

Introduction to
Industrial Ethernet, Part 5.

Part 4 was featured in Issue 6,
MAY–JUNE 2000. If you would like
a copy, please send your request to
info@cccontrols.com

the EXTENSION

A Technical Supplement to control NETWORK

© 2000 Contemporary Control Systems, Inc.

INTRODUCTION TO SUBNETTING

How to maximize network addresses.

By George Thomas,
Contemporary Controls

INTRODUCTION

In a previous article we discussed the Internet Protocol and the structure of IP addresses. An IP address identifies the source and destination of a directed or unicast message and is defined in RFC 761. IPv4 is the most common version of IP addressing requiring 32-bit addresses. Although IPv6, the 128-bit version, will be used in the future, this article will restrict the discussion to IPv4. IPv6 was developed because the explosive growth of the Internet will soon deplete the inventory of available addresses. At one time, 32-bit addresses seemed to provide more than enough addresses but there was much waste in initial assignments and the class structure of IP addresses was inefficient. In order to make more efficient usage of IP address, the concept of subnetting was introduced with RFC 950. This article introduces this concept.

Networks and Hosts

When we talk about a network we usually envision a cluster of workstations with one or more servers connected to a local area network. Each server and workstation would have a unique

address to distinguish it from the other computers. With IP addressing, servers and workstations are all termed hosts but each address not only identifies a host but the address of the network on which the host resides. This is because IP is an internetworking protocol that not only allows communication between hosts on the same network, but communication between hosts on different networks as well. The 32-bit IP address identifies a particular host along with the network on which the host resides. The structure of IP addressing is defined so that any host on the public Internet can be found by any other host.

The format of the 32-bit address is <netid, hostid> and it is usually shown as four bytes of data. Although each byte could be represented as a binary, decimal or hexadecimal number, the decimal-dot-decimal notation is the most popular. Therefore, the range of IP addresses can span 0.0.0.0 to 255.255.255.255. For example 193.5.8.254 is a valid IP address but it is difficult to determine which part of the address is the network ID and which part is the host ID. To understand the two you need to know about class addressing.

Class Addressing

IPv4 is called a classful system under RFC 761 with IP addresses being defined as belonging to one of five classes A, B, C, D or E. Classes A, B and C define different possible combinations of network and host addresses. Class D is reserved for multicasting. Multicasting is the ability of one host to communicate with many other hosts with one transmission and is beyond the scope of this article. Class E is reserved for future use. The classes of interest to subnetting are A, B and C.

With class A addresses, the first byte of the address identifies the network address while the three remaining bytes identify the host. With class B addresses, the first two bytes identify the network address while the remaining two identify the host address. With class C addresses, the first three bytes identify the network address while the last byte identifies the host. That seems simple enough but how do you know you are looking at either an A, B, C, D or E address?

The four-byte IP address is viewed from left to right with the first byte on the left. This is the most significant byte. The first few bits (most significant) of that byte identify the class of address. For a

class A address, the left most bit must be a zero. For a class B address, the first two bits must be a 10. For a class C address, the first three bits must be a 110. For a class D address, the first four bits must be a 1110. For a class E address, the first four bits must be a 1111. Therefore, it is only necessary to observe the first byte of the IP address to determine its class. Figure 1 shows the decimal value of the first byte for each class.

Reserved Addresses

There are some reserved IP address besides those identified as classes D and E. For example, the class A network address 0.X.X.X cannot be used since it is used to indicate “this” network. Class A address 127.X.X.X is reserved for loop back testing. With the host portion of the address, you cannot have an all 0s host, which refers to the network address where the hosts reside. Likewise you cannot use the all 1s host address because that indicates a broadcast which is a message to all hosts on the network. Therefore, with any host addressing on either a class A, B or C network, you lose 2 host addresses. Still with 4 billion possible addresses from a 32-bit address space, you would think there are plenty of addresses even with reserved addresses. The problem is that there was much waste when addresses were originally assigned. For example, a class A address can handle 16 million hosts per one network ID. That is an enormous amount of hosts for just one network. Even a class B address can handle 65 thousand hosts per network ID. A class C address can handle only 254 hosts per network ID which may be too little for some networks. A scheme was needed to obtain a better balance between network and host assignments and that is called subnetting.

SUBNETTING

Subnetting creates additional network IDs at the expense of host IDs and can be used with either A, B or C class addresses. If you look at figure 2, you will notice that a class B address uses 14 bits for network addressing and 16 bits for host addressing. By simply reassigning one of the host bits to a network bit, you would double the number of available network addresses but halve the number of host addresses. Carrying the argument further, move eight of the host bits (actually the complete third byte) to the network side. The result is 22 bits for network addressing and eight bits for host addressing which is quite similar to a class C address. These additional network addresses are called subnets and not networks because to the Internet, the original address is still a class B network address but locally the class B network address can be broken down to manageable subnets that function as actual network addresses. Why use subnets? Subnets are interconnected using routers, and routers improve network performance by reducing traffic and minimizing disruption due to broadcast messages. Large networks become more manageable when subnets are deployed.

MASKING

To create subnets you need a subnet mask that defines which bits will be used to create the new network address out of the 32-bit IP addresses. By “ANDing” the 32-bit IP address with a 32-bit mask, we create a 32-bit IP address that represents <netid, subnetid> becoming our new network address. What do these masks look like? If we start with a basic class A address and do not define any subnets, the mask would look like 255.0.0.0 which is called a natural or default

mask. Only those bits that are set as a 1 will be considered when defining a network address. In this case, all the bits in the first byte of the IP address will be considered. The natural mask for a class B address is 255.255.0.0 and for a class C address it is 255.255.255.0. In order to create more network addresses (subnets) we need to move the mask bits to the right (changing 0 bits into 1s) in order to convert host bits into network bits. The best way to understand the concept is to use an example.

Class A	1-126
Class B	128-191
Class C	192-223
Class D	224-239
Class E	240-254

Figure 1. The class of an IP address can be quickly identified by observing only the first byte.

Assume we begin with IP address 165.10.0.0. From figure 1 we know that this is a class B address with a network address of 165.10 with the capability of assigning up to 65,534 hosts. We do not want 65,534 hosts on one network but would like to have up to 500 hosts on each subnet. In order to have 500 hosts on one subnet, we need to have 9 bits of host addressing. Currently, we have 16 bits of host addressing since we possess a class B address. That means that we can reassign 7 of those bits to signify subnet bits. Therefore, the subnet mask would be 255.255.254.0. In binary it would be:

```
11111111.11111111.11111110.00000000
```

The natural mask for a class B address is 255.255.0.0 so in order to create subnets we moved mask bits to the right in order to convert more

host bits to network bits. It must be remembered that these mask bits must be contiguous from the left. For example, the above mask allows up to 510 host assignments. Remember that we cannot use either an all 0s or all 1s host address. The next jump would be to allow up to 1022 host addresses. What would be the subnet mask? It would be 255.255.252.0. The 1s are still contiguous from the left. This approach creates many subnets, but it is recommended that neither an all 0s nor all 1s subnet be used. This could cause a problem on some networks. How many mask bits can you have? You need to have some hosts on a network and two host addresses are unusable so the maximum number of mask bits is 30 leaving two valid host addresses.

NOTATION

Using the last subnet mask in the above example, we have 1022 host addresses. What if our computer actually had host address 768 on subnet 4? What would be our actual IP address? We cannot say it is 165.10.4.768 since with decimal notation no byte can be more than 255. The actual IP address would be 165.10.7.0 so you do need to know the subnet mask before determining the actual subnet address and host address.

There is a simpler way of representing the actual IP address and that is by using the Classless InterDomain Routing (CIDR) scheme. With this scheme the concept of A, B and C classes is eliminated, but the concept of subnetting is retained. In the above example, we use a total of 22 bits of contiguous 1s in our mask so we would display our IP address as 165.10.7.0/22. Although it is still not obvious that we are host 768 on subnet 4 of network 165.10, we can

	Address Identifier	Network Address	Host Address
Class A	0	7 bits of network address <small>First byte</small>	24 bits of host address <small>Last three bytes</small>
Class B	10	14 bits of network address <small>First two bytes</small>	16 bits of host address <small>Last two bytes</small>
Class C	110	21 bits of network address <small>First three bytes</small>	8 bits of host address <small>Last byte</small>
Class D	1110	Multicast address in the range of 224.0.0.0 - 239.255.255.255	
Class E	1111	Class E - Reserved for future use	

Figure 2. Address classes define the split between network and host IDs.

figure it out using this single notation which tells us exactly where the subnet mask separates the network and host addresses.

For example, in a previous article we mentioned that there were one A, 32 B and 256 C addresses that were strictly private and cannot be accessed through the Internet. These are as follows:

10.0.0.0 to 10.255.255.255

172.16.0.0 to

172.31.255.255

192.168.0.0 to 192.168.255.255

Notice that the first range is a single A address with 24 bits of host addressing, the second are B addresses with 16 bits of host addressing and the third are C addresses with 8 bits of host addressing. Using CIDR notation these same address ranges can be displayed as follows:

10.0.0.0/8

172.16.0.0/12

192.168.0.0/16

The natural mark for a class A address is 255.0.0.0 which means eight contiguous 1s from the left so 10.0.0.0/8 represents the natural mask for a class A address. This is what we would expect. A single class A network address with provisions for 24 bits of host

addressing. The natural mask for a class B address is 255.255.0.0 which, with CIDR notation, would be /16 but the above class B addresses have only 12 mask bits of contiguous 1s. This seems to violate our rule for subnetting and it does. With subnetting you move the bits to the right of the natural mask thereby consuming host bits. Instead we are moving the mask to the left of the natural mask (changing 1 bits to 0s) consuming network bits. This is called supernetting which requires contiguous network addresses and will be discussed shortly. By moving the mask to the left by four bits from the natural mask, we can gain more host addresses at the expense of 16 contiguous network addresses. Therefore, the notation 172.16.0.0/12 is short for indicating a range of contiguous network addresses from 172.16.0.0 to 172.31.0.0. The same is true for the last example which are C class addresses. The natural mask for a C address is /24. Instead the CIDR notation is a /16 meaning eight less mask bits thereby yielding a range of network addresses from 192.168.0.0 to 192.168.255.0.

SUPERNETTING

The inverse of subnetting is supernetting. Instead of moving mask bits to the right of the natural mask for subnetting, we move mask bits to the left for supernetting. With subnetting we create more network

addresses at the expense of host addresses. With supernetting we create more host addresses at the expense of network addresses. Supernetting is not for users since it would be difficult for users to be granted a range of contiguous network addresses. Supernetting is for Internet Service Providers (ISPs) who are attempting to obtain the most efficient allocation of IP addresses using the A, B, C class scheme.

SUMMARY

Although a 32-bit IP address offers an extremely large number of addresses, the A, B, C, class structure does not make efficient use of <netid, hostid> assignments. Subnetting improves the situation by allowing a finer split between network and host assignments while improving the performance and maintainability of large networks.

SUBNETTING A CLASS C ADDRESS

Subnet mask	CIDR	# Subnets	# Host
11111111.11111111.11111111.00000000	/24	0	254
11111111.11111111.11111111.11000000	/26	2	62
11111111.11111111.11111111.11100000	/27	6	30
11111111.11111111.11111111.11110000	/28	14	14
11111111.11111111.11111111.11111000	/29	30	6
11111111.11111111.11111111.11111100	/30	62	2

The natural mask for a class C address is 255.255.255.000 which provides for up to 254 host addresses. By moving the mask bits to the right (replacing 0s for 1s), subnets are created at the expense of host bits. Not shown are masks /25 and /31 since they are not allowed. Similar charts can be made for class A and class B addressing. Class A subnetting begins at /10 and class B at /18. Both end at /30.

REFERENCES

Illustrated TCP/IP, Matthew Naugle, 1998, Wiley Computer Publishing

Practical Networking With Ethernet, Charles E. Spurgeon, 1997, International Thomson Computer Press

International Standard ISO/IEC 8802-3 ANSI/IEEE std 802.3, 1996, The Institute of Electrical and Electronic Engineers, Inc.

TCP/IP Clearly Explained, Pete Loshin, 1997, Academic Press

TCP/IP Illustrated, Volume 1, The Protocols, W. Richard Stevens, 1994, Addison-Wesley Publishing Company