UnixWorld Online: Tutorial: Article No. 001

Subnet Addressing

By Ron Cooney

The issue of subnetting has long been a mystery to many system administrators. It seems that there are just too many bits, bytes, and masks involved to make it worthwhile. Besides, who needs all that trouble when it's all you can do to just maintain the network as it is? You can certainly find sympathy, but if your network plans call for Internet access then IP (Internet Protocol) addressing and possibly subnet masking are topics that you should be familiar with.

As networks grow to increasing numbers of segments, more network address numbers will be needed as each segment requires an individual number. The InterNIC (whose name signifies cooperation of Network Information Centers or NICs), however, is not eager to hand out unlimited network addresses because they are quickly running out of them. The Internet community is generally taking a firm stand on limiting address availability, due to the accelerated growth demand for Internet access. Network administrators will have to work with what they have to better accommodate network requirements and the shrinking supply of address numbers.

One way of accomplishing this address conservation is to take the address that is assigned to your network and expand the capacity with subnets. Subnetting allows you to increase the number of networks available to you without applying for another IP address number.

IP Addressing

The IP addressing framework allows one to address about 16,000,000 unique hosts for a Class A address, around 65,000 hosts for a Class B address, but only 254 hosts for a Class C address. However, there are no more Class A addresses available, and the InterNIC has stopped assigning Class B addresses. Class C addresses are the most numerous, but their limitation is that each can identify only 254 unique hosts.

The IP address is composed of 32 bits, which consist of two parts: the most significant bits (MSBs) identify a particular network and the remaining bits specify a host on that network. The most significant bits of the network portion actually determine the address class as shown in this table:

```
Address   MSB
Class     Pattern

Class A   0
Class B   10
Class C   110
```

Class A Address Example

A class A address could be diagramed:

```
  Network           Host
 +------+  +----------------------+
 |      |  |                      |
[0xxxxxxx][xxxxxxxxxxxxxxxxxxxxxxxx]
```

which shows the eight network bits followed by the 24 host bits.

These 32-bit IP addresses are almost always written as four dot-separated decimal numbers, one for each byte of the address. Thus, our class A address would have a range of address numbers from `1.0.0.0` through `126.0.0.0` (`0.`*x.x.x* and `127.`*x.x.x* are reserved). The number of host addresses per network is 16,777,214, which is two less than two raised to the 24th power because both host numbers `0.0.0` and `255.255.255` are reserved.

In practice, people don't really attach 16 million hosts to a network so administrators of a Class A site often divide the host address portion into a (sub)network and host portion. (Subnetting is now supported by most operating systems.) Each Class A network number can support up to 65,534 subnets (network numbers `0.0` and `255.255` are reserved) with each having 254 hosts (host numbers `0` and `255` are reserved). This is done by using the 16 high -order bits of the host portion for the subnet number and the lower eight bits for the host as diagramed here:

```
 Network        Subnet         Host
 +------+  +--------------+  +------+
 |      |  |              |  |      |

[0xxxxxxx][xxxxxxxxxxxxxxxx][xxxxxxxx]
```

Class B Addresses

The first two bits of a Class B address are 1 and 0, the next fourteen bits identify the network and the last sixteen the host, as diagramed:

```
     Network           Host
 +--------------+  +--------------+
 |              |  |              |
[10xxxxxxxxxxxxxx][xxxxxxxxxxxxxxxx]
```

Thus, Class B addresses include the network numbers in the range from `128.1.0.0` through `191.254.0.0` for a total of 65,534 host addresses.

As with the Class A address, we can divide the host portion of a Class B address into subnet and host parts. For instance, let's split our Class B network number on the byte boundary, that is, the eight MSBs of the host portion identifies the subnet and the remaining bits the host, as diagramed:

```
     Network       Subnet     Host
 +--------------+  +------+  +------+
 |              |  |      |  |      |
[10xxxxxxxxxxxxxx][xxxxxxxx][xxxxxxxx]
```

This arrangement allows 254 subnets each with 254 hosts.

Other Address Classes

The first three bits of a Class C address are 1, 1, and 0, the next 21 bits identify the network and the last eight the host, as diagramed:

```
         Network                 Host
 +----------------------+  +------+
 |                      |  |      |
[110xxxxxxxxxxxxxxxxxxxxx][xxxxxxxx]
```

Thus, Class C addresses include the network numbers in the range `192.0.1.0` through `223.255.254.0` for a total of 254 host addresses per network address.

Finally, we have Class D and Class E addresses. Class D address start at `224.0.0.0` and are used for multicast purposes. Class E addresses start at `240.0.0.0` and are currently used only for experimental purposes.

The Subnet Mask

A subnet mask (or number) is used to determine the number of bits used for the subnet and host portions of the address. The mask is a 32-bit value that uses one-bits for the network and subnet portions and zero-bits for the host portion.

Let's look at an example. Here we have a Class B address of `191.70.55.130` and apply some different subnet masks. A logical AND operation is performed between the IP address and the subnet mask as shown:

Here we use a mask that retains the default 16 network and host bits for a Class B address:

```
   191          70            55           130

1011 1111    1000 0110    0011 0111    1000 0010   IP address
1111 1111    1111 1111    0000 0000    0000 0000   Subnet mask

1011 1111    1000 0110    0000 0000    0000 0000   Result
```

Here we employ a mask that divides the host portion into a subnet and host that are each eight bits wide:

```
   191          70            55           130

1011 1111    1000 0110    0011 0111    1000 0010   IP address
1111 1111    1111 1111    1111 1111    0000 0000   Subnet mask

1011 1111    1000 0110    0011 0111    0000 0000   Result
```

This division allows 254 (256-2 reserved) subnets, each with 254 hosts.

This division on a byte boundary makes it easy to determine the subnet and host from the dotted-decimal IP address. However, the subnet-host boundary can be at any bit position in the host portion of the IP address. Here, we use a mask that allows more subnets (512-2 reserved), but with the trade-off of fewer hosts (128-2) per subnet:

```
   191          70            55           130
```

```
1011 1111    1000 0110    0011 0111    1000 0010  IP address
1111 1111    1111 1111    1111 1111    1000 0000  Subnet mask

1011 1111    1000 0110    0011 0111    1000 0000  Result
```

The subnet-host number tradeoff

Here's a table that let's you see at a glance the trade off between the number of subnets and hosts with different subnet masks for both Class B and Class C addresses. We've already subtracted two from the results in the last two columns to take the reserved network and host numbers into account:

Class B Subnetting:

# Mask Bits	Subnet Mask	# Subnets	# Hosts
2	255.255.192.0	2	16382
3	255.255.224.0	6	8190
4	255.255.240.0	14	4094
5	255.255.248.0	30	2046
6	255.255.252.0	62	1022
7	255.255.254.0	126	510
8	255.255.255.0	254	254
9	255.255.255.128	510	126
10	255.255.255.192	1022	62
11	255.255.255.224	2046	30
12	255.255.255.240	4094	14
13	255.255.255.248	8190	6
14	255.255.255.252	16382	2

Class C Subnetting:

# Mask Bits	Subnet Mask	# Subnets	# Hosts
2	255.255.255.192	2	62
3	255.255.255.224	6	30
4	255.255.255.240	14	14
5	255.255.255.248	30	6
6	255.255.255.252	62	2

The Subnet Advantage

Subnetting hides the internal network organization to external routers and thus simplies routing. For instance, a subnetted Class B address would require fewer routes than the equivalent number of Class C addresses. Shorter routing tables mean faster network transfers.

Subnetting allows address administration to be decentralized. Besides technical advantages, this approach may also provide political benefits for the organization. For instance, an administrator could assign a subnet to a department, which would then be responsible for their own network management.

Subnetting can help overcome distance limitations of physical networks by dividing up a physical network into individually addressed networks so they can be connected logically with routers.

Example: Subnetting a Class C Network

One of the first things a network administrator needs to do is define the requirements for the network. The logical place to start is to consider how many hosts are on the network.

Using the maximum number of hosts on one Ethernet segment is generally not good practice because it could create performance problems due to network congestion. If you only have one Class C address assigned to your network then what can you do? Refer to our table above that depicts the Class C address subnetting network number-host trade off.

Even though a Class C address can support up to 254 hosts, in my experience, 60-80 hosts is a good number for most LANs using office automation tools. I've seen overloaded Ethernet segments--with over 100 hosts--at client sites. My recommendation is that they segment their LAN in half or even further. Also, many hub cards come with 24 ports per card, which makes it easy to segment in 24-host multiples provided that the hub supports multiple segments on the backplane. Many do.

One reasonable approach would be to select six subnets each with 30 hosts. Although two subnets with 62 hosts is also feasible, it is not as flexible because there are only two subnets. The other alternatives that use more subnets probably don't provide enough hosts per subnet.

Subnets 0 and 7 are unusable because they are used for special addressing situations. For instance, a subnet of 7 (all one bits) is reserved for an all subnets-directed broadcast (a broadcast sent to all subnets of the specified subnetted network) when the host bits are all one. This leaves subnets 1 through 6 available for use.

In each subnet, the first host number (0) is reserved, and the resulting number is known as the network number. The last number in each subnet is reserved for the broadcast address, and cannot be used for a host address. Consequently, in this case there are only 30 host addresses available for each subnet.

Bibliography

Hunt, Craig. *TCP/IP Network Administration*. O'Reilly and Associates, Inc., Sebastopol, CA., 1992 (ISBN: 0-937175-82-X).

Markley, Richard W. *Data Communications and Interoperability*. Prentice Hall, Englewood Cliffs, N.J., 1990.

Stevens, W. Richard. *TCP/IP Illustrated, Volume 1: The Protocols*. Addison-Wesley Publishing Company, Reading, Mass., 1994 (ISBN: 0-201-63346-9).

Edited by Becca Thomas / Online Editor / UnixWorld Online / beccat@wcmh.com

Go To Content   Search

Last Modified: Tuesday, 22-Aug-95 15:52:44 PDT