IPv4 Addressing

IP addresses enable networked devices to communicate by providing unique identifiers for the devices and the networks on which they are located. Understanding how IP addresses are constructed and how they should be assigned is an essential part of network administration.

An IP address (in IPv4) is a 32-bit value that contains both a network identifier and a host identifier. The address is notated using four decimal numbers ranging from 0 to 255 (this is equivalent to binary values ranged from 00000000 to 11111111), separated by periods, as in 192.168.1.44 . This is known as dotted decimal notation. Each of the 8-bit values that make up an IP address is called an octet.

Each IP addresses represents a NIC, so devices that have more than one NIC such as routers, servers or even PCs can have more than one IP address for each device.

IP addresses are assigned either manually by a network administrator or automatically by a Dynamic Host Configuration Protocol (DHCP) Server. The IP address of a host is a "logical address" meaning it can be changed while the MAC address is a 48-bit "physical address" that is burned into the NIC and cannot change unless the NIC is replaced. The combination of the logical IP address and the physical MAC address helps to route packets to their proper destination.

IP Address Assignments

Unlike hardware addresses, which are hard-coded into network interface adapters at the factory, network administrators must assign IP addresses to the systems on their networks. It is essential for each network interface adapter to have its own unique IP address; when two systems have the same IP address, they cannot communicate with the network properly.

IP addresses consist of two parts: a network identifier and a host identifier. All of the network interface adapters on a particular network (or subnetwork) have the same network identifier but different host identifiers. For systems that are on the Internet, the Internet Assigned Numbers Authority (IANA) assigns network identifiers to ensure that there is no address duplication on the Internet. When an organization registers its network, it is assigned a network identifier. It is then up to the network administrators to assign unique host identifiers to each of the systems on that network.

Because IP addresses consist of four octets separated by dots, one, two, or three of these octets may be used to identify the network number. Similarly, the rest of these octets may be used to identify the host portion of an IP address. This process is organized by dividing IP addresses into classes.

Note: Although IANA is responsible for maintaining the network address assignments, virtually all of the IP addresses available using the current addressing scheme have already been assigned to Internet Service Providers (ISPs). When you are building a new network and want to obtain a registered network address, you now get one from an ISP, not directly from the IANA.

Classful IP Addressing

The most complicated aspect of an IP address is that the division between the network identifier and the host identifier is not always in the same place. A hardware address, for example, consists of three bytes assigned to the manufacturer of the network adapter and three bytes that the manufacturer itself assigns to each card. IP addresses can have various numbers of bits assigned to the network identifier, depending on the size of the network.

IANA defines five different classes of IP addresses, which provide support for networks of different sizes. In These five classes of IP addresses (A through E). Only the first three classes are used commercially (Classes A ,B and C). Class A addresses are reserved for

governments throughout the world and some large companies. Class B addresses for medium-sized companies. All other requestors are issued Class C addresses.

Class	1st Octet Decimal Range	Network / Host ID (N=Network, H=Host)
A	1 - 127*	N.H.H.H
B	128 – 191	N.N.H.H
C	192 – 223	N.N.N.H
D	224 – 239	Reserved for Multicasting
E	240 - 254	Experimental, used for research

Note: * Class A address 127 cannot be used and is reserved for loopback and diagnostic functions

Network fields and Host fields

The following table gives the number of bits allocated for network and host fields, number of networks, number of hosts per network in A, B, and C classes

Class	1st Octet Decimal Range	1st Octet High Order Bits	Network bits	Host bits	Number of Networks	Hosts per Network (usable addresses)
A	1 – 127	0	8	24	126 (2^7 - 2)	16,777,214 (2^{24} - 2)
B	128 - 191	1 0	16	16	16,384 (2^{14})	65,534 (2^{16} - 2)
C	192 - 223	1 1 0	24	8	2,097,152 (2^{21})	254 (2^8 - 2)

For the network fields, all the bits cannot be set to 0's or to 1's. So, network address 0 cannot be used in class A network, and since network address 127 is reserved as mentioned earlier, this reduces the total number of networks in class A to 126 instead of 128.

Regarded to the host fields, there are two important concepts these are:

1- Network address:

An IP address that ends with binary 0s (zeros) in all host bits is reserved for the network address. Therefore, as a Class A network example, host 113.1.2.3 is part of a network that has a network address 113.0.0.0 . A router uses a network's IP address when it forwards data on the Internet.

It is important to understand the significance of the network portion of an IP address, the network ID. Hosts on a network can only communicate directly with devices that have the same network ID. If they have different network numbers, even though they may share the same physical segment, they usually cannot communicate with each other directly, unless there is another device that can make a connection between the networks

2- Broadcast Address

A broadcast occurs when a host sends out data to all devices on its network. To ensure that all of the devices on the network pay attention to the broadcast, the sender must use a destination IP address that all of them can recognize and will pick up. Broadcast IP addresses end with binary 1s in the entire host part of the address (the host field).

For the network in the example (113.0.0.0), where the last 24 bits make up the host field (or host part of the address), the broadcast that would be sent out to all devices on that network would include a destination address of 113.255.255.255

All devices on a network recognize their own host IP address as well as the broadcast address for their network

Note: From the discussion above, it is important for network administrator to remember that the first address in each network is reserved for the actual network address and the final

address in each network is reserved for broadcasts. For this reason we have the number of usable hosts per network in classes A, B, and C equals the actual number minus two.

Subnetting :-

A subnet is simply a subdivision of a network address that can be used to represent one LAN on an internetwork. For example, class A network can have (2^{24} - 2) or over 16 million hosts. It is impractical to have all these hosts on the same physical network. It is common to subdivide the network into smaller groups called subnetworks. Most of the time subnetworks are simply referred to as subnets. Another example is the case of large ISP that might have a Class A address registered to it, and it might farm out pieces of the address to its clients in the form of subnets. In many cases, a large ISP's clients are smaller ISPs, which in turn supply addresses to their own clients.

Similar to the host number portion of Class A, Class B, and Class C addresses, subnet addresses are assigned locally, usually by the network administrator. Also, like other IP addresses, each subnet address is unique.

Subnet addresses include the Class A, Class B, or Class C network portion, plus a subnet field and a host field. The subnet field and the new host field are created from the original host portion for the entire network. The ability to decide how to divide the original host portion into the new subnet and host fields provides addressing flexibility for the network administrator.

To create a subnet address, a network administrator borrows bits from the host field and designates them as the subnet field. The minimum number of bits that can be borrowed is 2. If only 1 bit was borrowed to create a subnet, then there would only be a network number and the broadcast number. The maximum number of bits that can be borrowed can be any number that leaves at least 2 bits remaining, for the host number.

The primary reason for using subnets is to reduce the size of a broadcast domain by dividing the large broadcast domain into smaller broadcast domains and thus increase network efficiency.

Subnet Mask

The subnet mask, is not an address. The subnet mask determines which part of an IP address is the network field and which part is the host field. A subnet mask is 32 bits long and has 4 octets, just like an IP address.

When you configure a TCP/IP system, you assign it an IP address and a subnet mask. Simply put, the subnet mask specifies which bits of the IP address are the network identifier and which bits are the host identifier. For a Class A address, for example, the default subnet mask value is 255.0.0.0. When expressed as a binary number, a subnet mask's 1 bits indicate the network identifier, and its 0 bits indicate the host identifier. A mask of 255.0.0.0 in binary form is as follows:

11111111. 00000000 . 00000000 . 00000000

Thus, this mask indicates that the first 8 bits of a Class A IP address are the network identifier bits and the remaining 24 bits are the host identifier. The default subnet masks for the three main address classes are listed in the following table.

Class	Subnet Mask
A	255.0.0.0
B	255.255.0.0
C	255.255.255.0

However, you can create multiple subnets within a given address class by using a different mask. If, for example, you have a Class B address, using a subnet mask of 255.255.0.0 would allocate the first 16 bits for the network identifier and the last 16 bits for the host identifier. If you use a mask of 255.255.255.0, you allocate an additional 8 bits to the

subnetwork identifier, which you are borrowing from the host identifier. The third byte of the address thus becomes a subnet identifier.

However, the boundary between the network identifier and the host identifier does not have to fall in between two bytes. Suppose, for example, you have a Class C network address of 199.24.65.0 that you want to subnet. There are already 24 bits devoted to the network address, and you obviously can't allocate the entire fourth byte as a subnet identifier, or there would be no bits left for the host identifier. You can, however, allocate part of the fourth byte. If you use 4 bits of the last byte for the subnet identifier, you have 4 bits left for your host identifier. To do this, the binary form of your subnet mask must appear as follows:

11111111. 11111111 . 11111111 . 11110000

Number. of subnets

Number of subnets is proportional to the number of bits borrowed from the host field. If the number of bits borrowed is n , then the actual number of subnets is 2^n , but remember that the first and last subnets cannot be used , so the usable number of subnets is (2^n -2). So for the previous class C example the number of subnets is 14.

Number of hosts per subnet:-

As we borrow some bits of the original host field to the subnet field, number of bits remains as host bits decreases and the actual number of hosts per each subnet is 2^m, where m is the number of bits in the new host field. It is also important to remember that the usable number of hosts per subnet is (2^m-2). And again for the previous class C example the number of subnets is 14.

Expressing subnet in decimal

1. Express the subnet mask in binary form.
2. Replace the network and subnet portion of the address with all 1s.
3. Replace the host portion of the address with all 0s (zeros).
4. Convert the binary expression back to dotted-decimal notation.

So for the class C example in which binary form of the subnet mask is,

11111111 . 11111111 . 11111111 . 11110000

The decimal subnet mask is **255.255.255.240.**

Note: When you create subnets, you lose quite a few potential addresses. For this reason, network administrators must pay close attention to the percentage of addresses that they lose by creating subnets.

ANDing process

The lowest numbered address in an IP network is the network address (the network number plus 0s (zeros) in the entire host field). This also applies to a subnet, where in the lowest numbered address is the address of the subnet (the network address plus the subnet field plus 0s (zeros) in the remaining host field).

Hosts and routers use the "ANDing" process to determine if a destination host is on the same network or not. The ANDing process is done each time a host wants to send a packet to another host on an IP network. First the source host will compare (AND) its own IP address to its own subnet mask. The result of the ANDing is to identify the network or subnet where the source host resides. It will then compare the destination IP address to its own subnet mask. The result of the 2nd ANDing will be the network or subnet that the destination host is on. If the source network or subnet address and the destination network or subnet address are the same they can communicate directly. If the results are different then they are on different networks or subnets and will need to communicate through routers or may not be able to communicate at all.

Private IP Addresses

There are certain addresses in each class of IP address that are not registered. Network administrator can freely assign these addresses without obtaining them from an ISP or the IANA. Private addresses might be used by hosts that use a proxy server, to connect to a public (registered) network. They also might be used by hosts that do not connect to the Internet at all. By agreement, any traffic with a destination address within one of the private address ranges will not be routed on the Internet.

Private addresses are given in the following table. When building your own private network, you should use these addresses rather than simply choosing an address at random.

Class	Network Address
A	10.0.0.0 through 10.255.255.255
B	172.16.0.0 through 172.31.255.255
C	192.168.0.0 through 192.168.255.255