Network Addressing

IPv4 Addressing

The most widely deployed version of the internet protocol is version 4. Hosts on a network are assigned an IP address and a subnet mask. These two addresses together give the host a unique identifier on a network, as well as inform the host which network it is on. By knowing which network a host is on, decisions can be made about how to route packets within and outside of the network.

The IP Address

As you learned earlier, IP addresses are used by devices on a network to communicate with each other. An IP address consists of four numbers, separated by periods. Each number in an IP address is called a *dotted quad*. For example, the IP address for OST's website is 199.27.144.89. Each device on a particular network must have a unique IP address.

An IP address can be split into two parts, a *network prefix* and a *host identifier*. The host identifier does exactly what its name suggests: it identifies a host on a network. The network prefix identifies the network to which the IP address belongs. The split between network prefix and host identifier can vary and is determined by the *subnet mask*.

The Subnet Mask

The subnet mask, like an IP address, can be written using dotted quad notation. IP addresses and subnet masks share this notation because they are of the same length. That's no coincidence. The subnet mask is used to "mask off" parts of the IP address using *bit masking*. Bit masking is accomplished by comparing one binary number to another "bitwise," which means that each bit in one number is compared to the corresponding bit in the other number. Here's a table that outlines the results of various comparisons:

Bit 1	Bit 2	Result
0	0	0
1	0	0
0	1	0
1	1	1

Next we'll look at an example of how one 8-bit number (11110000) masks off another 8 bit number (10101100):

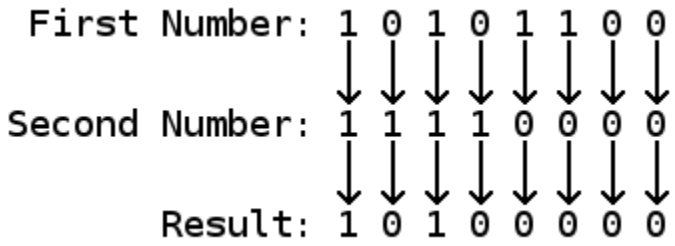

You can see that when our second number (the mask) is a 1, the first number is preserved as a result.

This is how the subnet masks off bits in the IP address to denote them as being part of the network prefix. Bits that are not a part of the network prefix are part of the host identifier. You can also write subnet masks using *prefix length notation*, which is a quicker way to specify the subnet mask. The prefix length is specified as /*n*, where *n* is the number of 1s in the subnet mask, starting from the bit farthest to the left. Below is a diagram that demonstrates the relationship between an IP address (199.27.144.89), a subnet mask (255.255.248.0), a prefix length notation (/21), and the resulting network prefix and host identifier.

	21 Bits	11 Bits
Address In Binary:	11000111.00011011.10010	000.01011001
Subnet Mask In Binary:	11111111.11111111.11111	000.00000000
Network Prefix:	11000111.00011011.10010	
Host Identifier:		000.01011001

The number of bits left when you subtract the prefix length from 32 (the longest prefix length possible) tells us how many hosts we can have on our network. In this example, we have 11 bits for the host identifier, so we can have up to 2048 (2^{11}) hosts—or can we?

The Network and Broadcast Addresses

For all networks with a subnet mask of /30 or smaller, two addresses must be reserved from the pool of available addresses, the *network address* and *broadcast address*. This means that in a /21 network like the one we had in our example, we can only have 2046 (2^{11}-2) hosts, not 2048. These two addresses serve a special purpose in the IPv4 protocol. The network address is used to refer to an entire network of hosts at one time. Most commonly, this address would be used in addition to a subnet mask in order to

specify a large range of hosts for some operation, such as traffic filtering. The network address is formed by picking the smallest address from the network.

The broadcast address also exists as a shorthand for all the hosts on a network, but there is one important distinction. You can use the network address to refer to an entire network, but you can't send packets to that address. However, the broadcast address can receive packets. Packets sent to the broadcast address go to all hosts on the specified network. The broadcast address is deerived from the largest address on the network. Here are some examples of various network and broadcast addresses for networks of varying sizes:

Host Range	Subnet Mask	Network Address	Broadcast Address	Number of Hosts
192.168.0.1 to 192.168.0.254	255.255.255.0 or /24	192.168.0.0	192.168.0.255	254
199.27.144.1 to 199.27.151.254	255.255.248.0 or /21	199.27.144.0	199.27.151.255	2046
10.0.0.9 to 10.0.0.14	255.255.255.248 or /29	10.0.0.8	10.0.0.15	6

Note That last entry shows that the network address doesn't always end in a 0, and the broadcast address does not always end in 255.

Private Networks

Devices on a network don't necessarily need a publicly addressable IP address; they only need to be accessible on your network, and not from the internet at large. In fact, in many situations, for security purposes it's better if devices are *not* accessible directly from the internet. Privately addressed devices can send packets to the internet using *network address translation* or *NAT*, but NAT is beyond the scope of this course. There are three networks that are designated as private networks. Anyone can use these networks without obtaining permission from one of the internet's regulatory bodies. Here's a summary:

Network	Address Range	Subnet Mask	Number of Hosts
192.168.0.0/16	192.168.0.0-192.168.255.255	255.255.0.0 or /16	65,536 addresses or 65,534 hosts
172.16.0.0/12	172.16.0.0-	255.240.0.0 or	1,048,576 addresses or 1,048,574

	172.31.255.255	/12	hosts
10.0.0.0/8	10.0.0.0-10.255.255.255	255.0.0.0 or /8	16,777,216 addresses or 16,777,214 hosts

Note When configured, your machine will be a part of the "172.16.0.0/12" network.

Subnets Revisited

The subnet mask determines which network includes a particular host. This, coupled with the fact that the length of a subnet mask is variable, allows us to take a network block and split it into several smaller network blocks. This can help with efficient IP address allocation on your network, as well as provide a method by which to segregate machines according to security policies. As an example, let's use a private network, 10.0.0.0/8. With over 16 million allocatable addresses, it's a prime candidate for being broken into smaller blocks of allocatable addresses. In fact, it's possible to have up to 65,536 /24 networks, just inside the 10.0.0.0/8 space! Let's take a look at the differences between two hosts on distinct /24 blocks in the 10.0.0.0/8 space:

```
          Host: 10.0.0.127/24
     In Binary: 00001010.00000000.00000000.01111111
   Subnet Mask: 11111111.11111111.11111111.00000000
Network Prefix: 00001010.00000000.00000000
       Host ID:                              01111111

          Host: 10.0.1.127/24
     In Binary: 00001010.00000000.00000001.01111111
   Subnet Mask: 11111111.11111111.11111111.00000000
Network Prefix: 00001010.00000000.00000001
       Host ID:                              01111111
```

As you can see, these two hosts have the same host id, but they're on different networks: 10.0.0.0/24 and 10.0.1.0/24. In fact, without some kind of routing in between them, these two hosts cannot communicate with each other. Now take these two IP addresses, but alter their subnet mask to /23. Now they are no longer two separate networks with 256 addresses each, but instead they comprise a single network with 512 addresses. If you're still a little fuzzy on this concept, draw yourself a chart like the one

above and see how the subnet mask affects the network address and and host id for each host.

Source: http://courses.oreillyschool.com/sysadmin2/network_addressing.html