IPv4 Address Structure

Version 4 IP addresses are 4-byte binary values. The "printable" version of an IP address is coded as "n.n.n.n", where "n" is the decimal value of each byte in the address. Therefore "n" can range between 0 and 255

Each IP address contains both a network number and a host number. The network number is used for "gross" routing and the host routes to a specific IP stack on the network.

The network portion of an IP address can be 1, 2, or 3 bytes, with the remainder being the host number.

Networks are grouped by "class". The class is determined by the number of bytes in the network number. A 1-byte network is a "Class A", 2 bytes is "Class B" and 3 bytes is "Class C".

You can tell how many bytes are in a network number by examining the left-most byte. Its value will determine class:

1 – 127		Class A
128 – 191	=	Class B
192 - 223		Class C

The following diagram shows the layout of each of the three network classes.

Class A:

01.....7	8.......................................31
Network	Host

The class A address begins with a bit of 0, followed by a 7-bit network ID and a 24-bit host ID. A Class A address supports a maximum of 128 networks, and each network can have 16 million hosts. To obtain a class A address, you'll have to explain why you need one. Class A addresses are easily recognizable because they have values of 1 through 127, inclusive, in the first field of the IP address.

Class B:

0.	2...................15	16......................31
10	Network	Host

The class B address begins with bits of 10, followed by a 14-bit network ID and a 16-bit host ID. A Class B address supports a maximum of 16,383 networks, and each network can have 65,535 hosts. Class B addresses have values of 128 through 191, inclusive, in the first field of the IP address.

Class C:

0..	3...............................23	24........31
110	Network	Host

The Class C address begins with bits of 110, followed by a 21-bit network ID and an 8-bit host ID. The Class C address is the most common type of address. A class C address supports a maximum of 2,097,151 separate networks, and each network can have 256 hosts. In practice this works well because most networks have less than 256 hosts. Class C addresses have values of 192 through 223, inclusive, in the first field of the IP address.

Sub Networking

Frequently, it is convenient to subdivide a network into "sub networks". This is done by reserving a portion of the host number to be used as a “sub network number”

Sub network numbers are not part of an IP address and they are not transmitted across the network. They are extremely useful, however, as a shorthand way of dealing with groups of addresses. At the theoretical level, anyone can be assigned any IP address, regardless of location. In practice, IP addresses are assigned to physical networks in blocks of consecutively-numbered addresses. This means that you can generally treat IP traffic by grouping it together, rather than specifying each individual address.

To work with “groups” of addresses we need to divide then by network number and then by sub network number

Masks

A subnet mask is simply an IP address that is used as a Boolean operator against another IP address to separate it into subnet and host numbers.

In a subnet mask, all non-zero bits identify subnet positions, leaving the zero bits for host positions.

Thus, a mask of 255.255.255.0 would reserve all but the final 8 bits of an IP address for the subnet mask.

Before a subnet mask is applied to an IP address, the network number is removed, leaving just the host bytes. This can be 1, 2, or 3 bytes, depending on the address class. Since only the host portion participates in sub networking, masks of 255.255.255.0, 0.255.255.0, and 0.0.255.0 will all produce the same result when applied to a Class B address (2 host bytes).

Zero-Host

A zero-host address is just as it sounds: an IP address in which the host number is 0. For our purposes, we will consider this to mean the portion of the host number NOT masked as a sub network number.

By specifying an address with a zero-host, we are specifying a pattern that will match any IP address with matching network and sub network numbers. In this manner, we can specify an entire sub network of addresses with a single value.

If we extend this concept to “zero-subnet”, we can refer to a complete network of addresses. [You should note that “zero-network” cannot be specified with a non-zero host.]

Finally, an address of 0.0.0.0 matches all IP addresses.

An Example

Consider the address **150.168.129.0**

By looking at the first byte, **192**, we know this is a "Class B" network. Please note that the "network number" is NOT 150. Since Class B networks reserve two bytes for the network number, its number is really **150.168 or x'96A8' or 38,568**.

The host number is **129.0 or x'8100' or 33,024**.

This address is a "complete" address and uniquely identifies a particular host on a particular network.

Using the same address, we now apply a "subnet mask" of "**255.255.255.0**" (this may also correctly be specified as "0.0.255.0" since only the host number participates in the process).

In view of the mask, the network number remains the same, the subnet number is **129 or x'81'** and the host number is **0**.

This completely changes the interpretation of the address. By using the subnet mask (255.255.255.0), the address in question is now a "zero-host" address. Since no network can have a host with number zero, this address now refers to ANY host on the specified network/sub network. This type of addressing is generally used as a shorthand notation for routing and selection for entire networks and sub networks.

Next, we change the subnet mask to "**255.255.240.0**". Again the network number remains "**150.168**". However, the subnet number is **128 or x'80'** and the host number is **256 or x'100'**. (The right-most two bytes are divided at the bit level.) Interpreted in this manner, it is NOT a zero-host address, but that of a specific destination.