary as 10000000 01000011 00100110 00010111 (hexadecimal 80-43-26-17) is written (and spoken) in dotted decimal as 128.67.38.23.



## NET & HOSTID:

- Each IP address has two components: a network identifier (NETID) and a host identifier (HOSTID).
- The NETID identifies the specific network to which the host is attached.
- The HOSTID uniquely identifies a host within that network. This distinction is important because routers route to a given NETID and don't care about the HOSTID. IP actually permits the boundary between the NETID and the HOSTID to shift. By extending the NETID and shrinking the HOSTID for a given network, one network can be partitioned into multiple subnetworks, a process known as subnetting. To indicate where the new boundary has been set for a given network, a subnet mask is used.

IP does not permit the NETID or HOSTID to be all ones or all zeros. All ones means broadcast and can be used for all networks or all hosts. For example, the IP address 128.17.255.255 has NETID 128.17 and HOSTID 255.255. It means all hosts on the network with NETID 128.17. To any IP device, 128.17.0.0 refers to the entire 128.17 network, regardless of HOSTID.

In classful addressing, an IP address of class A, B and C is divided into netid and hostid. The netid determines the network address while the hostid determines the host connected to that network.

	Byte 1	Byte 2	Byte 3	Byte 4
Class A	Netid	Hostid		
Class B	Netid		Hostid	
Class C	Netid			Hostid
Class D	Multicast Address			
Class E	Reserved for future use			

## Blocks and Hosts

Each class is divided into blocks. The numbers of block in a class can be calculated by the number of bits in the netid.

### CLASS A

Class A has 1 byte (8 bits) **netid** and from the binary notation we see that Class A address starts with **0** so there are total 7 bits that can be changed out of 8.

**Therefore total number of blocks in Class A =  $2^7 = 128$**

There are 3 bytes (24 bits) for **hostid** in Class A so total number of host in each block =  $2^{24} = 16,777,216$

So total number of addresses in Class A = No. of Blocks in Class A x No. of Hosts in each block of Class A

$$= 128 \times 16,777,216$$

$$= 2,147,483,648$$

This is 50% of the total addresses in IPv4.

1<sup>st</sup> block of Class A has the netid 0

And the host id is between

**0.0.0.0**

...

**0.255.255.255**



Similarly, 2<sup>nd</sup> block of Class A has the netid 1

And the host id is between

**1.0.0.0**

**1.255.255.255**

	Netid 0	Netid 1		Netid 127
Class A	0.0.0.0 to 0.255.255.255	1.0.0.0 to 1.255.255.255	...	127.0.0.0 to 127.255.255.255
	Block 1	Block 2		Block 128

## CLASS B

Class B has 2 bytes (16 bits) netid and from the binary notation we see that Class B address starts with 10, so there are total 14 bits that can be changed out of 16.

**Therefore total number of blocks in Class B =  $2^{14} = 16,384$**

There are 2 bytes (16 bits) for hostid in Class B so total number of host in each block =  $2^{16} = 65,536$

So total number of addresses in Class B = No. of Blocks in Class B x No. of Hosts in each block of Class B

= 16,384 x 65,536

= 1,073,741,824

This is 25% of the total addresses in IPv4.

1<sup>st</sup> block of Class B has the netid 128.0

And the host id is between

**128.0.0.0**

...

**128.0.255.255**

Similarly, 2<sup>nd</sup> block of Class B has the netid 128.1

And the host id is between

**128.1.0.0**

**128.1.255.255**

	Netid 128.0	Netid 128.1		Netid 191.255
Class B	128.0.0.0 to 128.0.255.255	128.1.0.0 to 128.1.255.255	...	191.255.0.0 to 191.255.255.255
	Block 1	Block 2		Block 16,384

## CLASS C

Class C has 3 bytes (24 bits) netid and from the binary notation we see that Class C address starts with 110, so there are total 21 bits that can be changed out of 24.

**Therefore total number of blocks in Class C =  $2^{21} = 2,097,152$**

There is 1 byte (8 bits) for hostid in Class C so total number of host in each block =  $2^8 = 256$

So total number of addresses in Class C = No. of Blocks in Class C x No. of Hosts in each block of Class C

=  $2,097,152 \times 256$

= 536,870,912

This is 12.5% of the total addresses in IPv4.

1<sup>st</sup> block of Class C has the netid 192.0.0

And the host id is between

**192.0.0.0**

...

**192.0.0.255**

Similarly, 2<sup>nd</sup> block of Class B has the netid 192.0.1

And the host id is between

**192.0.1.0**

**192.0.1.255**

	Netid 192.0.0	Netid 192.0.1		Netid 223.255.255
--	---------------	---------------	--	-------------------

Class C	192.0.0.0 to 192.0.0.255	192.0.1.0 to 192.0.1.255	...	223.255.255.0 to 223.255.255.255
	Block 1	Block 2		Block 2,097,152

## CLASS D

It consists of a single block. It is designed for multicasting.

Class D	224.0.0.0 to 239.255.255.255
	Single block of 268,435,456 addresses

## CLASS E

It also consists of a single block. It is reserved for future use.

Class E	240.0.0.0 to 255.255.255.255
	Single block of 268,435,456 addresses

- There are 128 blocks in Class A so only **128 organizations can be assigned Class A** address, but each block has **2,147,483,648 hosts** which mean the **organization needs to be really huge to consume all the addresses in the block.**
- There are **16,384 blocks in Class B** so total **16,384 organizations** can be assigned Class B address and there are **65,536 hosts in each block** which mean each **organization must be quite large to consume all the addresses in the block.**
- There are **2,097,152 blocks in Class C** so total **2,097,152 organizations** can be assigned Class C address but there are only **256 hosts in each block** which

Class	First Octet Range	Max Hosts	Format
A	1-126	16M	
B	128-191	64K	
C	192-223	254	
D	224-239	N/A	
E	240-255	N/A	

mean the organization must be quite small to use this class.

This is the reason why a lot of addresses are wasted because of classful addressing.

- Class A address can be allocated to only 128 organizations which sound quite ok but each of these organizations must have 16,777,216 machines (host) which is not impossible.
- While Class C address can be assigned to a lot of organizations (2,097,152 to be precise) but if an organization gets Class C then it cannot have more than 256 computers (host).

<i>Class</i>	<i>Number of Blocks</i>	<i>Block Size</i>	<i>Application</i>
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

## Private IP Addresses

Because the IP address space is relatively small, much work has been done to conserve that address space. Enter NAT (or NAPT). These technologies make it possible for a consumer (corporate or residential) to be allocated as few as a single address from their ISP and use it to support multiple systems within their network. The addresses actually used by the internal systems are translated by the NAT or NPAT device before being forwarded onto the Internet.

To prevent conflict, the IANA has allocated three address blocks for these private networks.

These addresses, which may not be used across the public Internet, include:

- **10.0.0.0** through **10.255.255.255**
- **172.16.0.0** through **172.31.255.255**
- **192.168.0.0** through **192.168.255.255**

A fourth block is reserved for Automatic Private IP Addressing (APIPA):

- **169.254.0.0** through **169.254.255.255**



Class	Binary	Dotted-Decimal	CIDR
A	11111111 00000000 00000000 00000000	255.0.0.0	/8
B	11111111 11111111 00000000 00000000	255.255.0.0	/16
C	11111111 11111111 11111111 00000000	255.255.255.0	/24

**Table: Default masks for classful addressing**

The mask can help us to find the netid and the hostid. The mask for a class A address has eight 1 s, which means the first 8 bits of any address in class A define the netid; the next 24 bits define the hostid.

## Private addresses:

Early network design, when global end-to-end connectivity was envisioned for communications with all Internet hosts, intended that IP addresses be uniquely assigned to a particular computer or device. However, it was found that this was not always necessary as private networks developed and public address space needed to be conserved.

Computers not connected to the Internet, such as factory machines that communicate only with each other via TCP/IP, need not have globally unique IP addresses. Three non-overlapping ranges of IPv4 addresses for private networks were reserved in RFC 1918. These addresses are not routed on the Internet and thus their use need not be coordinated with an IP address registry.

Today, when needed, such private networks typically connect to the Internet through network address translation (NAT).

### IANA-reserved private IPv4 network ranges










	Start	End	No. of addresses
24-bit block (/8 prefix, 1 × A)	10.0.0.0	10.255.255.255	16777216
20-bit block (/12 prefix, 16 × B)	172.16.0.0	172.31.255.255	1048576
16-bit block (/16 prefix, 256 × C)	192.168.0.0	192.168.255.255	65536

Any user may use any of the reserved blocks. Typically, a network administrator will divide a block into subnets; for example, many home routers automatically use a default address range of 192.168.0.0 through 192.168.0.255 (192.168.0.0/24).

## Public address

A public IP address, in common parlance, is a globally routable unicast IP address, meaning that the address is not an address reserved for use in private networks, such as those reserved by RFC 1918, or the various IPv6 address formats of local scope or site-local scope, for example for link-local addressing. Public IP addresses may be used for communication between hosts on the global Internet.

### Difference between Public IP addresses and Private IP addresses:

Public IP addresses	Private IP addresses
 Public address can be reached across the Internet.	 Private address can only be reached internally/through the LAN/Intranet // private address cannot be reached across the Internet.
	 NAT (Network Address Translation) is necessary for a private IP address to access the Internet directly.
 Public addresses are provided by ISP / assigned by InterNIC	 Private addresses are assigned by the router (of the network concerned).
 Public addresses are unique (to the Internet)	 Private addresses (are unique within their network, but) can be duplicated within other (discrete) networks.
 IP addresses from the private address space are never assigned as public.	 10.0.0.1 to 10.255.255.254 and 172.16.0.1 to 172.31.255.254 and 192.168.0.1.to 192.168.255.254 form the private address space

Indicate for each IP address whether it is valid or invalid and give a reason.

Address	Denary / Hexadecimal	Valid or Invalid	Reason
3.2A.6AA.BBBB	Hexadecimal		
2.0.255.1	Denary		
6.0.257.6	Denary		
A.78.F4.J8	Hexadecimal		

9608/11/M/J/16

### IPv6 addresses

An IPv6 address (in hexadecimal)

**2001:0DB8:AC10:FE01:0000:0000:0000:0000**



**2001:0DB8:AC10:FE01::** Zeroes can be omitted



1000000000000001:0000110110111000:1010110000010000:1111111000000001:

0000000000000000:0000000000000000:0000000000000000:0000000000000000

The rapid exhaustion of IPv4 address space prompted the Internet Engineering Task Force (IETF) to explore new technologies to expand the addressing capability in the Internet. The permanent solution was deemed to be a redesign of the Internet Protocol itself. This new generation of the Internet Protocol was eventually named Internet Protocol Version 6 (IPv6) in 1995.<sup>[3][4]</sup>

The address size was increased from 32 to 128 bits (16 octets), thus providing up to  $2^{128}$  (approximately  $3.403 \times 10^{38}$ ) addresses. This is deemed sufficient for the foreseeable future.

The intent of the new design was not to provide just a sufficient quantity of addresses, but also redesign routing in the Internet by more efficient aggregation of subnetwork routing prefixes.

This resulted in slower growth of routing tables in routers. The smallest possible individual allocation is a subnet for  $2^{64}$  hosts, which is the square of the size of the entire IPv4 Internet. At these levels, actual address utilization rates will be small on any IPv6 network segment.

The new design also provides the opportunity to separate the addressing infrastructure of a network segment, i.e. the local administration of the segment's available space, from the addressing prefix used to route traffic to and from external networks.

IPv6 has facilities that automatically change the routing prefix of entire networks, should the global connectivity or the routing policy change, without requiring internal redesign or manual renumbering.

The large number of IPv6 addresses allows large blocks to be assigned for specific purposes and, where appropriate, to be aggregated for efficient routing. With a large address space, there is no need to have complex address conservation methods as used in CIDR.

All modern desktop and enterprise server operating systems include native support for the IPv6 protocol, but it is not yet widely deployed in other devices, such as residential networking routers, voice over IP (VoIP) and multimedia equipment, and network peripherals.

## CIDR

### (Classless Inter-Domain Routing or supernetting)

CIDR (Classless Inter-Domain Routing, sometimes called *supernetting*) is a way to allow more flexible allocation of Internet Protocol (IP) addresses than was possible with the original system of IP address classes. As a result, the number of available Internet addresses was greatly

increased, which along with widespread use of network address translation (NAT), has significantly extended the useful life of IPv4.

Originally, IP addresses were assigned in four major address classes, A through D. Each of these classes allocates one portion of the 32-bit IP address format to identify a network gateway -- the first 8 bits for class A, the first 16 for class B, and the first 24 for class C. The remainder identify hosts on that network -- more than 16 million in class A, 65,535 in class B and 254 in class C. (Class D addresses identify multicast domains.)

To illustrate the problems with the class system, consider that one of the most commonly used classes was Class B. An organization that needed more than 254 host machines would often get a Class B license, even though it would have far fewer than 65,534 hosts. This resulted in most of the block of addresses allocated going unused.

The inflexibility of the class system accelerated IPv4 address pool exhaustion. With IPv6, addresses grow to 128 bits, greatly expanding the number of possible addresses on the Internet. The transition to IPv6 is slow, however, so IPv4 address exhaustion continues to be a significant issue.

**CIDR** reduced the problem of wasted address space by providing a new and more flexible way to specify network addresses in routers. CIDR lets one routing table entry represent an aggregation of networks that exist in the forward path that don't need to be specified on that particular gateway. This is much like how the public telephone system uses area codes to channel calls toward a certain part of the network. This aggregation of networks in a single address is sometimes referred to as a **supernet**.

Using CIDR, each IP address has a *network prefix* that identifies either one or several network gateways. The length of the network prefix in IPv4 CIDR is also specified as part of the IP address and varies depending on the number of bits needed, rather than any arbitrary class assignment structure. A destination IP address or route that describes many possible destinations has a shorter prefix and is said to be less specific. A longer prefix describes a destination gateway more specifically.

Routers are required to use the most specific, or longest, network prefix in the routing table when forwarding packets. (In IPv6, a CIDR block always gets 64 bits for specifying network addresses.)

A CIDR network address looks like this under IPv4:

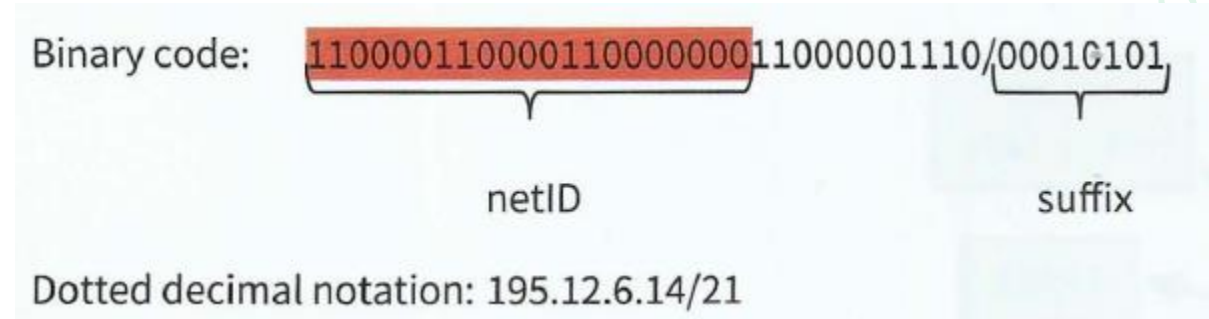
**192.30.250.0/18**

The "**192.30.250.0**" is the network address itself and the "**18**" says that the **first 18 bits are the network part** of the address, leaving the last **14 bits for specific host** addresses.

**Classless inter-domain routing (CIDR)**



The simple method used to achieve this is to add an 8-bit suffix to the address that specifies the number of bits for the netID. If, for instance, we define the suffix as 21, that means that 21 bits are used for the netID and there are 11 bits remaining (of a 32-bit address) to specify hostIDs allowing **2<sup>11</sup>** i.e. **2048**, hosts. One example of an IP address using this scheme is shown in Figure 2.06. The 21 bits representing the netID have been highlighted. The remaining 11 bits represent the hostID which would therefore have the binary value **11000001110**.



### Sub netting

A quite different approach, sub-netting, allows further structure in the addressing. To illustrate an example of this we can consider a medium-sized organisation with about 150 employees each with their own computer workstation. Let's assume that there are six individual department LANs and one head-office LAN. Figure 2.07 shows a schematic diagram of how the LANs would be connected to the Internet if the original scheme were used. The organization would need seven individual Class C netIDs. Each of these would point to one of the LAN gateways (which have to function as routers). Each netID would be associated with 256 hosts so an organisation with just 150 computer workstations would leave 1642 IP addresses unused and unavailable for use by any other organisation.

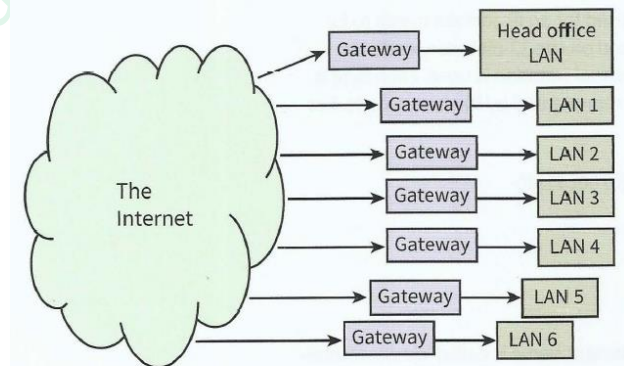


Figure 2.07 Connecting LANs using the original classful IPv4 scheme

The sub-netting solution for this organisation would require allocating just one Class C netID. For example, the IP addresses allocated might be **194.10.9.0** to **194.10.9.255** where the **netID comprises the first three bytes**, represented by the decimal values **194, 10 and 9**.

The sub-netting now works by having a defined structure for the **256 codes constituting the host ID**. A sensible solution for this organisation is to use the **top three bits as a code for the individual LANs** and the **remaining five bits as codes for the individual workstations**. Figure

2.08 shows a schematic diagram of this arrangement.

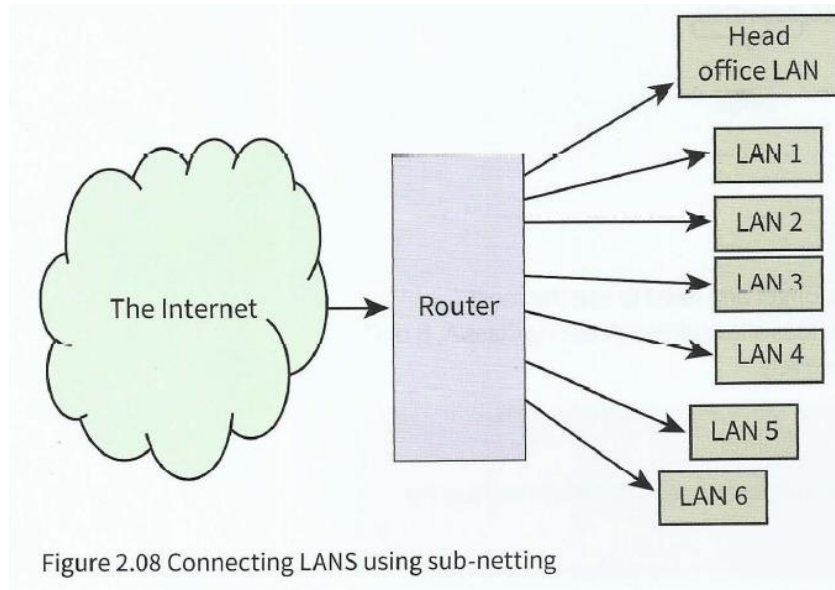


Figure 2.08 Connecting LANS using sub-netting

On the Internet, all of the allocated IP addresses have a netID pointing to the router. The router then has to interpret the hostID to direct the transmission to the appropriate host on one of the LANS. For example:

- hostID code 00001110 could be the address for workstation 14 on the head office LAN (LAN 000).
- hostID code 01110000 would be the address for workstation 16 on LAN 3 (LAN 011).

A **subnet mask** is a screen of numbers used for routing traffic within a **subnet**. Once a packet has arrived at an organization's gateway or connection point with its unique **network** number, it can be routed to its destination within the organization's internal gateways using the **subnet** number.

### What is a DNS server?

The Domain Name System (DNS) is a standard technology for managing public names of Web sites and other Internet domains. DNS technology allows you to type names into your Web browser like "*soundcloud.com*" and your computer to automatically find that address on the Internet. A key element of the DNS is a worldwide collection of *DNS servers*.

**DNS:** It's like your computer's GPS for the Internet. Its basic job is to turn a user-friendly domain name like "**google.com**" into an Internet Protocol (IP) address like **64.233.167.104** that computers use to identify each other on the network. Computers and other network devices on the Internet use an IP address to route your request to the site you're trying to reach. This is similar to dialing a phone number to connect to the person you're trying to call.

Thanks to DNS, though, you don't have to keep your own address book of IP addresses. Instead, you just connect through a domain name server, also called a DNS server or name server, which manages a massive database that maps domain names to IP addresses.

A **DNS server** is any computer registered to join the Domain Name System. A DNS server runs special purpose networking software, features a public IP address, and contains a database of network names and addresses for other Internet hosts. A DNS server is similar to looking up contacts on your phone, to call a contact, you simply look up that

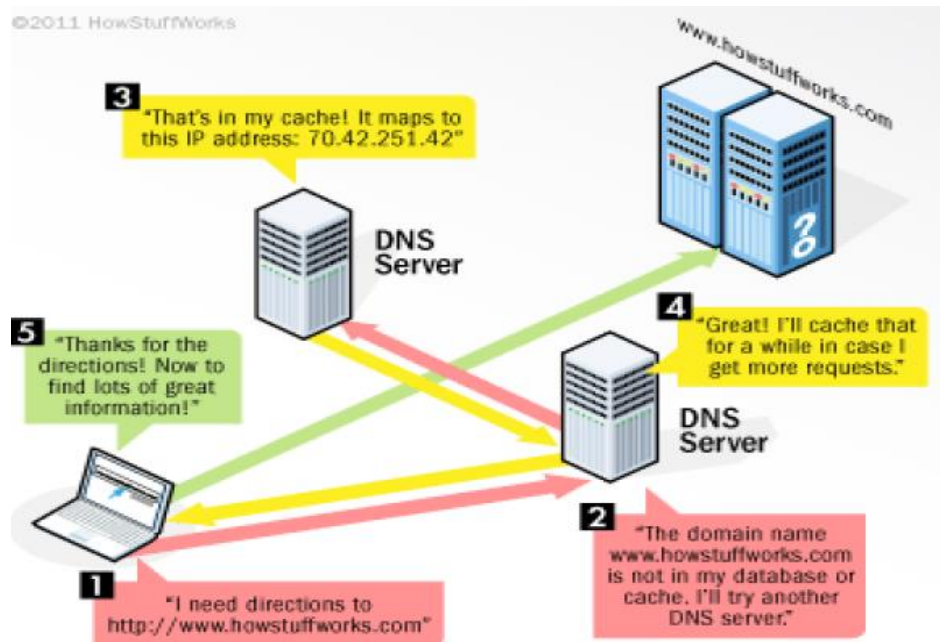
Person's name, but that name is of no use to the phone itself, it has to look up the contact number and dial that. Simply speaking, both systems translate the website/contact name into an IP address or phone number.

## DNS Root Servers

DNS servers communicate with each other using private network protocols. All DNS servers are organized in a hierarchy. At the top level of the hierarchy, so-called "*root servers*" store a complete database of Internet domain names and their corresponding IP addresses. The Internet employs 13 root servers that have become somewhat famous for their special role. Maintained by various independent agencies, the servers are aptly named A, B, C and so on up to M. Ten of these servers reside in the United States, one in Japan, one in London, UK and one in Stockholm, Sweden.

## How DNS Works

The DNS is a distributed system, meaning that only the 13 root servers contain the complete database of names and addresses. All other DNS servers are installed at lower levels of the hierarchy and maintain only certain pieces of the overall database.





Most lower level DNS servers are owned by businesses or Internet Service Providers (ISPs). For example, Google maintains various DNS servers around the world that manage the google.com, google.co.uk, and other domains.

Your ISP also maintains DNS servers as part of your Internet connection setup. DNS networking is based on the client / server architecture. Your Web browser functions as a DNS client (also called *DNS resolver*) and issues requests to your Internet provider's DNS servers when navigating between Web sites.

When a DNS server receives a request not in its database (such as a geographically distant or rarely visited Web site), it temporarily transforms from a server to a DNS client. The server automatically passes that request to another DNS server or up to the next higher level in the DNS hierarchy as needed.

Eventually the request arrives at a server that has the matching name and IP address in its database (all the way to the root level if necessary), and the response flows back through the chain of DNS servers to your computer.

### DNS and the World Wide Web

All public Web sites run on servers connected to the Internet with public IP addresses . The Web servers at About.com, for example, have addresses like 207.241.148.80. Although people can type address information like http://207.241.148.80/ into their Web browser to visit sites, being able to use proper names like http://www.about.com/ is much more practical.

The Internet utilizes DNS as a worldwide name resolution service for public Web sites. When someone types a site's name into their browser, DNS looks up the corresponding IP address for that site, the data required to make the desired network connections between Web browsers and Web servers.



Answers:

Address	Denary / Hexadecimal	Valid or Invalid	Reason
3.2A.6AA.BBBB	Hexadecimal	Invalid	<p>One point from:</p> <ul style="list-style-type: none"> <li>This is more than <u>32 bits</u></li> <li><u>6AA /BBBB</u> in Hex is bigger than <u>FF / 255</u> in denary</li> <li><u>6AA / BBBB</u> uses more than 8 bits / a byte</li> <li>The third / fourth group is bigger than <u>FF / 255</u> in denary</li> <li>The third / fourth group uses more than 8 bits / a byte</li> </ul>
2.0.255.1	Denary	Valid	There are 4 bytes, each 255 or below // All the values are in the range 0 - 255
6.0.257.6	Denary	Invalid	<u>257</u> is above 255 // The third group is above 255
0A.78.F4.J8	Hexadecimal	Invalid	J is not a valid hexadecimal digit // J8 is not a valid Hex number

One mark for each combination of valid or invalid and the reason.

Computer Science 9618