IP addresses and Subnetting

IP addresses & subnetting - an overview

The following gives an introduction to IP addresses and subnetting on local area networks. If you want to find out about the advantages of using private network IP addresses on you local area network, or what subnetting can do for you, the explanation is here. You can also find the recipe for how you calculate a subnet mask, a network address and broadcast address.

However, the course also offers a fast track to getting the advantages of subnetting on local area networks without having to do all the calculations yourself. If this is what you are looking for, you might want to jump directly to the last chapter in this course: **'The fast track to the advantages of subnetting'.**

IP addresses

Each computer on a TCP/IP based network (including the Internet) has a unique, numeric address called an IP address (IP stands for Internet Protocol), enabling data packages to be addressed to this specific recipient.

What is an IP address?

An IP address consists of four so-called octets separated by dots. The octet is a binary number of eight digits, which equals the decimal numbers from 0 to 255. To make IP addresses more easy to read and write, they are often expressed as four decimal numbers, each separated by a dot. This format is called "dotted-decimal notation".

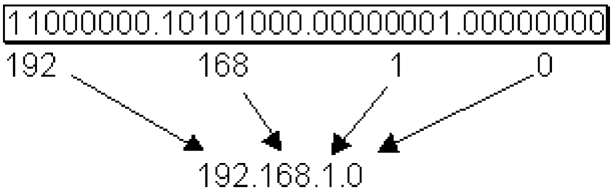

An IP address in its binary and dotted-decimal notation

In a local area network based on TCP/IP, an IP address must be assigned to each host (computer or device) in the network. The IP address must be unique to each host. (If two hosts were given the same address, the data to these hosts would be picked up randomly by one of them – be it the intended receiver or not – causing network irregularities.)

In addition, a device that serves as router to another network, contains two or more network adaptors and belongs to two or more networks. In this case, each adaptor must be assigned a unique IP address on each network.

Part of an IP address designates the network, while another part designates the individual host. The network number field is also referred to as the 'network prefix'.

NETWORK NUMBER
HOST NUMBER

The two parts of an IP address

Exactly where the network part ends and the host part begins is calculated by routers, using a so-called subnet mask as a deciphering key.

All hosts on a given network share the same network number, but each of them must have a unique host number:

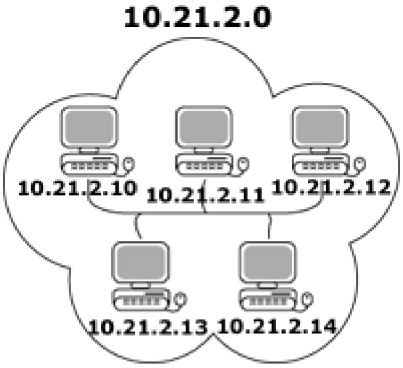

The host portion of the IP address is unique to each host

The network portion of an IP address is inherited down through a network hierarchy, as illustrated below.

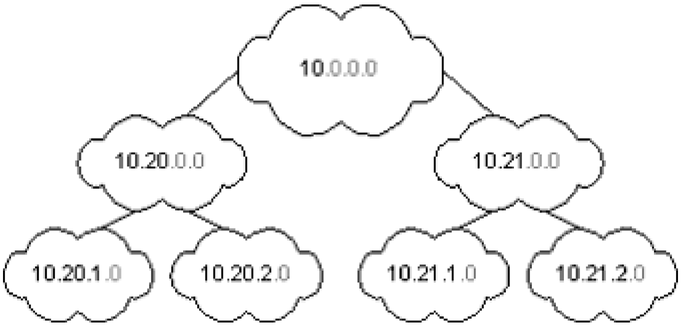

Each cloud symbolizes a network segment

to the top

Classes of IP addresses

In order to provide the flexibility required to support differently sized networks, IP addresses come in three *classes*, A, B, and C. Every class fixes the boundary between the network portion and the host portion of the IP address at a different point. This makes them appropriate for different size networks.

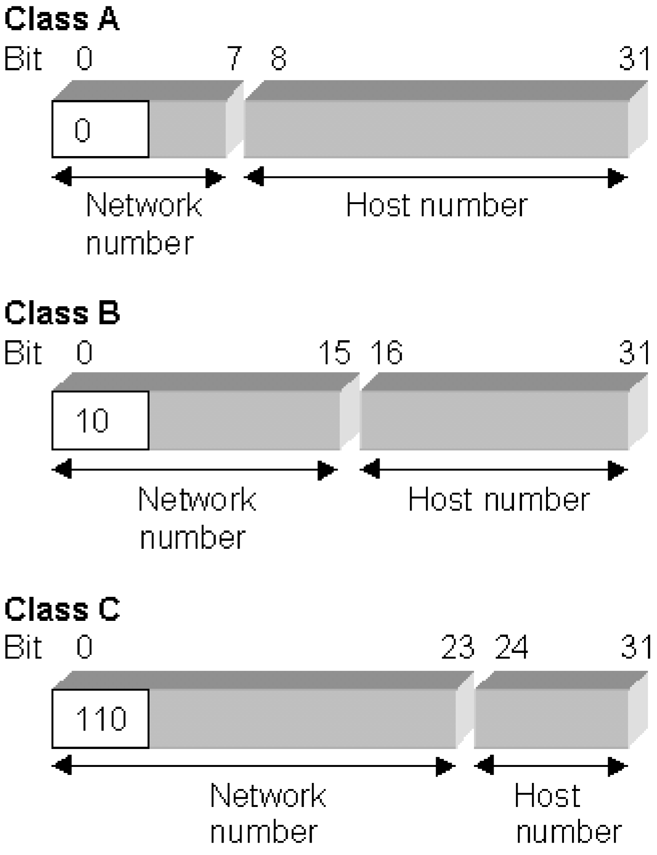

Class C addresses allow 254 hosts per network and are typically used by smaller and middle-sized companies. Class B networks allow a maximum of 16,384 hosts, while Class A networks allow more than 16 million hosts. As a consequence, Class A networks are only used by really large organisations.

Calculating the number of possible hosts requires a closer look at the IP classes in their binary form. (The binary system is a base-2 number system, just like the base-10 number system is known as the decimal number system). It is done as follows:

- In a Class C network only the last octet is used to designate the hosts. The maximum decimal number that you can write using eight bits is 256 (2^8). The host calculation now requires that 2 is subtracted, because two host addresses must be reserved for a network address and a broadcast address (for a further explanation of network and broadcast addresses, see the section on ‘Subnets’). Ergo the maximum number of hosts on a Class C network is 256-2=254.
- A class B network allows a maximum of 16,384 hosts (2^{16}-2) per network (two octets designate the hosts).
- A class A network allows up to 16,777,214 (2^{24}-2) hosts per network (three octets are used to designate the hosts).

The table below shows the range of dotted-decimal values that can be assigned to each of the three address classes. An x represents the host number field of the address which is assigned by the network administrator.

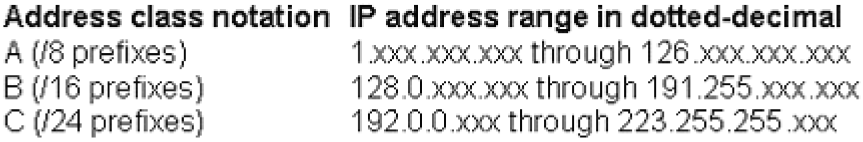

Address class notation	IP address range in dotted-decimal
A (/8 prefixes)	1.xxx.xxx.xxx through 126.xxx.xxx.xxx
B (/16 prefixes)	128.0.xxx.xxx through 191.255.xxx.xxx
C (/24 prefixes)	192.0.0.xxx through 223.255.255.xxx

Class A networks are also referred to as ‘/8’s (pronounced slash eight’s or just eight’s) since they have an 8-bit network prefix (one octet is used to designate the network). Following the same convention, Class B networks are called ‘/16s’ and Class C networks ‘/24s’.

to the top

Globally routable and private network IP addresses

There are two *types* of IP addresses – those which are globally routable (included in the routing tables on the Internet), and those which have been set aside for private networks. It is generally recommended that organisations use IP addresses from the blocks of private network addresses for hosts that require IP connectivity within their company network, but do not require external connections to the global Internet.

The system with non-routable IP addresses was introduced to help prevent a future shortage of IP addresses due to the explosive growth of the Internet. Because addresses belonging to these address blocks are not routed through the Internet routing system, the same numbers can be used at the same time by many different organisations.

The three blocks of IP addresses which have been reserved for private networks are:

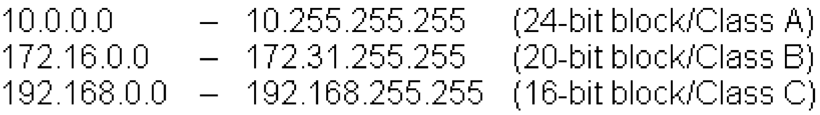

10.0.0.0	–	10.255.255.255	(24-bit block/Class A)
172.16.0.0	–	172.31.255.255	(20-bit block/Class B)
192.168.0.0	–	192.168.255.255	(16-bit block/Class C)

There are no official rules for when to use which of the three private network IP address blocks, but generally the one of the most suitable size is used. For obvious reasons there is no need to use 10.x.x.x if it is unthinkable that your LAN will ever grow to more than 254 hosts. However, when using private addresses the network administrator can be liberal on the usage of the addresses when assigning them to the different parts of a network, as the strict rules that govern public IP address assignment do not apply.

Hosts with private network IP addresses cannot communicate directly with the Internet, because the Internet refuses to receive and transmit data with such origin or destination address. For a host with a private network IP address to be allowed to communicate with the Internet, it must have its data stream to the Internet handled by an intermediary host, which can act as an ‘Internet representative’ for the private host. The intermediary host must have ways to relay data between the global Internet and the host on the private network. Therefore it must have a globally routable IP address that it uses when communicating with the Internet, and a private network IP address that is uses for communication with the private host. There are a number of different types of intermediary hosts that fit this description. The most common types of intermediary hosts are proxy servers, firewalls and firewalls with NAT (Network Address Translation).

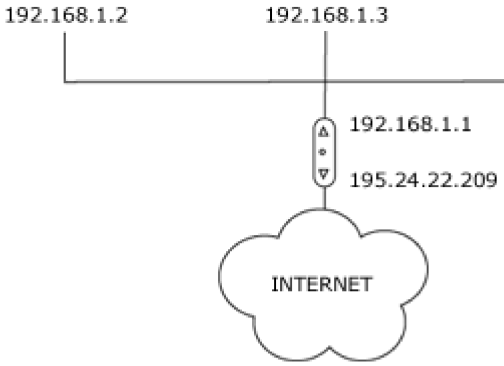

A NAT router translating private network IP addresses to globally routable IP addresses

An advantage of using private network addresses is that it makes it easier for organisations to change their Internet service provider without having to renumber their IP addresses. If private network addresses are not used, renumbering when changing ISP is necessary because globally routable IP addresses are "owned" by the Internet service provider that the company has "leased" the IP addresses from. It is possible to buy and own IP addresses, but this only applies to very large organisations that need in the magnitude of 40,000 globally routable IP addresses.

Using private network IP addresses also gives a company a measure of security. Globally routable IP addresses are advertised in the routing tables on the Internet, making the system vulnerable to hackers. When private IP network addresses are used, however, the intermediary host (such as a firewall with NAT) will work as a barrier against unwanted visits from the Internet.

The current version of IP, IP version 4, defines a 32-bit address, which means that there are only 2^{32} (4,294,967,296) addresses available globally. Over the past few years, the number of available IP addresses on the Internet has started to run out, as the number of companies and people wishing to go on-line has exploded. As a consequence, a new generation of IP addresses (IPv6) is currently in the works. The current IP system will not become obsolete overnight, however, as the two systems will coexist for some time after the new version has been implemented.

to the top

Subnetting

What is subnetting?

A subnet is a segment of a network. Subnetting is a technique that allows a network administrator to divide one physical network into smaller logical networks and, thus, control the flow of traffic for security or efficiency reasons.

Dividing a network into several subnets can serve a number of purposes: to reduce network traffic by decreasing the number of broadcasts (if used in combination with a switch), to exceed the limitations in a local area network, for instance the maximum number of allowed hosts, or to enable employees to be able to dial in to the network from home, without opening the entire network up to unwanted visits from the Internet.

Subnets are created by using a so-called subnet mask to divide a single Class A, B, or C network number into smaller pieces, thus allowing an organisation to add subnets without having to obtain a new network number through an Internet service provider. Subnets can again be subnetted into sub-subnets.

Subnets were originally invented to help solve the lack of IP addresses on the Internet.

*Please note: There is a fast track to getting the advantages of subnetting on local area networks without having to go through the process of calculating a subnet mask, etc. The recipe can be found in the last chapter: '**The fast track to the advantages of subnetting**'.*

to the top

How does subnetting work?

An IP address consists of a network portion and a host portion. A subnet is created by borrowing bits from the part of the IP address which normally designates the host and using them to designate one or more smaller, secondary networks (subnets) within the original network. The network prefix and subnet number in combination are called the extended network prefix (in every day talk often, somewhat confusingly, referred to as the network number!).

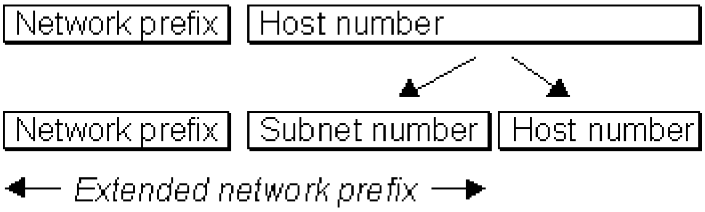

to the top

Subnet masks

A 32-bit subnet mask is used as a deciphering key to determine how an IP address is to be divided into extended network prefix and host part. It is used by routers and network devices to determine where traffic should be routed to.

Like IP addresses, subnet masks consist of four numbers of 8 bits, separated by dots. They are usually written in the corresponding decimal notation.

The typical subnet masks used for Class A, B and C addresses are as follows:

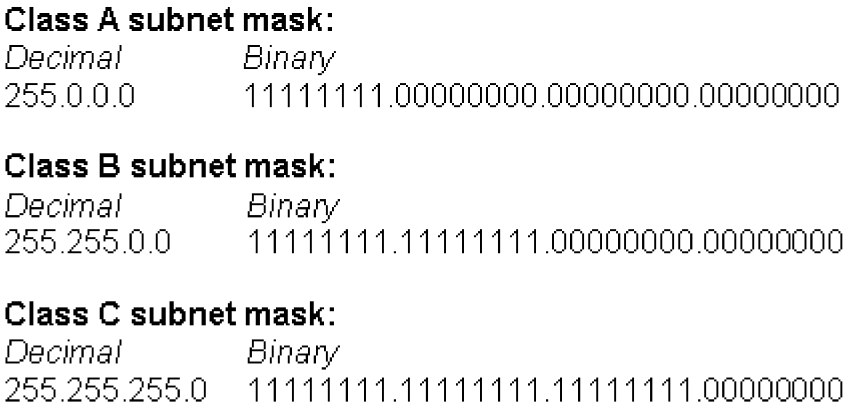

Class A subnet mask:

Decimal	*Binary*
255.0.0.0	11111111.00000000.00000000.00000000

Class B subnet mask:

Decimal	*Binary*
255.255.0.0	11111111.11111111.00000000.00000000

Class C subnet mask:

Decimal	*Binary*
255.255.255.0	11111111.11111111.11111111.00000000

All the 0's in the subnet mask specify that this part in a corresponding IP address is the host portion, while the 1's indicate that the corresponding bits in the IP address constitute the network portion.

The three subnet masks above set the change from network to host portion at the end of a whole octet – Class A after one octet, Class B after two octets, and Class C after three. However, a subnet masks does not have to follow the address classes, but can specifiy a host portion that is not a whole octet.

The subnet mask 255.255.255.240 (11111111.11111111.11111111.1111*0000*) for instance, marks the breaking point four bits into the last octet.

The purpose of having subnet masks defining networks is that the technical devices that the network is build from will be able to determine if traffic should be routed out of the network or kept within it. Using a mask saves the routers from having to handle the entire 32-bit address, because they can simply look at the bits selected by the mask (and thus not worry about the host portion of the address).

Internet routers use only the network number of the destination address to route traffic to a subnetted environment. Subnetting, thus, also has the advantage that it keeps the size of the routing tables on the Internet down because Internet routers only need to know the one common network address for all the individual computers and devices on the different subnets. The route from the Internet to any subnet of a network is the same, no matter which subnet the destination host is on, namely that of the mother network. From there, the local network router(s) divides the communication out into individual subnets and to the individual hosts on these subnets.

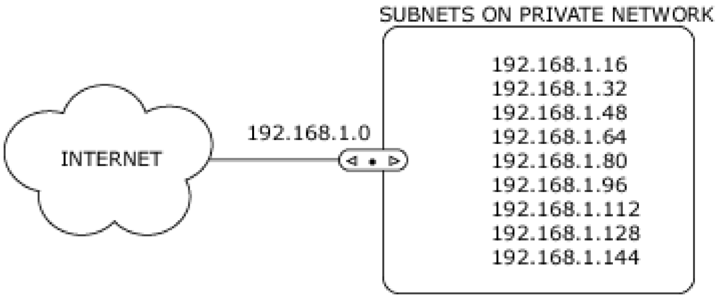

Subnetting keeps the size of the routing tables on the Internet down, as the Internet routers only use the network number of a subnetted environment *to route traffic to any of the subnets*

A router within a subnetted environment uses the extended network prefix to route traffic between the individual subnets. The extended network prefix is composed of the network prefix and the subnet number.

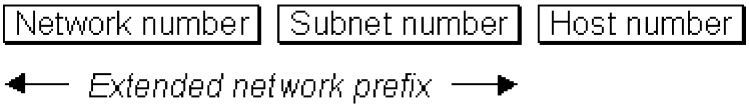

to the top

Calculating a network number using a subnet mask

The network number is the part of the IP address that all hosts on a network share. Network numbers are entered in routing tables and used by routers to direct traffic between networks. The network number, or *extended network prefix,* of an IP address is found by using the subnet mask to mask off the host portion of the IP address.

An example:
You choose the IP address 192.168.1.1 and the subnet mask 255.255.255.0. The above IP address and subnet mask written in their binary notation looks as follows:

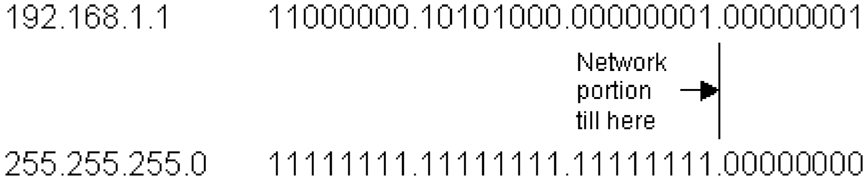

Every bit in the IP address is compared to the corresponding bit in the subnet mask: a '1' in the subnet mask indicates that the corresponding bit in the IP address is part of the network portion, while a '0' in the subnet mask illustrates that the corresponding bit in the IP address is part of the host portion.

In the above example, the host portion is thus all the bits in the first three octets, which in decimal numbers is written 192.168.1.0.

Subnet masks written in binary notation always consist of a continuous string of 1's followed by a continuous string of 0's, e.g.

11111111.11111111.11111111.*00000000* or
11111111.1111*0000.00000000.00000000*

As a consequence, the host range that a subnet mask defines will always be either 2 (2^1 – corresponding to a situation where only the last bit defines hosts), 4(2^2 – corresponding to a situation where the last two bits define hosts), 8 (2^3), 16(2^4), 32(2^5), 64(2^6), 128(2^7) or 256(2^8).

In reality, 2 must be subtracted from all the numbers of hosts above to get the actual number of IP addresses available to use for hosts, because two addresses, namely the address which has all-0's in the host bits (this network) and the address which has all 1's in the host bits (broadcast), can not be assigned to hosts. As a consequence, it is not possible to make a network that consists of fewer than four IP addresses (2 hosts + the broadcast and network addresses).

In the above example, based on the IP address 192.168.1.1 and the subnet mask 255.255.255.0, the network address (all host bits set to 0) was 192.168.1.0. The broadcast address for this network would be 192.168.1.255 as illustrated below.

Calculating a broadcast address using a subnet mask

The broadcast address is the address where all the bits in the host portion are set to 1. The broadcast address is used when you want to communicate data to all the hosts on a network. Here follows an example of how it can be calculated:

In our example above, the last 8 bits were hosts. As a consequence, the broadcast address for the network 192.168.1.0 with the subnet mask 255.255.255.0 is 11000000.10101000.00000001.**111111111** (host bits set to 1) or in decimal notation: 192.168.1.255

Note: If you know the IP address segment your network consists off, the lowest IP address is the network number, while the highest IP address is the broadcast address.

Prefix length notation (CIDR notation)

For the sake of convenience, prefix length notations (CIDR notation, Classless Inter-Domain Routing notation) is often used instead of writing the subnet mask. This means that the IP address above (192.168.1.1) with the subnet mask 255.255.255.0 can also be expressed as 192.168.1.1/24.
The /24 indicates the network prefix length, which is equal to the number of continuous one-bits in the subnet mask.

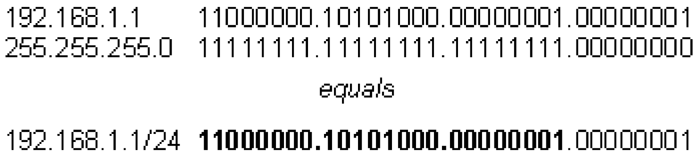

192.168.1.1 11000000.10101000.00000001.00000001
255.255.255.0 11111111.11111111.11111111.00000000

equals

192.168.1.1/24 **11000000.10101000.00000001**.00000001

to the top

Calculating a subnet mask

When subnetting a network, you first need to determine two things:

- how many subnets do you need to create?
- how many host addresses do you need on each net (you should always add some extra host addresses to be used for future growth).

Once you have determined the required number of subnets and hosts, the next step is to calculate a corresponding subnet mask, which will support the desired network structure.

In the following you will find two examples of how the subnetting of a Class C network can be planned and the required subnet mask calculated.

Example A:
Imagine that you are setting up a network on the network number 192.168.1.0/24. You need a local area network which is going to connect a number of workstations, servers and others devices, totalling more than 80. To allow some slack, you set the number of required hosts to be 90. Now, the calculation of the subnet mask can begin. The calculation is best understood if the numbers are looked at in their binary form (see example below).

The first step is to determine the lowest number of bits required to identify 90 hosts. Since IP addresses of hosts can only be created along binary boundaries, the number of hosts must be created in blocks of powers of two – 2 (2^1), 4 (2^2), 8 (2^3), 16 (2^4) and so on. In other words, we must first determine what the lowest power is that we can lift 2 to and get a number equal to or greater than 90. Since 2^7 equals 128 and 2^6 equals 64, we need 7 bits to designate 90 hosts. This means that the host portion of the IP address must be the last 7 bits. An IP address consists of 32 bits all in all. The network portion must thus consist of 32-7=25 bits. As every '1' in a subnet mask indicates that the corresponding bit in the IP address belongs to the network portion and every '0' indicates that the corresponding bit in the IP address is part of the host portion, the corresponding subnet mask must consist of a series of 25 1's, followed by 7 0's (as illustrated below). Written in decimal notation, the subnet mask is 255.255.255.128.

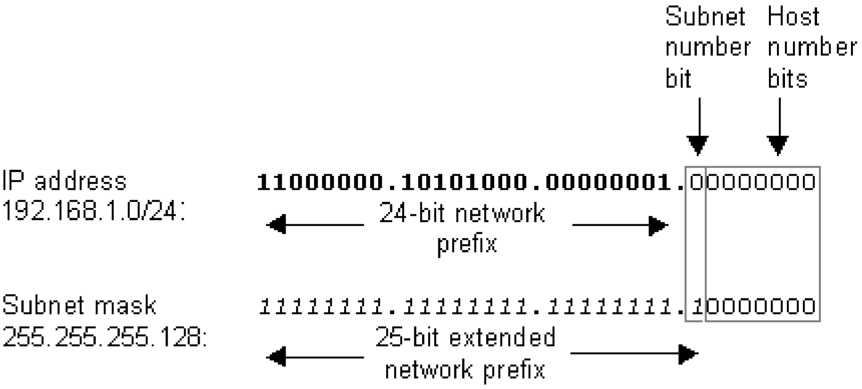

The number of subnets that can be created using this subnet mask is calculated as follows: The original network prefix was 24 bit long (192.168.1.0/24), and the extended network prefix (network prefix + subnet prefix) masked off by the subnet mask is 25 bits long. As a consequence, one bit is available to designate subnets. In other words, it is possible to create 2 (2^1) subnets of this given size using this subnet mask, should we wish to do so.

Example B:
Now pretend that through an estimation of the number of subnets and hosts that the subnet you are setting up will have to support, you have come to the conclusion that you need to define ten subnets. The largest subnet is required to support 10 hosts. You have again chosen to create the subnet on the network number 192.168.1.0/24. Now, the calculation of the subnet mask can begin.

The first step is to determine the number of bits required to define the ten subnets. Since a network address can be subnetted only along binary boundaries, subnets must be created in blocks of powers of two 2 – 2, 4, 8, 16 and so on. Thus, it is impossible to define an IP address block so that it contains exactly ten subnets. In this case, the network administrator must define a block of 16 (2^4) and have six unused subnet addresses for future growth.

Since we need to raise 2 to the power of four (2^4) to get 16, four bits are required to designate the sixteen subnets in the block. In this example, you are subnetting a Class C network (/24) so it will need four more bits (/28) as the extended network prefix. A 28-bit extended network prefix can be expressed in dotted-decimal notation as 255.255.255.240. This is illustrated below:

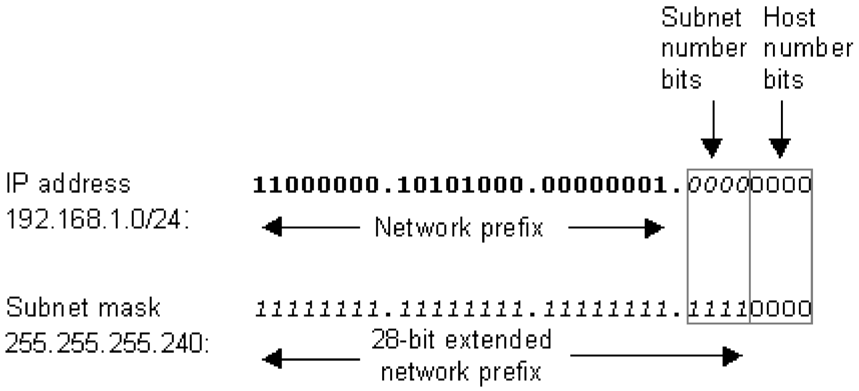

A 28-bit extended network prefix leaves 4 bits to define host addresses on each subnet. This means that each subnet with a 28-bit prefix represents a continuous block of 2^4 (16) individual IP addresses. However, since the all-0's ('this network') and the all-1's host addresses ('broadcast') must not be allocated, there are 14 (2^4-2) assignable host addresses on each subnet. We needed a maximum of 10 hosts on each subnet, so the result is satisfactory.

Defining subnet numbers

In example B above, with the ten subnets, the subnets will be numbered 0 through to 9. The 4-bit binary representation of the decimal values 0 though 9 are: 0 (0000), 1 (0001), 2 (0010), 3 (0011), 4 (0100), 5 (0101), 6 (0110), 7 (0111), 8 (1000), 9 (1001), 10 (1010).

To find the subnet number of each subnet, place the binary representation of the subnet number, e.g. 0001, into the bits in the base network address that is used to designate the subnet (see illustration below). For example, to define subnet number 8, the network administrator places the binary representation of 8 (1000) into the 4-bits in the base network address that are used to designate the subnet.

The ten subnet numbers for the example are given below. The italicised portion of each address identifies the extended network prefix, while the bold digits identify the 4 bits representing the bits in the address that are used to designate the subnet:

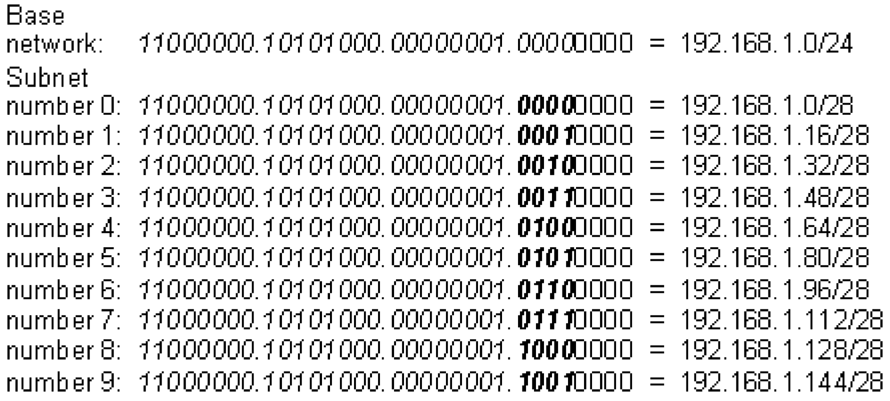

Base network:	*11000000.10101000.00000001.0000*0000	=	192.168.1.0/24
Subnet number 0:	*11000000.10101000.00000001.**0000***0000	=	192.168.1.0/28
number 1:	*11000000.10101000.00000001.**0001***0000	=	192.168.1.16/28
number 2:	*11000000.10101000.00000001.**0010***0000	=	192.168.1.32/28
number 3:	*11000000.10101000.00000001.**0011***0000	=	192.168.1.48/28
number 4:	*11000000.10101000.00000001.**0100***0000	=	192.168.1.64/28
number 5:	*11000000.10101000.00000001.**0101***0000	=	192.168.1.80/28
number 6:	*11000000.10101000.00000001.**0110***0000	=	192.168.1.96/28
number 7:	*11000000.10101000.00000001.**0111***0000	=	192.168.1.112/28
number 8:	*11000000.10101000.00000001.**1000***0000	=	192.168.1.128/28
number 9:	*11000000.10101000.00000001.**1001***0000	=	192.168.1.144/28

An easy way to ensure that the subnets are calculated correctly is to ensure that they are all multiples of the subnet number 1 address. In this case, all subnets are multiples of 16.

to the top

The fast track to the advantages of subnetting

There is a fast track to getting the advantages of subnetting on local area networks without having to go through the process of calculating a subnet mask, etc. The fast track involves using a standard class subnet mask in combination with addresses from the IP address blocks set aside for private networks. For instance by using 192.168.0.0 to designate your local network 1, 192.168.1.0 to designate your local area network 2 and 192.168.2.0 to designate your local area network 3. The standard subnet mask to use on networks with up to 254 hosts is 255.255.255.0.

The three blocks of IP addresses which have been reserved for private networks and the corresponding standard subnet masks are:

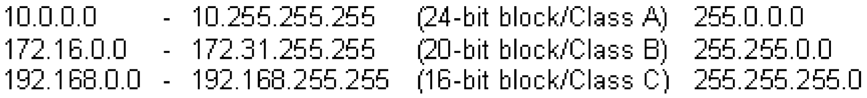

10.0.0.0	-	10.255.255.255	(24-bit block/Class A)	255.0.0.0
172.16.0.0	-	172.31.255.255	(20-bit block/Class B)	255.255.0.0
192.168.0.0	-	192.168.255.255	(16-bit block/Class C)	255.255.255.0

If you do not want to use this method, a list of precalculated subnet masks together with the number of hosts available on the networks they create can be found below for easy reference.

List of subnet masks

The list below can be used as a fast track when subnetting. It describes the relationship between the number of host IP addresses required and the corresponding subnet mask. The example above, for instance, required subnets with 10 host addresses on each. The nearest number that is equal to or greater than 10 is 16. The subnet mask corresponding to 16 hosts is listed in the table below. It is 255.255.255.240.

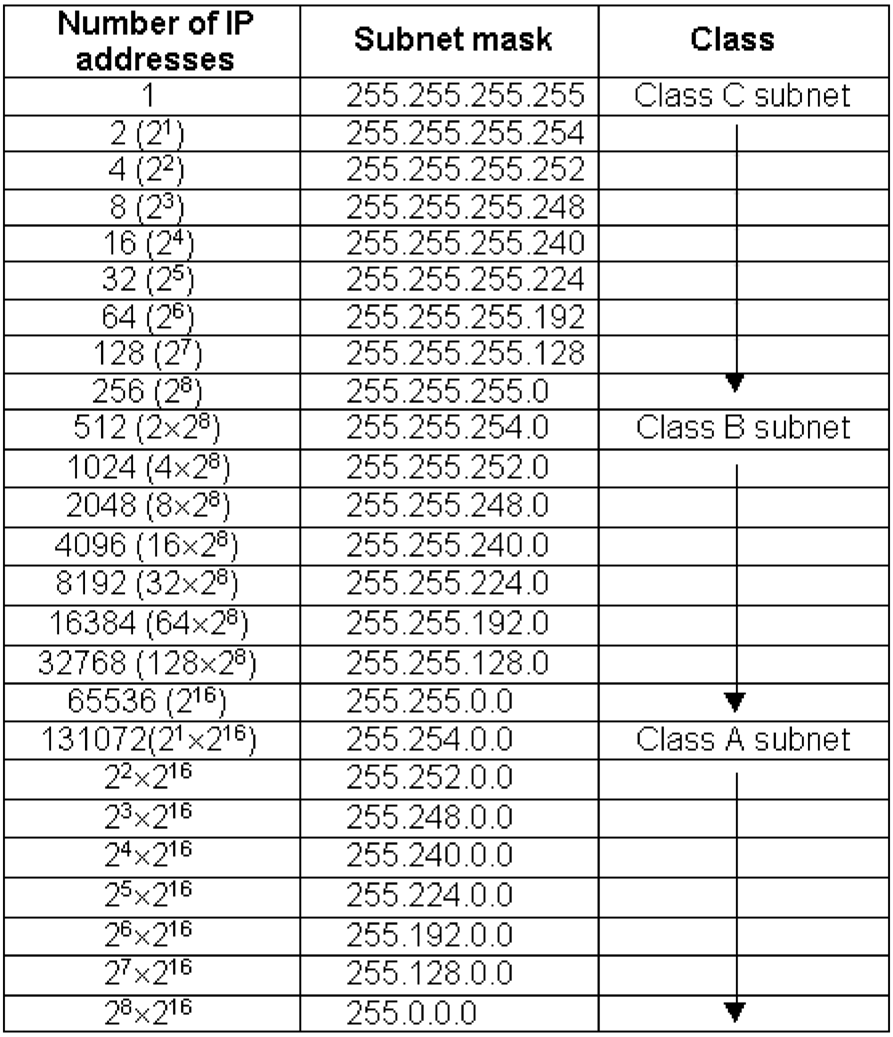

Number of IP addresses	Subnet mask	Class
1	255.255.255.255	Class C subnet
2 (2^1)	255.255.255.254	
4 (2^2)	255.255.255.252	
8 (2^3)	255.255.255.248	
16 (2^4)	255.255.255.240	
32 (2^5)	255.255.255.224	
64 (2^6)	255.255.255.192	
128 (2^7)	255.255.255.128	
256 (2^8)	255.255.255.0	↓
512 (2×2^8)	255.255.254.0	Class B subnet
1024 (4×2^8)	255.255.252.0	
2048 (8×2^8)	255.255.248.0	
4096 (16×2^8)	255.255.240.0	
8192 (32×2^8)	255.255.224.0	
16384 (64×2^8)	255.255.192.0	
32768 (128×2^8)	255.255.128.0	
65536 (2^{16})	255.255.0.0	↓
131072($2^1 \times 2^{16}$)	255.254.0.0	Class A subnet
$2^2 \times 2^{16}$	255.252.0.0	
$2^3 \times 2^{16}$	255.248.0.0	
$2^4 \times 2^{16}$	255.240.0.0	
$2^5 \times 2^{16}$	255.224.0.0	
$2^6 \times 2^{16}$	255.192.0.0	
$2^7 \times 2^{16}$	255.128.0.0	
$2^8 \times 2^{16}$	255.0.0.0	↓

to the top

Test your knowledge

1.

What is a network number?

○ The part of an IP address that all hosts on a network share

○ The part of an IP address which networks share

○ The part of an IP address which no hosts on a network share

2.

What is a host number?

○ The part of an IP address which all hosts on a network share

○ The part of an IP address which networks share

○ The part of an IP address which no hosts on a network share

3.

How many hosts can you set up on a Class C network (without subnetting)?

○ 254

○ 256

○ 16,384

4.

A '/8' is also referred to as?

A class A network

○

○ A class B network

○ A class C network

5.

What is a private network IP address?

○ The IP address of a secret server on the Internet

○ An IP address which is included in the routing tables on the Internet

○ An IP address which is NOT included in the routing tables on the Internet

6.

You are setting up a LAN with 20 hosts. Which of the following private network IP address blocks does it make the most sense to choose your IP addresses from?

○ 10.0.0.0-10.255.255.255

○ 172.16.0.0-172.31.255.255

○ 192.168.0.0-192.168.255.255

7.

What is subnetting?

○ The division of a physical network into two or more physical networks

○ The division of a logical network into two or more physical networks

○ The division of a physical network into one or more logical networks

8.

How does subnetting work?

○ Bits from the network portion of the IP address are borrowed to designate the subnetwork

○ Bits from the host portion of the IP address are borrowed to designate the subnetwork

○ An additional cable is attached to the servers LAN port

9.

What is a subnetmask?

○ A deciphering key used to determine which part of an IP address constitutes the Host and Network portions respectively

○ The network number of a subnet

○ The network of a number before it is subnettet

10.

What does the 0's in a subnet mask (written in its binary form) mean?

○ They indicate that this part in the corresponding IP address is the network portion

○ They indicate that this part in the corresponding IP address is the host portion

○ They indicate that the network has no subnets

Evaluate Reset