

NetInfo Reference

## Internet Addressing

### SA4padPointRule2.eps ↗

In this topic

The Internet ;IPAddressing.rtf;The Internet;↗  
Network Traffic ;IPAddressing.rtf;Network Traffic;↗  
Internet Addresses ;IPAddressing.rtf;Internet Addresses;↗  
Routing ;IPAddressing.rtf;Routing;↗  
Assigning Internet Addresses ;IPAddressing.rtf;Assigning Internet Addresses;↗

### 1PointDashedRule3Black.eps ↗

As part of setting up a network, you'll have to assign an Internet address to each NeXT computer connected to a network. An Internet address is, simply, the means used by one networked computer to identify another. However, in order to be sure that your Internet addresses are unique, you must apply to a central organization to have a network address assigned. If you have a large network, you may want to build your network to include subnets.

This section provides some background information about the Internet, then explores the components of an Internet address, describes how to use subnetting, and provides contact information for obtaining an Internet address.

### 281558\_PointDashedRule3Black.eps ↗

## The Internet;↗The Internet

The Internet is a huge network of networks connecting university, government, and commercial research labs around the world. The individual networks that make up the Internet share an address structure and communications protocol suite. This protocol suite is TCP/IP, which stands for Transmission Control Protocol/Internet Protocol.

NeXT computers support TCP/IP and can easily communicate with other computers

running TCP/IP. The TCP/IP software is independent of the hardware used for communication. While NeXT computers (and many others) are commonly connected in an Ethernet network, TCP/IP runs on many other types of media, including ISDN (Integrated Services Digital Network).

114750\_PointDashedRule3Black.eps ↯

## **Network Traffic; ↯ Network Traffic**

Information transmitted on a network is generally called *traffic*. Individual pieces of information are called *packets*. The amount of traffic on a network is proportional to the number of computers using the network, the types of packets they send, the types of services provided by the network, and the frequency of transmission, among other factors. The capacity of a network to handle traffic is limited by the type of communication hardware as well as the network configuration.

Packets can generally be grouped into a few categories:

349528\_smBlksquare.tiff ↯ Broadcast↯Data sent to all computers on a network. In TCP/IP based networks, this is accomplished by sending data to the broadcast address of a given network (see the next section, <<sup>a</sup>Internet Addresses<sup>o</sup>>). When a networked NeXT computer boots, the first packet it sends is a request asking for its Internet address, sent to the broadcast address.

180483\_smBlksquare.tiff ↯ Multicast↯Data sent to a specific group of hosts.

700123\_smBlksquare.tiff ↯ Unicast↯Data sent to a single destination. This is the most common type of packet.

339920\_smBlksquare.tiff ↯ Error↯Data re-sent due to an error. When a packet can't be completely delivered to the destination in a set period of time, it's retransmitted. Retransmission is a normal part of network operation, but it can cause problems

since error packets can interfere with other traffic.

**Note:** In Ethernet-based networks, all packets are actually sent to every host on the network. Each host is responsible for rejecting packets that aren't addressed to it. Ethernet networks allow only one computer to transmit at a time. If two or more computers transmit at the same instant, a collision occurs. All computers involved in a collision will retransmit their packets after a random length of time. A large number of collisions on a network can indicate that it's overloaded.

285773\_PointDashedRule3Black.eps ↵

## **Internet Addresses; ↵Internet Addresses**

An *Internet address* or *IP address* is simply a 32-bit integer. Addresses are used to uniquely identify individual computers, groups of computers, or entire networks. The address itself is generally written in dotted notation like this: x.x.x.x where each  $x$  is an integer from 0-255 (called an *octet*), and each octet represents one byte of the address.

Valid Internet addresses start at 0.0.0.1 and range through 255.255.255.254. Although all Internet addresses contain 4 octets, there are several abbreviations. For example, 127.1 is the same as 127.0.0.1, while 127. is shorthand for 127.0.0.0.

### **Address Components**

Internet addresses have two components: the network part and the host part. The network part, or network address, identifies a network, while the host part identifies a host on that network. All computers on the same physical network must have the same network address.

Some networks only have a few hosts, and therefore only need a few host addresses,

while other networks need many host addresses. Internet addresses are divided into classes, determined by the first (high-order) bits of the address. Each class allocates a specific number of bits for the network part of the address, with the remaining bits used for the host address. The following table describes the primary classes:

<b>HO Bits</b>	<b>Allocation</b>	<b>Class</b>	<b>Address Range</b>
0	7/24	A	1.0.0.0-127.255.255.255
10	14/16	B	128.0.0.0-191.255.255.255
110	21/8	C	192.0.0.0-223.255.255.255

Class A addresses allocate 7 bits to the network portion of the address and 24 to the host portion; class B addresses allocate 14 bits to the network portion and 16 to the host portion; and class C addresses allocate 21 bits to the network portion and 8 to the host portion. The fewer the number of bits used in the network part, the greater the number of bits allocated to the host portion. The number of bits allocated to the host portion determines the number of hosts that can be attached to one network.

You can determine the class of a given Internet address by looking at the value of the first octet. For example, 192.42.172.0 is a class C network address, since 192 is in the range of addresses for that class. With 8 bits allocated to the host portion of the address, class C networks have 254 host addresses available (two addresses are reserved, as described in the next section).

## **Special Addresses**

There are several types of addresses in addition to host addresses:

660513\_smBlksquare.tiff ↪ Network Address ↪ Specifies an entire network, as opposed to an individual host. The bits in the host portion of a network address are all set to zero, as in 189.52.0.0 (a class B network address). You'll see network addresses in the output of several commands, including **netstat**.

140277\_smBlksquare.tiff ↪ Broadcast Address ↪ Used to send broadcast packets. The bits in the host portion of a broadcast address are all set to one, as in 189.52.255.255. All hosts with the same network address will receive information sent to the broadcast address. The address 255.255.255.255 is defined as the broadcast address of the current network.

580589\_smBlksquare.tiff ↪ Loopback Address ↪ Reserved for non-network interprocess communication and testing. All addresses beginning with the octet 127 are loopback addresses and are not considered to be on a network. On NeXT computers, as well as the majority of UNIX computers using TCP/IP, the special name **localhost** is given to the address 127.0.0.1.

**Important:** Don't assign any of these addresses to hosts on your network.

576708\_PointDashedRule3Black.eps ↪

## Routing; ↪ Routing

When talking about the network part of an Internet address, the term *network* means a set of hosts that can directly communicate with each other. For example, assume there are three different hosts connected with Ethernet cables that have been assigned the addresses 192.42.172.1, 192.42.172.2, and 192.42.178.1. By examining the first octet, you can see that each is a class C address. For two class C addresses to be on the same logical network, the first three octets (24 bits) must match, since only the last octet (8 bits) is used for the host portion. Therefore, the hosts assigned the first two addresses are on the same network, and can communicate with each

other, while the third is on a separate network, unable to communicate with either of the first two.

Hosts on different logical networks can communicate with each other if they have access to a *router*. A router attaches to two or more networks and routes traffic between them. Routers maintain tables of paths to destinations as well as values that approximate the delay of choosing one path over another.

## **Subnetting**

A large organization that has (or is planning) hundreds or thousands of hosts will find that the 254 host addresses provided by a class C network are insufficient. A class B network address allows for tens of thousands of host addresses. A network with a large number of hosts will probably be made up of several smaller networks in different locations.

Routers can be used to connect these smaller networks, but remember that routers only transfer data between different logical networks (hosts with addresses that have different network portions). In order for routing to work successfully in this situation, you need to have each smaller network be a different logical network, even though you've been assigned a single network address.

By using a *subnet mask*, you can extend the network portion of the Internet address, effectively providing multiple logical networks. A subnet mask (or *netmask*) determines which bits of an Internet address will be used as the network portion, and which as the host portion. For example, assume you've been assigned the class B network address 140.211.0.0. You've got several sites distributed among different buildings and different cities, each with up to 200 hosts. You set a subnet mask of 255.255.255.0, which assigns the first 16 bits to the network portion and 8 bits to the host portion. This gives you the potential for 254 different subnets, each with 254 hosts. You can now, for

example, assign one site the network address 140.211.128.0, and another the network address 140.211.192.0.

To see how this works, we need to examine the addresses in binary form:

<b>Decimal</b>	<b>Binary</b>
307955_PointRule2.eps ↵	
255.255.255.0	11111111.11111111.11111111.00000000
427825_PointDashedRule3Black.eps ↵	
140.211.128.0	10001100.11010011.10000000.00000000
658435_PointDashedRule3Black.eps ↵	
140.211.192.0	10001100.11010011.11000000.00000000
687477_PointDashedRule3Black.eps ↵	

Each bit in a subnet mask set to 1 indicates that the corresponding bit of the Internet address should be considered part of the network portion, while bits set to 0 indicate the host portion. In this example, the first 16 bits of the address have been assigned to you, and the subnet mask extends the network portion to include the next 8 bits, leaving 8 bits for the host portion.

Notice that the two network addresses have been chosen by adjusting the leftmost bits of the third octet. This is a common practice that allows for future expansion. If you find that a subnet needs more than 254 host addresses, you can assign a new subnet mask that leaves more bits available for the host portion. If you had already used some of those bits for the network portion, you would have to reassign host addresses (not a pleasant task). By using the leftmost bits of the third octet for network addresses, you leave the rightmost bits of the address available for possible future use in the host portion.

You can use the **ifconfig** program to see the Internet address, subnet mask, and broadcast address currently in use. For example, enter the following in a shell window:

```
ifconfig en0
```

On NeXT computers, the name of the Ethernet network interface is **en0**. The output will look something like this:

```
en0: flags=23<BROADCAST,UP,NOTRAILERS>
      inet 140.211.192.27 netmask ffffffff broadcast 140.211.192.255
```

The second line of output includes the Internet address of the local computer, the netmask, and the broadcast address for the current network. Notice that the netmask is displayed as a string of hexadecimal numbers, two for each octet. The netmask value `ffffff00` is the same as `255.255.255.0`. For more information, see the UNIX manual page for **ifconfig**.

You can use the **netstat** program to see a list of active network interfaces. Enter the following in a shell window:

```
netstat -r
```

The **r** option is used to display routing table information. The output will look similar to this:

```
Routing tables
Destination      Gateway          Flags           Refs      Use  Interface
localhost        localhost       UH              4         1520  lo0
140.211.192.0    140.211.192.27 U                0          0    en0
```

The destination information is a network address or host name, and the gateway information is the Internet address or host name used to get to the destination. In this example, packets destined for the `140.211.192.0` network will be sent to `140.211.192.27`. For more information, see the UNIX manual page for **netstat**.

## **Assigning Internet Addresses; Assigning Internet Addresses**

If you're absolutely sure your network will never be on the Internet, you can choose your own Internet addresses, since you don't have to worry about your addresses being distinct from other networks. (Being <sup>a</sup>on the Internet<sup>o</sup> just means that you have some kind of connection that allows you to send packets to another computer on the Internet.) If your network will be on the Internet, or you aren't sure, have an address assigned at the outset, to avoid having to reassign addresses later.

To request an address, contact the Network Information Center (NIC):

Government Systems, Inc.  
Attention: Network Information Center  
14200 Park Meadow Drive, Suite 2000  
Chantilly, Va. 22021  
Telephone: 1-800-365-3642  
Fax: (703) 802-4535

There may be a fee for obtaining a network address. Anyone may request one. The class of address you will be assigned depends on the number of hosts you plan to have on your network and the size of your organization.