



Introduction to IP subnetting

Networking Fundamentals

What you will learn

At the core of the lesson

You will learn how to:

- Locate your local IP address
- Create subnetting, subnet masks, a host ID range, number of usable host IDs, and the broadcast ID without using a lot of math
- Divide a network into two or more networks (this is called IP subnetting)
- Use the Classless Inter-Domain Routing (CIDR) notation to specify subnet address ranges
- Describe how Amazon's Virtual Private Cloud (VPC) handles subnetting
- Explain AWS' VPC architectural components






Refresher: What is an IP Address?

What is an IP address?



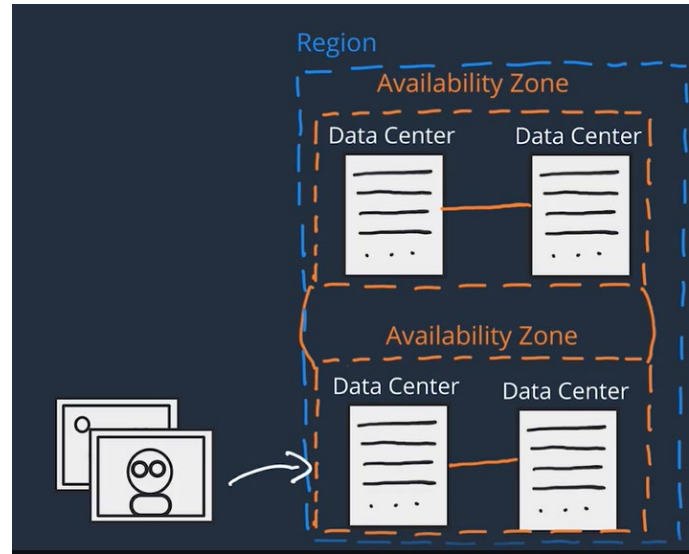
Networking is how you connect computers and other devices around the world and allow them to communicate with one another. In this course, you've already seen a few examples of networking. Currently, we're looking at systems from personal computers to smaller network infrastructures. Each has an IP address.



Think about them as delivery systems that can deliver mail to a rural address (only needing one IP address) or to many addresses in a large town (like a mass mailer sent to everyone in three zip code areas).


What does an IP address do in larger networks?

Later, we will begin focusing on AWS-specific Global Infrastructure. These are huge areas that have buildings of servers, that reside in rows in data centers. The data centers reside in availability zones, which reside in regions.



Availability Zones

But for now, remember that every device has an original IP address.





What is an IP address in an IPv4 IP?

If you recall from an earlier lesson, an IP address is a unique address that identifies a device on the internet. This could be a:

- Any device on the internet
- A local network
- Or even a printer or other local device

IPv4 IP addresses are expressed as sets of four numbers (called 8-bit binary bits and are also referred to as octets). Each octet is assigned a decimal value. The left-most bit is assigned a value of 128. Each subsequent bit is assigned 64, 32, 16, 8, 4, 2, 1 and 1, going from left to right.

Each of those four sets can range from 0-255. Here is an example:

190.111.25.37

An IP address in Decimal, Binary, and Decimal Calculations

Each bit in the octet can be either a 1 or a 0. If the value is 1, it is counted as its decimal value, and if it is 0, it is ignored.

Just by using the set of 8 bits and manipulating the 1's and 0's, you can obtain any value from 0 to 255 for each octet. If all the bits are 0, the value of the octet is 0. If all the bits in the octet are 1, the value is 255, which is $128+64+32+16+8+4+2+1$.

Decimal Value	Binary Value	Decimal Calculation
10	00001010	$8+2=10$
192	11000000	$128+64=192$
205	11001101	$128+64+8+4+1=205$
223	11011111	$128+64+16+8+4+2+1=223$



How to read a binary and decimal IP address

The phenomenon of logical addressing works on the Layer-3 of the OSI reference model and network components like routers and switches are the host devices that are most popularly used. An IP Address is a 32-bit logical address that distinctively classifies a host of the network. The host can be a computer, Mobile handset, or even a tablet. The 32 bits binary IP address is made up of two distinctive parts: **The Network address and the Host address.**

It also has 4 octets as each octet is having 8 bits. This octet is converted into decimal and is separated by a format referred to as dotted-decimal format. The range of an octet in binary is from 00000000 to 11111111 and in decimal from 0 to 255. Binary IP addresses are separators are grouped by a colon between each octet.

Example of an IP Address format:

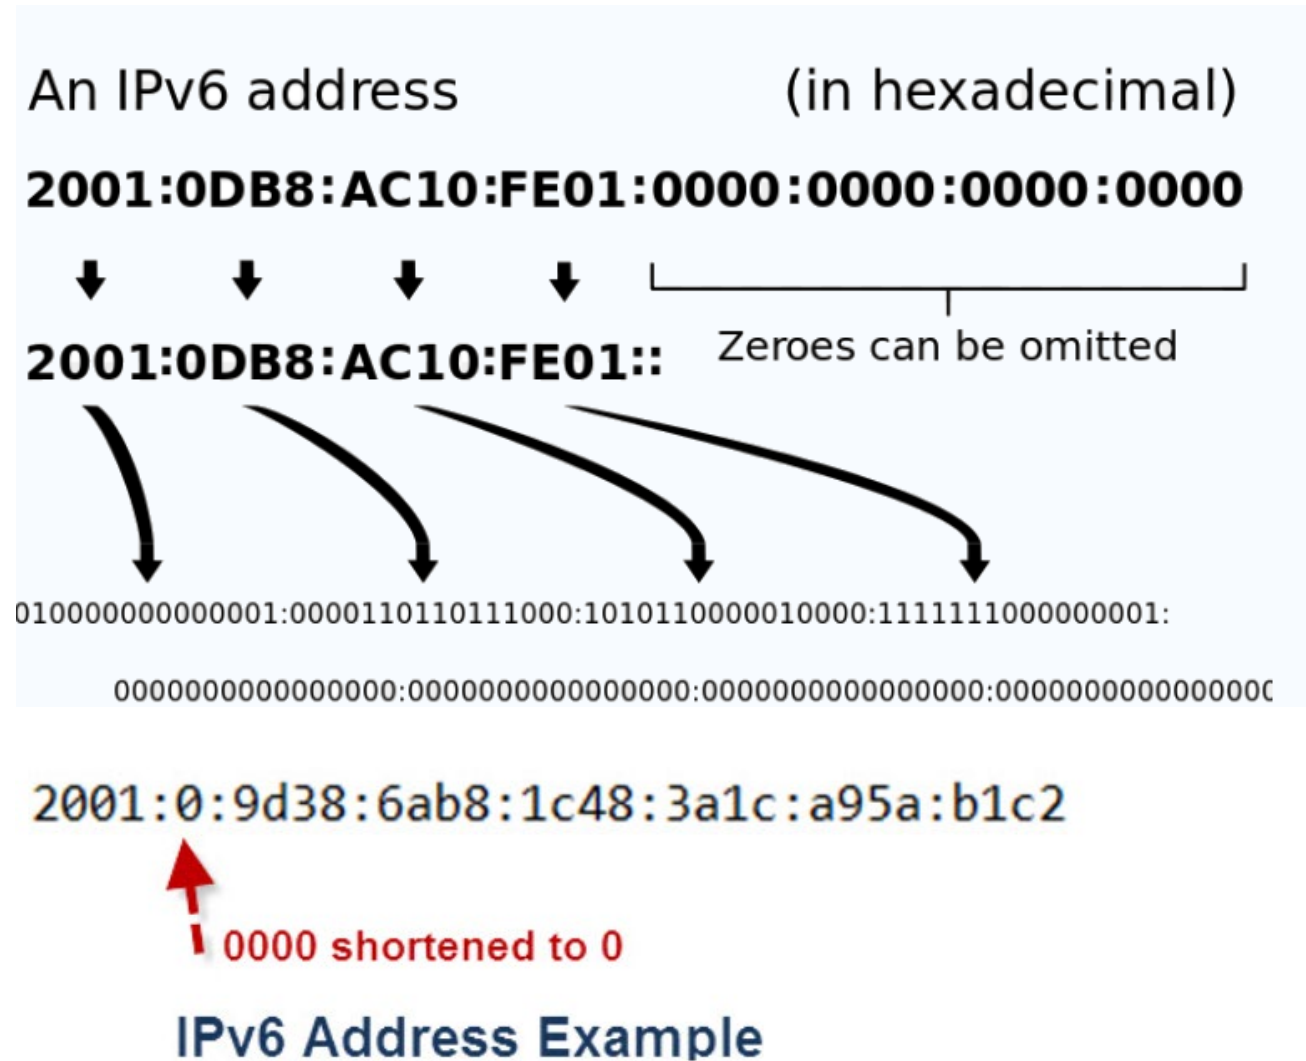
Decimal: 192.168.1.64 (IPv4 protocol that uses a 32-bit address.)

Binary: 11000000:10101000:00000001:01000000 (IPv6 protocol that uses 128-bit hexadecimal IP address).

What an IP address in an IPv6 IP?

Because we use so many IP addresses now and we're running out of them in IPv4, we developed IPv6. These are longer, numeric labels that help identify and locate the network.

- IPv6 uses 128 bits, whereas IPv4 uses 32 bits.
- IPv6 uses hexadecimal.
- Because it uses hex hexadecimal (4 bits) it can consists of 32 hexadecimal numbers.
- IPv6 IP addresses are grouped in 4 sections, separated by colons.
- Zeroes are omitted from an address.





How do IP addresses work and change?

Internet Protocol works the same way as any other language, by communicating using set guidelines to pass information. All devices find, send, and exchange information with other connected devices using internet protocols. By speaking the same language, any computer in any location can talk to one another.

The use of IP addresses typically happens behind the scenes. The process works like this:

1. Your device indirectly connects to the internet by first connecting to your local service provider, which gives you initial access.
2. From home, this is your ISP. From work, it will be your company network.
3. This is the point where your IP address is recognized or assigned, depending on whether you have a fixed or dynamic IP address.



How do IP addresses work and change?

4. Your internet activity goes through the internet provider, and they route it back to you, using your IP address. Since they are giving you access to the internet, it is their role to assign an IP address to your device.
5. Your IP address can change. For example, turning your modem or router on or off can change it. Or you can contact your ISP, and they can change it for you.
6. When you are out (traveling or shopping), and you take your device with you, your home or work IP address does not come with you. Why? This is because you will be using another network (Wi-Fi at a hotel, airport, or coffee shop, etc.) to access the internet. All of these internets will assign different (and temporary) IP address.




How to locate an IP address

How to find an IP address

You should always be able to find the IP address on any device. The best way to do it is to use google to look up the instructions for the device you are using. Here are the instructions for finding your local IP address for the PC.

In Windows :

1. To the right of the Search button, in the Search text box, enter “command prompt” (without quotes), and press Enter.



2. In the Command Prompt box, enter “ipconfig” (no quote marks), and press Enter.
3. Scroll down to see a list of Wireless Lan Adapter Wi Fi information, including IP addresses.

How to find an IP address (Continued)

Here is an example of what you'll see after scrolling down the results of your Ipconfig search.

```
Wireless LAN adapter Wi-Fi:
```

```
Connection-specific DNS Suffix  . :  
IPv6 Address. . . . . : 2600:100f:b011:b30d:d159:1d49:e48f:6645  
Temporary IPv6 Address. . . . . : 2600:100f:b011:b30d:b17b:6e32:c2d5:7e98  
Link-local IPv6 Address . . . . . : fe80::d159:1d49:e48f:6645%18  
IPv4 Address. . . . . : 192.168.1.2  
Subnet Mask . . . . . : 255.255.255.0  
Default Gateway . . . . . : fe80::8846:f3ff:fe79:196f%18  
                             192.168.1.1
```

Notice that you can see:

- IPv4 address (192.168.1.2)
- IPv6 address (2600:100f:b011:b30d:b17b:6e32:c2d57e98)
- Subnet (255.255.255.0)



Key Takeaways

- What is the main command used to find any IP address?
- How many IPvx IP addresses will you see when you search for one?
- What are the IPvx *numbers you can see*?
- What is one other item you might see in the example shown in the section you just completed?




What is a subnet?

What is a subnet?

A subnet is a segmented piece of a larger network and is often thought of as a subnetwork. More specifically, subnets are a logical partition of an IP network into multiple, smaller network segments.

Organizations use a subnet to subdivide large networks into smaller, more efficient subnetworks. They split larger networks into groupings of smaller, interconnected networks to help minimize traffic. Using subnets, traffic takes the most efficient routes, increasing network speeds.

Reasons for using subnets




Why are subnets used?

A subnet is a segmented piece of a larger network and is often thought of as a subnetwork. More specifically, subnets are a logical partition of an IP network into multiple, smaller network segments.

Organizations use subnets to subdivide large networks into smaller, more efficient subnetworks. They split larger networks into groupings of smaller, interconnected networks to help minimize traffic.

In this example, splitting the network into two allows less printer traffic on each printer so jobs print faster. If this weren't done, you would have 6 people standing in line waiting for their 100-page reports to print. That would cost a loss of productivity, efficiency, and more.



Note: A subnet is a logical organization of connected devices whose main purpose is to relieve network traffic congestion.




More reasons to use subnets

Use subnets to:

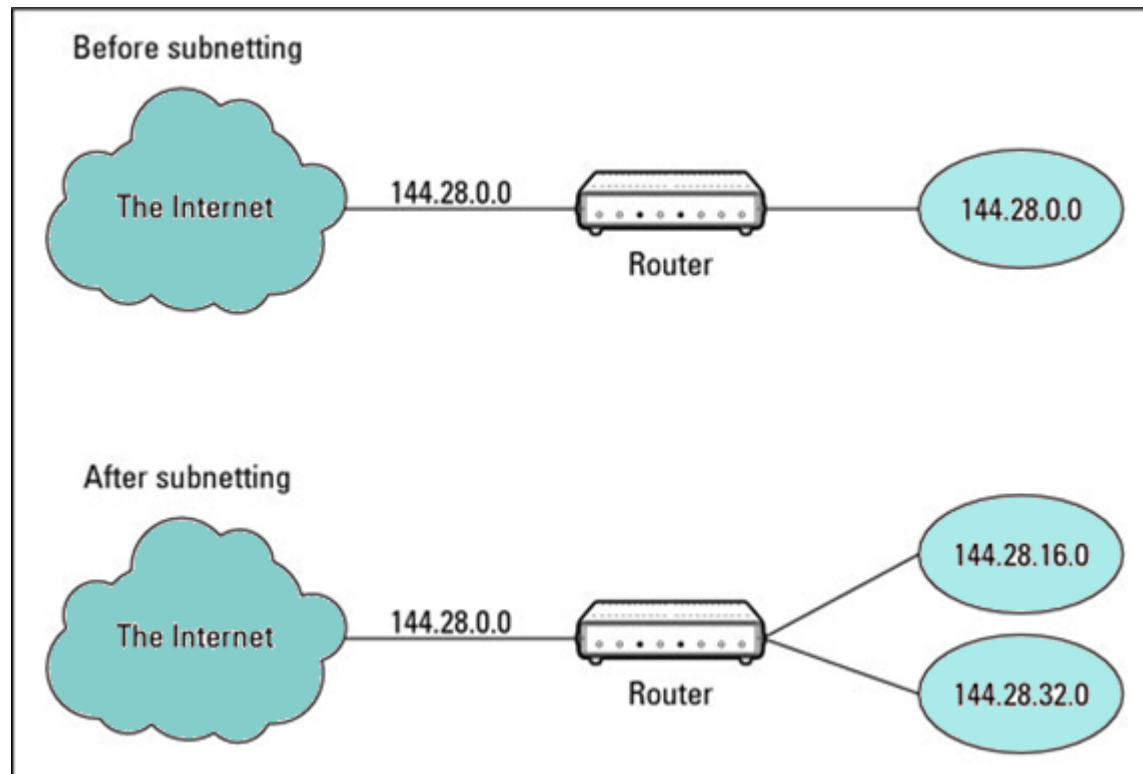
- Optimize network performance
- Maximize the efficiency of IP addressing
- Extend the life of IPV4 because of its scarcity
- Reduce network traffic by eliminating collision and broadcast traffic
- Allows the efficient application of network security policies at the interconnection between subnets
- Facilitate spanning of large geographical distances (especially for AWS's needs)
- Prevents the allocation of large numbers of unused IP network addresses

What is a subnet?

- 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. In our previous example, there is a network (the company), which is divided into logical subnets (department a and b), each of which has its own hosts (users and printers).




Example showing an IP before and after subnetting



1. Before applying subnetting, the router only moves traffic to one network.
2. After subnetting is applied, notice that the router splits the traffic into two networks.
3. The Router contains a table called the router table that contains a map showing where the router should direct traffic. (Not seen here.)

More about after subnetting

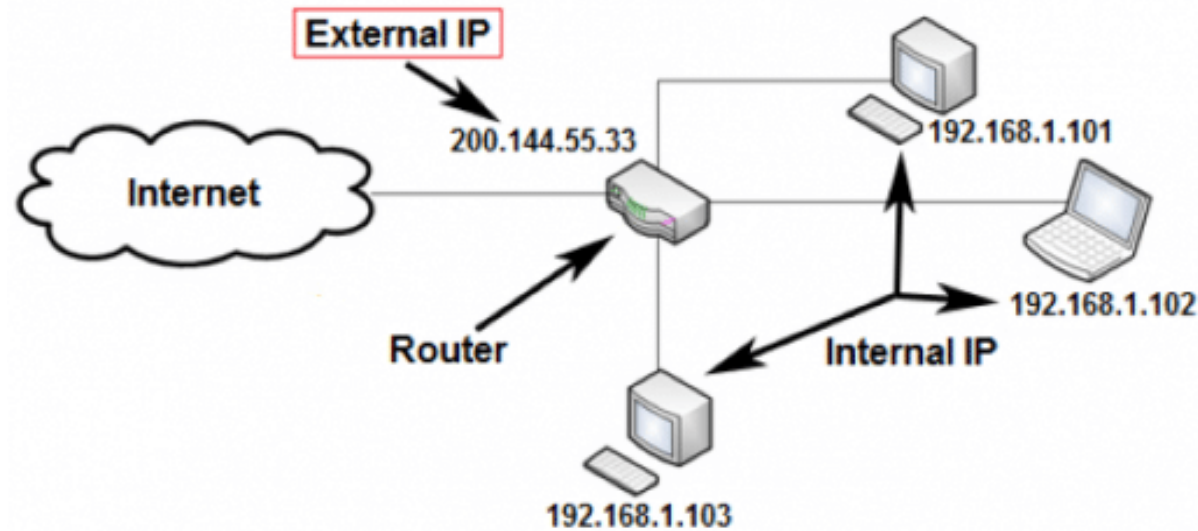


Q: What do you notice when you look at the IP addresses after the subnetting has been applied?

A: The first two sections of network are the network identifier (144.28.0.0) and the subnets are identified as 144.28.16.0 and 32.0. But to the outside world, the IP address is only seen as 144.28.0.0.

Note: you can use up to 3 of the 4 octets for subnets, depending on the class, which will be explained in an upcoming section.

Larger internet configuration with subnets



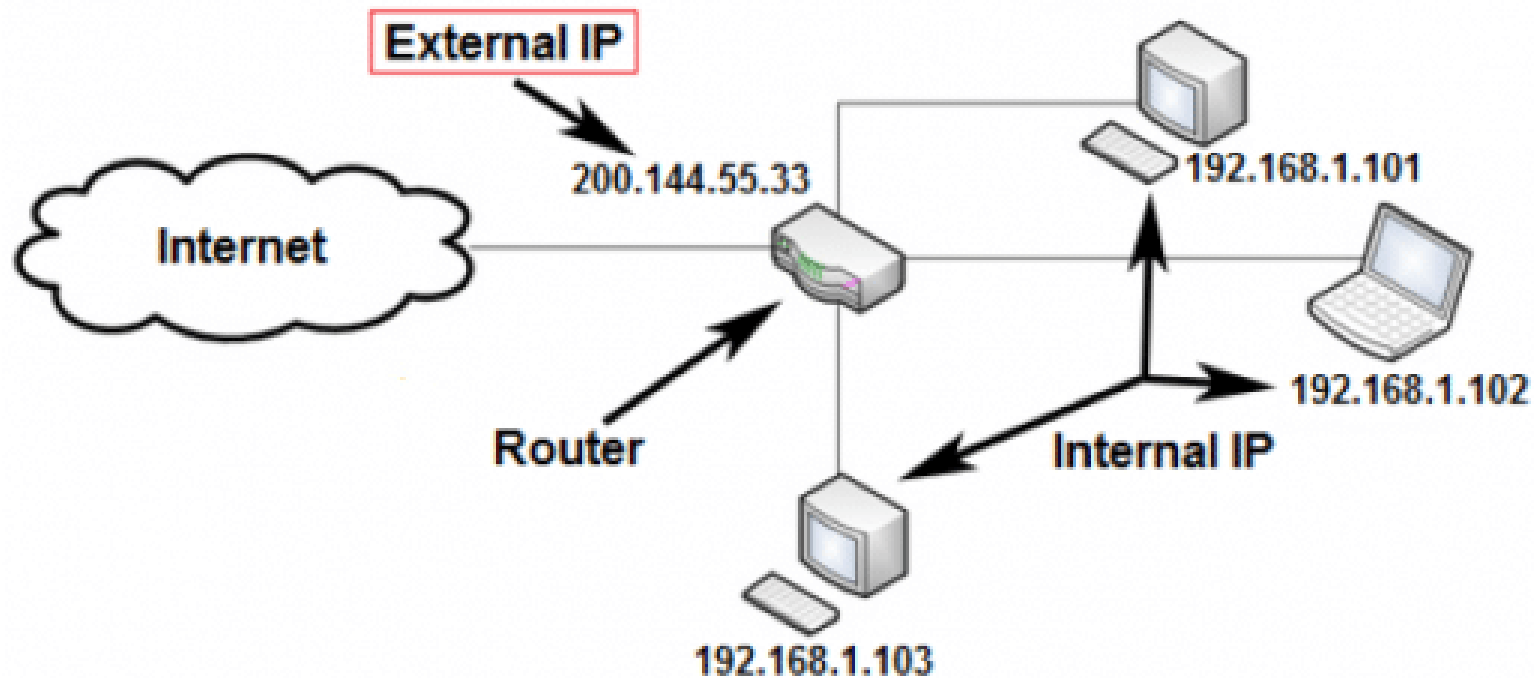
Here is a larger configuration with additional devices. Imagine a large corporation where there are many external IP addresses connected to routers, switches, and printers, computers, smart tv monitors, and many, many more devices and you can start to understand the incredible power that you can harness once you understand and use subnets in your network configurations.

In this example, the external IP is 200.44.55.33 but the subnets, (after the router) all begin with 192.168.1 and only the last octet is different. The subnets are 101-103.

Larger internet configuration with subnets

Q: At what point do you think the subnet is split from its primary IP address?

A: the first two sections of network are the network identifier (144.28.0.0) and the subnets are identified as 144.28.16.0 and 32.0. But to the outside world, the IP address is only seen as 144.28.0.0.



Note: you can use up to 3 of the 4 octets for subnets, depending on the class, which will be explained in an upcoming section.



What you need to know to calculate subnets

Subnets can be complicated. However, the following slides will show you a simple method for creating them. You will have to know the following items:

- Subnet mask
- Network ID
- Host ID
- Broadcast ID

But how can you find those items?

Remember earlier when we learned to use ping? You can use it again to see everything available to you when using that command.

For this example, lets start with the following IP address: **192.168.4.0/24**



Calculating subnets (contd)

Now let's create 3 subnets for IP address: **192.168.4.0/24** for each department:

- Reception/front desk
- Mail room
- Guest or public use

Next, list this information for each subnet:

- Network ID
- Subnet mask
- Host ID range
- # of usable host IDs
- Broadcast ID

Wow! This sounds really complicated, right?. But don't worry. We've got tricks to show you the easy way to do these tasks.

Calculating subnets (contd)

Let's create a table like the following example:

Subnet	1	2	4	8	16	32	64	128	256
Host	256	128	64	32	16	8	4	2	1
Subnet Mask	/24	/25	/26	/27	/28	/29	/30	/31	/32

Having a table makes it easy to solve most of the subnetting issues.

- Row 1 – Has 9 numbers. Each of the numbers is a double of the previous number.
- Row 2 – Also has 9 numbers is a reverse of the first row.
- Row 3 - We use the shorthand format for this row, starting with 24 and going up to 32.

OK, we have our subnet table, but now what? How do we get all those fields we need?

Calculating subnets (contd)

1. Use the network ID you were assigned earlier **192.168.4.0/24** to create three subnets (Reception, Mail room, and guest use).
2. Find the Network ID for the subnet:
3. Identify the number that will come the closest to containing 3 subnets. You can't find the number 3 in the first row, but 4 will work.

Subnet	1	2	4	8	16	32	64	128	256
Host	256	128	64	32	16	8	4	2	1
Subnet Mask	/24	/25	/26	/27	/28	/29	/30	/31	/32

Calculating subnets (contd)

If you circle all of the numbers in this column, you will get most of your answers from here.

- 4 is the number of subnets.
- 64 is the number of total host ids, including network ID and broadcast ID
- /26 is the new subnet mask for all of the four subnets we are creating.

OK, So far it sounds simple, but now what? How do we get the rest of the items we're supposed to get and what do we do with this information?

Subnet	1	2	4	8	16	32	64	128	256
Host	256	128	64	32	16	8	4	2	1
Subnet Mask	/24	/25	/26	/27	/28	/29	/30	/31	/32

Creating a table for your subnets

Create a table with the information we have so far:

Network ID	Subnet Mask	Host ID Range	# of Usable Hosts	Broadcast IDs
192.168.4.0				
192.168.4.64				
192.168.4.128				
192.168.4.192				

1. The first network ID is always the original ID but the last number is 0.
2. The second ID is zero plus 64. Look back at the rows you circled (4,64,/26).
3. The third Network ID is 64+64, so it would be 128.
4. The 4th Network ID is still calculated, even though we are only planning to use 3. It would be 128+64, which is 192. Notice that you are adding the prior two line's Network IDs to create additional network IDs.

Finding the subnet mask

The subnet mask for this exercise is /26 for the entire group of Network IDs.

Network ID	Subnet Mask	Host ID Range	# of Usable Hosts	Broadcast IDs
192.168.4.0	/26			
192.168.4.64	/26			
192.168.4.128	/26			
192.168.4.192	/26			

You didn't even need to break a sweat to figure out the subnets. See, it isn't as hard to do this as people make it out to be. They just don't know the Amazon way.

Finding the # of usable hosts

The first host ID is reserved for the the Network ID. And the last Host ID is reserved for the Broadcast ID. So, the # of usable hosts is $64 - 2 = 62$. Use 62 for all of the Usable Hosts.

Network ID	Subnet Mask	Host ID Range	# of Usable Hosts	Broadcast IDs
192.168.4.0	/26		62	
192.168.4.64	/26		62	
192.168.4.128	/26		62	
192.168.4.192	/26		62	

You didn't even need to break a sweat to figure out the subnets. See, it isn't as hard to do this as people make it out to be. They just don't know the Amazon way.

Finding the Broadcast ID

- Remember that the last host ID is reserved for the broadcast ID. So, the first Broadcast ID is 63.
- We don't add anything to the first Network ID, so the second line is $0+128-1=127$.
- The next Broadcast ID is $128+64-1=191$.
- The last subnet's Broadcast ID is calculated as $191+128-1=255$.

Tip: You can also just add 64 to each ID after the first one to calculate the Broadcast IDs.

Network ID	Subnet Mask	Host ID Range	# of Usable Hosts	Broadcast IDs
192.168.4.0	/26		62	192.168.4. 63
192.168.4.64	/26		62	192.168.4. 127
192.168.4.128	/26		62	192.168.4. 191
192.168.4.192	/26		62	192.168.4. 255

Finding the Host ID range

The last piece of the puzzle that we need to calculate is the Host ID Range. This range can be any IDs between its Network ID (192.168.4.0) and its broadcast ID (192.168.4.255).

So, the first Host ID range is between 192.168.4.1-192.168.4.62.

Remember that the first and last Host IDs are reserved.

Network ID	Subnet Mask	Host ID Range	# of Usable Hosts	Broadcast IDs
192.168.4.0	/26	192.168.4.1-192.168.4.62	62	192.168.4.63
192.168.4.64	/26		62	192.168.4.127
192.168.4.128	/26		62	192.168.4.191
192.168.4.192	/26		62	192.168.4.255

Finding the Host ID range

Let's calculate the other Host ID ranges:

- The second range is between 64 and 127.
- The third range is between 128 and 191.
- The last range is between 192 and 255.

How easy was that? You just learned to figure out how to create subnets, a subnet mask, host ID ranges, the number of usable hosts, and broadcast IDs.

Network ID	Subnet Mask	Host ID Range	# of Usable Hosts	Broadcast IDs
192.168.4.0	/26	192.168.4.1-192.168.4.62	62	192.168.4.63
192.168.4.64	/26	192.168.4.65-192.168.4.126	62	192.168.4.127
192.168.4.128	/26	192.168.4.129-192.168.4.190	62	192.168.4.191
192.168.4.192	/26	192.168.4.193-192.168.4.254	62	192.168.4.255


Now what?

The original objective was to create three subnets:

- ✓ Reception/front desk
- ✓ Mail room
- ✓ Guest or public use

Then we needed to list this information for each subnet:

- ✓ Network ID
- ✓ Subnet mask
- ✓ Host ID range
- ✓ # of usable host IDs
- ✓ Broadcast ID



Adding the Subnet use to the table

Let's add a column to our table to assign which subnet is used for what subnet so you don't confuse them. Notice there is a new column 1 that indicates the use of each subnet and the extra one that isn't being used at present. Documentation is important when you work with subnets. There are lots of IP addresses to track.

Subnet Use	Network ID	Subnet Mask	Host ID Range	# of Usable Hosts	Broadcast IDs
Reception/Front desk	192.168.4.0	/26	192.168.4.1-192.168.4.62	62	192.168.4.63
Mail room	192.168.4.64	/26	192.168.4.65-192.168.4.126	62	192.168.4.127
Guest/Public Use	192.168.4.128	/26	192.168.4.129-192.168.4.190	62	192.168.4.191
Not used at present	192.168.4.192	/26	192.168.4.193-192.168.4.254	62	192.168.4.255

Adding the Subnet use to the table

Let's add a column to our table to assign which subnet is used for what subnet so you don't confuse them. Notice there is a new column 1 that indicates the use of each subnet and the extra one that isn't being used at present. Documentation is important when you work with subnets. There are lots of IP addresses to track.

Subnet Use	Network ID	Subnet Mask	Host ID Range	# of Usable Hosts	Broadcast IDs
Reception/Front desk	192.168.4.0	/26	192.168.4.1-192.168.4.62	62	192.168.4.63
Mail room	192.168.4.64	/26	192.168.4.65-192.168.4.126	62	192.168.4.127
Guest/Public Use	192.168.4.128	/26	192.168.4.129-192.168.4.190	62	192.168.4.191
Not used at present	192.168.4.192	/26	192.168.4.193-192.168.4.254	62	192.168.4.255

Routing tables

Let's add a column to our table to assign which subnet is used for what subnet so you don't confuse them. Notice there is a new column 1 that indicates the use of each subnet and the extra one that isn't being used at present. Documentation is important when you work with subnets. There are lots of IP addresses to track.

Subnet Use	Network ID	Subnet Mask	Host ID Range	# of Usable Hosts	Broadcast IDs
Reception/Front desk	192.168.4.0	/26	192.168.4.1-192.168.4.62	62	192.168.4.63
Mail room	192.168.4.64	/26	192.168.4.65-192.168.4.126	62	192.168.4.127
Guest/Public Use	192.168.4.128	/26	192.168.4.129-192.168.4.190	62	192.168.4.191
Not used at present	192.168.4.192	/26	192.168.4.193-192.168.4.254	62	192.168.4.255

Adding routes to the route table

A routing table is a set of rules, often viewed in table format, that is used to determine where data packets traveling over an Internet Protocol (IP) network will be directed. All IP-enabled devices, including routers and switches, use routing tables. The table we created is a sample of the type of information you will track in the table.

Subnet Use	Network ID	Subnet Mask	Host ID Range	# of Usable Hosts	Broadcast IDs
Reception/Front desk	192.168.4.0	/26	192.168.4.1-192.168.4.62	62	192.168.4.63
Mail room	192.168.4.64	/26	192.168.4.65-192.168.4.126	62	192.168.4.127
Guest/Public Use	192.168.4.128	/26	192.168.4.129-192.168.4.190	62	192.168.4.191
Not used at present	192.168.4.192	/26	192.168.4.193-192.168.4.254	62	192.168.4.255



How routing works

You can scan your own computer's network connection properties to see its properties and begin to see how routing works at the local level.

In Windows 10:

1. Go to Settings > Network & Internet > Status > View hardware and connection properties.
2. The next screen displays details for your different network connections.
3. Scroll through the settings to find devices that are connected to the internet.

How routing works

Here is a sample of what you might see once you find a device that is connected to the internet. If it doesn't say it is connected, keep looking through all the devices that you might be connected to but just are not turned on at present time.

🏠 View hardware and connection properties

Name:	Wi-Fi	IPv4 address:	160.16.1.10/23
Description:	Intel(R) Wi-Fi 6 AX200 160MHz	IPv6 address:	fe80::d159:1d49:e48f:6645%18/64
Physical address (MAC):	b8:9a:8a:b8:fa:0b	Default gateway:	160.16.0.1
Status:	Operational	DNS servers:	160.16.0.1
Maximum transmission unit:	1500	DNS domain name:	
Link speed (Receive/Transmit):	78/86 (Mbps)	DNS connection suffix:	
DHCP enabled:	Yes	DNS search suffix list:	
DHCP servers:	160.16.0.1	Network name:	██████████
DHCP lease obtained:	Monday, December 27, 2021 5:17:57 PM	Network category:	Public
DHCP lease expires:	Tuesday, December 28, 2021 5:17:57 PM	Connectivity (IPv4/IPv6):	Connected to Internet / Connected to unknown network



How to access a router as an admin

You must be an administrator to access a router. If you run a large internet, you probably will need additional security to make changes to your route tables. Some companies only designate people with specific skills with the permission levels to make traffic changes to routers.

In this lesson, you will learn some very basic skills to access a router. Later, once you have a better understanding of network security, how Amazon Web Services handles VPC access, and other information, then you'll be ready for more advanced instructions.

There are several reasons you may need to access your router as an administer. One basic reason is to change the default username and password. You access the router through a web browser using either an Ethernet cable or a wireless connection.



How to access a router as an admin

1. Identify the IP address of the router.
2. Use a web browser with either Chrome, Firefox, or Internet Explorer and type in the router's IP address (in the format of <http://192.168.4.0>).
3. Use your admin login information to access the admin settings.
(**Note:** Routers are shipped with default usernames and passwords—usually, the word **admin**, but it could be different for your router. Some routers might not have a password or username.)

How subnetting works

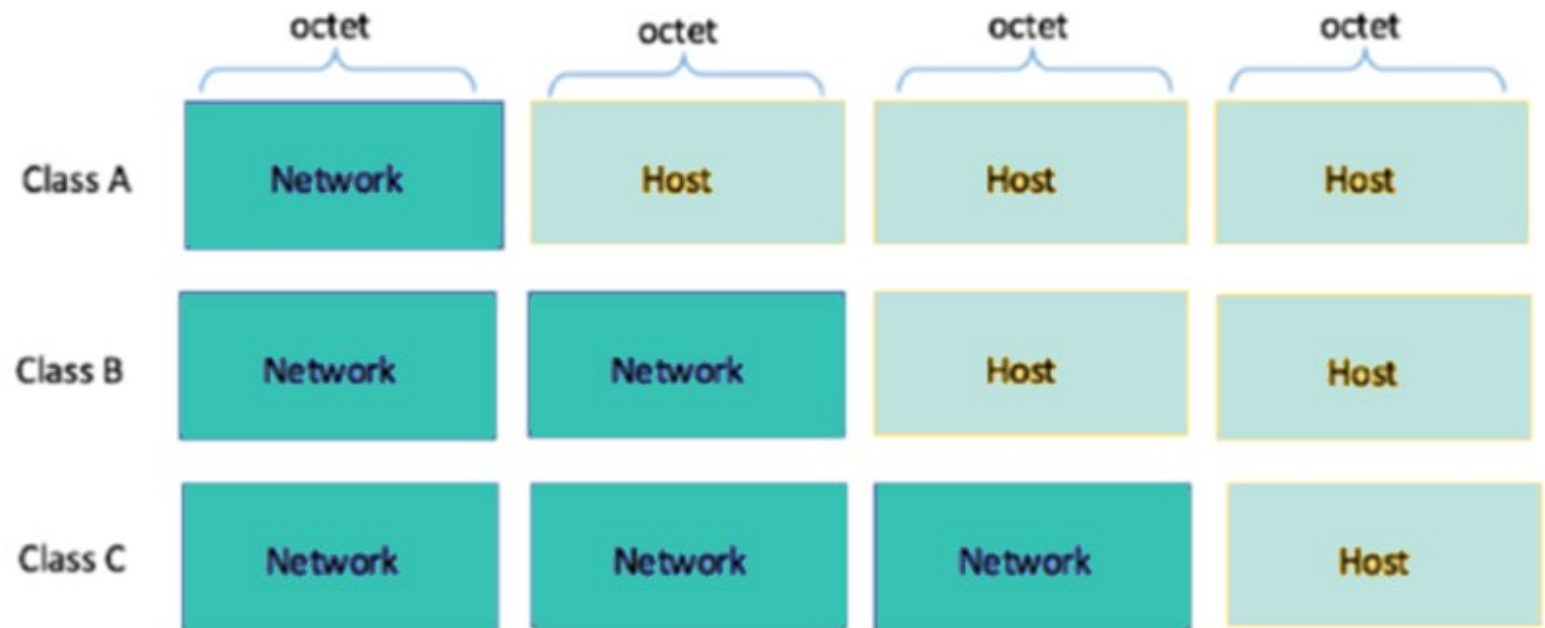
IP subnetting is a method for dividing a single, physical network into smaller subnetworks (subnets).

Subnetting in an IPv4 address gives you 32-bits to divide into two parts:

- a network ID
- a host ID

Note: Here is a table indicating the number of octets that can be used by class.

Depending on the number of bits you assign to the network ID, subnetting allows for either a greater number of total subnetworks or more hosts (devices that can be part of each subnet).





Key Takeaways

- A subnet is a logical organization of connected devices
- A subnet gives you a lot of capacity for growth
- A host device will have the last two octets as 0.0 but the subnet last next to last will have numbers ranging from 8, 16, or 24 depending upon its class.
- IPv6 uses 128 bits for its IP addresses and 32 hexadecimal numbers




What is a subnet mask?

What is a subnet mask?

Every device has an IP address with two pieces: the client or host address and the server or network address. IP addresses are either configured by a Dynamic Host Configuration Protocol (DHCP) server or manually configured (static IP addresses).

The subnet mask splits the IP address into the host and network addresses, thereby defining which part of the IP address belongs to the device and which part belongs to the network. It also covers up the subnet so that it isn't seen outside of allowed traffic.



The Router subnet is split up at the router. Notice the external IP address is completely different until it gets to the router and then it changes completely to another IP address.

Note: Routers can connect to different subnets and switches can connect to the same subnet.

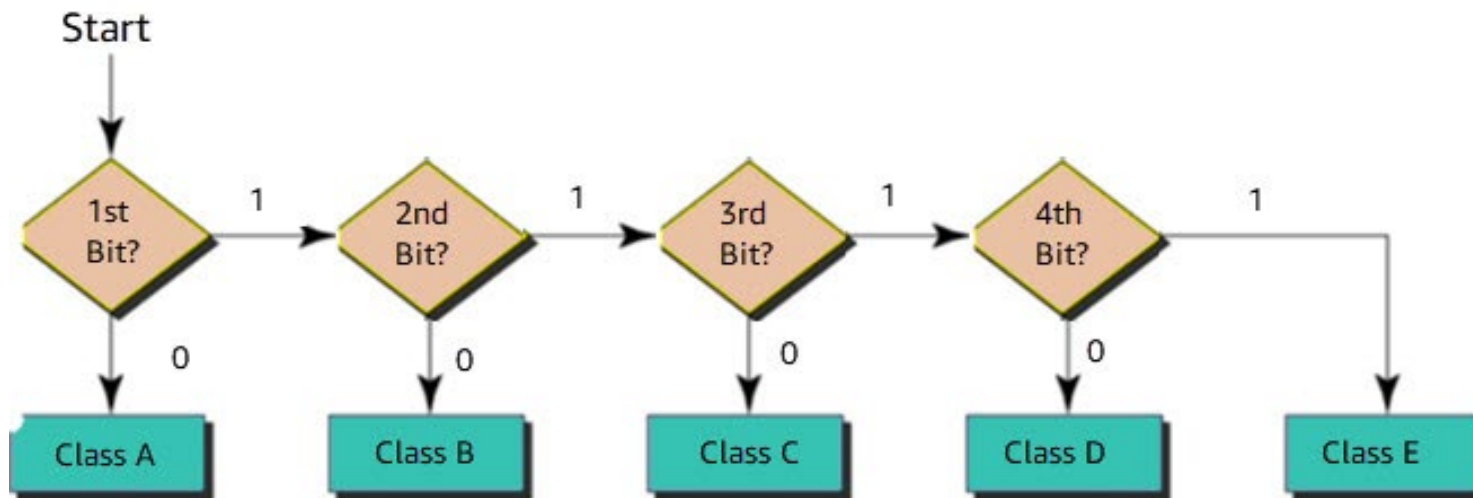
What is a subnet mask?

As you saw in the previous section, the IP address, subnet mask and gateway or router, together create an underlying structure called the Internet Protocol, that most networks use to facilitate inter-device communication.

When organizations require more subnetworking, as their need for additional devices grow, subnetting divides the host element of the IP address to further increase the number of available subnets. The goal of subnet masks are to create more subnets. The phrase “mask” is applied because the subnet mask essentially uses its own 32-bit number to mask the IP address. The subnet mask is used by the router to cover up the true network address. It shows which bits are used to identify the subnet but cannot be seen by just anyone. It allows for privacy and makes it more difficult for hackers to penetrate.

How to use a submask

Every network has its own unique address. Like here, class B network has network address 172.20.0.0, which has all zeroes in the host portion of the address. Example IP address: 11000001. Here 1st and 2nd bits are 1, and the 3rd bit is 0; hence, it is class C.



Identifying which class the IP address belongs to

This example shows how IP addresses should be deconstructed, which makes it simple for Internet routers to find the right Network to route data into. However, in a Class A network there could be millions of connected devices, and it could take some time for the router to find the right device.

Methods of subnet masking

There are two methods of submasking: Straight and Short-cut.

Straight:

Use the binary notation method for both the address and the mask. Then apply the AND operation to get the block address.

Cheat sheet for additional network calculations

Calculating the Network Address:

- The network address is the logical AND of the respective bits in the binary representation of the IP address and network mask.
- Align the bits in both addresses, and then perform a logical AND on each pair of the respective bits. Then convert the individual octets of the result back to decimal.
- Logical AND truth table:

INPUT		OUTPUT
A	B	A AND B
0	0	0
0	1	0
1	0	0
1	1	1

```
128.42.5.4      in binary: 10000000 00101010 00000101 00000100
255.255.248.0   in binary: 11111111 11111111 11111000 00000000
-----
                  10000000 00101010 00000000 00000000
```

As you can see, the network address of 128.42.5.4/21 is 128.42.0.0

Methods of subnet masking

There are two methods of submasking: Straight and Short-cut.

Straight:

Use the binary notation method for both the address and the mask. Then apply the AND operation to get the block address.

Class	Default subnet mask	No. of networks	No. of host per network
A	255.0.0.0	256	16,777,214
B	255.255.0.0	65,536	65,534
C	255.255.255.0	16,77,216	126

Subnet Cheat Sheet

CIDR	SUBNET MASK	WILDCARD MASK	# OF IP ADDRESSES	# OF USABLE IP ADDRESSES
/32	255.255.255.255	0.0.0.0	1	1
/31	255.255.255.254	0.0.0.1	2	2*
/30	255.255.255.252	0.0.0.3	4	2
/29	255.255.255.248	0.0.0.7	8	6
/28	255.255.255.240	0.0.0.15	16	14
/27	255.255.255.224	0.0.0.31	32	30
/26	255.255.255.192	0.0.0.63	64	62
/25	255.255.255.128	0.0.0.127	128	126

Subnet Cheat Sheet

CIDR	SUBNET MASK	WILDCARD MASK	# OF IP ADDRESSES	# OF USABLE IP ADDRESSES
/24	255.255.255.0	0.0.0.255	256	254
/23	255.255.254.0	0.0.1.255	512	510
/22	255.255.252.0	0.0.3.255	1,024	1,022
/21	255.255.248.0	0.0.7.255	2,048	2,046
/20	255.255.240.0	0.0.15.255	4,096	4,094
/19	255.255.224.0	0.0.31.255	8,192	8,190
/18	255.255.192.0	0.0.63.255	16,384	16,382
/17	255.255.128.0	0.0.127.255	32,768	32,766

Subnet Cheat Sheet

CIDR	SUBNET MASK	WILDCARD MASK	# OF IP ADDRESSES	# OF USABLE IP ADDRESSES
/16	255.255.0.0	0.0.255.255	65,536	65,534
/15	255.254.0.0	0.1.255.255	131,072	131,070
/14	255.252.0.0	0.3.255.255	262,144	262,142
/13	255.248.0.0	0.7.255.255	524,288	524,286
/12	255.240.0.0	0.15.255.255	1,048,576	1,048,574
/11	255.224.0.0	0.31.255.255	2,097,152	2,097,150
/10	255.192.0.0	0.63.255.255	4,194,304	4,194,302
/9	255.128.0.0	0.127.255.255	8,388,608	8,388,606

Subnet Cheat Sheet

CIDR	SUBNET MASK	WILDCARD MASK	# OF IP ADDRESSES	# OF USABLE IP ADDRESSES
/8	255.0.0.0	0.255.255.255	16,777,216	16,777,214
/7	254.0.0.0	1.255.255.255	33,554,432	33,554,430
/6	252.0.0.0	3.255.255.255	67,108,864	67,108,862
/5	248.0.0.0	7.255.255.255	134,217,728	134,217,726
/4	240.0.0.0	15.255.255.255	268,435,456	268,435,454
/3	224.0.0.0	31.255.255.255	536,870,912	536,870,910
/2	192.0.0.0	63.255.255.255	1,073,741,824	1,073,741,822
/1	128.0.0.0	127.255.255.255	2,147,483,648	2,147,483,646
/0	0.0.0.0	255.255.255.255	4,294,967,296	4,294,967,294

What is a subnet mask?

The IP address, subnet mask and gateway or router, together create an underlying structure called the Internet Protocol, that most networks use to facilitate inter-device communication.

When organizations require more subnetting, as their need for additional devices grow, subnetting divides the host element of the IP address to further increase the number of available subnets. The goal of subnet masks are to create more subnets. The phrase “mask” is applied because the subnet mask essentially uses its own 32-bit number to mask the IP address.

The subnet mask is used by the router to cover up the network address. It shows which bits are used to identify the subnet.

Every network has its own unique address. Like here, class B network has network



Subnet masks hide your private servers

For instance, if you pinged to find the root IP address of your company website, and found it to be 192.124.145.26, what parts would you be able to use to create a subnet so you could still use your host server, but direct incoming and outgoing traffic to information housed on different subnets?

If you recall from an earlier slide, the first two sections of network are the network identifier (144.28.0.0) and the subnets are identified as 144.28.16.0 and 32.0. But to the outside world, the IP address is only seen as 144.28.0.0 for security.



What is a subnet mask?

Every device has an IP address with two pieces: the client or host address and the server or network address. IP addresses are either configured by a Dynamic Host Configuration Protocol (DHCP) server or manually configured (static IP addresses).

The subnet mask splits the IP address into the host and network addresses, thereby defining which part of the IP address belongs to the device and which part belongs to the network.

Every device has an IP address made up of two pieces: You can use a subnet mask to define the range of IP addresses that can be used within a network or subnet. A subnet mask separates an IP address into two parts: 1) Network bits, and 2) Host bits.

By adjusting a subnet mask, you can set the number of available IP addresses within a network.



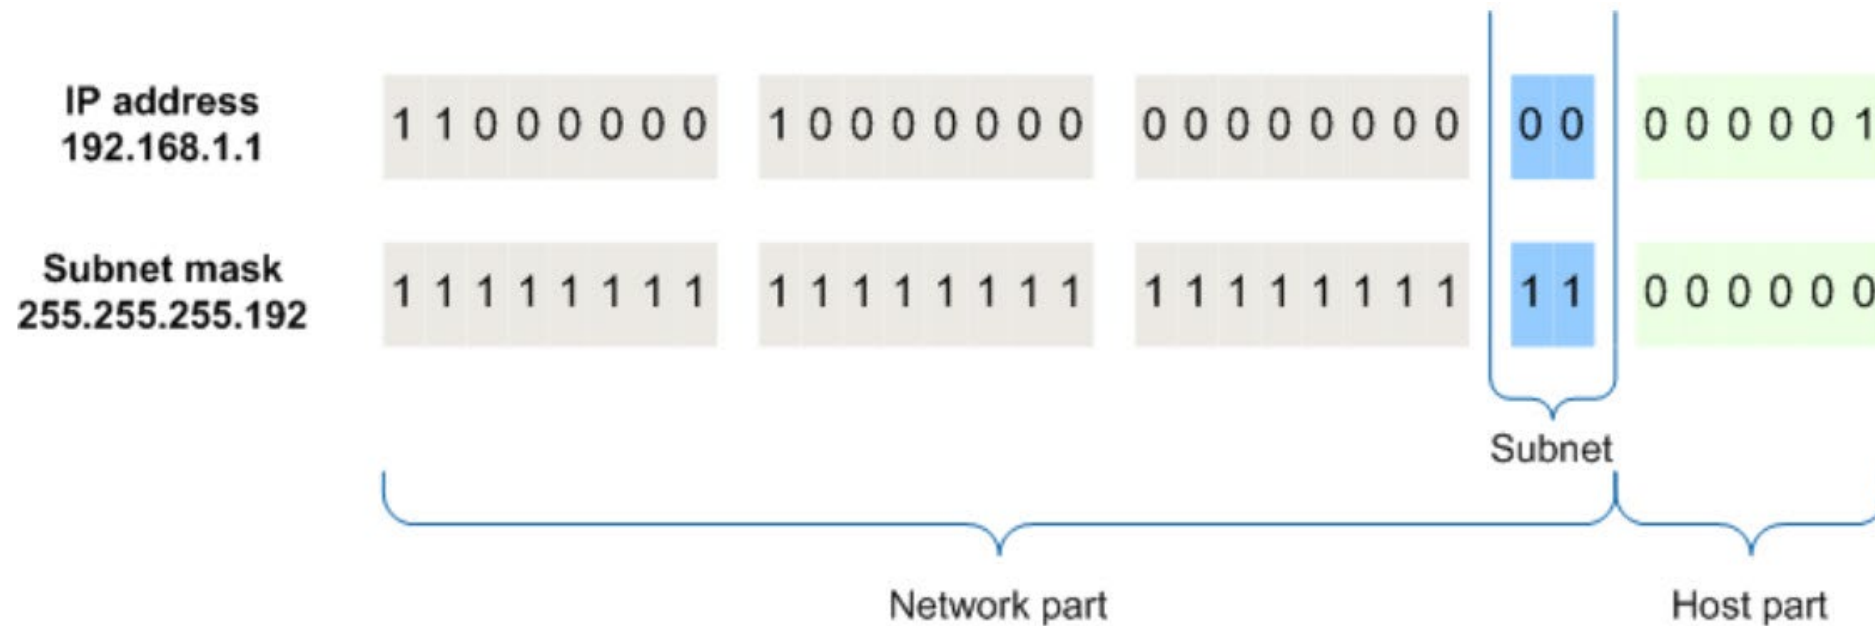
What is a subnet mask?

A common subnet mask for a home network would be: **255.255.255.0**.

This subnet mask lets you have up to 254 usable IP addresses in your home. You can use those address for all types of connected devices. Today, with all the technology we own, our phones, each laptop, computer, printer, smart tv's, gaming equipment, and more each need their own IP address. Using subnet masking allows you to have so many devices in your home without any IP conflicts. With such a large number of devices, you might need a high-speed internet to avoid speed conflicts.

Understanding how network bits and host bits are used

Subnet masks break an IP address up into network bits and host bits. When a device sees the network and host bits of another device's IP address, it can figure out if the other device is part of the same network (home, business, etc.), or is somewhere else online.



An IP address and subnet mask. Source: [IPv4](#)



Calculating the netmask length (prefix)

Convert the dotted-decimal representation of the netmask to binary. Then, count the number of contiguous 1 bits, starting at the most significant bit in the first octet (for example, the left-hand-side of the binary number).

255.255.248.0

in binary it equals: 11111111 11111111 11111000 00000000

There are 21 1's which equal “/21”

The prefix of 128.42.5.4 with a 255.255.248.0 netmask is /21.



Subnet mask key takeaways

- A subnet is a logical organization of connected devices.
- Classes are used to divide a host into subnets.
- A host device will have the last two octets as 0.0 but the subnet last next to last will have numbers ranging from 8, 16, or 24 depending upon its class.



What is a Classless Inter-Domain Routing (CIDR) notation?

Classes you can and cannot use in subnetting

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	IPv4 Bits for Network ID Sizes
0-126	Class A	34.126.35.125	8
128-191	Class B	134.23.45.123	16
192-223	Class C	212.11.123.3	24
224-239	Class D	225.2.3.40	Used for multicast and cannot be used for regular internet traffic
240-255	Class E	245.192.1.123	Reserved and cannot be used on the public internet



What is CIDR?

Both are IP addressing schemes that improve the allocation of IP addresses.

The general rule is that subnets are used at the organizational level but CIDRs are used at the ISP level and higher.

Subnets: When you place a mask over the subnet, you instantly create an entire subnetwork that is a subordinate network of the internet. The subnet mask signals to the router which part of the IP address is assigned to the hosts (individual participants of the network) and which determines the network.

CIDRs: This is a scheme to add suffixes to the integrate them directly into the IP address. Using CIDRs, you can not only create subnets, but also supernets. It also lets you subdivide a network more precisely, but lets you combine several networks.



CIDR notations

Previously, you were looking at a single IP address. But what if you want to send data to a range of IP addresses as in this example 192.168.1.0 and 192.168.1.255? How can you do that?

Classless Inter-Domain Routing (CIDR) notation is a compressed method of specifying a range of IP addresses. This method determines how many IP addresses are available to you. Here is an example of how to use a range of IP addresses.


192.168.1.0/24

When you add a slash after the fourth integer and another number, that number specifies how many of the bits in the IP address are fixed.

Earlier, you saw the /26 when you set up 3 subnets. This is another way of showing the range.

CIDR notations

In this example, the first 24 bits of the IP address are fixed. The rest are flexible.



32 total bits subtracted by 24 fixed bits leaves 8 flexible bits. Each of these flexible bits can be either 0 or 1, because they are binary. That means that you have two choices for each of the 8 bits, providing 256 IP addresses in that IP range.

This is slightly different than the easy way you learned to calculate the subnets earlier in this lesson.



Key takeaways

- Understand what CIDR notations are.
- How many characters of an IP address are fixed?
- Describe which classes are unusable.



What is Amazon's VPC?

What is Amazon's VPC?

Amazon Virtual Private Cloud (Amazon VPC) provides features that let you increase and monitor for Virtual Clouds hosted on AWS.

While you can learn about the specific features at <https://aws.amazon.com/vpc/features/>, this section will discuss how subnets are used with AWS VPC.



Amazon Virtual Private Cloud
(Amazon VPC)

One specific area of support enabled by AWS VPC is IP addressing.

IP addresses allow resources in your VPC to communicate with each other and with resources over the internet. Amazon VPC supports both the IPv4 and IPv6 addressing protocols.

Core Amazon VPC concepts

Amazon VPC is the networking layer for Amazon EC2. The following are the key concepts for VPCs:

- **CIDR block:** Classless Inter-Domain Routing. An internet protocol address allocation and route aggregation methodology.
- **Internet gateway:** A gateway that you attach to your VPC to enable communication between resources in your VPC and the internet.
- **Route table:** A set of rules, called routes, that are used to determine where network traffic is directed. This is not the document table we created earlier but some of the contents are similar.
- **Subnet:** A range of IP addresses in your VPC.
- **Virtual private cloud (VPC):** A virtual network dedicated to your AWS account.



Core Amazon VPC concepts contd

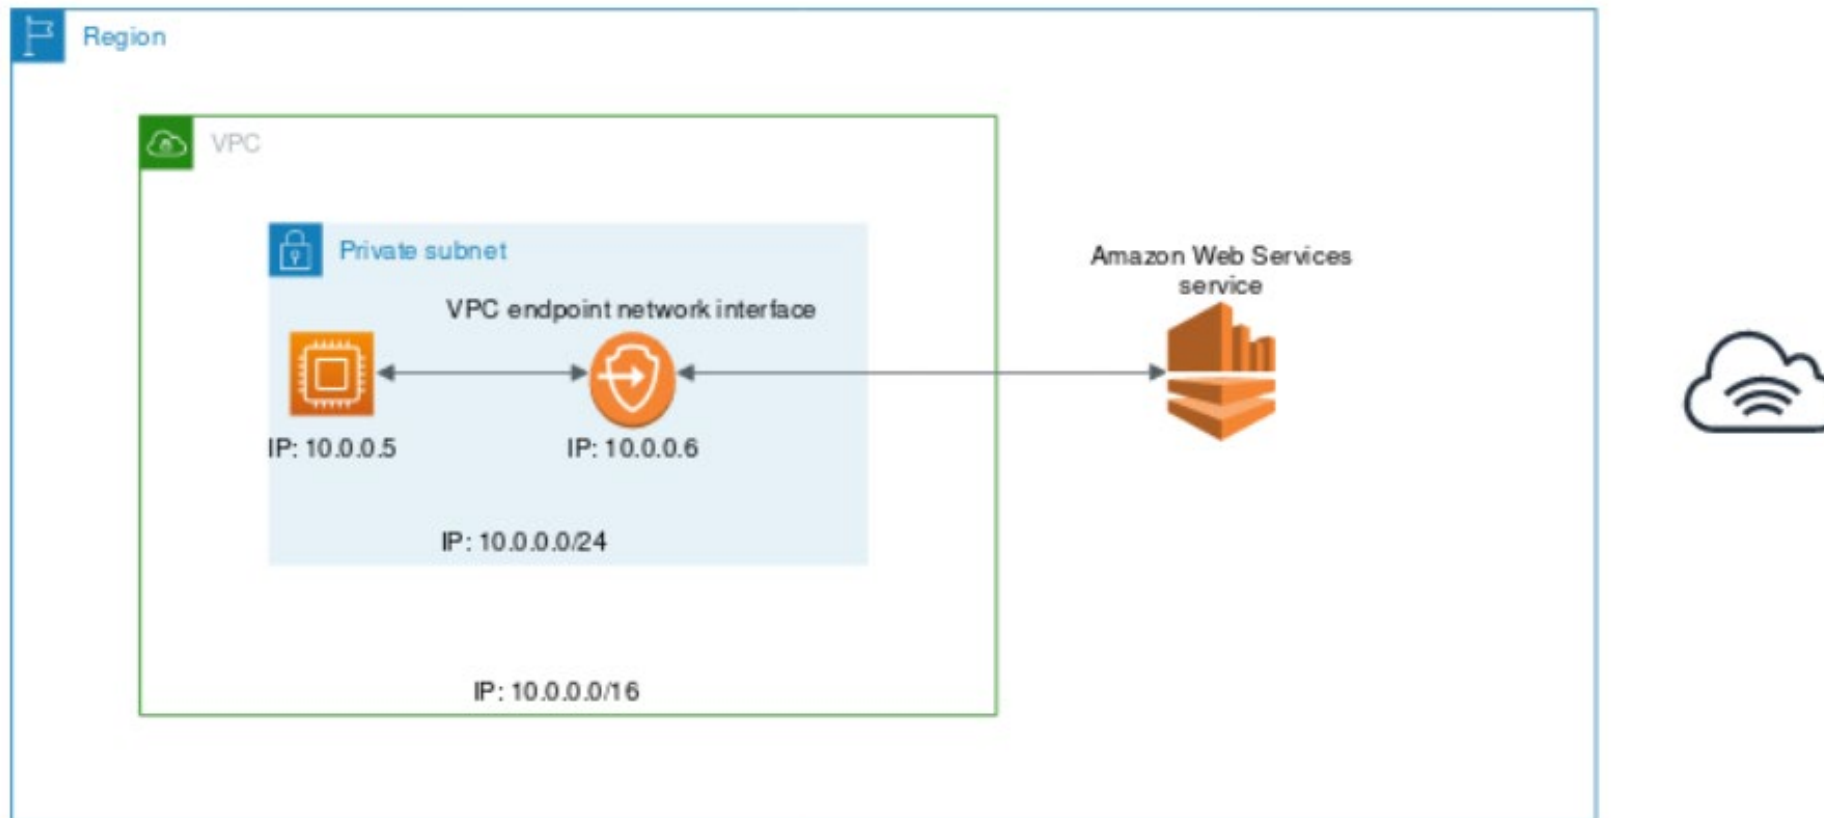
VPC endpoint: Use this service to privately connect your VPC to supported AWS services and VPC endpoint services powered by PrivateLink without requiring an internet gateway, NAT device, VPN connection, or AWS Direct Connect connection.

PrivateLink establishes the connectivity between virtual private clouds (VPCs) and services hosted on AWS or on-premises, without exposing traffic between your VPC and the service to the internet. (See the image in the following slide.)

When you use PrivateLink, you create an endpoint for the service in your VPC. The elastic network interface in your subnet connects with a private IP address that serves as an entry point for traffic destined to the service. The following diagram shows the basic architecture to securely connect your VPC to an AWS service that supports AWS PrivateLink.

Core Amazon VPC concepts contd

Instances in your VPC do not require public IP addresses to communicate with resources in the service. Traffic between your VPC and the other service does not leave the Amazon network. For more information, see [AWS PrivateLink and VPC endpoints](#).





Amazon's VPC Capabilities

In a VPC, you can create the following subnets:

- IPv4-only
- dual-stack
- IPv6-only subnets

You can also launch Amazon EC2 instances in these subnets.



Amazon's VPC Capabilities contd

Amazon also gives you multiple options to assign public IP addresses to your instances.

You can use the Amazon-provided:

- Public IPv4 addresses
- Elastic IPv4 addresses
- or an IP address from the Amazon provided IPv6 CIDRs

Apart from this, you have the option to bring your own IPv4 or IPv6 addresses within the Amazon VPC that can be assigned to these instances. You can read more about IP addressing in your VPC [in the VPC User Guide](#).

Note: To that your instances can communicate with the internet, you must also attach an internet gateway to your VPC. For more information, see [Internet gateways](#).



Key takeaways

- How does the Amazon VPC method of handling subnetting differ from other cloud providers?
- Why does it handle subnetting in this manner?
- Does Amazon VPC support one or both methods of IPv?
- What is AWS PrivateLink used for?



How Amazon VPC differs from other cloud providers



Amazon VPC differs from cloud providers

Subnet basics at Amazon VPC

A *subnet* is a range of IP addresses in your VPC. That part is the same as anywhere else. However, you can launch AWS resources, such as EC2 instances, into a specific subnet. When you create a subnet, you specify the IPv4 CIDR block for the subnet, which is a subset of the VPC CIDR block. This is different from the way subnets are handled outside of Amazon.

Amazon's VPC is scalable infrastructure that exists in the networking layer for Amazon's Elastic Compute Cloud (EC2) product. That's where this difference in handling subnets is accomplished.

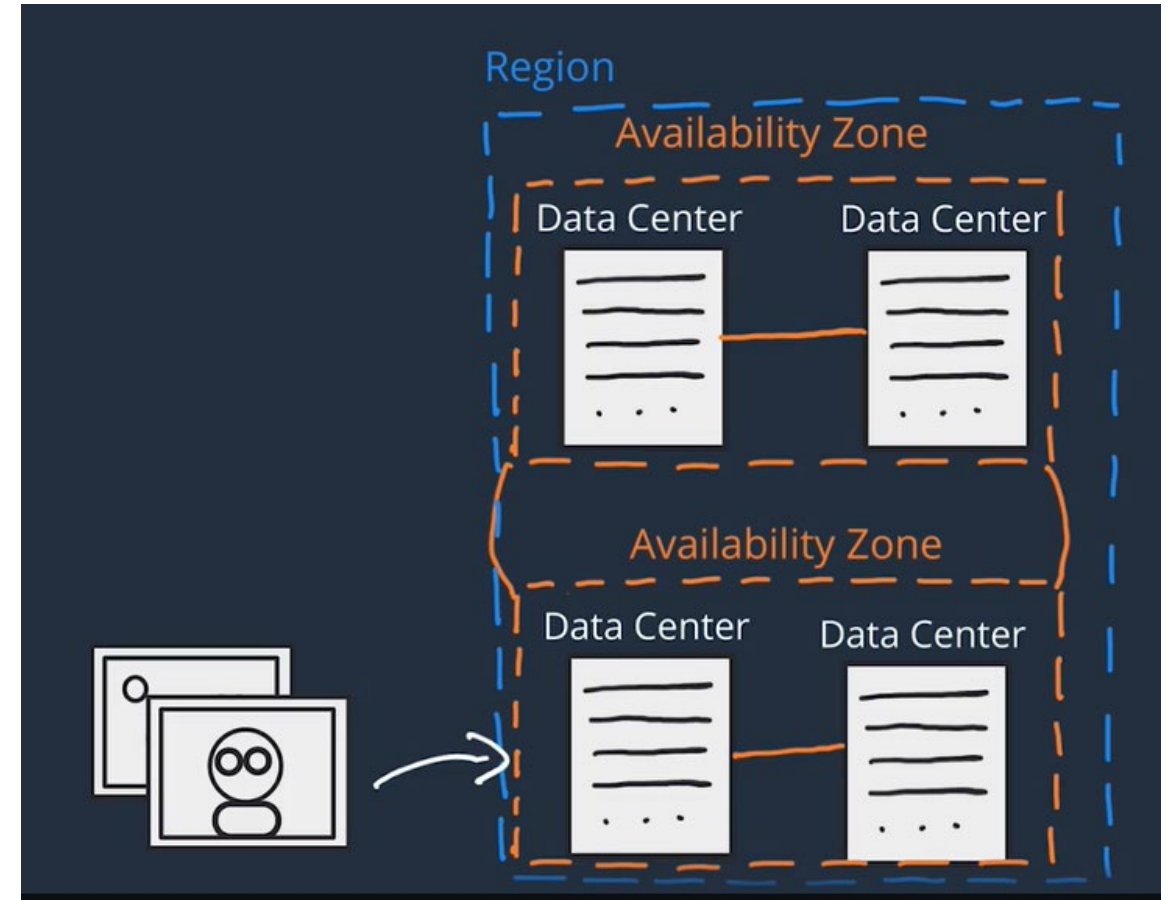
In the Amazon VPC, each subnet must reside entirely within one Availability Zone and cannot span zones. By launching instances in separate Availability Zones, you can protect your applications from the failure of a single zone.

The Amazon VPC architecture

To understand the differences between how subnets are used outside of Amazon and inside using VPC, let's start by understanding the potential scale and architecture of Amazon's network.

Amazon's infrastructure is made up of the following components:

- **Regions:** Physical locations around the world where multiple data centers are clustered.
- **Availability Zones:** A group of data centers are called availability zones. These are physically separated within a geographic area, although all are within 100 km (60 miles) of each other.





The Amazon VPC architecture contd

Data Centers: These are where multiple servers are housed. A data center can contain hundreds or even thousands of servers in large warehouses.

Unlike other cloud providers, who often define a region as a single data center, the multiple AZ design of every AWS Region offers advantages for customers.

Each AZ has independent power, cooling, and physical security and is connected via redundant, ultra-low-latency networks. AWS customers focused on high availability can design their applications to run in multiple AZs to achieve even greater fault-tolerance. AWS infrastructure Regions meet the highest levels of security, compliance, and data protection.



The Amazon VPC architecture contd

AWS provides a more extensive global footprint than any other cloud provider, and to support its global footprint and ensure customers are served across the world, AWS opens new Regions rapidly. AWS maintains multiple geographic Regions, including Regions in North America, South America, Europe, China, Asia Pacific, South Africa, and the Middle East.



IP address ranges and size relationships

The higher the number after the /, the smaller the number of IP addresses in your network. For example, a range of 192.168.1.0/24 is smaller than 192.168.1.0/16.

Note: When working with networks in the AWS Cloud, you choose your network size by using CIDR notation. In AWS, the smallest IP range you can have is /28, which provides 16 IP addresses. The largest IP range you can have is a /16, which provides 65,536 IP addresses.



Amazon VPC components

Amazon VPC comprises a variety of objects that will be familiar to customers with existing networks:

- **A Virtual Private Cloud:** A logically isolated virtual network in the AWS cloud. You define a VPC's IP address space from ranges you select.
- **Subnet:** A segment of a VPC's IP address range where you can place groups of isolated resources.
- **Internet Gateway:** The Amazon VPC side of a connection to the public Internet.
- **NAT Gateway:** A highly available, managed Network Address Translation (NAT) service for your resources in a private subnet to access the Internet.



Amazon VPC components contd

- **Virtual private gateway:** The Amazon VPC side of a VPN connection.
- **Peering Connection:** A peering connection enables you to route traffic via private IP addresses between two peered VPCs.
- **VPC Endpoints:** Enables private connectivity to services hosted in AWS, from within your VPC without using an Internet Gateway, VPN, Network Address Translation (NAT) devices, or firewall proxies.
- **Egress-only Internet Gateway:** A stateful gateway to provide egress only access for IPv6 traffic from the VPC to the Internet.



Why use Amazon VPC?

Amazon VPC lets you to build a virtual network in the AWS cloud - no VPNs, hardware, or physical datacenters required.

You can:

- Define your own network space
- Control how your network and the Amazon EC2 resources inside your network are exposed to the Internet
- Leverage the enhanced security options in Amazon VPC to provide more granular access to and from the Amazon EC2 instances in your virtual network.



What happens when I use the wizard?

Your AWS resources are automatically provisioned in a ready-to-use default VPC. You can choose to create additional VPCs by going to the Amazon VPC page in the AWS Management Console and selecting "Start VPC Wizard".

Ready
Set
Go!

You'll be presented with four basic options for network architectures. After selecting an option, you can modify the size and IP address range of the VPC and its subnets.

If you select an option with Hardware VPN Access, you will need to specify the IP address of the VPN hardware on your network. You can modify the VPC to add or remove secondary IP ranges and gateways or add more subnets to IP ranges.



Key takeaways

- What is the largest architectural area in Amazon's VPC?
- How does Amazon VPC handle IP address ranges and size relationships?
- What architectural area is the largest in Amazon's VPC?
- What are some of the key things that are done when you use the Amazon VPC to create your VPC network environment?



Thank you

© 2022 Amazon Web Services, Inc. or its affiliates. All rights reserved. This work may not be reproduced or redistributed, in whole or in part, without prior written permission from Amazon Web Services, Inc. Commercial copying, lending, or selling is prohibited. Corrections, feedback, or other questions? Contact us at <https://support.aws.amazon.com/#/contacts/aws-training>. All trademarks are the property of their owners.

