

Networks and Subnetting

When TCP/IP is configured on a computer or other network device, each connection point on the device that will communicate TCP/IP is called an **interface**. This includes each ethernet port, token ring port, AUI port, wireless network adapter, or serial line connection that is used for TCP/IP networking.

When an IP address and netmask pair are assigned to an interface they make known to the TCP/IP device what other IP addresses are reachable as locally connected to that interface. For example, a computer connected to an ethernet LAN, with an address of 205.217.146.198 and netmask of 255.255.255.0, defines that all IP addresses from 205.217.146.1 through 205.217.146.254 would be local addresses on that ethernet segment, if they exist.

To determine if two addresses are on the same local segment, a router uses the netmask and performs a logical AND operation on each of the two addresses. If the resulting network address is the same, the two addresses are on the same network segment. If the resulting network address is different, then the two addresses are on different segments.

For example, a computer has only one interface and it is connected to an ethernet LAN. The interface has the IP address 192.168.10.30 and a netmask of 255.255.255.252. By combining the netmask and the address together, a network address of 192.168.10.28 is derived.

Interface	192.168.10.30	11000000	10101000	00001010	00011110
Netmask	255.255.255.252	11111111	11111111	11111111	11111100

Network Address	192.168.10.28	11000000	10101000	00001010	00011100

To determine if the destination address 192.168.10.37 is on the same network, the netmask is combined with the address, yielding a network address of 192.168.10.36.

Destination	192.168.10.37	11000000	10101000	00001010	00100101
Netmask	255.255.255.252	11111111	11111111	11111111	11111100

Network Address	192.168.10.36	11000000	10101000	00001010	00100100

The two network addresses are not the same, so the two addresses are not on the same network. To reach the destination address from that interface, a datagram would need to be passed to a gateway system on the local network for delivery.

Classical IP networks have default netmasks:

- Class A - 255.0.0.0
- Class B - 255.255.0.0
- Class C - 255.255.255.0

Subnetting is the process of taking a classical IP network (Class A, B, or C) and using a mask to break it into smaller pieces. This is usually performed by a network administrator to provide services to various organizations and divisions within a network. It might be done because of physical requirements, political necessities, or because of hardware differences. Offices could be in different geographic areas, departments might need or want their networks separated, or an organization might have ethernet, token ring, or fast ethernet networks that need to be connected. Subnetting might also be performed by an ISP to make efficient

use of IP addresses or other reasons.

There are specific rules for using subnetting to break up classical IP networks. The primary rule is that when the netmask is represented in binary, all ones must be contiguous to the left, and all zeroes must be contiguous to the right. This results in a limited number of valid netmasks.

Another important subnetting rule is that the highest and lowest numbered subnets are not valid and should not be used. The default netmask for a classical IP network divides the address into a network portion and a host portion. For example, the Class C default netmask 255.255.255.0 assigns the first 24 bits (3 bytes) as the network address, and the last 8 bits (1 byte) as the host portion. Subnetting adds additional One bits to the netmask, in the host portion, which are sometimes referred to as the **subnet bits** or **subnet address**. A classical IP network that is subnetted has additional One bits, the most significant bits of what would normally be the host portion, which are used to extend the network into a subnet:

Example Class C Subnet Masks

Mask	Network	Subnet	Host
255.255.255.192	11111111 11111111 11111111	11111111	11000000
255.255.255.248	11111111 11111111 11111111	11111111	11111000

Example Class B Subnet Masks

Mask	Network	Subnet	Host
255.255.192.0	11111111 11111111	11000000	00000000
255.255.248.0	11111111 11111111	11111000	00000000
255.255.255.128	11111111 11111111	11111111	10000000

Example Class A Subnet Masks

Mask	Network	Subnet	Host
255.192.0.0	11111111	11000000 00000000	00000000
255.255.248.0	11111111	11111111 11111000	00000000
255.255.255.128	11111111	11111111 11111111	10000000

Just as the highest and lowest numbered host address within a network (the network address and announce address) are reserved, the highest and lowest subnet numbers in a network are reserved. Many types of equipment and software will allow a network administrator to assign and use these addresses, but some equipment will reject them as invalid. Some software, especially diagnostic programs, will have problems talking to systems if these reserved networks are used. Many systems interpret the subnet address with all ones as a subnet announce address and the subnet address with all zeroes as a subnet group address.

When using subnet masks to break up a Class C network, there are only five valid netmasks:

Netmask	Number of Bits	Usable Networks	Usable Hosts per Network
255.255.255.192	26	2	62
255.255.255.224	27	6	30
255.255.255.240	28	14	14
255.255.255.248	29	30	6
255.255.255.252	30	62	2

This table usually brings up several questions:

Why can't you use a netmask of 255.255.255.128 (25 bits)?

The 255.255.255.128 subnet mask would break a Class C network into two subnets, the first with a subnet address of all zeroes and the second with a subnet address of all ones. These two subnet addresses are reserved and should not be used.

If the 255.255.255.192 (26 bit) netmask breaks a Class C network of 254 addresses into two pieces, why do you get only two 62 address networks (124 addresses)? What happened to the rest of the addresses?

When you use a net mask of 192, you have 4 possible subnets:

```
0 00000000
64 01000000
128 10000000
192 11000000
```

The first and last subnets are reserved because the first, 0, has a subnet address that is all zeroes, and the last, 192, has a subnet address that is all ones. This leaves only subnets 64 and 128. These have only 62 usable addresses each because the first and last addresses in each network are the reserved network address and announce address.

Why can't you use a netmask of 255.255.255.254 or 255.255.255.255?

A subnet mask of 255.255.255.254 breaks a Class C network into subnets that have only two addresses each. These addresses would be the network address and announce address, which would be unusable as host addresses. The netmask 255.255.255.255 would define a network of only one device. About the only time this mask could be useful is defining a loopback, an IP device that can only talk to itself.

I have seen other netmasks used, or have used them successfully myself - Why would they be invalid?

It is possible to configure networks using other netmasks. If the addresses involved come from Class A or Class B ranges, then the restrictions on the valid masks are different: for Class B networks the subnet mask 255.255.255.128 is valid, 255.255.128.0 is not. Also, newer equipment or software often does not enforce subnet masking restrictions. However, many common pieces of equipment and software do enforce these restrictions, especially diagnostic programs and more expensive routers. Even if you are able to subnet your Class C network using a 255.255.255.128 netmask, or using addresses in the first or last subnet, there is a distinct possibility that the next piece of equipment or software added to your network could **"break"** the network.