IPv4 Address

What is IPv4 address?

- 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.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.
- On the other hand, if a device operating at the network layer has *m* connections to the Internet, it needs to have *m* addresses.
- 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.

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 2^N because each bit can have two different values (0 or 1) and N bits can have 2^N values.

- IPv4 uses 32-bit addresses, which means that the address space is 2^{32} 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

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

Notations

- 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

Notations:

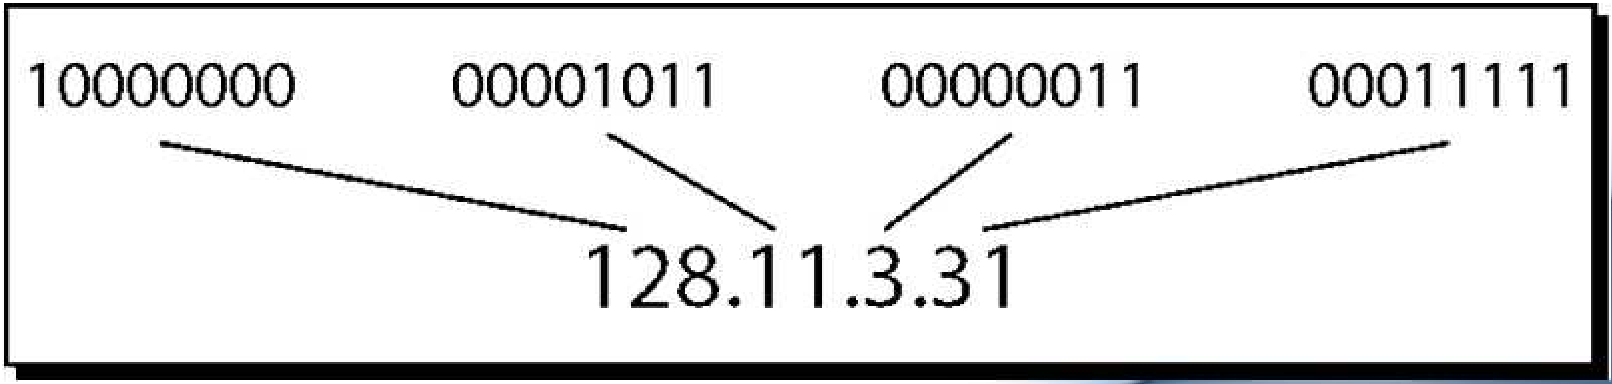

Classful Addressing

- IPv4 addressing, at its inception, used the concept of classes. This architecture is called classful addressing. This scheme is becoming obsolete.
- 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. 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.

Classfull Addressing

	First byte	Second byte	Third byte	Fourth byte
Class A	0			
Class B	10			
Class C	110			
Class D	1110			
Class E	1111			

a. Binary notation

	First byte	Second byte	Third byte	Fourth byte
Class A	0–127			
Class B	128–191			
Class C	192–223			
Class D	224–239			
Class E	240–255			

b. Dotted-decimal notation

Classes and Blocks

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

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. 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 contiguous 1s followed by contiguous 0s.

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

- This notation is also called slash notation or Classless Interdomain Routing (CIDR) notation.

Subnetting and Supernetting

- Why is it Required?
- Has it solved any Problem?

Subnetting

- During the era of classful addressing, subnetting was introduced.

- If an organization was granted a large block in class A or B, 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 1s in the mask.

Supernetting

- The time came when most of the class A and class B addresses were depleted; however, there was still a huge demand for midsize blocks. The size of a class C block with a maximum number of 256 addresses did not satisfy the needs of most organizations. Even a midsize organization needed more addresses. One solution was 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. For example, 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 2^{32} 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 Addressing

- 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.

- For 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 3 restrictions on the 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 (1, 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 0s.
- 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 */n* notation completely define the whole block (the first address, the last address, and the number of addresses).
- First Address The first address in the block can be found by setting the 32 - *n* rightmost bits in the binary notation of the address to 0s.