Cisco – Subnet Zero and the All–Ones Subnet

Table of Contents

Subnet Zero and the All–Ones Subnet

Introduction

Subnetting breaks down a given network address into smaller subnets. It allows for the more efficient use of available IP address space, in tandem with other technologies like Network Address Translation (NAT) and Port Address Translation (PAT), thereby alleviating the problem of address depletion to a great extent. Subnetting has guidelines regarding the use of the first and the last subnets, known as subnet zero and the all–ones subnet, respectively. This document discusses subnet zero and the all–ones subnet and their uses.

Before You Begin

Conventions

For more information on document conventions, see the Cisco Technical Tips Conventions.

Prerequisites

There are no specific prerequisites for this document.

Components Used

This document is not restricted to specific software and hardware versions.

The information presented in this document was created from devices in a specific lab environment. All of the devices used in this document started with a cleared (default) configuration. If you are working in a live network, ensure that you understand the potential impact of any command before using it.

Subnet Zero

If a network address is subnetted, the first subnet obtained after subnetting the network address is called subnet zero.

Consider a Class B address, 172.16.0.0. By default the Class B address 172.16.0.0 has 16 bits reserved for representing the host portion, thus allowing 65534 (2^{16}–2) valid host addresses. If network 172.16.0.0/16 is subnetted by borrowing three bits from the host portion, eight (2^3) subnets are obtained. The table below is an

example showing the subnets obtained by subnetting the address 172.16.0.0, the resulting subnet mask, the corresponding broadcast addresses, and the range of valid host addresses.

Subnet Address	Subnet Mask	Broadcast Address	Valid Host Range
172.16.0.0	255.255.224.0	172.16.31.255	172.16.0.1 to 172.16.31.254
172.16.32.0	255.255.224.0	172.16.63.255	172.16.32.1 to 172.16.63.254
172.16.64.0	255.255.224.0	172.16.95.255	172.16.64.1 to 172.16.95.254
172.16.96.0	255.255.224.0	172.16.127.255	172.16.96.1 to 172.16.127.254
172.16.128.0	255.255.224.0	172.16.159.255	172.16.128.1 to 172.16.159.254
172.16.160.0	255.255.224.0	172.16.191.255	172.16.160.1 to 172.16.191.254
172.16.192.0	255.255.224.0	172.16.223.255	172.16.192.1 to 172.16.223.254
172.16.224.0	255.255.224.0	172.16.255.255	172.16.224.1 to 172.16.255.254

In our example above, the first subnet, subnet 172.16.0.0/19, is called subnet zero.

Note that the class of the network subnetted and the number of subnets obtained after subnetting have absolutely no role in determining subnet zero. It is simply the first subnet obtained when subnetting the network address. Further, when you write the binary equivalent of the subnet zero address, all the subnet bits (bits 17, 18, and 19 in this case) are zeros. Subnet zero is thus also known as the all−zeros subnet.

The All−Ones Subnet

When a network address is subnetted, the last subnet obtained is called the all−ones subnet.

With reference to the example above, the last subnet obtained when subnetting network 172.16.0.0, subnet 172.16.224.0/19, is called the all−ones subnet.

Note that the class of the network subnetted and the number of subnets obtained after subnetting have absolutely no role in determining the all−ones Subnet. Also, when you write the binary equivalent of the subnet zero address, all the subnet bits (bits 17, 18, and 19 in this case) are ones, hence the name.

Problems with Subnet Zero and the All−Ones Subnet

Traditionally, it was strongly recommended that subnet zero and the all−ones subnet not be used for addressing. According to RFC 950 , "It is useful to preserve and extend the interpretation of these special (network and broadcast) addresses in subnetted networks. This means the values of all zeros and all ones in the subnet field should not be assigned to actual (physical) subnets." This is the reason why network engineers required to calculate the number of subnets obtained by borrowing three bits would calculate $2^3−2$ (6) and not

2^3 (8). The "−2" takes into account that subnet zero and the all−ones subnet are not used traditionally.

Subnet−Zero

Use of subnet zero for addressing was discouraged because of the confusion inherent in having a network and a subnet with indistinguishable addresses.

With reference to our example above, consider the IP address 172.16.1.10. Now, if you calculate the subnet address corresponding to this IP address, the answer you would arrive at would be subnet 172.16.0.0 (subnet zero). Note that this subnet address is identical to network address 172.16.0.0, which was subnetted in the first place, so whenever you perform subnetting, you get a network and a subnet (subnet zero) with indistinguishable addresses. This was formerly a source of great confusion.

Prior to Cisco IOS® Software Release 12.0, Cisco routers, by default, did not allow an IP address belonging to subnet zero to be configured on an interface. However, if a network engineer working with Cisco IOS version older than 12.0 finds it safe to use subnet zero, the **ip subnet−zero** command in the global configuration mode can be used to overcome this restriction. As of Cisco IOS version 12.0, Cisco routers now have **ip subnet−zero** enabled by default, but if the network engineer feels that it is unsafe to use subnet zero, the **no ip subnet−zero** command can be used to restrict the use of subnet zero addresses.

In versions prior to Cisco IOS version 8.3, the **service subnet−zero** command was used.

The All−Ones Subnet

Use of the all−ones subnet for addressing has been discouraged in the past because of the confusion inherent in having a network and a subnet with identical broadcast addresses.

With reference to our example above, the broadcast address for the last subnet (subnet 172.16.224.0/19) is 172.16.255.255, which is identical to the broadcast address of the network 172.16.0.0, which was subnetted in the first place, so whenever you perform subnetting you get a network and a subnet (all−ones subnet) with identical broadcast addresses. In other words, a network engineer could configure the address 172.16.230.1/19 on a router, but if that is done, he can no longer differentiate between a local subnet broadcast (172.16.255.255 (/19)) and the complete Class B broadcast (172.16.255.255(/16)).

Although the all−ones subnet can now be used, misconfigurations can cause problems. To give you an idea of what can happen, consider the following:

Cisco − Subnet Zero and the All−Ones Subnet

Note: See Host and Subnet Quantities for details.

Routers 2 through 5 are access routers that each have several incoming asynchronous (or ISDN) connections. We have decided to break up a network (195.1.1.0/24) into four pieces for these incoming users. Each piece is given to one of the access routers. Also, the asynchronous lines are configured **ip unnum e0**. Router 1 has static routes pointing at the correct access router, and each access router has a default route pointing at Router 1.

The Router 1 routing table looks like this:

```
C  195.1.2.0/24   E0
        S  195.1.1.0/26   195.1.2.2
        S  195.1.1.64/26  195.1.2.3
        S  195.1.1.128/26 195.1.2.4
        S  195.1.1.192/26 195.1.2.5
```

The access routers have the same connected route for the Ethernet, the same default route and several host routes for their asynchronous lines (courtesy of Point−to−Point Protocol (PPP)).

```
    Router 2 routing table:                  Router 3 routing table:

        C  195.1.2.0/24   E0                     C  195.1.2.0/24   E0
        S  0.0.0.0/0       195.1.2.1             S  0.0.0.0/0       195.1.2.1
        C  195.1.1.2/32    async1                C  195.1.1.65/32   async1
        C  195.1.1.5/32    async2                C  195.1.1.68/32   async2
        C  195.1.1.8/32    async3                C  195.1.1.74/32   async3
        C  195.1.1.13/32   async4                C  195.1.1.87/32   async4
        C  195.1.1.24/32   async6                C  195.1.1.88/32   async6
        C  195.1.1.31/32   async8                C  195.1.1.95/32   async8
        C  195.1.1.32/32   async12               C  195.1.1.104/32   async12
        C  195.1.1.48/32   async15               C  195.1.1.112/32   async15
        C  195.1.1.62/32   async18               C  195.1.1.126/32   async18

     Router 4 routing table:                  Router 5 routing table:

        C  195.1.2.0/24   E0                     C  195.1.2.0/24   E0
        S  0.0.0.0/0       195.1.2.1             S  0.0.0.0/0       195.1.2.1
        C  195.1.1.129/32   async1               C  195.1.1.193/32   async1
        C  195.1.1.132/32   async2               C  195.1.1.197/32   async2
        C  195.1.1.136/32   async3               C  195.1.1.200/32   async3
        C  195.1.1.141/32   async4               C  195.1.1.205/32   async4
        C  195.1.1.152/32   async6               C  195.1.1.216/32   async6
        C  195.1.1.159/32   async8               C  195.1.1.223/32   async8
        C  195.1.1.160/32   async12              C  195.1.1.224/32   async12
        C  195.1.1.176/32   async15              C  195.1.1.240/32   async15
        C  195.1.1.190/32   async18              C  195.1.1.252/32   async18
```

What if we have misconfigured the hosts on the asynchronous lines to have a 255.255.255.0 mask instead of 255.255.255.192 mask? Everything works fine.

Now let's take a look at what happens when one of these hosts, say 195.1.1.24, does a local broadcast (NetBIOS, WINS). The packet looks like this:

```
    s: 195.1.1.24 d: 195.1.1.255
```

The packet is received by Router 2. Router 2 sends it to Router 1, which sends it to Router 5, which sends it Router 1, which sends it Router 5, and so on, until the Time To Live (TTL) expires.

Let's look at another example, say, host 195.1.1.240:

```
        s: 195.1.1.240  d: 195.1.1.255
```

This packet is received by Router 5. Router 5 sends it to Router 1, which sends it to Router 5, which sends it Router 1, which sends it Router 5, and so on, until the TTL expires. If this situation occurs, you might think you were under a packet attack. Given the load on Router 5, this wouldn't be an unreasonable assumption.

In this example, we've created a routing loop. Because Router 5 is handling the all−ones subnet, it gets blasted. Routers 2 through 4 see the "broadcast" packet only once. Router 1 is hit, too, but what if it's a Cisco 7513, which can handle this situation? In that case, you need to configure your hosts with the correct subnet−mask.

To protect against misconfigured hosts, create a loopback interface on each access router with a static route 195.1.1.255 to the loopback address. You could use the Null0 interface, but this causes the router to generate Internet Control Message Protoco (ICMP) "unreachable" messages.

Using Subnet Zero and the All−Ones Subnet

It should be noted that even though it was discouraged, the entire address space including subnet zero and the all−ones subnet has always been usable. The use of the all−ones subnet was explicitly allowed and the use of subnet zero is explicitly allowed since Cisco IOS version 12.0. Even prior to Cisco IOS version 12.0, subnet zero could be used by entering the ip subnet−zero global configuration command.

On the issue of using subnet zero and the all−ones subnet, RFC 1878 states, "This practice (of excluding all−zeros and all−ones subnets) is obsolete! Modern software will be able to utilize all definable networks." Today, the use of subnet zero and the all−ones subnet is generally accepted and most vendors support their use, though, on certain networks, particularly the ones using legacy software, the use of subnet zero and the all−ones subnet can lead to problems.

Related Information

- **Technical Support − Cisco Systems**
- **Tools and Utilities**