8 Steps to Understanding IP Subnetting

Introduction

Understanding IP subnetting is a fundamental requirement for almost any techie - whether you're a coder, a database administrator or the CTO. However, as simple as the concepts are, there is a general difficulty in understanding the topic.

Here we'll break this topic into eight simple steps and help you put the pieces together to fully understand IP subnetting.

These steps will give you the basic information needed in order to configure routers or understand how IP addresses are broken down and how subnetting works. You'll also learn how to plan a basic home or small office network.

A basic understanding of how binary and decimal numbers work is required. In addition, these definitions and terms will get you started:

- IP Address: A logical numeric address that is assigned to every single computer, printer, switch, router or any other device that is part of a TCP/IP-based network
- Subnet: A separate and identifiable portion of an organization's network, typically arranged on one floor, building or geographical location
- Subnet Mask: A 32-bit number used to differentiate the network component of an IP address by dividing the IP address into a network address and host address
- Network Interface Card (NIC): A computer hardware component that allows a computer to connect to a network

Step 1 - Why We Need Subnets

To understand why we need subnets (short for subnetwork), let's start right from the beginning and recognize that we need to talk to "things" on networks. Users need to talk to printers, email programs need to talk to servers, and each of these "things" needs to have some sort of address. This is no different from a house address, but with one minor exception: the addresses need to be in numerical form. It is not possible to have a device on a network that has alphabetical characters in its address like "23rd Street." Its name can be alphanumeric - and we could translate that name to a numeric address - but the address itself must be numbers alone.

These numbers are called IP addresses, and they have the important function of figuring out not only the address of "things," but how communication can occur between them. It is not enough to just have an address. It is necessary to figure out how a message can be sent from one address to another.

This is where a little organization comes into play.

It is often necessary to group things on a network together for both organizational and efficiency's sake. For example, let's say you have a group of printers in your company's marketing department and a different bunch in the sales offices. You want to limit the printers that each user sees to those of each department. You could accomplish this by organizing the addresses of these printers into unique subnets.

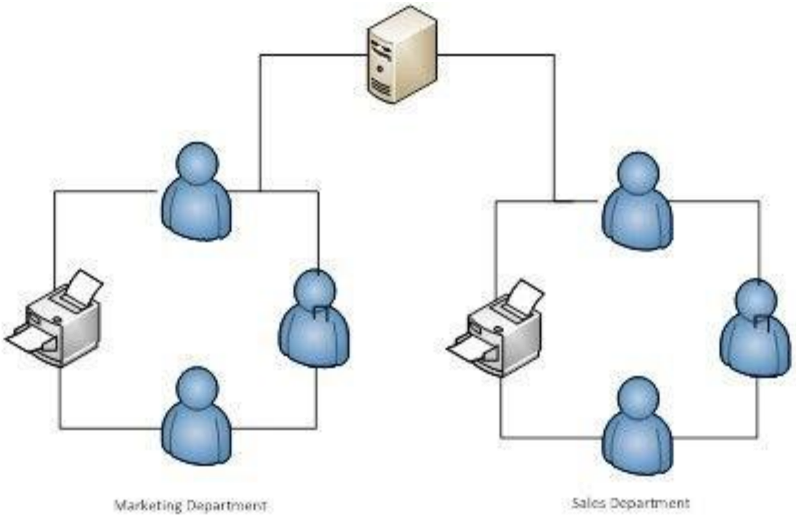

A subnet then, is a logical organization of connected network devices.

Each device on each subnet has an address that logically associates it with the others on the same subnet. This also prevents devices on one subnet from getting confused with hosts on the other subnet.

In terms of IP addressing and subnets, these devices are referred to as hosts. So, in our example, there is a network (the company), which is divided into logical subnets (marketing and sales departments), each of which has its own hosts (users and printers).

Step 2 - Understanding Binary Numbers

Just the sound of "binary numbers" sends pangs of fear through many people with different shades of arithmophobia (the irrational fear of numbers and arithmetic). Have no fear - or at least put your fear to rest. Binary numbers are just a different way to count. That is all. The concept is as easy as one plus one.

Appreciate that we use the *decimal numbering system* in our everyday lives, where our numbers are based on 10s of things - probably because we have 10 toes and 10 fingers. All the decimal system has are symbols that represent quantities. We call the straight vertical line a "1" and the round circle a "0".

That does not change with *binary numbering systems*.

With the decimal system, we can represent larger and larger numbers by tacking numbers together. So, there are single-digit numbers, like 1, double-digit numbers, like 12, triple-digit numbers, like 105, and so on and so on. As numbers get larger, each digit represents a progressively greater value. There is a 1's place, a 10’s place, a 100’s place and so on.

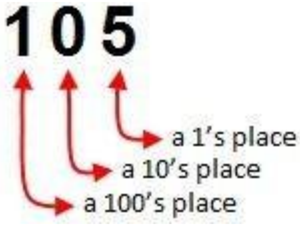

With this number, we have a 5 in the 1’s place, a 0 in the 10’s place and a 1 in the 100’s place. Hence,

1 x 100 + 0 x 10 + 5 x 1 = 105

Binary numbering systems are based on the same concept except that because the binary system only has two numbers, 0 and 1, it takes a lot more groupings to represent the same number. For example, the binary equivalent of 105 is 01101001 (actually, it would be usually written as 1101001 because just like in the decimal numbering system, leading zeros are dropped. However, we'll keep that first zero in place in order to explain the next concept).

Once again, as binary numbers get larger, each digit represents a progressively greater value, but now the binary system has a 1's place, a 2's place, a 4's place, 8's place, a 16's place, a 32's place and so on.

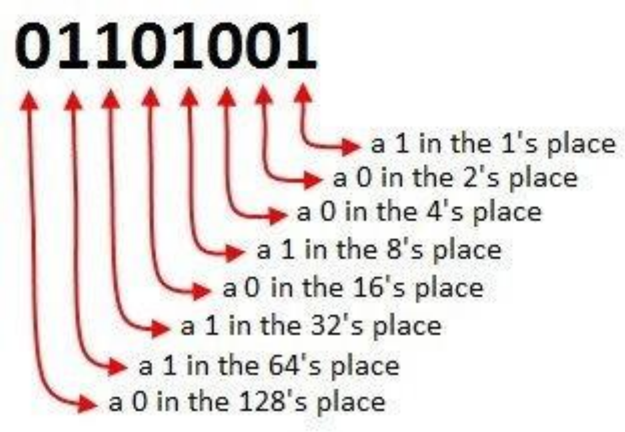

Hence,

0 x 128 + 1 x 64 + 1 x 32 + 0 x 16 + 1 x 8 + 0 x 4 + 0 x 2 + 1 x 1

equals:

0 + 64 + 32 + 0 + 8 + 0 + 0 + 1 = 105

Step 3 - IP Addresses

The "IP" in IP addresses refers to the Internet Protocol, where protocol is loosely defined as "rules of communication". Imagine using a two-way radio in a police car. Your conversations would probably end with "over" to indicate you are finishing a particular part of the conversation. You might also say "over and out" when you are finished the conversation itself. These are nothing more than the rules of talking over a two-way radio - or the protocol.

So, IP addressing must be understood as part of the rules for conversations over the Internet. But it has grown so popular that it is also used on most any network connected to the Internet, making it safe to say IP addressing is relevant for most networks as well as the Internet.

So what is an IP address? Technically, it is the means whereby an entity on a network can be addressed. It is made up solely of numbers, and these numbers are conventionally written in the particular form of XXX.XXX.XXX.XXX, which is referred to as dotted decimal format.

Any one of the numbers between the dots can be between 0 and 255, so example IP addresses include:

- 205.112.45.60
- 34.243.44.155

These numbers can also be written in binary form by taking each of the decimal values separated by dots and converting to binary. So a number like 205.112.45.60 could be written as:

11001101.01110000.00101101.00111100

Each of these binary components is referred to as an octet, but this term is not often used in subnetting practice. It does seem to come up in classrooms and books, so know what it is (and then forget about it).

Why is each number limited to 0 to 255? Well, IP addresses are limited to 32 bits in length and the maximum number of combinations of binary numbers you could have in an octet is 256 (mathematically calculated as 28). Hence, the largest IP address you could have would be 255.255.255.255, given that any one octet could be from 0 to 255.

There is one more aspect of an IP address that is important to understand **-** the concept of a class.

Each IP address belongs to a class of IP addresses depending on the number in the first octet. These classes are:

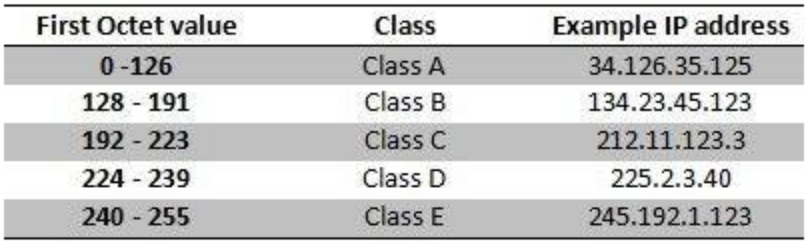

First Octet value	Class	Example IP address
0 -126	Class A	34.126.35.125
128 - 191	Class B	134.23.45.123
192 - 223	Class C	212.11.123.3
224 - 239	Class D	225.2.3.40
240 - 255	Class E	245.192.1.123

Notice that the number 127 is not included. That's because it is used in a special, self reflecting number called a loopback address. Think of this as an address that says, "this is my address." Note that only the first three classes **-** A, B and C **-** are used by network administrators. These are the commonly used classes. The other two, D and E, are reserved.

You define the class of an IP address by looking at its first octet value, but the structure of an IP address for any one class is different. Each IP address has a network address and a host address. The network part of the address is the common address for any one network, while the host address part is for each individual device on that network. So, if your phone number is 711-612-1234, the area code (711) would be the common, or network, component of the telephone system, while your individual phone number of (612-1234) would be your host address.

The network and host components of class IP addresses are:

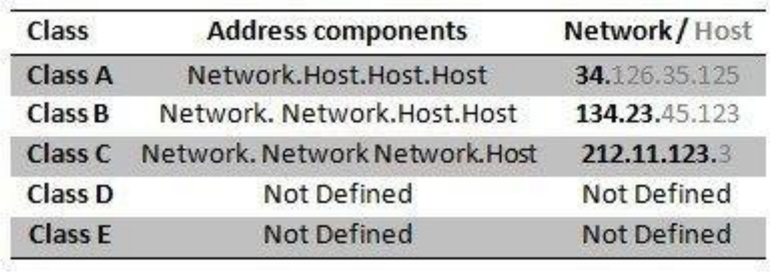

Class	Address components	Network / Host
Class A	Network.Host.Host.Host	**34.**126.35.125
Class B	Network. Network.Host.Host	**134.23.**45.123
Class C	Network. Network Network.Host	**212.11.123.**3
Class D	Not Defined	Not Defined
Class E	Not Defined	Not Defined

The technical numbers behind class addressing are as follows:

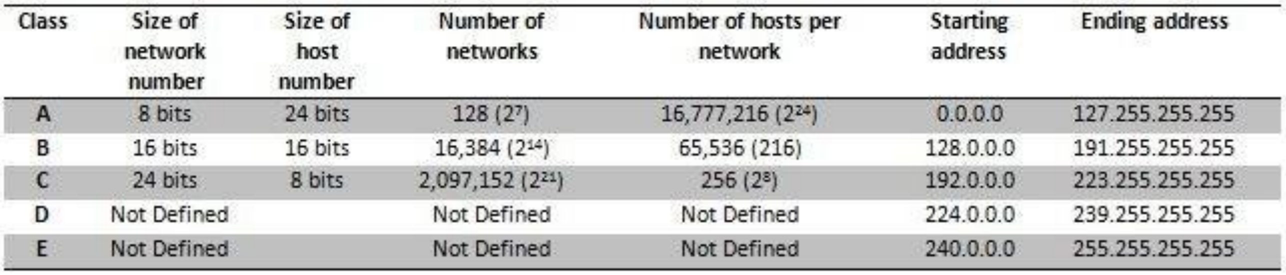

Class	Size of network number	Size of host number	Number of networks	Number of hosts per network	Starting address	Ending address
A	8 bits	24 bits	128 (2^7)	16,777,216 (2^{24})	0.0.0.0	127.255.255.255
B	16 bits	16 bits	16,384 (2^{14})	65,536 (216)	128.0.0.0	191.255.255.255
C	24 bits	8 bits	2,097,152 (2^{21})	256 (2^8)	192.0.0.0	223.255.255.255
D	Not Defined		Not Defined	Not Defined	224.0.0.0	239.255.255.255
E	Not Defined		Not Defined	Not Defined	240.0.0.0	255.255.255.255

Step 4 - Subnetting and the Subnet Mask

To subnet a network is to create logical divisions of the network. Subnetting, therefore, involves dividing the network into smaller portions called subnets. Subnetting applies to IP addresses because this is done by borrowing bits from the host portion of the IP address. In a sense, the IP address then has three components - the network part, the subnet part and, finally, the host part.

We create a subnet by logically grabbing the last bit from the network component of the address and using it to determine the number of subnets required. In the following example, a Class C address normally has 24 bits for the network address and eight for the host, but we are going to borrow the left-most bit of the host address and declare it as identifying the subnet.

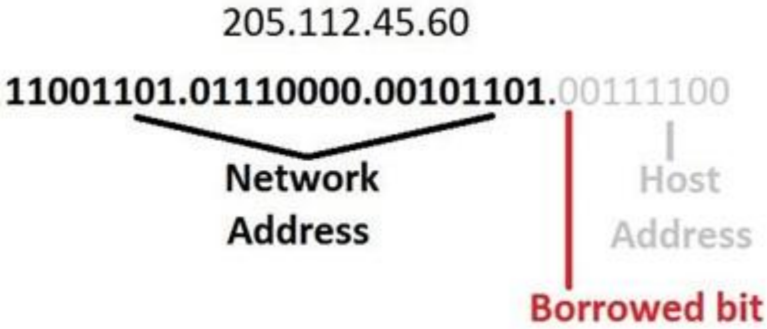

If the bit is a 0, then that will be one subnet; if the bit is a 1, that would be the second subnet. Of course, with only one borrowed bit we can only have two possible subnets. By the same token, that also reduces the number of hosts we can have on the network to 127 (but actually 125 useable addresses given all zeros and all ones are not recommended addresses), down from 255.

So how can you tell how many bits should be borrowed, or, in other words, how many subnets we want to have on our network?

The answer is with a subnet mask.

Subnet masks sound a lot scarier than they really are. All that a subnet mask does is indicate how many bits are being “borrowed” from the host component of an IP address. If you can’t remember anything about subnetting, remember this concept. It is the foundation of all subnetting.

The reason a subnet mask has this name is that it literally masks out the host bits being borrowed from the host address portion of the IP address.

In the following diagram, there is a subnet mask for a Class C address. The subnet mask is 255.255.255.128 which, when translated into bits, indicates which bits of the host part of the address will be used to determine the subnet number.

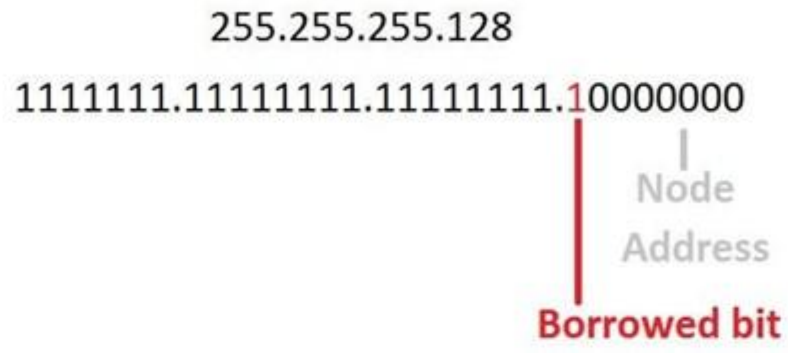

Of course, more bits borrowed means fewer individually addressable hosts that can be on the network. Sometimes, all the combinations and permutations can be confusing, so here are some tables of subnet possibilities.

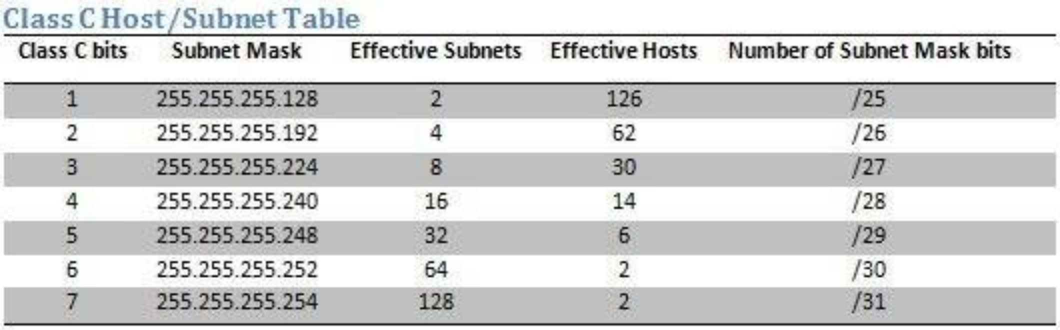

Class C Host/Subnet Table

Class C bits	Subnet Mask	Effective Subnets	Effective Hosts	Number of Subnet Mask bits
1	255.255.255.128	2	126	/25
2	255.255.255.192	4	62	/26
3	255.255.255.224	8	30	/27
4	255.255.255.240	16	14	/28
5	255.255.255.248	32	6	/29
6	255.255.255.252	64	2	/30
7	255.255.255.254	128	2	/31

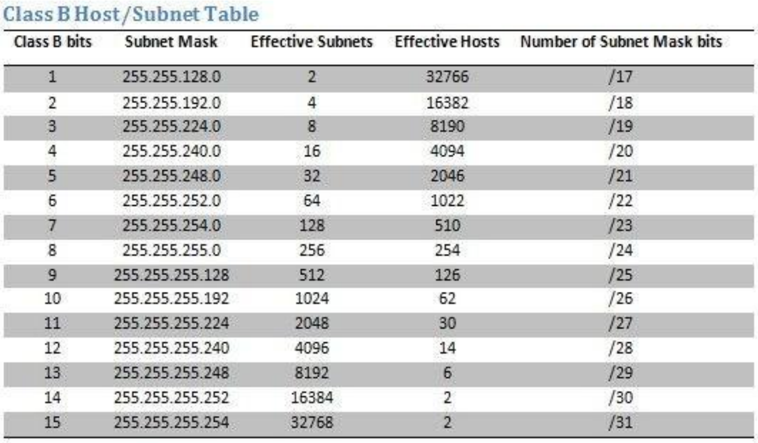

Class B Host/Subnet Table

Class B bits	Subnet Mask	Effective Subnets	Effective Hosts	Number of Subnet Mask bits
1	255.255.128.0	2	32766	/17
2	255.255.192.0	4	16382	/18
3	255.255.224.0	8	8190	/19
4	255.255.240.0	16	4094	/20
5	255.255.248.0	32	2046	/21
6	255.255.252.0	64	1022	/22
7	255.255.254.0	128	510	/23
8	255.255.255.0	256	254	/24
9	255.255.255.128	512	126	/25
10	255.255.255.192	1024	62	/26
11	255.255.255.224	2048	30	/27
12	255.255.255.240	4096	14	/28
13	255.255.255.248	8192	6	/29
14	255.255.255.252	16384	2	/30
15	255.255.255.254	32768	2	/31

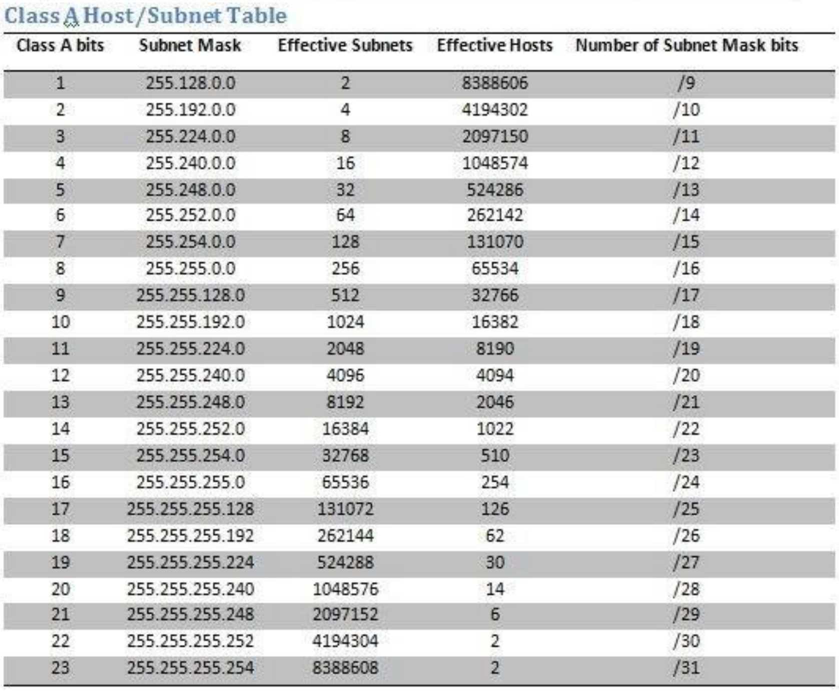

Class A Host/Subnet Table

Class A bits	Subnet Mask	Effective Subnets	Effective Hosts	Number of Subnet Mask bits
1	255.128.0.0	2	8388606	/9
2	255.192.0.0	4	4194302	/10
3	255.224.0.0	8	2097150	/11
4	255.240.0.0	16	1048574	/12
5	255.248.0.0	32	524286	/13
6	255.252.0.0	64	262142	/14
7	255.254.0.0	128	131070	/15
8	255.255.0.0	256	65534	/16
9	255.255.128.0	512	32766	/17
10	255.255.192.0	1024	16382	/18
11	255.255.224.0	2048	8190	/19
12	255.255.240.0	4096	4094	/20
13	255.255.248.0	8192	2046	/21
14	255.255.252.0	16384	1022	/22
15	255.255.254.0	32768	510	/23
16	255.255.255.0	65536	254	/24
17	255.255.255.128	131072	126	/25
18	255.255.255.192	262144	62	/26
19	255.255.255.224	524288	30	/27
20	255.255.255.240	1048576	14	/28
21	255.255.255.248	2097152	6	/29
22	255.255.255.252	4194304	2	/30
23	255.255.255.254	8388608	2	/31

Note that this combination of IP addresses and subnet masks in the charts are written as two separate values, such as Network Address = 205.112.45.60, Mask = 255.255.255.128, or as an IP address with the number of bits indicated as being used for the mask, like 205.112.45.60/25.

Subnet masks work because of the magic of Boolean logic. To best understand how a subnet mask actually does its thing, you must remember that a subnet mask is only relevant when getting to a subnet. In other words, determining what subnet an IP address lives on is the only reason for a subnet mask. It’s devices like routers and switches that make use of subnet masks.

Step 5 - Public Vs. Private IP Addresses

Technically, if all the possible combinations of IP addresses were available, there would be about 4,228,250,625IP addresses for use. This would have to include all public uses *and* private uses - which would then mean, by definition, there would be nothing but public IP addresses.

However, not all addresses are available. Some are used for special purposes. For example, any IP address ending in 255 is a special broadcast address.

Other addresses are used for special signaling, including:

- Loopback (127.0.0.1) when a host is referring to itself
- Multicast routing mechanisms
- Limited broadcasts sent to every host, but limited to the local subnet
- Directed broadcasts first routed to a specific subnet, and then broadcast to all hosts on that subnet

The concept of a private address is similar to that of a private extension in an office phone system. Someone who wants to call an individual in a company would dial the company’s public phone number, through which all employees can be reached. Once connected, the caller would enter in the extension number of the person to whom they wished to speak. Private IP addresses are to IP addresses what extension numbers are to phone systems.

Private IP addresses allow network administrators to extend the size of their networks. A network could have one public IP address that all traffic on the Internet sees, and hundreds - or even thousands - of hosts with private IP addresses on the company subnet.

Anyone can use a private IP address on the understanding that all traffic using these addresses must remain local. It would not be possible, for example, to have an email message associated with a private IP address to move across the Internet, but it is quite reasonable to have the same private IP address work well in the company network.

The private IP addresses that you can assign for a private network can be from the following three blocks of the IP address space:

- 10.0.0.1 to 10.255.255.255: Provides a single Class A network of addresses
- 172.16.0.1 to 172.31.255.254: Provides 16 contiguous Class B network addresses
- 192.168.0.1 to 192.168.255.254: Provides up to 216 Class C network addresses

A typical network setup using public and private IP addresses with a subnet mask would look like:

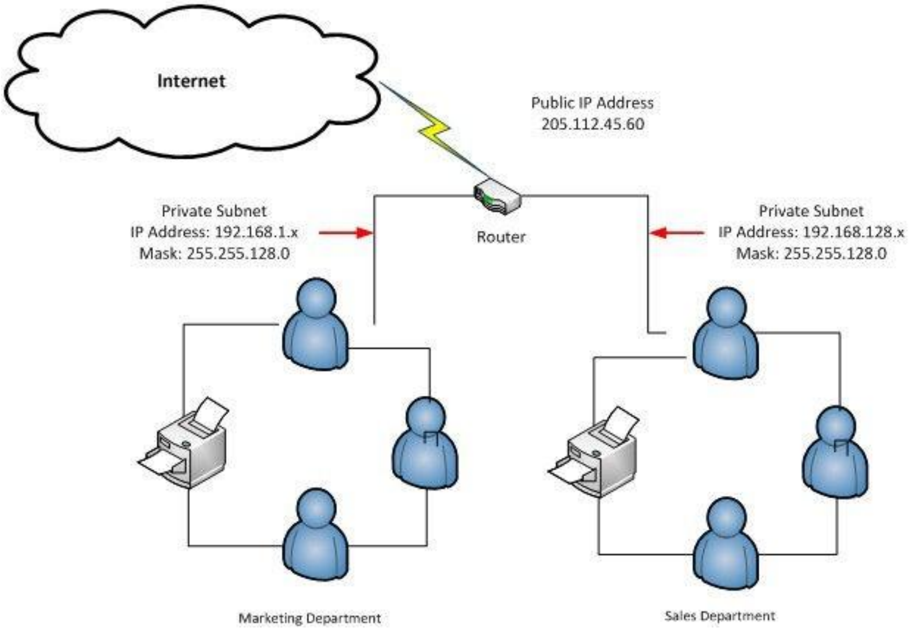

Step 6 - CIDR IP Addressing

Having spent a whole bunch of time learning about IP addresses and classes, you might be surprised that in reality they are not used anymore other than to understand the basic concepts of IP addressing.

Instead, network administrators use Classless Internet Domain Routing (CIDR), pronounced "cider", to represent IP addresses. The idea behind CIDR is to adapt the concept of subnetting to the entire Internet. In short, classless addressing means that instead of breaking a particular network into subnets, we can aggregate networks into larger supernets.

CIDR is therefore often referred to as supernetting, where the principles of subnetting are applied to larger networks. CIDR is written out in a network/mask format, where the mask is tacked onto the network address in the form of the number of bits used in the mask. An example would be 205.112.45.60/25. What is most important to understand about the CIDR method of subnetting is the use the network prefix (the /25 of 205.112.45.60/25), rather than the classful way of using the first three bits of the IP address to determine the dividing point between the network number and the host number.

The process for understanding what this means is:

1. The “205” in the first octet means this IP address would normally contain 24 bits to represent the network portion of the address. With eight bits to an octet, the arithmetic is 3 x 8 = 24, or looking at it the other way around, “/24” means no bits are being borrowed from the last octet.
2. But this is “/25,” which indicates it is “borrowing” one bit from the host portion of the address.
3. With only one bit, there can only be two unique subnets.
4. So this is the equivalent of a net mask of 255.255.255.128, where there is a maximum of 126 host addresses addressable on each of the two subnets.

So why did CIDR become so popular? Because it’s a much more efficient allocator of the IP address space. Using CIDR, a network admin can carve out a number of host addresses that’s closer to what is required than with the class approach.

For example, say a network admin has an IP address of 207.0.64.0/18 to work with. This block consists of 16,384 IP addresses. But if only 900 host addresses are required, this wastes scarce resources, leaving 15,484 (16,384 – 900) addresses unused. By using a subnet CIDR of 207.0.68.0/22 though, the network would address 1,024 nodes, which is much closer to the 900 host addresses required.

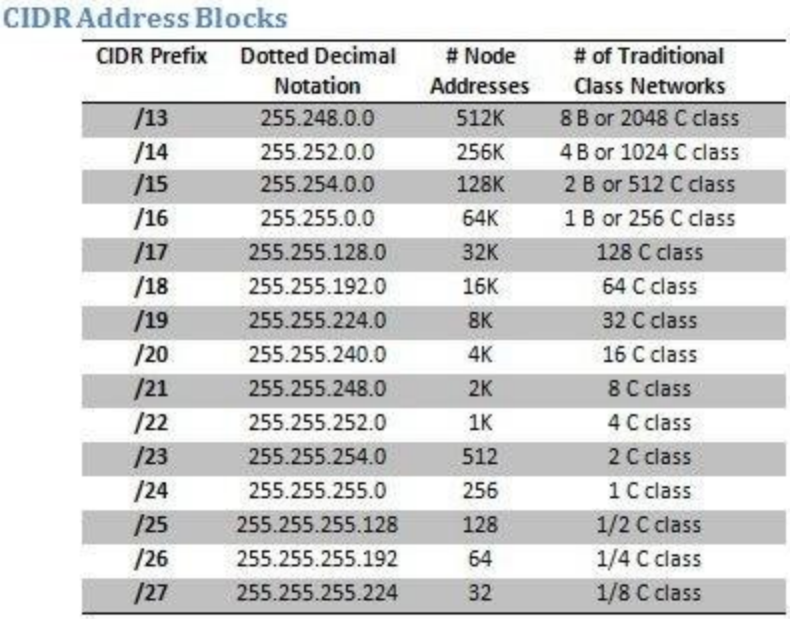

CIDR Address Blocks

CIDR Prefix	Dotted Decimal Notation	# Node Addresses	# of Traditional Class Networks
/13	255.248.0.0	512K	8 B or 2048 C class
/14	255.252.0.0	256K	4 B or 1024 C class
/15	255.254.0.0	128K	2 B or 512 C class
/16	255.255.0.0	64K	1 B or 256 C class
/17	255.255.128.0	32K	128 C class
/18	255.255.192.0	16K	64 C class
/19	255.255.224.0	8K	32 C class
/20	255.255.240.0	4K	16 C class
/21	255.255.248.0	2K	8 C class
/22	255.255.252.0	1K	4 C class
/23	255.255.254.0	512	2 C class
/24	255.255.255.0	256	1 C class
/25	255.255.255.128	128	1/2 C class
/26	255.255.255.192	64	1/4 C class
/27	255.255.255.224	32	1/8 C class

Step 7 - Variable Length Subnet Masking

When an IP network is assigned more than one subnet mask, it is said to a have a variable length subnet mask (VLSM). This is what is required when you are subnetting a subnet. The concept is very straightforward: Any one subnet can be broken down into further subnets by indicating the proper VLSM.

What must be appreciated about VLSM is how RIP 1 routers work. Originally, the IP addressing scheme and RIP 1 routing protocol did not take into consideration the ability to have different subnet masks on the same network. When a RIP 1 router receives a packet destined for a subnet, it has no idea of the VLSM that has been used to generate the packet address. It just has an address to work with without any knowledge of what CIDR prefix was originally applied - and therefore no knowledge of how many bits are used for the network address and how many are for the host address.

A RIP 1 router would handle this by making some assumptions. If the router has a subnet of the same network number assigned as the local interface, then it assumes the incoming packet has the same subnet mask as the local interface, otherwise it assumes there is no subnet involved and applies a classful mask.

The relevance of this is that RIP1 only allows a single subnet mask, making it impossible to get the full benefit of VLSM. You must use a newer routing protocol like Open Shortest Path First (OSPF) or RIP2, where the network prefix length or mask value is sent along with route advertisements from router to router. With these in use, it is possible to use VLSM to its full potential and have more than one subnet or sub-subnets.

Step 8 - IPv6 to the Rescue

Obviously, the 32-bit IP address has a limited number of addresses and the explosion of interconnectivity has proved that there are just not enough IPv4 addresses to go around. The answer to future growth lies in the IPv6 addressing scheme. This is more than just the big brother to IPv4 in that it not only adds a significant number of addresses to the IP addressing scheme but eliminates the need for CIDR and the network mask as used in IPv4.

IPv6 increases the IP address size from 32 bits to 128 bits. A 128-bit number supports 2^{128} values, or

340,282,366,920,938,463,463,374,607,431,768,211,456 possible IP addresses. This number is so big there is not even a name for it.

Even the text representation of IPv6 is different from that of IPv4, although it does have a similar-looking dotted decimal look. You will see an IPv6 address written one of three ways:

- Preferred
- Compressed
- Mixed

Preferred IPv6 Addressing Notation
The preferred form is written using hexadecimal values to refer to the 128-bit numbers in each address segment separated by a colon. It would be written like X:X:X:X:X:X:X:X, where each X consists of four 16-bit values. An example would be:

2001:0db8:85a3:0000:0000:8a2e:0370:7D34

Each of the eight sections of an IPv6 number separated by the colons is written as a hexadecimal number which, when translated to decimal value, would range between 0 and 65,535. So where IPv4 text representations of addresses use decimal numbers, IPv6 uses hexadecimal. It really does not matter though - both boil down to binary numbers, which we covered in detail in Section 2.

The following illustration shows how the text representation of an IPv6 address written in hexadecimal is translated into decimal and binary values.

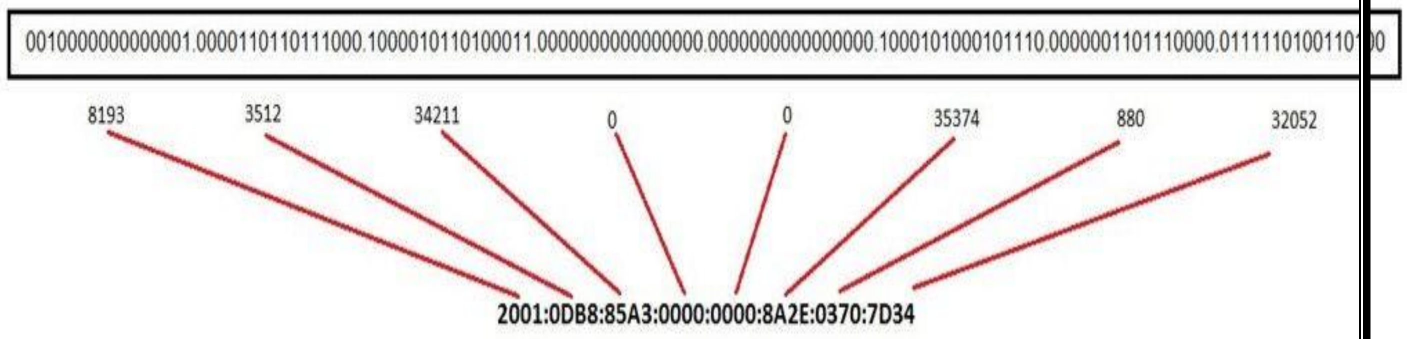

Compressed IPv6 Addressing Notation
The compressed form simply substitutes zero strings with double colons to indicate the zeros are "compressed". For example, the above address in compressed notation would become:

2001:0db8:85a3::8a2e:0370:7D34

There are some rules to follow when doing this zero substitution. First, a substitution can only be done on one "section," or a full 16-bit group; second, the double colon can only be used one time in any given address. There is one other slightly confusing consideration: a double colon automatically suppresses neighboring leading or trailing zeros in an address. Therefore, the above address only indicates one set of double colons as a compressed IPv6 address even though there are two sets of zeros.

Mixed IPv6 Addressing

The mixed addressing notation is useful in environments using both IPv4 and IPv6 addresses. A mixed address would look like X:X:X:X:X:X:X:X:D:D:D:D, where "X" represents the hexadecimal values of the six highest-order 16-bit components of an IPv6 address, and"D" represents an IPv4 value that would plug into the four lower-order values of an IPv6 address.

IPv6 Routing and Prefix Notation

IPv6 does not use subnet masks but does have a means of communicating with subnets that is similar to CIDR. IPv6 routing is based on a prefix length as well, where the prefix length represents the bits that have fixed values or are the bits of t

Conclusion

Whew! We have covered a lot of ground. Let's recap what we've learned:

- For components to communicate on a network, each needs a unique address. For computer networks using the Internet Protocol, these addresses are numeric and are commonly referred to as IPs .
- To make efficient use of IP addresses we also need logical groupings of devices. A subnet then, is a logical organization of connected network devices.
- Binary numbers look very confusing but it's really just because we use the base10 numbering system day to day. The concept of binary numbering is the same.
- Think of the Internet Protocol as simply the rules of communication.
- IP addresses are written in the form of XXX.XXX.XXX.XXX, where each IP address belongs to a certain class depending on the first octet.
- Subnetting involves dividing the network into smaller portions called subnets. In a sense, the IP address then has three components - the network part, the subnet part and, finally, the host part.
- All a subnet mask does is indicate how many bits are being "borrowed" from the host component of an IP address.
- Some IP addresses are used for special purposes.
- Public versus private IPs are similar in theory to public telephone numbers versus private extensions.
- CIDR is used to adapt the concept of subnetting to the entire Internet. It's sometimes referred to as supernetting.
- Variable length subnet masking (VLSM) is another concept that essentially refers to subnetting a subnet.
- IPv6 is the future. It not only adds to the number of available IP addresses but also eliminates the need for CIDR and network masks in IPv6.
- There are three ways to write an IPv6 address: Preferred, compressed and mixed.