

4

Layer 3 Network Layer and the Internet Protocol (IP)

In this section you're going to learn about the Network Layer of the OSI model, and the Internet Protocol (IP). The Network Layer of the OSI model is the layer that's responsible for end-to-end delivery of network data, as opposed to the point-to-point delivery that happens in the Data Link Layer. Almost all networks, including the Internet use a protocol called the Internet Protocol or IP at this layer. The Internet Protocol defines several things, including IP addresses and the rules for packing data from the upper layers into packets. There's quite a bit to IP, so we'll be breaking this section up into the following:

1. History of IP
2. IP Address Basics
3. IP Packets
4. Network Mask (Netmask) Basics
5. Subnetting and Netmask Basics - Delivery of IP Packets
6. IP Address Assignment and Classes
7. Subnetting – Advanced
8. Non-routable IP Addresses
9. DNS Basics
10. Default Router (Gateway)
11. IP Configuration Basics
12. IPV6

At the end of this section you will be able to:

1. Define what IP addresses are, how they are used, and differentiate them from MAC addresses.
2. Given two IP addresses and a netmask, determine whether the IP addresses are on the same or different network segments.
3. Describe the role of the default router, and why its IP address must be part of the TCP/IP configuration.
4. Given a netmask, determine the number of networks and number of hosts.
5. Identify the sections of an IP packet.
6. List the IP address classes.
7. Describe how non-routable IP addresses function.

8. Describe how non-routable IP addresses freed up most of the IPv4 address space.
9. Describe the purpose of DNS in the network stack.
10. Configure TCP/IP settings on a Windows based computer.
11. Use Wireshark to analyze IP information.
12. Troubleshoot problems with TCP/IP configuration.

Resources

There are several resources that you should use to learn this material.

1. These Lecture Notes
2. Videos linked in Canvas.
3. Labsim Section 4

If you feel like any of the concepts presented in this section are not clear or need additional resources, you can always do your own research on the Internet or contact the instructor for help.

Lecture Notes

Introduction

In this section you're going to learn about the Network Layer and the Internet Protocol (IP) which is the main protocol used at this layer. You'll start with a quick history lesson, and then learn about the Network Layer and IP. Learning about how everything works in the Network Layer can be a little tricky because there are lots of interrelated parts, and understanding any single part requires learning about the other parts first. Rather than trying to learn all the components at the same time we're going to take cover everything three times. The first time through you'll get an overview of the components and how they work together. Next, we'll swing back around cover each separately, providing you with the basics of each component. And finally, we'll cover each component one more time, but this time covering any important details.

History of IP

Let's start with a quick look backwards at the history of the Internet Protocol (IP). You don't really need to know the names and dates, but learning about the history of the IP is important because it will help you understand how it became the main protocol used on the Internet. You'll also learn about the IETF which is the organization in charge of IP. You'll learn about how the IETF came to be in charge of maintaining the standards for all protocols used on the Internet, including IP.

In the early days of computer networking, there were several competing protocols for transmitting data over networks. However, the development of the IP proved to be a significant milestone in the history of networking, as it provided a standard way to move data from end-to-end in a network. Unlike protocols

that relied on MAC addresses to move data from point-to-point, IP allowed data to be transmitted across multiple networks, enabling the creation of the global internet we know today.

The Internet Protocol was developed in the early 1970s by a team of researchers led by Vinton Cerf and Bob Kahn^{1,2,3}. At the time, there was a need for a new networking protocol that could connect different computer networks together, as the existing protocols were not suitable for this purpose. Cerf and Kahn's work on IP was part of a larger project known as the ARPANET, which was a research project



Bob Kahn and Vinton Cerf, developers of the Internet Protocol. From <https://northernviriniamag.com/culture/culture-features/2021/07/23/meet-mclean-residents-bob-kahn-and-vint-cerf-they-invented-the-internet/>

funded by the US Department of Defense. The ARPANET was the precursor to the modern internet, and the development of IP was a critical step in the creation of the internet as we know it today.

The IP protocol defines two main items, the format for network packets and a system of addresses called IP Addresses. You'll learn the details about both these items later in this section, but for now here's a quick overview. The IP network packet contains two parts: the header and the data. The header contains information about the source and destination of the packet, as well as other information. The data contains the actual information being transmitted, such as an email message or a web page. The other key component defined by the IP protocol is the system of IP addresses. An IP address is a unique identifier assigned to every device connected to the Internet and they are used to move network data from end-to-end, as opposed to MAC addresses which are used to move data from point-to-point. IP

¹ Cerf, V. G., & Kahn, R. E. (1974). A protocol for packet network intercommunication. *IEEE Transactions on Communications*, 22(5), 637-648

² Cerf, V. G., & Kahn, R. E. (1978). The current state of the art of packet switching: An overview. *IEEE Communications Magazine*, 16(2), 129-138

³ Internet Society. (n.d.). Brief history of the Internet. Retrieved from <https://www.internetsociety.org/internet/history-internet/brief-history-internet/>

addresses are a 32-bit number that is usually represented in a dotted decimal notation, such as 192.168.0.1.

IP was designed to be a simple, lightweight protocol that could route packets of data across different networks. However, it had some limitations, such as a lack of reliability and no built-in error checking. To address these issues, Cerf and Kahn developed the Transmission Control Protocol (TCP), which was added to IP to create the TCP/IP protocol suite. TCP/IP was developed in the 1970s and adopted as the protocol standard for ARPANET (the predecessor to the Internet) in 1983. The development of TCP/IP was a significant breakthrough in the history of computer networking, as it allowed different computer networks to communicate with each other using a common set of protocols. Today, TCP/IP is the foundation of the internet, and it is used by virtually every device that is connected to the internet.

Sadly, unlike other tech moguls like Gates, Zuckerberg, and Musk, most people won't recognize the names Cerf and Kahn even though their work made today's Internet possible. They are highly recognized and respected by the tech industry and were awarded the National Medal of Technology and Innovation in 1997.

IETF – Who's (not) the Boss

While Cerf and Kahn were the original developers of IP, they eventually turned the responsibility for managing the protocol to the Internet Engineering Task Force (IETF). The IETF is a standards organization that is open to anyone who wants to participate. The IETF develops standards through a process of consensus-building and publishes the resulting standards in a series of documents known as Request for Comments (RFCs).

The Internet Engineering Task Force (IETF) was created in 1986 as a result of the merging of two groups: the Internet Configuration Control Board (ICCB) and the Defense Data Network (DDN) Network Information Center (NIC). The ICCB had been responsible for managing the early development of the internet, while the DDN NIC had been responsible for managing ARPANET, the defense network that the internet had grown out of. The merger of these two groups resulted in the creation of the IETF, which was tasked with developing and promoting voluntary technical standards and protocols for the Internet. The IETF was initially composed of a small group of researchers and engineers who were working on the development of the internet, but it has since grown into a large, international organization with thousands of members from around the world.

The IETF is not "in charge" of internet technical standards in the traditional sense. Rather, it is a community-driven organization that operates through a process of open, collaborative discussion and consensus building. The IETF doesn't have any legal authority to enforce its standards or protocols, but its work is widely respected and adopted by industry and governments around the world. The only way for the Internet to work is for everyone that uses it to cooperate and follow the IETF's suggested standards. This means that anyone that manages a network segment connected to the Internet, and all the organizations that route packets between these network segments must agree to cooperate. If you take a minute to think about this you might be pleasantly surprised that everyone does cooperate and that the Internet does work, especially in today's highly polarized social climate where it seems that no one wants to cooperate for the greater good.

In any case, the IETF is responsible for developing and maintaining many of the protocols that are used on the internet, including IP. The IETF publishes the standards in documents called RFCs, which stands for Request for Comments. RFCs are technical documents that contain information on how to implement various protocols, technologies, and standards that make up the Internet. The RFCs are not meant for general use or explanation; instead, they serve as a blueprint for building and implementing various technologies. They include detailed technical specifications, diagrams, and examples to help developers create and maintain interoperable systems. Some of the most well-known RFCs include:

RFC 791 - Internet Protocol (IP)

RFC 822 - Standard for the Format of ARPA Internet Text Messages (Email)

RFC 1034 and 1035 - Domain Name System (DNS)

RFC 1945 - Hypertext Transfer Protocol (HTTP/1.0)

RFC 8446 - Transport Layer Security (TLS 1.3)

These RFCs are just a few examples of the many standards and protocols that make up the Internet's infrastructure. You can find a complete list on several Internet sites, but the official repository run by the IETF can be found at <https://www.rfc-editor.org/>. While the IETF has serious responsibilities, they also have a well-developed although slightly geeky, sense of humor. There are several RFCs that are meant to be jokes and can be fun to read.⁴

One thing about the RFCs that many people find curious is the name "Request for Comments", instead of something more authoritative like "Standards" or "Technical Specifications and Requirements". The story behind the name "RFC" comes from the early days of the ARPANET, the precursor to the Internet. In 1969, Steve Crocker wrote the first RFC, which was a memo describing how to implement a simple host-to-host protocol for sending and receiving messages on the ARPANET. According to Dr. Crocker, the term "Request for Comments" was originally used as a bit of a joke. He and his colleagues were concerned that calling their publications "standards" might come across as overly presumptuous or arrogant, given that the technology they were working on was still very much in its infancy.

As Dr. Crocker recounted in a 2012 interview with the Internet Society⁵:

"The first RFC was simply an internal note, but when we decided to publish it, we needed a name. I personally was very uncomfortable with the term 'standard' because we were just a bunch of graduate students, and it seemed presumptuous to call what we were doing a 'standard.' I just suggested 'Request for Comments' partly as a joke, but partly also to say, 'Look,

⁴ <https://tangentsoft.com/rfc/humorous.html>

⁵ <https://www.internetsociety.org/wp-content/uploads/2017/09/rfc-retrospective.pdf>

we're just putting this out there. We're not claiming it's perfect. We're not claiming it's a standard. We're asking for feedback."

The term "Request for Comments" caught on and has since become synonymous with the process by which new Internet standards are developed and refined. Today, the RFC series includes over 9,000 documents, and remains an essential part of the Internet's infrastructure.

When the Internet was first starting the IETF developed a few "rules" that any network had to follow to connect to the Internet. These rules include the following:

1. The Internet Protocol must be used.
2. Every device on the network must have a unique IP address, which for smaller networks can be obtained from an ISP, and for larger networks can be obtained from the official internet registry, the Internet Assigned Numbers Authority (IANA), or a regional internet registry (RIR). You'll learn the details about obtaining IP Addresses later in this section.
3. The Domain Name System (DNS), which is used to translate human-readable domain names such as google.com into IP addresses, must be used. To connect a network to the internet, the network must have at least one DNS server that is able to resolve their DNS names into IP addresses. registry (RIR). You'll learn the details about DNS later.
4. To connect a network to the internet, the network must have a routing protocol in place that allows it to exchange routing information with other networks and routers on the Internet. You'll learn the details about routing later.

The main rule is that you must use IP as the network protocol. This isn't a rule in the sense that it's a legal standard, it's just that you must use the Internet Protocol get your network data through the Internet routers.

Once again, meeting these requirements is essential for ensuring that the network can communicate with other networks and devices on the Internet. Connecting a network to the Internet without meeting the necessary requirements can have various consequences, the main one being that devices on the network will be unable to share information with other devices on the Internet. If your network doesn't follow the IETF standards and starts causing problems with other networks, the router(s) connecting your network to the Internet will shut down your connection. Unlike the rest of human society, the Internet realizes that the only way to work is to work together and cooperate. For this reason, all countries use the same set of IP numbers and recognize the Internet Assigned Numbers Authority (IANA) as the organization responsible for allocating IP addresses and managing the global IP address space. While some countries may have their own regulations and policies related to the use of IP addresses and the management of their own national networks, the underlying technical standards and protocols used for communication on the internet are global and universal, they're not limited to the United States.

If a country, such as China, were to decide to ignore the Internet Assigned Numbers Authority (IANA) and set up their own set of IP addresses, it would create a parallel system that would not be compatible with the existing global internet and devices and networks within that country would not be able to communicate with devices and networks on the global internet. Such a move could have significant geopolitical and economic consequences, as it would likely lead to fragmentation and isolation of the internet. This could impact the ability of businesses and individuals to engage in international commerce, research, and communication. So, even though other countries may have the technical capability to create its own set of IP addresses, doing so would likely have significant negative consequences and would not be compatible with the global nature of the Internet.

While you don't need to memorize any of the history of the Internet Protocol, there are a few things you should take away from this section. The first is the concept of how the Internet and the groups that control IP and the other protocols used on the Internet protocols came about. That is, they weren't formed by a company, and they weren't part of an international effort like the ISO. Instead, the Internet grew out of a US government research project. But even though the Internet started in the US it's now a global network that only maintains a global span because of a cooperative effort among all parties that use it.

The second point is a technical point, and this is that the Internet Protocol defines two main items, the format for network packets and a system of addresses called IP Addresses. So, let's quit talking about IP's history and talking about what it is in a vague sense, and start learning about IP Addresses and IP network packets.

Overview of the Network Layer Process and Components

Now let's start learning about what actually happens in the Network Layer, and the protocols and components used in the process. In this section you're going to get the 40,000 foot view of what happens in the Network Layer. The main thing to note is that the Network Layer is responsible for transmitting data between networks. That is, the Data Link Layer, ethernet frames, and MAC addresses can be used to send network data between devices on the same network segment, they aren't able to send data between devices on different networks. An analogy for this would be with snail mail, where your local postal carrier can deliver mail between houses and buildings in your town, but they can't deliver mail between houses and buildings in different counties or different states.

The main component used in the Network Layer is a protocol called the Internet Protocol, which defines three main things. It defines a format for the packets, called IP packets, that will be used to hold the data being transferred, a format for network addresses which are called IP addresses, and something called netmask which is used to interpret how much of each IP address describes a network and how much of the address describes a specific device on the network.

The IP packet is much like an ethernet frame, in the sense that it's used to send network data and has a header section that contains things like source and destination addresses, and a data section. But when you look at the specifics IP packets and ethernet frames there are some obvious differences. One of the big differences is that IP packets will use IP addresses instead of MAC addresses.

IP addresses are made of 4 numbers, called octets, with dots or periods separating the numbers. Each octet number can range from 0 to 255, although the numbers 0 and 255 serve special purposes. An example IP address would be 12.234.1.98. Each device on a network connected to the Internet must

have a unique IP address, but unlike MAC addresses which are assigned by vendor, the IP addresses are assigned to networks, so all the devices on the same network segment will have similar addresses.

Each IP address actually contains two pieces of information, the network number, and a unique number assigned to each host on the network. This is kind of like MAC addresses, where the first 6 numbers describe the vendor and the last 6 describe the specific device, but there are two huge differences. The first, as you just learned, is that the first part of the IP address describes a network instead of a vendor, and the second is that the portion of each IP address that describes the network can vary, and you can't tell how many of the octets describe the network without another piece of information. The network portion of an IP address may be the first octet, it may be the first two octets, or it may be the first three octets. Or in advanced cases, the network part of the IP address may be a fraction of one of the octets. The only way to know how to divide an IP address into the network portion and the host portion is by using a third piece of information called the netmask.

Another network component used by IP and the Network Layer is a router, which Microsoft calls the default gateway instead of the default router for some reason. A router is a networking device that forwards data packets between different computer networks. Hopefully you remember that switches and hubs connect devices to create a single network segment, while routers connect multiple networks or network segments. On most networks, there's a single router that connects the network to the Internet or the organizations larger network, and in Microsoft terminology this device is called the default gateway. Since the Network Layer is tasked with moving network data between networks, one of the key components in the process is the default gateway, as this is the device that will handle any network packets that are being sent to devices on other networks. Whenever you configure the IP settings on a computer, you must provide the IP address of the default gateway.

The last main component used by the Network Layer is the Domain Name System or DNS. This is the system that converts or resolves names like google.com to IP addresses. DNS is used because as humans are much better at remembering names than we are at remembering long strings of numbers. You can think of DNS as being like a phone book, but instead of converting human names to phone numbers, DNS resolves names to IP addresses. Any network that connects to the Internet needs to have access to two DNS servers that it can use to resolve names to IP addresses, and any computer connected to a network needs the IP addresses of the two DNS servers.

Here's a summary of the protocols, components, and services used by the Network Layer:

1. IP Protocol – defines the format for IP packets.
2. IP Addresses – used to identify networks, and hosts on those networks.
3. Netmask – used to determine how to divide an IP address into the network portion and the host portion.
4. Default Gateway – IP address of the router used to transmit packets to other networks.
5. DNS Servers – Resolve host names to IP addresses.

Now let's look at an overview of the process used by the Network layer, and how it uses the protocols, components, and services to package and deliver data across the network.

1. The Network Layer receives a request to send data from the upper layers in the OSI network stack. This request includes the data to send, and the DNS name of the host to send the data to.

2. The Network Layer will build an IP Packet by:
 - a. Placing the data in the data portion of the IP Packet
 - b. Building the IP packet header, placing the source and destination IP addresses in the correct locations.
 - i. The source IP will be the IP address of this computer, which must be configured.
 - ii. The destination IP address will be the IP address associated with the DNS name. The Network Layer will put the current IP packet on hold and create a new IP packet containing a DNS request. The DNS system will respond with the IP address of the destination computer.

3. The Network Layer will start to build an ethernet frame by:
 - a. Placing the IP Packet in the data section of the ethernet frame.
 - b. Placing the MAC address of this computer in the source MAC portion of the ethernet frame header.
 - c. Decide the destination for the ethernet frame by comparing the network portions of the source and destination IP addresses using this process:
 - i. Use the netmask to determine the network portion of each IP address. The netmask must be configured on each computer that uses IP.
 - ii. Compare the two network portions. If the network portions of the source and destination IP addresses are identical the ethernet frame will be sent directly to the destination device. If the network portions of the source and destination IP addresses are different the ethernet frame will be sent to the default router.
 - d. While the IP address of the next hop in the network transmission is known, the MAC address may not be known. If not, use ARP to discover it.

Ok ... that's the overview, and as you can see there's a lot going on. It may be a bit fuzzy now, but after you step into the details of the various protocols and components it should start to make more sense.

IP Address Basics

One of the things defined by the IP protocol, and used at the Network layer are IP addresses. IP addresses are used to uniquely identify a computer or device on the network and allow end-to-end delivery. This might seem just like MAC addresses as they're both unique addresses used to identify devices on a network, but there are a couple of significant differences.

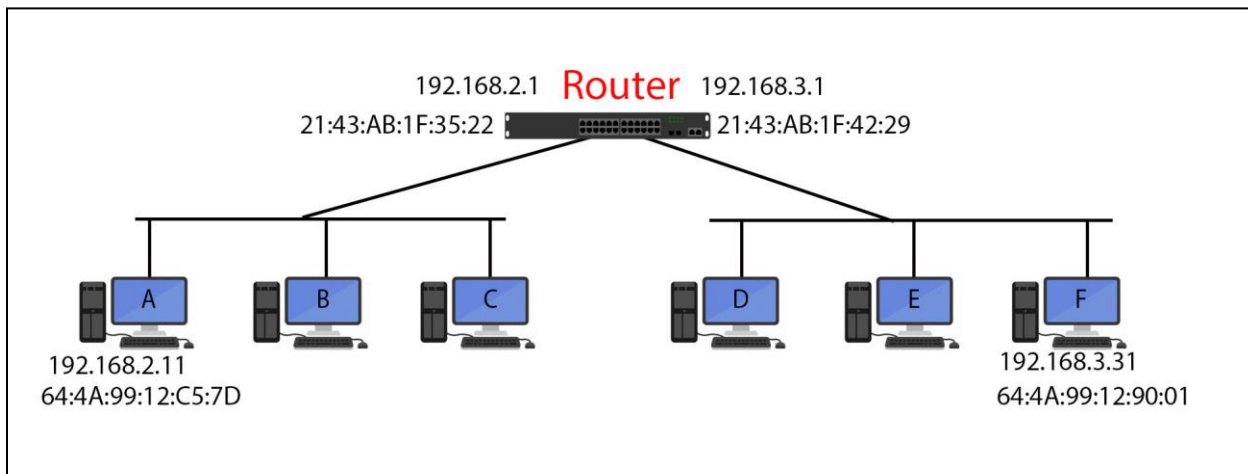
The first difference is that MAC addresses are burned onto a network card and cannot be changed. (Ok, this isn't technically correct, as the MAC address is burned onto an EEPROM that can be reprogrammed to change the MAC address, but this isn't a simple process.) IP addresses are configured for each computer and can be easily changed. You can think of this as the relationship between you and your phone number. You are like the MAC address, it doesn't matter where you go or what you do, you will always be glorious you. Your phone number is like an IP address, as it can always be changed, it's just a number that people use to call you, but it's not you. Of course, it might be painful for your friends if you change your number, but the point is it can be changed, while no matter what phone number you have you will be the same person.

You should note that just because IP addresses can be changed, it doesn't mean they are frequently changed. But if you connect your computer to a different network, you will be assigned a new IP

address. For example, if you connect your computer to your home network it will have one IP address. But, if you later take your computer to school or on the road and stay in a hotel and connect to a different network, it will get a different IP address. You'll learn the details of this later.

The second difference between IP addresses and MAC addresses is that IP addresses are used for end-to-end delivery, while MAC addresses are used for point-to-point delivery. Any computer connected to a network will have both a MAC address and an IP address. The MAC address is used to transmit data from point-to-point, while the IP address is used to transmit data from end-to-end.

Let's look at an example of what this means, using the following figure. Say that computer A wants to send some data across the network to computer B. In this case computer A and computer F are on different networks, with a router connecting the two network segments. When we say end-to-end, we mean from computer A to computer F. But for the packets to get from end-to-end they need to go through the router, which means there will be a few point-to-point hops. The first point-to-point hop will be from computer A to the router, and the second point-to-point hop will be from the router to computer F. You'll learn the details of how this happens later, the main point to take away is that IP addresses are used for getting network data from end-to-end, while MAC addresses are used to move network data from point-to-point.



IP Address Versions

There are two versions of IP addresses, called IPv4 and IPv6. The IPv4 addresses were the original addresses and were originally just called IP addresses. As you'll learn there are only so many IPv4 addresses, over ~4.2 billion, which is a lot, but not enough if every computer and device connected to the Internet needs one. When it appeared that the Internet would need more addresses than IPv4 could provide work was started on IPv6, which mainly increased the address space to $\sim 3.4 \times 10^{38}$ but also made a few other improvements. Once IPv6 was released, we started using IPv4 to refer to the old version. But if you just mention IP addresses in most cases you can assume that this means IPv4, and if IPv6 is being used it will be called out by name.

Most of the rest of the discussion in this section will be based on IPv4, but once again I'm just going to call it IP. While IPv6 is being used in some situations, most organizations are still using IPv4. And the processes used by IPv4 will still be applicable in IPv6, so I think it will be simpler and less confusing to concentrate on IPv4. You will learn about IPv6 at the very end of this section, and that will be the only time, besides now, that we use IPv6 and not IPv4.

If you're the curious type and want to know why the only two versions are 4 and 6, you can look the explanation up on the Internet read <https://www.ipxo.com/blog/what-happened-to-ipv5/>

Anatomy of an IP Address.

Now it's time to learn what an IP address looks like, and how this information is used. IP addresses are 4 part numbers, with each of the numbers separated by a dot or period. Each individual number can range from 0-255, although you never see a 0 in the first number, and 255 has a special meaning so it isn't used in a normal device address.

173.143.229.111
0-255 . 0-255 . 0-255 . 0-255

Here are a few examples of valid and invalid IP addresses. See if you can find the problems with the invalid examples.

Valid

182.34.29.100

192.168.1.1

12.13.14.15

254.1.31.254

Invalid

182.34.29

192.168.1.1.2

12.13.14.355

255.1.31.254

173.143.229.111

1010 1100 . 1000 1111 . 1110 0101 . 0110 1111

An IP address is really 4 8-bit binary numbers, but they're easier for humans to read and understand if we convert the binary numbers to decimal numbers.

The reason the numbers range from 0-255 is because they're actually 8-bit binary numbers that range from 0000 0000 to 1111 1111. But because binary is so awkward for humans to deal with, we convert the numbers to decimal, which results in the 0-255 range. To work with IP addresses, you will need to be able to convert numbers between binary or base 2, and decimal or base 10. If you don't know how to do this, you can use the video tutorials I've created or find other tutorials on the Internet. You can also make use of a calculator, which will do the conversion for you, but if you're planning on a career in cyber security or in any aspect of computer science, you'll find that this is a necessary skill, so I suggest you take the time to learn how to work in binary on your own without the use of a calculator.

Each number is also referred to as an octet, since it's 8 bits. Using this terminology, we would say that an IP address is made up of 4 octets. You'll often hear portions of an IP address referred to as the first octet, which would be the number before the first dot, or the first two octets which would be the first two numbers.

Sidebar on Counting in Binary

If you want to earn your Junior Geek badge one of the things you need to be able to do is count in binary or the base 2 number system. Before we jump into binary, we'll do a quick review of base 10 or the decimal number system you're already familiar with to provide you with a perspective.

In base 10 the places in a number are assigned values. That is, the right most numeral is in what is referred to as the 1's place. Moving left, the next numeral is in what's referred to as the 10's place, followed by the 100's place, followed by the 1000's place, etc. Each place value is calculated by raising the number base, in this case 10, to a power starting with 0 on the right and increasing the power by 1 for each position as you move left. That is, the value of the first place is 10^0 , the value of the second place (moving left) is 10^1 , the value of the third place is 10^2 , and the value of the fourth place is 10^3 .

1000's place

10^3

100's place

10^2

10's place

10^1

1's place

10^0

Sidebar on Counting in Binary (Continued)

Remember that any number to the 0th power is 1, and any number to the 1st power is itself.

Calculating the total value of a number is done by multiplying the numeral in each place by the place value, and then summing the result of all the multiplications. For example, if we had the number 4378 the total value would be:

$$(4 \times 1000) + (3 \times 100) + (7 \times 10) + (8 \times 1)$$

<u>4</u>	<u>3</u>	<u>7</u>	<u>8</u>
1000's place	100's place	10's place	1's place
10^3	10^2	10^1	10^0

The point of this isn't to tell you what 4378 is. I know you know what it is, even without calculating $(4 \times 1000) + (3 \times 100) + (7 \times 10) + (8 \times 1)$. We can all pretty much look at numbers in base 10 and tell exactly what they are. The point of this to remind you how number systems work, as binary or base 2 works exactly the same with a few differences. The first difference is that each place is going to be a power of 2 instead of a power of 10. The second difference is that there are only two numerals in binary, 0 and 1. The following diagram shows the values of each place in a 4-digit binary number.

<u> </u>	<u> </u>	<u> </u>	<u> </u>
8's place	4's place	2's place	1's place
2^3	2^2	2^1	2^0

Converting between binary and decimal is also done using the same general process used for decimal numbers. That is, first multiply the number in each place by its place value, and then sum the results of all the multiplications. For example, 0110_2 would be $(0 \times 8) + (1 \times 4) + (1 \times 2) + (0 \times 1) = 6_{10}$.

<u>0</u>	<u>1</u>	<u>1</u>	<u>0</u>
8's place	4's place	2's place	1's place
2^3	2^2	2^1	2^0

Since IP addresses contain 8 bits in each place, you'll need to have 8 places, with the left most place having the value of 2^7 or 128 and the right most having the value 2^0 or 1. The following shows all the place values for an 8-bit number.

Sidebar on Counting in Binary (Continued)

Since IP addresses contain 8 bits in each place, you'll need to have 8 places, with the left most place having the value of 2^7 or 128 and the right most having the value 2^0 or 1. The following shows all the place values for an 8-bit number.

128	64	32	16	8	4	2	1
2^7	2^6	2^5	2^4	2^3	2^2	2^1	2^0

How to view a device's IP address

Now let's show you how to view a device's IP address. Every device connected to an IP based network needs to have an IP address. And, since the Internet is an IP based network, any device connected to the Internet will have an address, even your phone or home computer will have one if they're connected to the Internet through your home network.

I'll show you how to view the IP address on a Windows based computer, a Linux computer, and an Android phone, since I have those types of devices.

Windows: On a Windows computer the `ipconfig` command is used to display network settings, including the IP address. To run the `ipconfig` command, follow this process:

1. Start the Command Prompt application by going to the Windows start bar (or search area) in the lower left, and type **cmd**.
2. Windows will display the Command Prompt app. Click on this to open a command window.
3. Place the cursor in the Command Prompt window and type **ipconfig**. If you get more than one window's worth of data type **ipconfig | more**. The **|** character is typically located above the <enter> key. This will display something that looks like the following:
4. The IP address for the computer will be displayed in the IPv4 Address field. In this figure the IP address is 10.90.98.162

```
Command Prompt
Connection-specific DNS Suffix . :
C:\WINDOWS\system32>ipconfig | more
Windows IP Configuration

Ethernet adapter Ethernet 2:

    Media State . . . . . : Media disconnected
    Connection-specific DNS Suffix . :

Ethernet adapter Ethernet:

    Connection-specific DNS Suffix . : arda.cbc
    Link-local IPv6 Address . . . . . : fe80::5266:6cfe:aa04:b194%18
    IPv4 Address. . . . . : 10.90.98.162
    Subnet Mask . . . . . : 255.255.252.0
    Default Gateway . . . . . : 10.90.96.1
```

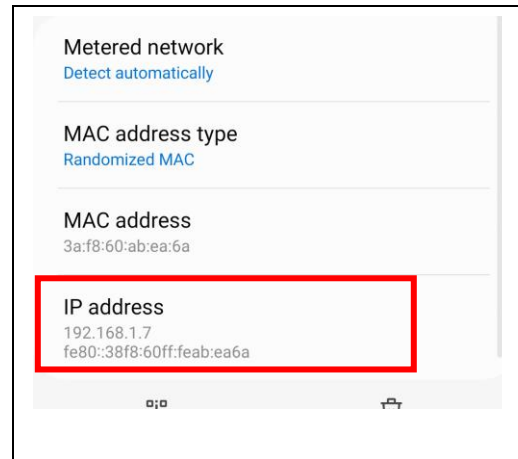
Linux: You don't really need to know all the Linux commands, unless you're taking the Linux class. Although you may have to do this in a few of the Labsim labs. In any case, on a Linux computer the `ip a` command is used to display network settings. To run the `ip a` command, follow this process:

1. Get to the Linux command line, by either logging in to an account that starts with the command line, or logging in to an account that uses a GUI and then starting a Command Window.
2. Type the `ip a` command.
3. The IP address for the computer will be displayed under the `eth0` interface, as the first part of the `inet` field. In this figure the IP address is 192.168.34.19

```
[root@balrog ~]# ip a
1: lo: <LOOPBACK,UP,LOWER_UP> mtu 65536 qdisc noqueue state UNKNOWN group default qlen 1000
    link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
    inet 127.0.0.1/8 scope host lo
        valid_lft forever preferred_lft forever
    inet6 ::1/128 scope host
        valid_lft forever preferred_lft forever
2: eth0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc mq state UP group default qlen 1000
    link/ether 00:15:5d:32:8e:a6 brd ff:ff:ff:ff:ff:ff
    inet 192.168.34.19/24 brd 192.168.34.255 scope global noprefixroute eth0
        valid_lft forever preferred_lft forever
    inet6 fe80::215:5dff:fe32:8ea6/64 scope link
        valid_lft forever preferred_lft forever
```

Android: This is another process you don't need to know, but I thought I'd include it in case you're curious. To see the IP address on an Android phone, follow these steps:

1. First, ensure the phone is connected to a network.
2. Go to Settings / Connections / Wi-Fi
3. Select the Gear icon by the Current Network
4. Select View More at the bottom of the screen.
5. Scroll down until you see the IP address.



Apple Devices: If you have an Apple Mac computer the process is very similar to the Linux computer since the MacOS is a version of Linux. If you have an Apple phone, I feel sorry for you, just kidding but not really, you can find tutorials on the Internet that will show you how to find the IP assigned to your phone. Just remember, your portable devices will only have an IP address if they're connected to a network.

Ok, that's it for the basics of IP addresses. You'll learn more about how they're used and how they're assigned a little later in this section.

IP Packets

Another thing specified by the IP protocol and implemented at the Network Layer of the OSI stack is the structure for the packets that will contain the data to be sent across the network.

Snail mail or UPS packages make a good analogy for explaining the IP rules for creating network packets. With snail mail, anytime you want to send someone a letter you must put it in an envelope and address the envelope correctly. There are sets of rules for how many pages you can place in a single envelope, and for where and how the sender and recipient addresses must be written⁶. The addressing rules state that the recipient address must be in the center on the front of the envelope, and the sender's address must be in the upper left corner. The addressing rules also state that the format of each address must be the sender's name on the first line, street address on the second line, with the city, state, and zip code on the third line. There is also a rule about how much you can fit in a single envelope, or to be more correct there's a rule that says you can only send up to 1 ounce with a single stamp, which is typically 4-6 pages of normal weight paper and a standard envelope. For our analogy we'll assume that you're always going to use a normal size envelope, and let's say you can only put 5 pages in each envelope. So, if you want to send more than 5 pages, you'll have to break your message up into multiple envelopes, and number the envelopes so the person you're sending them to knows how to reassemble them in the correct order.

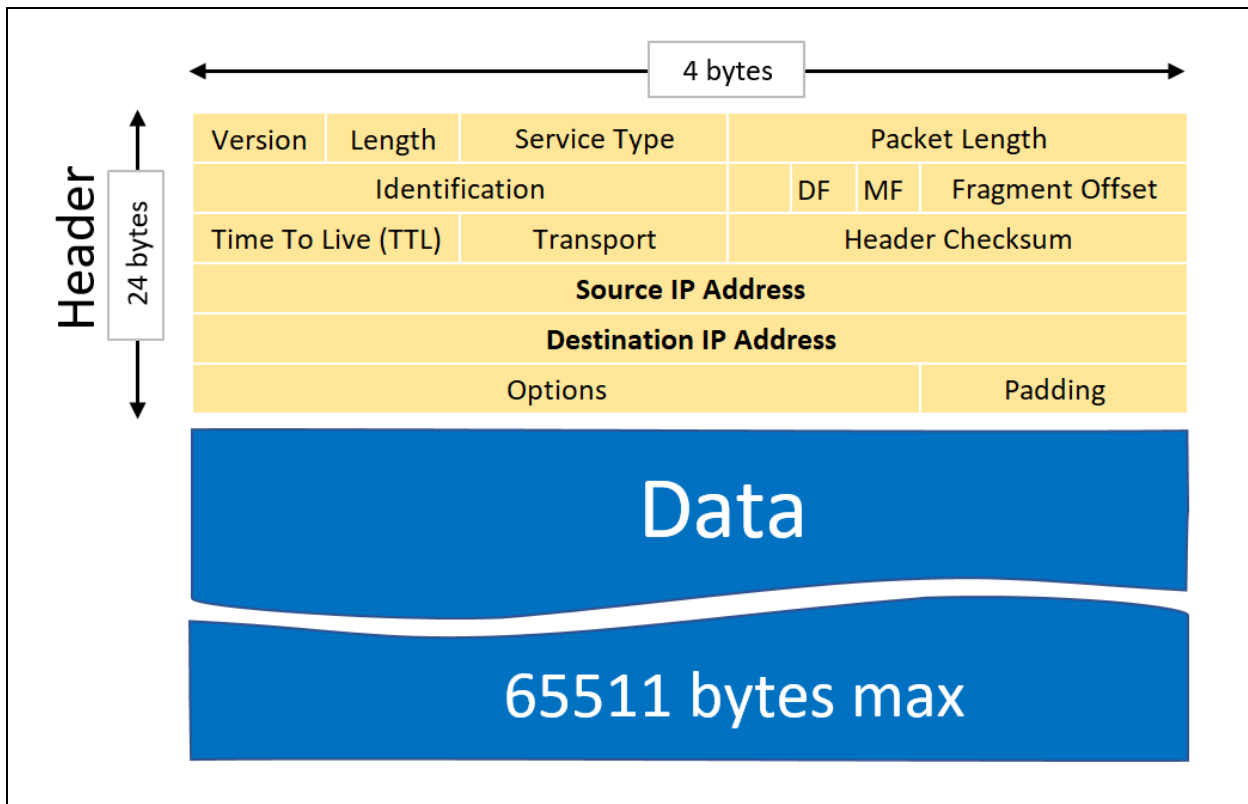
With IP, the task of building network packets falls to the Network Layer, but instead of placing letters inside envelopes these layers take the data and place it inside network packets. The network packets are

⁶ <https://mystampguide.com/mail-pages-paper-with-one-stamp/>

really just strings of binary data, with the packet organized into a header section, which is like the outside of the snail mail envelope, and the data section which is like the pieces of paper containing the actual message you want to send and put inside an envelope.

These IP packets are used to send data from end-to-end across a network, and the protocol specifies rules for addressing the packets and how to place the data inside the packet. When you first learn about IP packets you may have a feeling of Deja vu because if we speak in general terms an IP packet will be much like an ethernet frame. They're both used to send network data, and they both have a header with the sender and recipient's address. However, ethernet frames and IP packets are only the same in general terms, their main purpose is different, and the specifics of their structure is completely different.

Let's start by inspecting the structure of an IP packet, which is shown in the following figure. Each packet can be up to 65535 bytes in total, with a header section which is 24 bytes. This means the packet can hold up to 65511 bytes of data, which is the 65535 total minus the 24 bytes for the header.



Each packet starts with a header section which is organized as shown in the figure. The header stores things like the Packet Length which is a number that shows how much data is being sent and a Header Checksum which is used to check and ensure that the data in the header hasn't been mangled or changed during transmission. While all the header information is important, the items we're going to concentrate on now are the Source IP Address and the Destination IP Address. At this point you don't

need to understand or memorize all the fields in the header, you just need to be able to describe how an IP packet contains a header as well as the actual data, and how the header contains the Source and Destination IP addresses.

Now let's look at what the source and destination addresses are, and how the Network Layer gets them. The Source IP address, as the name implies, is the IP address of the computer sending the data. Any computer connected to the Internet must be configured with an IP address, so it will be readily available to the Network layer.

The Destination IP address is the IP address of the computer the data is being sent to. While a computer will know its own IP address, it probably won't know the IP address of the destination computer. It will most likely know the destination's computer's name, but it won't know the IP address. The reason it won't know the Destination IP is because the application sending the data down the network stack to the Network layer, will most likely pass a computer's name in the form of a DNS name. For example, if the user is running a web browser and clicks on a link, the browser will pass a request for data down the network stack, asking for a URL like <http://google.com> or <http://ebay.com>. Or, if the user is sending email, the email client will send data down the network stack asking to send data to an email address like joe@gmail.com or elon@tesla.com.

The computer names are called Domain Name Service (DNS) names. DNS is a system that was developed to resolve DNS names to IP addresses, because humans are much better at remembering names than they are remembering numbers. This is the same reason that your phone allows you to create a contact list, so you can simply say call Mom instead of having to remember her number.

When the Network layer receives a request to send data to a DNS name, the first thing it has to do is resolve the name to an IP address. It does this by placing your current IP packet on hold, building a different network packet with the DNS request, and sending this packet to the DNS server. The DNS server will then return a network packet containing the IP address associated with the DNS name. You'll learn all the steps in the DNS process later, but for now, let's just assume that it works. At this point, the Network layer can resume building the header for your original IP packet, placing the IP address of the destination computer in the packet's header.

Once the IP packet header is complete the network layer will add your data to the data section. The data can be as small as a single byte or as large as 65511 bytes.

65511 might seem like a weird number, because it is, but there's a reason why this is the maximum amount of data. Since the IP packet can hold a varying amount of data, one of the fields in the header is used to say how much data the packet holds. This way devices that receive the packet will know how much data to look for when they read the packet. The field in the header that holds the length of the packet is 2 bytes or 16-bit binary number. This means the largest the number can be is 2^{16} or 65535. You might think that this means that the packet could hold 65535 bytes of data, but the length also includes the header. And since the header is 24 bytes, the maximum amount of data that can be placed in the packet must be reduced by 24 bytes, which means the max amount of data is 65511 bytes.

If the original amount of data is larger than 65511 bytes, the Network layer will break it up into multiple packets. The Network layer will keep track of the sequence of packets and store the sequence number for each packet in the Fragment Offset field in each packet's header. For example, if the data required 5

packets, the Fragment Offset in the first packet would be set to 1, the Fragment Offset in the second packet would be set to 2, etc. This way the Network stack on the recipient computer can put the data back together in the correct order.

Subnetting and Netmask Basics - Delivery of IP Packets

Once the IP packet is built, the next thing the Network layer needs to do is determine whether the destination IP address is on the same network segment or a different network segment. It needs this information to tell the Data Link Layer whether to send the packet directly to the destination, or if the Data Link Layer should send the network packet to the default gateway/router. That is, the IP packet contains the end-to-end addressing, but not the point-to-point addressing. The Network layer will decide the next point in the delivery process using another piece of information called the netmask.

You just learned about IP addresses and how they are used to identify computers and devices on a network. Even though each IP address looks like a single number it actually has 2 parts. Part of the address describes which network a device is on, and the other part identifies the computer or device itself. We generically call these two parts the network portion of the IP address and the host portion. Determining which part of each IP address is the network portion and which part is the host portion requires using another number called the netmask. In this section you'll learn about network and host portions of an IP address, and how the netmask is used to specify which part of an IP address identifies the network and which part identifies the host.

Let's start by looking at how an IP address is used to specify two things, a network segment, and a specific host on the network segment. In some ways this is like street addresses or land line phone numbers. In a street address, most of the address identifies a street in specific city or town, while the rest of the address identifies a specific building on that street. Or, with a complete land line phone number, there will be several parts, a country code, an area code, the prefix or exchange number, and finally the line number. When you dialed a phone number with the phone system in the 1930s and 1940s, the country code, area code and exchange would get your phone call connected to a human operator in a specific geographic location. You would then ask the operator to connect your call to a specific line, and the operator would physically move some a cable to patch your call from the incoming line to the line you're calling⁷. The human operators were eventually replaced by mechanical switching devices, but the concept and the parts of the phone number remains the same, at least for land line numbers.

1	(314)	159	-	2653
country	area code	prefix or exchange		line

⁷ <https://www.youtube.com/watch?v=r46zXIN3Nus>

The point of all this is that snail mail addresses and land line phone numbers contain two pieces of information. Part of each mail address gets the mail to a street and part of the address identifies a specific building, while part of each phone number gets the call to an exchange, and the rest of the phone number identifies a specific phone connected to the exchange.

IP addresses are similar in that they contain two pieces of information, part of the IP address identifies a specific network segment and part of the address identifies a specific device on that network segment. However, there's one big difference between how IP addresses, and street addresses and phone numbers divide their information.

The big difference is that phone numbers and street addresses are consistent in the way they divide the overall number or address, while the portion of the IP address that describes the network and the part that describes the host can vary. In mail addresses, the part that identifies the specific house is always fixed, and in the same location.

742 Evergreen Terrace

Springfield OT 87342

The same thing for phone numbers, the part of a land line phone number that identifies a specific phone line will always be the same part of any phone number. It's different with IP addresses, where the part of the IP address that specifies a specific computer on a network segment can vary.

With IP addresses the network portion may be the first octet, or it may be the first two octets, or it may be the first three octets. For example, if we look at the IP address 22.56.79.115 the network portion could be 22, in which case the part that identifies a specific computer would be 56.79.115. Or the network portion could be 22.56, in which case the host portion would be 79.115. Or the network portion could be 22.56.79, in which case the host portion would be 115. The network portion is always made up of numbers on the left, but you can't tell just by looking at the IP address. To make this determination you'll also need another number called the netmask.

22.56.79.115

22.56.79.115

22.56.79.115

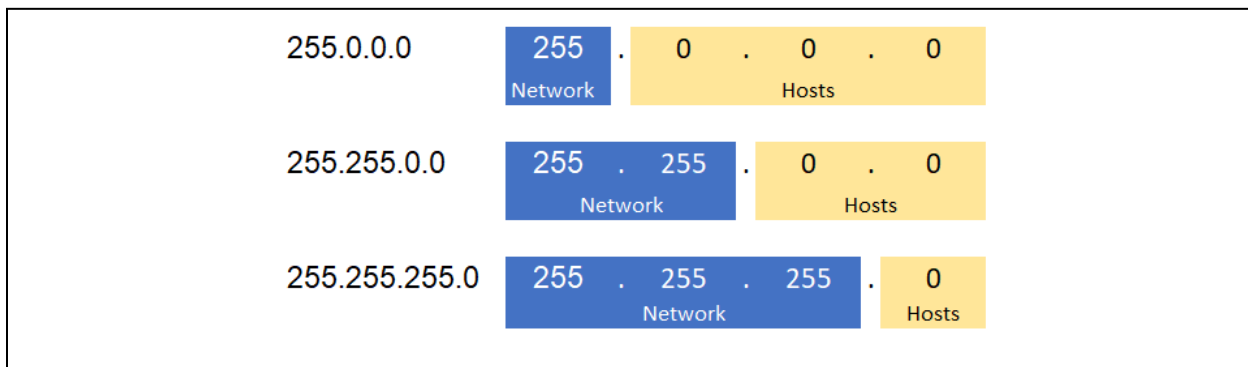
The netmask is also made up of 4 numbers, separated by dots or periods. However, in a netmask the only numbers used are 255 or 0. (This is a small lie, as other numbers can be used, which you'll learn about later, when you learn about advanced subnetting.) For example, 255.255.255.0 would be a valid netmask. To apply a netmask and find the network portion of an IP address, find the octets in the netmask that contain the number 255. These octets will be the ones in the IP address that describe the network. Any octets in the netmask that contain a 0 will correspond to the host portion of an IP address.

For example, if we know an IP address is 22.56.79.115, and the netmask is set to 255.255.255.0, we know that the first three numbers describe the network. Which means in this example the network

portion of the IP address would be 22.56.79. This also means that the host portion of the IP address would be the last octet, or in this case 115.

Here's another example using the same IP address, 22.56.79.115, but this time with a netmask of 255.255.0.0. Since the netmask has 255s in the first two octets, the network portion of the IP address is 22.56, while the host portion of the IP address is 79.115.

When netmasks are configured the network portion of the IP address must always be made up of octets on the left, and the host portion must always be made from octets on the right. This means that the only valid netmasks will be 255.255.255.0, or 255.255.0.0, or 255.0.0.0. This is one of those occasions where I'm not telling the truth. There are other valid netmasks, but they're a bit more complicated. You'll learn about them later, but for now we'll stick with these basic netmasks.



One thing to note is that the 255s will always be to the left part, and the 0's will always be to the right. That is, you won't ever see a netmask like 0.255.255.0 or 255.0.255.0. The netmask 0.255.255.255 would not be valid because it builds the host portion of the IP address from the left, not the right. The netmask 255.0.0.255 would also not be valid as it splits the network portion and puts part of it on the right.

Hopefully you can see that the netmask is a critical piece of information for any IP network implementation, and that building and using a network with more than a single segment would be impossible without knowing the netmask. For this reason, the netmask, along with the IP address, the DNS server IP address, and the default router or default gateway are the 4 pieces of information that must be configured on any device that wants to connect to an IP based network.

You'll learn more about netmasks later, but for now, you just need to know the following:


1. IP addresses are used to deliver network packets from end-to-end.
2. Each device connected to a network must have a unique IP address.
3. IP addresses are made up of 4 numbers or octets, with each number ranging from 0-255.
4. The only way to know how an IP address is divided into the network portion and the host portion is to use the netmask.

- The numbers assigned to each octet in a netmask must be 255 or 0. Any octet that contains a 255 will describe the network and must be on the left, and any octet that contains a 0 describes the host and must be on the right.

Using the Netmask

In this section you'll learn about using the netmask to determine whether two IP addresses are on the same network segment or different segments. This is an important part of the overall networking process, as it allows the Network Layer to tell the Data Link Layer where to send the ethernet frames. That is, this is the process for determining the next point in the point-to-point transmission.

Let's use the following diagram to illustrate this decision-making process, and how the netmask is used. The figure shows two network segments connected by a router, with computers A, B, and C on Network 1, and computers D, E, and F on Network 2. As you've learned, if computer A wants to send data to computer B, the data link layer will address the ethernet frames using computer A's MAC address as the source MAC and computer B's MAC address as the destination MAC. However, if computer A wants to send data to computer E, the data link layer on computer A will have to address the ethernet frames to the router's MAC. The question or problem that the netmask solves, is how does computer A's network stack decide whether to send the ethernet frames directly to another device on the same network segment, or whether to send the frames to the router?



In general terms this decision is made by comparing the network portion of the IP addresses of the source and destination devices. If the network portions show the devices are on the same network, then the ethernet frames will be sent directly, but if the network portions show the devices are on different networks the ethernet frames must be forwarded through the router, or possibly several routers.

But once again, how do we know which portion of the IP address is the host portion? This is where we must know the netmask, as it tells us exactly how to divide the IP address into the network portion and the host portion.

Let's go through a few examples using a netmask of 255.255.255.0, which specifies that the network portion of the IP address is the first three octets.

If Computer A wants to send network packets to Computer E, it will compare its IP address of 192.168.2.11 with Computer E's IP address which is 192.168.3.31. Since the netmask is 255.255.255.0 Computer A's Network layer will compare the first three octets of the IP addresses. For Computer A this is 192.168.2 and for Computer E it's 192.168.3. Since the network numbers are different, Computer A will tell its Data Link Layer to send the ethernet frame to the router, instead of sending the ethernet frame directly to Computer E.

Now let's look at what happens if Computer A wants to send network packets to Computer C. In this case Computer A will compare its IP address of 192.168.2.11 with Computer C's IP address which is 192.168.2.31. Since the netmask is 255.255.255.0 Computer A's Network layer will compare the first three octets of the IP addresses. For Computer A this is 192.168.2 and for Computer C it's 192.168.2. Since the network numbers are the same, Computer A will tell its Data Link Layer to send the ethernet frame directly to Computer C.

It's important to note that there's a difference between the IP Network Packet addressing and the ethernet frame addressing. The IP packet will always be addressed with the end-to-end IP addresses, while the ethernet frames will be addressed with the MAC addresses of the next point in the delivery process. Hopefully this makes sense, but if not, you'll get a detailed demonstration of this a little later.

Before we end this section, here are a few examples, where you can check to see if you are able to apply netmasks to determine whether two IP addresses are on the same network segment or different segments.

Example 1:

IP address 1: 192.168.1.10

IP address 2: 192.168.1.20

Netmask: 255.255.255.0

In this example, because the netmask is 255.255.255.0 the first three octets describe the network, and both IP addresses have the same numbers in the first three octets, 192.168.1. This means that both IP addresses belong to the same network segment.

Example 2:

IP address 1: 192.168.1.10

IP address 2: 192.168.2.20

Netmask: 255.255.255.0

In this example, because the netmask is 255.255.255.0 the first three octets describe the network. When we compare the first three octets, we find that the last number is different between the two IP addresses, 1 for the first address and 2 for the second address. This means that the IP addresses belong to different network segments.

Example 3:

IP address 1: 192.168.1.10

IP address 2: 192.168.4.20

Netmask: 255.255.0.0

In this example, the netmask is 255.255.0.0, which means that the first two octets are the network portion, and the last two octets are the host portion. This means that both IP addresses belong to the same network segment because their first two octets, 192.168, are identical.

Example 4:

IP address 1: 192.168.1.10

IP address 2: 10.0.0.20

Netmask: 255.255.255.0

In this example, the IP addresses belong to different network segments because their first octets are different 192 vs. 10. Because the first octet is different, we don't really need to even use the netmask. But if we go through the process, we find that the netmask is still 255.255.255.0, which means that the first three octets are the network portion, and the last octet is the host portion. Since the first three octets is different, the two computers are on different network segments.

Example 5:

IP address 1: 12.168.1.10

IP address 2: 12.230.76.20

Netmask: 255.0.0.0

In this example, the netmask is 255.0.0.0, which means that the first octet is the network portion, and the last three octets are the host portion. When we compare the first octets, we find that they're both set to 12, which means the two computers are on the same network segment.

Example 6:

IP address 1: 211.68.19.10

IP address 2: 211.68.19.156

Netmask: 255.255.255.0

In this example, because the netmask is 255.255.255.0 the first three octets describe the network. Both IP addresses have the same numbers in the first three octets, 211.68.19. This means that both IP addresses belong to the same network segment.

Example 7:

IP address 1: 10.0.0.1

IP address 2: 10.0.1.1

Subnet mask: 255.255.0.0

In this example, the subnet mask is 255.255.0.0, which means that the first two octets of the IP address represent the network portion. When we compare the first two octets of the two IP addresses, we see that they are both 10.0, which means that they are on the same network segment.

Classless Inter-Domain Routing (CIDR)⁸

Another way to specify the netmask uses something called Classless Inter-Domain Routing or CIDR notation. This notation combines an IP address and the netmask in a single string, with the IP address listed first, followed by a slash, then a single number which specifies number of network bits. For example, consider 22.56.79.115/24. In this case the /24 on the end specifies that the first 24 bits of the address describe the network. Since 24 bits is equivalent to 3 bytes, this it would be equivalent to a netmask of 255.255.255.0.

Here's the basic subnet masks and their equivalent CIDR notations for an IP address 13.45.234.17.

	Netmask	CIDR	Network Address
13.45.234.17	255.0.0.0	13.45.234.17/8	13.45.234.0
13.45.234.17	255.255.0.0	13.45.234.17/16	13.45.0.0
13.45.234.17	255.255.255.0	13.45.234.17/24	13.0.0.0

The nice thing about the CIDR notation is that it conveys the same information as the IP address and a separate subnet mask in a much more compact form. That is, it's much easier to write "13.45.234.17/24" than it is to write "13.45.234.17 and 255.255.255.0"

Here are a few examples of comparing IP addresses using CIDR notation:

Example 1:

IP address 1: 199.16.1.100/24

IP address 2: 199.16.2.100/24

⁸ https://en.wikipedia.org/wiki/Classless_Inter-Domain_Routing

In this example, the /24 means that the subnet mask is 255.255.255.0, which means that the first three octets of each IP address represent the network portion. When we compare the first three octets of the two IP addresses, 199.16.1 and 199.16.2, we see that they are different. This means that the two IP addresses are on different network segments.

Example 2:

IP address 1: 212.57.1.100/16

IP address 2: 212.57.19.200/16

In this example, the /16 means that the subnet mask is 255.255.0.0, which means that the first two octets (24 bits) of the IP address represent the network portion. When we compare the first two octets of the two IP addresses, we see that they are the same 212.57, which means that the two computers are on the same network segment.

Example 3:


IP address 1: 10.20.30.142/8

IP address 2: 10.168.231.11/8

In this example, /8 means the subnet mask is 255.0.0.0, which means that the first octet of each IP address represents the network portion. In both IP addresses the first octet is 10, which means that they are on the same network segment.

Complete Delivery Demonstration

Here's a complete walkthrough of the steps involved in delivering network packets from end-to-end, as the packets go through a router. The demonstration will use the IP addresses and MAC addresses shown in the following figure:



Let's assume that Computer A wants to send a message to Computer F, and the netmask is 255.255.255.0. Here are the steps involved in making this network transmission:

1. Computer A builds an IP packet to hold the data it wants to transmit. This is addressed using Computer A's IP address 192.168.2.11 as the source IP, and Computer F's IP address 192.168.3.31 as the destination IP.
2. Computer A's Network Layer determines whether to ask its Data Link Layer to send the ethernet frame directly to Computer F, or if the ethernet frame should go to the Router. It does this by using the netmask to find the network portion of 192.168.2.11 and 192.168.3.31. Since the netmask is 255.255.255.0 it uses the first three octets and compares 192.168.2 with 192.168.3. Since the network numbers are different, it asks the Data Link Layer to send the ethernet frame to the Router.
3. The Data Link Layer on Computer A builds an ethernet frame to send to the Router. It uses its own MAC address of 64:4A:99:12:C5:7D as the source MAC. If Computer A doesn't know the Router's MAC address it uses ARP to find it. Once it has the Router's MAC it adds 21:43:AB:1F:35:22 to the ethernet frame as the destination MAC.
4. The Data Link Layer hands the ethernet frame to the Physical Layer for transmission on the network media.
5. The Physical Layer on the Router sees the transmission and hands it up to its Data Link Layer.
6. The Data Link Layer on the Router checks the ethernet frame and sees that the destination MAC address matches its destination MAC. Since ethernet is addressed to the Router, it unpacks the IP packet data and hands it to the Router's Network Layer.
7. The Network Layer on the Router checks the destination IP address and sees that it doesn't match its own IP address. The router knows that this means it is supposed to route the packet to another network. Using its routing tables, the Router determines that the packet should be sent out across its other network interface. The Router then builds another ethernet frame, adding the IP packet as data. The Router sets the source MAC as the MAC address of its second NIC, which is 21:43:AB:1F:42:29. If the Router doesn't know Computer F's MAC address it uses ARP to find it. Once it has Computer F's MAC address 64:4A:99:12:90:01 it puts it in the ethernet frame as the destination MAC address.
8. The Router hands the ethernet frame to its Physical Layer, which places the transmission on the network media.
9. Computer F's NIC sees the transmission, and hands the ethernet frame to its Data Link Layer, which checks the destination MAC address and sees that it matches its own MAC address. Since it's a match, the Data Link Layer unpacks the IP packet from the ethernet frame and hands it up to the Network Layer.

10. The Network Layer on Computer F checks the destination IP address and sees that it matches its IP address, so it knows it should further process the network transmission.

This demonstration shows how everything you've learned up to this point is used to move network data from end-to-end, by making more than one point-to-point transmission.

Default Gateway (Router) Basics

The next component to learn about is the default gateway or default router. You'll first learn what routers do, and then learn about the function of the default gateway and how it's used by the Network Layer.

Let's start with learning what routers do in a network. The first concept to understand about routers is that they are used to connect different networks. That is, if you only have one network segment, you can connect the devices using a hub or a switch. These devices will pass network traffic between the devices on the network segment, but they only know about MAC addresses and don't know how to pass network traffic to other networks or network segments. Routers on the other hand, are built specifically for passing network packets between networks. Routers look at the destination IP address and use this information to move the network packets from one network to another.

Every network connected to the Internet needs at least one router. Most home networks and small networks will only have one, while larger networks may have more than one for redundancy.

When you add a device to a network and configure the network settings, one of the mandatory items you need to configure is the IP address of the default router. Note that Microsoft decided to call this the default gateway, which is really a confusing mistake, since there are network devices called gateways that perform a completely different function. But, regardless of what it's called, the default gateway is the device that will take network packets from one network and deliver them to a different network.

Another thing to note is that you must configure the IP address of the default gateway, you can't put in a DNS name or the MAC address. The reason you can't use the DNS name is a short logic puzzle. If you could use the DNS name, then before sending any network packets to the router you would have to first send network packets to a DNS server to resolve the router's name to an IP address. But, sending the network packets for the DNS request requires sending those packets through the router itself, which you can't do if you don't know its IP address. But if you already have the router's IP address, you'll be able to use it to send packets to the DNS server.

The reason we don't configure the MAC address, which is the piece of information we really want, is because using the IP address makes the system much more flexible. If we use MAC addresses, then any time the router, or the NIC in the router is changed, we will have to change the network settings of every device on the network. But, if we use the IP address, we can swap out the router or the router's NIC without making any changes to the network settings on any of the network devices. Note that this flexibility is only possible because of ARP. That is, if we swap out the router or the router's NIC, all devices on the network will need to use ARP to discover the new MAC.

The last thing to mention about the default router/gateway is that this might be a little confusing if your only exposure to networks has been your home network, because in a home network the router that

connects your home network to the Internet typically also works as a switch, and maybe even as a modem. Instead of having a separate box for the home network switch and home network router, most wireless routers used in home networks act as a switch and a router. This will act as a switch, connecting the devices in your home network so you can do things like access your home printer or media server, but it will also act your network's router, moving network packets between the devices on your home network and the Internet.

You'll learn more details about routers later in the class, but for now make sure that you know the basic function of the default gateway, and why any computer connected to a network must know the IP address of its default gateway.

DNS Basics

Another important component used by the Network Layer is DNS (Domain Name System). DNS is a protocol combined with a distributed system of servers that are used to translate human-friendly domain names such as `google.com` or `whitehouse.gov` into IP addresses. DNS lets us humans use easy to remember names instead of trying to remember IP addresses when we want to send email or browse the web. This is just like remembering someone's name as opposed to remembering their phone number. I don't know about you, but I'd much rather ask my phone to call Saul than trying to remember his phone number is 505-503-4455. Since the Network Layer must have IP addresses to build and send IP packets the DNS system is essential to the functioning of the Internet.

The Domain Name System (DNS) has its roots in the early days of the internet. In the 1970s, the ARPANET (a precursor to the internet) used a central file to map human-readable hostnames to numerical IP addresses. The file was named `HOSTS.TXT` and it was maintained by the Stanford Research Institute (SRI) and had to be manually updated and distributed to all the computers on ARPANET. Whenever a new system was added, the computer owner would call SRI and ask them to add the new computer and its IP address to the file. This system worked fine in the early days when there weren't many computers on the network, but as more and more systems joined ARPANET using a single centralized file became more and more difficult, and soon it became obvious that this process would not scale easily.

In the early 1980s, the need for a more scalable and automated system for resolving domain names to IP addresses led to the development of the DNS. John Postel oversaw the ARPANET list, and he gave the task of selecting a new system to Paul Mockapetris from 5 candidate systems. Mockapetris ended up developing his own system, creating the basic architecture, and naming conventions that are still used today.

The first actual implementation of Mockapetris' DNS was created in 1984 by four students at the University of California, Berkeley, who wrote the Berkeley Internet Name Domain (BIND) software. BIND quickly became the de facto standard for DNS software and became formalized in 1987 with the publication of RFC 1035, which defined the structure and operation of the DNS system. In the 1990s BIND and DNS were adopted by Microsoft. Since then, DNS has undergone numerous changes and improvements to address issues such as security and scalability. Today, the DNS system is an essential part of the Internet infrastructure, consisting of a highly distributed network of servers and caches that work together to efficiently resolve domain names to IP addresses, enabling users around the world to access online resources quickly and easily.

You've probably been using DNS names your entire life, but just in case you don't know the basics of their format here's a little background information. Each domain name consists of at least two strings separated by dots, with the rightmost string indicating the top-level domain or TLD.


The top-level domain names used to categorize domain names based on their purpose or geographic location. There are two main categories of TLDs, generic top-level domains (gTLDs) and country-code top-level domains (ccTLDs). The first TLDs were created in 1984 and included seven gTLDs: .com, .edu, .gov, .mil, .net, .org, and .arpa. Generic top-level domains (gTLDs) are TLDs that are not restricted to a particular country or region, and are intended for use by individuals, organizations, and businesses worldwide. In contrast, country-code top-level domains (ccTLDs) are two-letter codes that are assigned to specific countries or territories, such as .us for the United States, .uk for the United Kingdom, and .de for Germany.

The distinction between gTLDs and ccTLDs is important because it affects how domain names are registered and used. While gTLDs are available for registration to anyone, ccTLDs may have specific registration requirements or restrictions based on the country or region they represent.

The TLDs were managed by the Internet Assigned Numbers Authority (IANA) until 1998, when they were transferred to the Internet Corporation for Assigned Names and Numbers (ICANN). ICANN is a non-profit organization that is responsible for managing and coordinating the DNS system, including the allocation of TLDs. Today, there are hundreds of TLDs, including many new gTLDs that were added as part of a major expansion of the DNS system in 2013. The expansion of TLDs was intended to provide greater choice and flexibility for domain name registration, and to support the growth of the Internet in new regions and industries. The expansion included the addition of hundreds of new gTLDs, such as .app, .blog, .club, and .xyz, as well as many new ccTLDs.

The actual DNS system is composed of several components. On the client side, or the side asking to resolve a DNS name to an IP address, there's the client and something we call a resolver. On the server side there are three main components, DNS root servers, Top Level Domain (TLD) servers, and DNS servers. In addition, there's the DNS protocol that defines the format for the DNS requests and responses⁹.

To help you understand how all the DNS components work together let's look at the process that's followed to resolve a name to an IP address. During this explanation we'll use the following diagram which shows the steps involved in the DNS process.



⁹ <https://www.catchpoint.com/blog/how-dns-works>

1. The DNS Client is any device or software that needs to translate a domain name into an IP address. For example, if you want to view a web page at msdn.microsoft.com, the network stack on your computer will need to know the IP address of msdn.microsoft.com. The first step in this process is the Network Layer on your computer will build a DNS Request using the DNS protocol, and place this inside an IP packet. This IP packet will be sent to the DNS Resolver for your network. The IP address of the DNS Resolver is one of the items that must be configured on every computer. You should note that when you look at the network settings on most computers, the DNS Resolver will be called the DNS Server. This choice of names is unfortunate because we use the name DNS Server in another place to describe something that provides a completely different service, so I'm going to stick with the name DNS Resolver.

In Step 1 of the diagram, the DNS Client sends the DNS Request to the DNS Resolver.

2. The DNS Resolver is the local system that takes responsibility for returning an IP address to the client. By local we mean that this computer is typically physically located near or in the client's network. The DNS Resolver handles making DNS requests for an entire network, and when it gets a response, it places the DNS name and IP address in its cache. So, when it receives a new DNS request from a client it may already have a copy of the requested DNS name and IP address. If it does find a copy in its cache, it returns the IP address to the client, which would be Step 8 in the diagram.

If the DNS Resolver doesn't have the requested data in its cache, it builds an IP packet with a DNS Request and sends it to one of the DNS Root Servers. Each DNS Resolver has a built-in list of IP addresses for the DNS Root Servers, so it knows the IP address to use as the destination address in the IP packet. The following figure shows a list of DNS Root Servers.

HOSTNAME	IP ADDRESSES	OPERATOR
a.root-servers.net	198.41.0.4, 2001:503:ba3e::2:30	Verisign, Inc.
b.root-servers.net	199.9.14.201, 2001:500:200::b	University of Southern California, Information Sciences Institute
c.root-servers.net	192.33.4.12, 2001:500:2::c	Cogent Communications
d.root-servers.net	199.7.91.13, 2001:500:2d::d	University of Maryland
e.root-servers.net	192.203.230.10, 2001:500:a8::e	NASA (Ames Research Center)
f.root-servers.net	192.5.5.241, 2001:500:2f::f	Internet Systems Consortium, Inc.
g.root-servers.net	192.112.36.4, 2001:500:12::d0d	US Department of Defense (NIC)
h.root-servers.net	198.97.190.53, 2001:500:1::53	US Army (Research Lab)
i.root-servers.net	192.36.148.17, 2001:7fe::53	Netnod
j.root-servers.net	192.58.128.30, 2001:503:c27::2:30	Verisign, Inc.
k.root-servers.net	193.0.14.129, 2001:7fd::1	RIPE NCC
l.root-servers.net	199.7.83.42, 2001:500:9f::42	ICANN
m.root-servers.net	202.12.27.33, 2001:dc3::35	WIDE Project

List of DNS Root Servers from IANA.org.

It may look like there are only 13 Root Servers but there are actually dozens of Root Servers at each IP address. If there were only 13, they would easily be overwhelmed as they have to respond to hundreds of thousands of requests each minute. Each location has multiple servers

and uses something called Round Robin DNS to balance the load between the servers. You'll learn about Round Robin DNS later, when you take a dive into the details of DNS.

The important thing to take away is that every DNS Resolver has this list of Root Servers and will build and send the packet shown as Step 2 in the diagram.

3. At this point the request the DNS server ecosystem starts a process to determine which DNS server will ultimately handle the request. Remember there isn't one central server to handle all requests, each network has its own server(s), so somehow the request needs to be sent to the correct server. In the DNS ecosystem there are DNS Root servers, Top Level Domain Servers, and DNS Servers, which have the actual DNS name and IP address tables for each network. The initial point of contact for any request will be one of the DNS Root Servers. It doesn't matter what the DNS name in the request is, all requests are initially sent to a Root Server.

Each DNS Root Server has a list of the IP addresses for the Top Level Domain (TLD) Servers, and it will return the IP address of the appropriate TLD Server to the DNS Resolver. The Root Server determines the appropriate TLD server by looking at the last portion of any DNS name which will be .com, .blog, etc. or a country code. In our example, the client is looking for the IP address for msdn.microsoft.com, so the Root Server will return the IP address for the .com TLD server, which we'll say is: 192.5.6.30.

This will be sent back to the DNS Resolver in the network packet shown as Step 3 in the diagram.

4. The DNS Resolver will receive this packet with the IP address for the appropriate TLD Server, and now know the IP address of the TLD server for any DNS names ending in .com. The DNS Resolver will build another network packet using the TLD's IP address as the destination IP. This is Step 4 in the diagram.
5. The TLD Servers won't know the IP address for every domain name, but they can provide the information we need to make the next step in finding the IP address we want. Each TLD Server has a list of all the next level domains in its domain, along with their IP address. For example, the .com TLD servers will have lists of every DNS name that ends in .com, while the .edu TLD servers will have a list of every DNS name that ends in .edu.

In our example we're looking for msdn.microsoft.com, so the next portion of the DNS name is microsoft. When our example request gets to the TLD Server, the TLD server will find the IP address associated with microsoft. For our example let's assume that this IP address is: 20.81.111.85. The TLD server will build another response to send back to the DNS Resolver on our network, telling it that Microsoft's DNS Server's IP address is 20.81.111.85. This is shown as packet 5 in the diagram.

6. The DNS Resolver will receive this packet, and now knows the IP address of the DNS Server that holds the information it really wants. In the DNS ecosystem these servers are called Authoritative servers because they have the authority to hand out IP addresses for any DNS

names in their zone of authority, as opposed to the Root and TLD servers, which play a different role in the process. The DNS Resolver will build another network packet to send to this last authoritative DNS Server. In the case of our example the DNS Resolver will build a network packet to send to Microsoft's DNS Server's at IP address 20.81.111.85. This is shown as step 6 in the diagram.

7. Each authoritative DNS Server has a list of DNS names and their corresponding IP addresses for the DNS names in its zone of authority. The zone of authority typically corresponds to a DNS name such microsoft.com. That is, any DNS name such as msdn.microsoft.com or xbox.microsoft.com would all fall into the microsoft.com zone of authority and be served by Microsoft's DNS server, while docs.google.com and sites.google.com will all fall under the google.com zone of authority and be served by Google's DNS server. However, sometimes DNS servers will cover multiple DNS names in their zone of authority. For example, if you register a DNS name for a web site you use with a web hosting service such as sites.google.com or GoDaddy, you won't run your own DNS server. Instead, you'll use one of Google's DNS Servers or GoDaddy's DNS Servers, and they will cover your respective zone of authority.

In the case of our example, the Microsoft DNS Server will look up the IP address for msdn.microsoft.com and build a network packet to send it back to the DNS Resolver. For example, let's say the IP address for msdn.microsoft.com is 13.107.238.70. This will be placed in a DNS response packet and sent back to the DNS Resolver.

8. The DNS Resolver will receive the DNS response, put the IP address for msdn.microsoft.com in its cache, then forward the DNS response with msdn.microsoft.com's IP address back to the client. The DNS Resolver caches the IP addresses to speed up the process in case another client on its network makes a DNS request for the same DNS name. The amount of time an IP address remains in cache is set by the authoritative DNS server using a setting in the DNS response called the Time To Live or TTL. Typically, the TTL is set to ~3600 seconds, or 1 hour, but this can vary as each DNS administrator can change it.
9. The client will receive the DNS response from the DNS Resolver, and now have the IP address it needs to build the network packet it was originally working on.

This system of using multiple levels of DNS servers might seem a little complicated, but it was adopted for several important reasons including scalability, redundancy, localized control and access, and security.

Scalability - A centralized server model would be impractical for the size and complexity of the DNS system. With millions of domain names and billions of requests every day, a single server would not be able to handle the volume of traffic and would quickly become overwhelmed. By using a distributed system of servers, the load can be spread across multiple servers, making the system more scalable and resilient.

Redundancy- A distributed system of servers provides redundancy and fault tolerance. If a single server fails or is taken offline, the other servers in the system can still handle requests and provide the necessary information, and there are redundant servers at every level of the DNS eco system. This ensures that the DNS system remains available and functional even in the face of hardware failures or other disruptions.

Local Control - Because each network controls its own DNS authoritative server, they can change or add to the DNS names and addresses within their own domains at any time. If a centralized system were used instead, adding or changing names and IP addresses would require filling out forms and waiting for someone else to make the changes.

Localized Access – The DNS Resolvers are typically located close to the clients they serve. This means DNS queries can be resolved more quickly and efficiently, especially if the DNS name and IP address are already cached, reducing latency, and improving performance.

Security - A distributed system of servers can be more secure than a centralized server model. If a single server is compromised, it could potentially compromise the entire system. By distributing the data and workload across multiple servers, the impact of a security breach can be minimized. Of course, the flip side of this is that each DNS Resolver and authoritative server needs to be secured by someone with the experience to ensure the servers are properly secured.

Another great thing about the DNS system is that even though resolving a name to an IP address requires several network transactions, the entire process happens very quickly. In fact, in most cases the amount of time it takes to resolve a DNS to an IP address is measured in milli seconds, or thousandths of a second.

Configuring Network Settings on a Windows Computer

Now that you've learned the basics about IP addresses, netmasks, the default gateway/router, and DNS, and how they work together at the Network Layer, let's look at configuring and checking these items on a Windows based computer.

There are three ways to configure the network settings on a Windows based computer. You can do the configuration manually, use a network service called Dynamic Hardware Configuration Protocol (DHCP) to do it automatically, or configure some of the settings using Automatic Private IP Addressing (APIPA).

Manual Configuration

Manually configuring the network settings on a computer or phone requires knowing a couple of things. The first is the network settings themselves. You'll need to know the IP address, the netmask, the IP address of the default gateway, and the IP address for the DNS resolver(s), of which there are typically two, a primary and a secondary which will act as the backup in case the primary is not available.

The second thing you'll need to know is where to go to get the values for the network settings. Remember you can't just make these up, you'll need to get them from someone who got them from IANA. If you're configuring the settings manually, then you're probably going to be working as a network administrator, where someone above you in the network food chain will provide you with the information you need. This might be a head network administrator, or if you're the head net admin

you'll get your IP addresses from your network's ISP. You may choose to manually configure the network settings on devices on your home network, but you'll need to have some advanced skills to do this as you'll also need to reserve any IP addresses you use from being used by your DHCP server, which is typically your home router.

Once you have the network settings, the last thing you'll need to know is how to configure them on your device. You'll learn how to do this for Windows based computers below, but the process will be different in every OS and on every device. If you need to configure the settings on a different OS or device, you can undoubtedly find help on the Internet.

DHCP

Dynamic Host Configuration Protocol (DHCP) is a network protocol and process that automatically assigns IP addresses, netmasks, and the IP addresses of the default gateway and DNS servers to devices on a network. Using DHCP to automatically configure the network settings makes it much easier to manage a large network.

DHCP works by having a DHCP server on the network that manages a pool of available IP addresses. When a new device connects to the network, it sends a broadcast message looking for a DHCP server and asking for an IP address. The DHCP server receives this message and responds with an offer of an available IP address. The device then accepts the offer and the DHCP server assigns the IP address to the device for a set amount of time, known as the lease time.

DHCP is particularly useful in large networks because it reduces the amount of manual configuration that needs to be done. It's not difficult to configure the network settings on any single device, but on a large network with hundreds or thousands of devices, configuring and managing the network settings on all the devices would take a lot of time.

DHCP also makes connecting to home networks or networks outside the home much easier. Think about your home network and the devices you connect. Typically, the only thing you, or anyone that connects to your home network, needs to do is connect to your wireless router, you don't need to set or change the IP address, netmask, default gateway, and DNS server. This is because your computer is set to DHCP, and when you connect the computer to your home network your router acts as a DHCP server and gives your computer the IP address to use, along with the correct netmask, default gateway IP address, and IP addresses for the primary and secondary DNS servers. The same thing happens when you travel with your computer and connect to a network at a business, motel, or airport. That is, you don't have to configure any of the network settings, instead your computer sends out a DHCP broadcast when you initially connect to the network and gets the network configuration settings from a DHCP server.

Another benefit of DHCP is that it allows for efficient use of available IP addresses, as devices are only assigned an IP address when they need it, and the address can be released when it is no longer needed.

You'll learn more about DHCP later in the class when you learn how to set up and configure a DHCP server.

APIPA

Automatic Private IP Addressing (APIPA) is a feature in Microsoft Windows operating systems that provides a fallback mechanism for devices when they cannot obtain an IP address from a DHCP server.

APIPA is designed to allow devices on a local network to communicate with each other even when there is no DHCP server available. When a device is unable to obtain an IP address from a DHCP server, it automatically assigns itself an IP address in the range of 169.254.0.1 to 169.254.255.254. This address is known as an APIPA address. Devices with APIPA addresses can communicate with each other on the local network, but they can't communicate with devices on other networks or access the Internet. APIPA does not provide a default gateway or DNS server, so devices with APIPA addresses cannot access resources outside of their local network.

APIPA isn't meant to be a replacement for DHCP, it's just meant to be used to connect to a network "in case of emergency", which means in situations where a DHCP server is not available. It will be enough to get a computer connected to the local network segment, but even then, the computer may not be able to communicate with other network devices such as file servers or printers, if the other devices don't have IP addresses that use the same network number. And this will probably be the case if your network has a DHCP server that unexpectedly becomes unavailable. That is, devices that connected to the network using settings from the DHCP server will use one set of network numbers, while devices that try to connect to the network when the

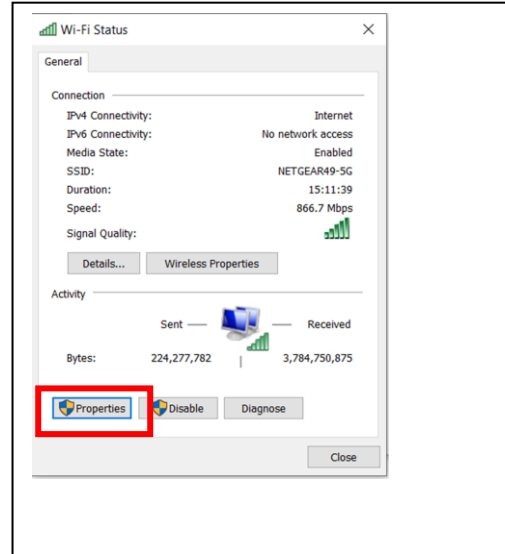
In general, APIPA should only be used as a last resort in emergencies when a DHCP server is not available, and it should not be relied upon as a long-term solution for network addressing.

Manual Configuration and DHCP Configuration in Windows

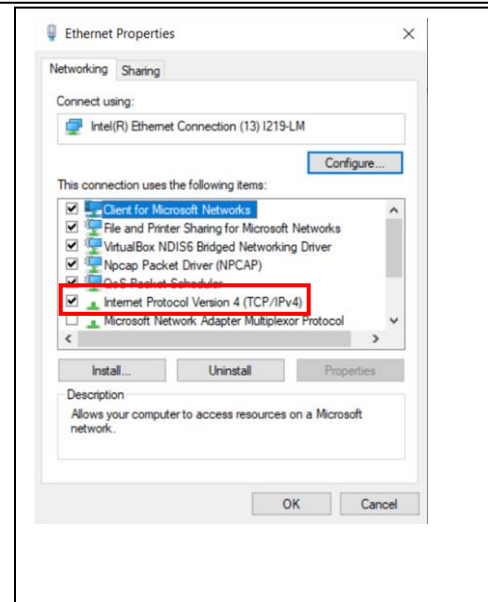
Here's the process for manually configuring the network settings or selecting DHCP on a computer running Microsoft Windows:

1. Open the **Control Panel** from the **Start** menu.
2. Click on **Network and Sharing Center**. Note that there are several ways to get the Network and Sharing Center, but these steps are the quickest.
3. In the left-hand menu, click on **Change adapter settings**.

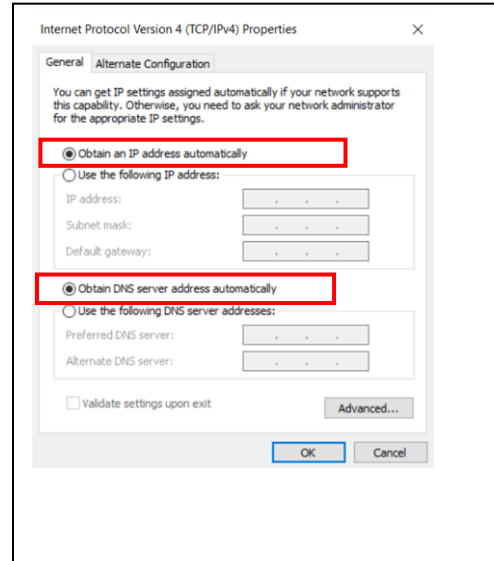
4. If you're checking a wired network connection double-click the network adapter you want to configure, or right-click on the network adapter and select **Properties**. If you're checking a wireless connection, double-click the wireless adapter, then select the **Properties** button at the bottom of the dialog box.



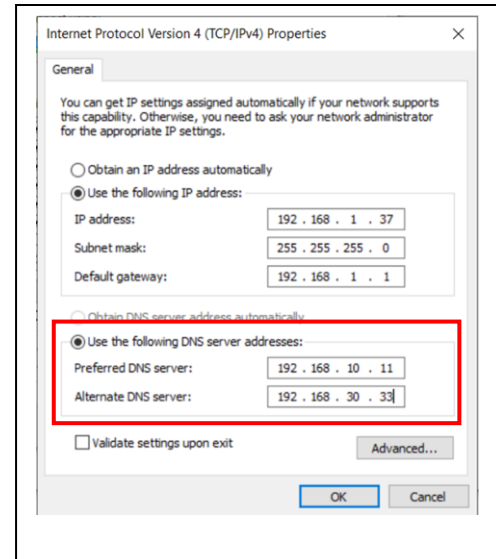
5. Find **Internet Protocol Version 4 (TCP/IPv4)** in the list of items in the middle of the dialog box, then select it and click **Properties** or double-click it.



6. In the new dialog box, ensure that the **General** tab is selected. This is where you will set the IP address, Subnet mask, Default gateway, and DNS server values. You can see that there are two sections in the dialog box for entering these values, one for the IP address information and one for the DNS server addresses. If these are set to **Obtain an IP address automatically** and **Obtain DNS server address automatically** it means the computer is set to use DHCP. Note that if you select DHCP, you won't see any values for the IP address settings of the DNS servers. If the computer uses DHCP and it's connected to the network, it will have appropriate values for these settings, but they will remain blank in this dialog box.



7. To manually set the network values, select **Use the following IP address** and **Use the following DNS server addresses**. Note that for DNS, Microsoft uses the terms **Preferred DNS server** and **Alternate DNS server** for what we called DNS Resolvers in the explanation of the overall DNS system.
8. Once the network settings are configured, click **OK** to exit the Properties window.



Basic Network Troubleshooting

Next, let's look at troubleshooting possible problems with the IP configuration. As you've seen there are a lot of settings that could be wrong, and several components such as DHCP, the default gateway, or DNS that could be having issues, so having a good troubleshooting process is important. In this section you'll learn about a few commands and tools used for troubleshooting, and a special IP address called the localhost or loopback address that's also used for troubleshooting the network stack. We'll start by going through a troubleshooting process that checks every possible problem, and then finish by looking at what steps of the process would be used to check specific problems.

In our general troubleshooting process, we'll do a set of checks that starts with checking things on the local computer, then work our way out to check network components like DNS or the default gateway. Hopefully you can see the sense in checking to make sure the IP settings are correct on a computer before we start checking DNS servers. It's kind of like if we were having trouble sending snail mail. It would be smarter to make sure we put a stamp on the envelope first, before we start checking for problems with mail sorting machines in Spokane and Seattle.

Here are the steps I suggest you use for troubleshooting network connections:

1. **Check the physical connections.** Ensure that all cables are securely connected and that the network interface is enabled. Look for LED lights on the network interface and if you have access to it, check for lights on the switch or router to verify that end of the cable is also connected and working.

2. **Check the IP settings.** Verify that the computer has a valid IP address, subnet mask, default gateway, and DNS server address. This is done by using the `ipconfig /all` command on Windows.

```
Ethernet adapter Ethernet:

    Connection-specific DNS Suffix . : 
    Description . . . . . : Intel(R) Wi-Fi 6 AX201 160MHz
    Physical Address. . . . . : 84-5C-F3-CA-82-B7
    DHCP Enabled. . . . . : Yes
    Autoconfiguration Enabled . . . . : Yes
    Link-local IPv6 Address . . . . . : fe80::f631:51fb:eb16:9e62%4(Preferred)
    IPv4 Address. . . . . : 192.168.1.6(Preferred)
    Subnet Mask . . . . . : 255.255.255.0
    Lease Obtained. . . . . : Wednesday, May 3, 2023 3:56:44 PM
    Lease Expires . . . . . : Friday, May 5, 2023 6:32:53 AM
    Default Gateway . . . . . : 192.168.1.1
    DHCP Server . . . . . : 192.168.1.1
    DHCPv6 IAID . . . . . : 59006195
    DHCPv6 Client DUID. . . . . : 00-01-00-01-28-2E-19-CC-34-73-5A-DE-7A-3B
    DNS Servers . . . . . : 192.168.1.1
```

3. **Use ping to check the network stack.** This test will check to see if the network stack is functioning by asking the computer to talk to itself through the computer’s network interface. If this test fails, it’s a sign that some part of the computer’s network software is not configured right or not working as expected. If the test is successful, it means that the network stack is functioning correctly.

The ping tool is a network utility used to test the connectivity and reachability of a host on an Internet Protocol (IP) network. When you start ping, you give it the name or IP address, and then ping sends a series of small packets of data, called Internet Control Message Protocol (ICMP) echo requests, to the target host. Ping waits for a response for each request, which should be an ICMP echo reply. The ICMP echo requests don’t have any data, there just used to “knock on the door” of network devices and see if anyone answers or not.

When you run the ping command, the program creates a series of ICMP packets and sends them to the target IP address. The ping tool measures the round-trip time for each packet to travel from the source host to the target host and back again. It also displays information about the number of packets sent, received, and lost, as well as the minimum, maximum, and average round-trip time (RTT) for each packet.

The ping tool is commonly used for testing and troubleshooting network connectivity issues, such as identifying network latency, packet loss, or connectivity problems, but in this case, we're having a computer ping itself, just to test to see if the network stack is operating. This is accomplished by telling ping to target a special IP address, 127.0.0.1, which all network software use to refer to the computer the software is running on. We also use the names localhost or loopback to refer to 127.0.0.1.

When packets are sent to the loopback address, they are processed by the network stack in the same way as packets that are sent to any other IP address. The network stack is the set of protocols and services that are responsible for sending and receiving network data on a computer.

When an application like ping sends data to the loopback address (127.0.0.1), the data is first processed by the upper layers of the network stack, then passed down to the network layer, where it is encapsulated in an IP packet and the source and destination IP addresses are set to 127.0.0.1. The IP packet is then passed down to the data link layer, where it is encapsulated in an ethernet frame where the source and destination MAC addresses are set to the computer's MAC address. The ethernet frame is then sent out on the loopback interface, which is a virtual network interface that represents the loopback address. The packet is not actually sent out on any physical network but is instead looped back to the computer's own network stack.

When the packet is received by the loopback interface, it is processed by the network stack in the reverse order that it was sent. The data link layer unpacks the ethernet frame and hands the IP packet to the network layer. The Network layer unpacks the IP packet and passes the data back up to the upper layers in the network stack.

Sending packets to 127.0.0.1 or the loopback interface allows applications on a computer to communicate with itself, testing the network stack as if the packets were being sent and received from the network.

<pre>C:\WINDOWS\system32>ping 127.0.0.1 Pinging 127.0.0.1 with 32 bytes of data: Reply from 127.0.0.1: bytes=32 time<1ms TTL=128 Reply from 127.0.0.1: bytes=32 time<1ms TTL=128 Reply from 127.0.0.1: bytes=32 time<1ms TTL=128 Reply from 127.0.0.1: bytes=32 time<1ms TTL=128 Ping statistics for 127.0.0.1: Packets: Sent = 4, Received = 4, Lost = 0 (0% loss), Approximate round trip times in milli-seconds: Minimum = 0ms, Maximum = 0ms, Average = 0ms</pre>	<pre>C:\WINDOWS\system32>ping 127.0.0.1 Pinging 127.0.0.1 with 32 bytes of data: Request timed out. Request timed out. Request timed out. Request timed out. Ping statistics for 127.0.0.1: Packets: Sent = 4, Received = 0, Lost = 4 (100% loss),</pre>
---	--

Pinging 127.0.0.1. A successful ping is shown on the left, and an unsuccessful ping is shown on the right.

4. **Check local network connections and default gateway.** If the computer can ping itself, you can be assured that the network stack is functioning. The next thing to do is to try and connect to other devices on the same network segment. The network connection we typically test next is the one to the default gateway. Testing this connection makes sense since we'll need the default gateway in later tests, so we might as well test it now.

The first step in performing this test is to find the IP address of the computer's default gateway, which can be found using the `ipconfig /all` command, which you did in step 2. Next, run ping using the default gateway's IP address as the target for ping. For example, if the default gateway's IP address is 192.168.1.1 you would run the command:

```
ping 192.168.1.1
```

As in the previous step, if this is successful the default gateway will send back replies, but if it isn't successful ping will display a request timed out message. If this test succeeds it means your computer can make network connections with other devices on the local network segment and reach the default gateway. If this test fails, it means that there's an issue making connections to the local network segment, or there are problems with the default gateway. To narrow this down you can:

- A. Check other devices on the local network segment and see if you can ping those devices from your computer. If you can ping another device, it's a good sign that you either have the wrong IP address for the default gateway, or there's an issue with the default gateway.

- B. Log in to another device on the same network segment and ping the default gateway. If you can ping the default gateway from another device, it's an indicator that there's a hardware issue on your computer, the IP address of the default gateway is incorrect on your computer, the netmask on your computer is incorrect, or possibly the host firewall is blocking the network traffic.
5. **Check the DNS server network connection.** If the computer can successfully ping the default gateway, you can be assured that the computer can make connections to other devices on the same network segment, and probably make connections to computers on the Internet or computers on other networks. If you're still having problems connecting to computers on other networks the next thing to test is the connection to the DNS server. This will help eliminate the possibility that the problem is in resolving a DNS name to an IP address.

The first step in performing this test is to find the IP address of the computer's DNS server, which can be found using the `ipconfig /all` command, which you did in step 2. Let's assume the IP address for the DNS server is 192.168.24.21. Next, run ping using the DNS server's IP address as the target for ping. For example:

```
ping 192.168.24.21
```

If the DNS server replies to the Ping packets, it means the test is successful and you can proceed to the next step. If this test fails, it means that there's an issue making network connections to the DNS server, or there are problems with the DNS server. To narrow this down you can try connecting to the DNS server from another device on the local network segment. If you can ping the DNS server from a different computer, it's a good sign that you either have the wrong IP address for the DNS server, or there's an issue with the DNS server. If all the computers on the local network segment have issues connecting to the DNS server, you should get in touch with the organization that manages the server and make them aware of the issue.

6. **Ask the DNS server to Resolve a DNS Name.** If the computer can successfully ping the DNS server, the next thing to test is to ensure the DNS server returns an IP address if you send it a DNS name. On windows systems this is done using the `nslookup` utility, while on Linux systems and Macs the `dig` utility program is used. To run `nslookup` on Windows, start the Command Prompt application, then type `nslookup` followed by a DNS name. For example:

```
C:\WINDOWS\system32>nslookup google.com
Server: UnKnown
Address: 192.168.1.1
```

```
Non-authoritative answer:
Name: google.com
Addresses: 2607:f8b0:400a:806::200e
          142.251.211.238
```

If the nslookup query returns an error message saying that it can't find the DNS name you asked for, chances are high that you mistyped the name. If this happens you should try a common name that's easy to type, such as google.com.

```
C:\WINDOWS\system32>nslookup nosuchhostmadeupname.com
Server: UnKnown
Address: 192.168.1.1
```

```
*** UnKnown can't find nosuchhostmadeupname.com: Server failed
```

If the nslookup query times out, it's an indication that you have the wrong IP address for the DNS server. At this point you should double check the DNS server's IP address or try running nslookup from another computer. If you can't connect to the DNS server from any computers on the network segment you will have to get in touch with the system administrator for the DNS server.

```
C:\WINDOWS\system32>nslookup google.com
DNS request timed out.
  timeout was 2 seconds.
Server: UnKnown
Address: 192.168.1.14
```

```
DNS request timed out.
  timeout was 2 seconds.
DNS request timed out.
  timeout was 2 seconds.
DNS request timed out.
  timeout was 2 seconds.
DNS request timed out.
  timeout was 2 seconds.
*** Request to UnKnown timed-out
```

7. **Check the Entire Internet.** If you're able to get the DNS server to resolve DNS names, and still can't connect to the Internet or a specific computer on the Internet the next step is to check all the routers between your network and the Internet, or between your network and the specific computer. You'll learn the details about routers in a later section, but there are typically several routers between any home network and the Internet or between any local network segment and the Internet. Any of these routers could be having trouble, preventing network packets from being transmitted. To check this, we use a program called `tracert` which is short for traceroute.

The traceroute program allows you to see the path your network packets take between your computer and a specific network host, reporting back on each successive router that handles the packets. Each time a network packet is sent to another router it's called a hop. Traceroute

works by sending ICMP packets, like those used by ping, with ever increasing Time To Live (TTL) values. The TTL is a setting that's used to keep network packets from wandering around the network forever. This is done by setting a max TTL when a packet is first sent, then decreasing it by one each time it makes another hop to a new router. The first packet that traceroute sends out will have TTL=1. The first router that handles the packet, which will be the default router, will send an ICMP response back to traceroute, then decrement the TTL, which in this case will set it to 0. Since the TTL is 0, the router will not forward the packet to the next router, it will simply delete it. Note that traceroute actually sends 3 packets to each router, not one, and will display the round-trip times for each packet.

When the traceroute program receives the ICMP response, it will generate a second ICMP request but this time set the TTL=2. This packet will go to the first router, which will decrement the TTL, making it 1, then send the packet to the next router on the path to the final destination. This second router will send back an ICMP response and decrement the TTL. Since the TTL is 0, the second router will drop the network packet.

This process continues, with the traceroute program sending successive packets and receiving successive ICMP response packets from each router as the packets take one additional hop on each iteration.

```
C:\WINDOWS\system32>tracert google.com
```

```
Tracing route to google.com [142.251.211.238]  
over a maximum of 30 hops:
```

```
 1  1 ms  1 ms  1 ms  192.168.1.1  
 2  4 ms  3 ms  4 ms  fdr01.knwc.wa.nwestnet.net [50.46.181.118]  
 3  3 ms  44 ms  4 ms  cr1-knwcwaxa-a-be500.bb.as20055.net [64.52.98.0]  
 4 10 ms 10 ms 10 ms lr1-umtloraz-b-be-12.bb.as20055.net [198.179.53.18]  
 5 10 ms  9 ms  8 ms lr1-yakmwafp-a-be-17.bb.as20055.net [204.11.64.133]  
 6  *    *    *    Request timed out.  
 7 13 ms 11 ms 32 ms google-sttlwawb-a.pni.as20055.net [107.191.239.0]  
 8 14 ms 11 ms 11 ms 142.251.229.137  
 9 11 ms 11 ms  9 ms 216.239.43.121  
10 10 ms  9 ms  9 ms sea30s13-in-f14.1e100.net [142.251.211.238]
```

```
Trace complete.
```

The output from tracert shows the number of hops in the first column, the round-trip times for each of the three ICMP packets sent to each router, the DNS name of the router if it's available, and the IP address of the router. The DNS names for most routers will be hard to decipher because they're typically strings of characters the ISPs and backbone providers use to identify the router, and not something meant for normal user consumption. Sometimes you can tease the location of the router from the name, but not always. This might give you an idea of where

the router is physically located. For example, I'm guessing that **lr1-yakmwafp-a-be-17** is in Yakima Washington, and **google-sttlwawb-a** is in Seattle Washington, but these are just educated guesses.

Another thing to note about the tracert output is that many routers on the Internet are configured not to respond to ICMP packets. When tracert encounters one of these routers, it will display an asterisk (*) in place of the RTT value for that router. This can make it difficult to pinpoint the exact location of the problem.

If you're having problems connecting to the Internet and use tracert as a diagnostic tool, it can be difficult to pinpoint the problem if the problematic router lies outside your network. If the problem with your connection is caused by a router within your LAN, then tracert will be able to identify it and provide valuable information for troubleshooting. However, if the problem is caused by a router outside of your LAN, then the results from tracert may not be as helpful, especially if the router doesn't respond to the ICMP requests. So, while the bad news may be that you can't connect to the Internet because someone else's router is having problems, the good news is that it's someone else's router having problems, not your router. And, typically these types of router issues will be fixed relatively quickly as they will impact many different networks and users, and there will be lots of unhappy people notifying the router owner of the problem.

There are a few different versions of the tracert tool that you can run on Windows There's the command line version, or you can download and run a version with a graphical user interface (GUI), or you can find online versions of traceroute, although most online tools show the route back from their web site to a DNS name you provide as opposed to showing the route from your computer to their web site. The GUI based tools are nice because they may have extra features, like showing you where each router is located on a map.

As you troubleshoot network connections, here are a few other things you should keep in mind, and may have to check:

1. Check for any recent changes or updates that could have caused the issue.
2. Disable any firewalls or antivirus software temporarily to see if they are blocking the connection.
3. Check for any known issues or outages with the network or specific websites/services being accessed.

Network Layer and IP Basics – Summary

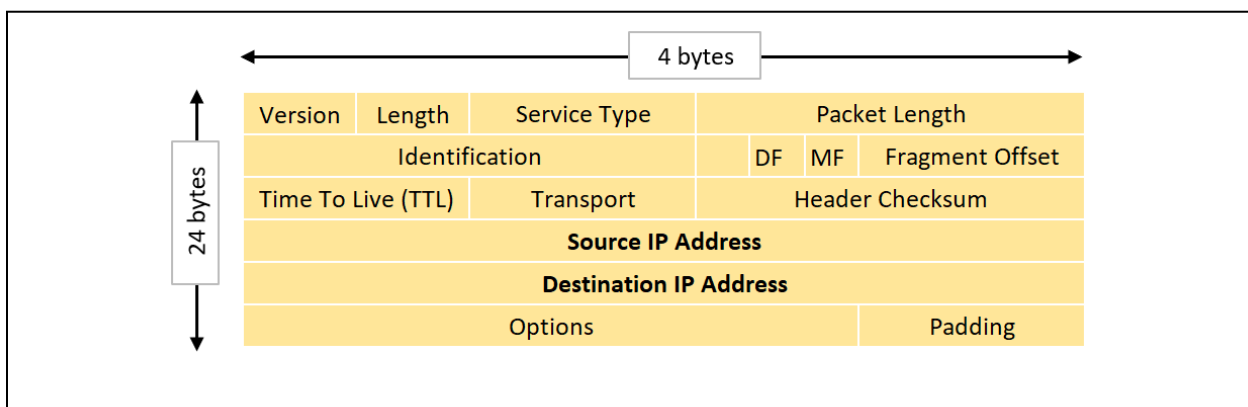
This wraps up our tour of the basics of the Network Layer of the OSI Network Model, and the Internet Protocol. During this initial look at these items, you learned the following:

1. The Network Layer is responsible for end-to-end delivery of network packets, as opposed to point-to-point delivery which is done by the Data Link Layer.
2. The main protocol used at the Network Layer is the Internet Protocol, which defines several things including the format of IP packets, and IP addresses and netmasks.

3. IANA is the group responsible for the overall management of IP addresses, but you will most likely receive an IP address from your ISP.
4. IP addresses are built from 4 octets, with each octet ranging from 0 to 255.
5. An IP address describes two things, a network, and a host number on that network. The netmask is used to determine how to divide and read an IP address.
6. Every computer or device that uses IP must be configured with the IP address of a default gateway/router, which will handle transmitting packets to other network segments.
7. The DNS system consists of many servers whose main role is to resolve DNS names to IP addresses. Every computer on an IP network must be configured with the IP address of a primary and secondary DNS server.
8. Configuring IP on a computer requires setting the IP address, netmask, default gateway/router IP address, and IP addresses of the primary and secondary DNS servers. This configuration can be done manually but is typically done with DHCP.
9. Using DHCP requires setting up a server that will manage and distribute the IP settings for devices on a network.

IP Packet Details

Now that you have a basic understanding of the components used by the Network Layer and the Internet Protocol, and how these components work together to build and deliver network packets, it's time to loop back around and take a deeper dive into the components. The first thing we'll look at is the IP packet format, and all the fields in the IP packet header. Note that it isn't critical that you know this information unless you end up working as a network administrator or plan on taking a certification exam.



Here are the fields in the header, along with a description of what the field contains:

Version (4 bits) - This field indicates the version of the IP protocol being used. This can be 4 or 6, but if it's 6 the rest of the header fields will be different than what's listed here.

Header Length (4 bits) - This field indicates the length of the IP packet header in 32-bit words. This field is necessary because the length of the IP packet header can vary because there are a few optional items that may or may not be included in fields at the end of the header. If the options field is not present, the header length will be fixed at 20 bytes. However, if options are included, the header length will increase in multiples of 4 bytes, up to a maximum of 60 bytes.

Type of Service (TOS) (8 bits) - This field was meant to be used to indicate the priority and type of service requested for the packet. The TOS field was later replaced by the Differentiated Services (DiffServ) field in IPv4, and the Traffic Class field in IPv6, and most routers ignore the TOS information, so its use is limited and it's not a reliable means of specifying QoS for IP packets in modern networks.

The TOS field is divided into two subfields: the Precedence subfield and the TOS bits subfield. The Precedence subfield is a 3-bit field that is used to specify the priority of the packet. The TOS bits subfield is a 5-bit field that is used to specify the type of service required by the packet, such as low delay, high throughput, or high reliability.

In practice, the TOS field is not widely used and is often ignored by routers. This is because routers primarily use the destination IP address and the protocol type (TCP, UDP, etc.) to determine how to handle a packet. Additionally, modern networks typically use QoS mechanisms such as DiffServ and traffic shaping to prioritize traffic, rather than relying on the TOS field. However, some legacy networks may still make use of the TOS field for QoS purposes.

Total Length (16 bits) - The Total Length field in an IPv4 header is a 16-bit field that specifies the total length of the IP packet in bytes, including both the header and data sections. The minimum value for the Total Length field is 20 bytes, which corresponds to a packet with no data, only the IP header. The maximum value for the Total Length field is 65,535 bytes, which is the maximum size of an IP packet.

Identification (16 bits) - This field is a number that's used to identify fragments of an original packet. When a packet is too large to be sent in a single transmission, it must be divided into smaller fragments that can be reassembled at the destination. The identification field is set to a unique value for each original packet, and all fragments of that packet will have the same value in this field, which tells the recipient they're all part of the same fragmented packet. The receiving device will use this information to reassemble the fragments, also using the Fragment Offset field (see below), to assemble them in the correct order.

The identification field has a minimum value of 0 and a maximum value of 65,535. When the maximum value is reached, the value wraps around to 0 and continues counting.

Flags (3 bits) - This field is used to indicate whether the packet can be fragmented or not by the routers as they forward the packet. Some routers and network links use a smaller Maximum Transmission Unit (MTU), so this field is used to tell them whether the packet can be fragmented, or if the router should just return an error if it can only transmit the packet by breaking it up.

The three bits in this field are:

Reserved: This is always set to zero.

Don't Fragment (DF): If this bit is set to 1, it indicates that the packet should not be fragmented. If a router receives a packet with the DF flag set and the packet needs to be fragmented to be forwarded, the router will drop the packet and send an ICMP "Destination Unreachable: Fragmentation Needed" message back to the sender. This message includes the MTU of the next hop, so the sender can reduce the size of the packet before sending it again.

More Fragments (MF): If this bit is set to 1, it indicates that there are more fragments to follow. If this bit is set to 0, it indicates that this is the last fragment.

Fragment Offset (13 bits) - This field is used along with the identification and Flag fields to determine the position of the current packet within the original packet and reassemble the fragments of the original packet. Here's an example of using the fragment offset field in an IPv4 header:

Suppose that a sender has a large data packet to send to a receiver, but the packet is too large to be sent in a single transmission unit. The sender can break up the packet into smaller fragments, each with its own IP header. The first fragment will have a "more fragments" flag set to 1, indicating that there are more fragments to follow. Each fragment, except the last one, will have a "fragment offset" field indicating the position of that fragment within the original packet. The offset is specified in units of 8 bytes, so if the first fragment is 1500 bytes long, the fragment offset of the second fragment would be $1500/8 = 187.5$, rounded down to 187.

When the receiver receives the fragments, it will use the identification field in the IP header to identify which fragments belong to the same original packet. The receiver will then use the fragment offset fields to reconstruct the original packet by concatenating the fragments in the correct order. The last packet sent will have the "more fragments" flag set to 0 to indicate that this is the last fragment.

Time to Live (TTL) (8 bits) - This field is used to limit the lifetime of a packet. This is necessary because it's possible for packets to get "lost" and wander in circles between routers. The number of hops a packet can make is limited by the TTL field, with each router subtracting one from the TTL as it moves the packet from one network to the next. The TTL field is typically set to 64 and decremented by one by each router that processes the packet, and the packet is discarded if the TTL field reaches zero.

Protocol (8 bits) - This field indicates the protocol used in the data portion of the IP packet. Common protocol values include TCP, UDP, ICMP, and IGMP. You'll learn more about TCP, UDP, and IGMP in later sections.

Header Checksum (16 bits): This field is used to detect errors in the IP packet header. The checksum is calculated over the entire IP packet header, and if any errors are detected the packet is discarded.

It is calculated by taking the 16-bit one's complement of the one's complement sum of all the 16-bit words in the header, with the checksum field itself being set to zero before the calculation is performed. To calculate the checksum, the header is first divided into 16-bit words, and then the sum of all these words is computed, including any carryover from the most significant bit of the previous addition. Once the sum has been computed, the one's complement of the result is taken, and this value is used as the checksum value.

When the packet arrives at the destination, the same calculation is performed on the received packet header, and the calculated checksum is compared with the value in the header. If the two values match

then the header is considered valid, and the packet is forwarded. If the checksums don't match, the packet is discarded.

The reason why a checksum is used in the IPv4 header instead of a hash is that a checksum can be computed quickly, with relatively low computational overhead. A hash function, on the other hand, is typically more computationally expensive to compute and would slow down the packet processing.

Source IP Address (32 bits) - This field contains the IP address of the sender of the packet.

Destination IP Address (32 bits) - This field contains the IP address of the intended recipient of the packet. Normally IP packets are sent to a single IP address, in which case we call the transmission a unicast. But there are also ways to send IP packets to multiple devices using something called broadcasts, and multicasts. In IPv4, there are two special addresses that can be used in the destination field for multicast or broadcast:

Broadcast Address: The broadcast address is used to send a packet to all hosts on a particular network segment. The IPv4 broadcast address is typically the highest address in the subnet, where all host bits are set to 1. For example, in a subnet with a netmask of 255.255.255.0, the broadcast address would be 192.168.1.255.

Multicast Address: The multicast address is used to send a packet to a group of hosts that belong to a specific multicast group. The IPv4 multicast address range is from 224.0.0.0 to 239.255.255.255 and is divided into two parts: well-known multicast addresses and dynamically allocated multicast addresses. Well-known multicast addresses are reserved for specific purposes, such as routing protocols, while dynamically allocated multicast addresses are assigned as needed by applications.

Options (variable length) - This field is used to provide additional information about the IP packet. Options are rarely used, but they can include things like security-related information or hints for routers. If included, each option has a two-byte "Option Type" field that identifies the type of option and a one-byte "Option Length" field that specifies the length of the option in bytes. The remainder of the option field contains the actual option data.

Some examples of options that can be included in the options field include:

"Timestamp": This option includes a timestamp that can be used to measure the round-trip time of packets.

"Record Route": This option instructs routers to record their IP addresses in the packet as it passes through them, allowing the sender to trace the route the packet took.

"Security": This option provides security-related information about the packet.

"Strict Source Route": This option specifies a strict path that the packet must follow through the network, allowing the sender to specify the exact route that the packet should take.

Note that the options field is not commonly used in typical IPv4 traffic, and its use is generally limited to specialized applications or situations where additional information or control over the routing of packets is needed.

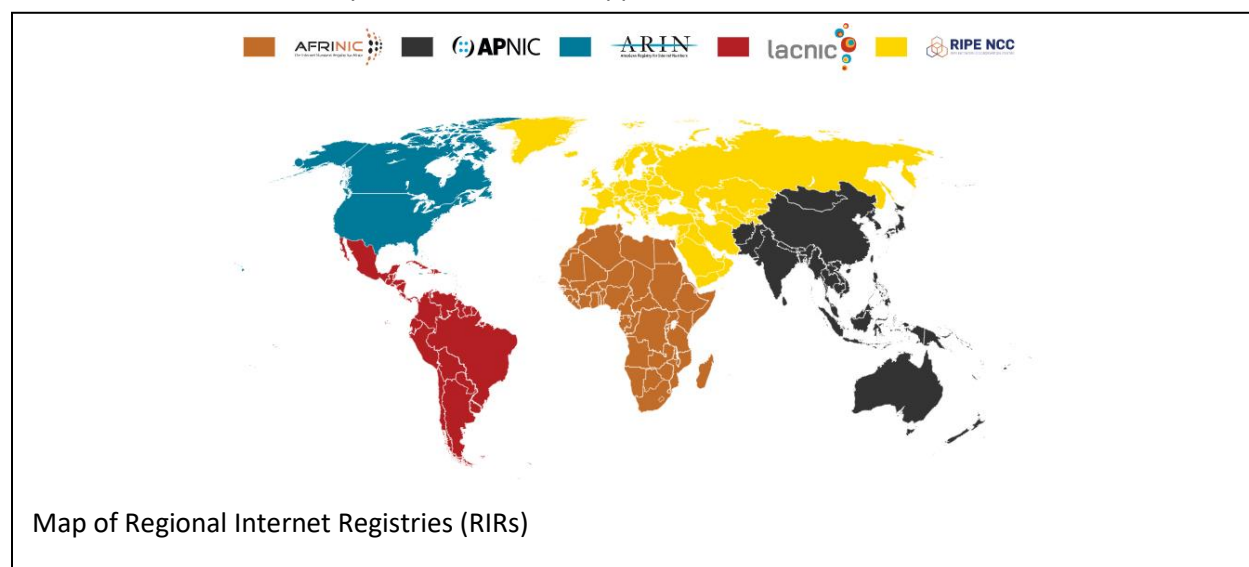
IP Address Details - How are IP Addresses Assigned

Now let's return and look at the details of IP addresses. We'll start by looking at how IP addresses are assigned. Any computer that wants to communicate on the Internet needs an IP address, but if you want an IP address, or a block of IP addresses, where do you go to get them?

There are two ways or three ways that IP Addresses are assigned. Like phone numbers, you can't just make them up, otherwise there would be the possibility of two devices having the same IP address, and for the Internet to function each device must have a unique address. The process for getting an IP address is a little like getting a phone number. That is, when you buy a new phone, you also need to get a phone number from whichever phone company you select. They will assign you a phone number from their pool of available numbers.

When you connect a computer to the Internet, that computer will be assigned an IP address by the network provider. For home networks and smaller businesses, the network provider is typically an ISP. Larger businesses and organizations will typically connect to the network run by a parent organization. For example, CBC connects to a network called the K20 network, which is run by the state of Washington, so the college receives their bucket of IP addresses from K20 network admins. Large companies like Microsoft or Boeing will also have a hierarchical network structure, where the company gets a block of IP addresses, which it then subdivides to its networks in different subsidiaries and locations. So, no matter what the situation you will get your IP address(es) from someone above you in the network hierarchy.

But where does your ISP or the person in charge of a large network get their block of addresses? The group in charge of handing out IP addresses is the IANA or Internet Assigned Numbers Authority. The IANA has the responsibility for managing the overall IP address pool, but they farm out the job of actually assigning numbers to groups called the Regional Internet Registries (RIRs)¹⁰. The RIRs are groups within a geographic region that have been given large blocks of IP addresses to distribute by the IANA. If you're a large company or organization that needs a large number of IP addresses to connect computers and devices to the Internet, you must submit an application to one of the RIRs. The number of available



¹⁰ <https://www.internetsociety.org/resources/deploy360/2015/short-guide-ip-addressing/>

IP address blocks is limited, so the RIRs won't hand them out to just anyone. Before assigning a block of addresses the RIR will perform some background checks to authenticate your organization and your need for the addresses, as well as ensure that your organization has met the technical requirements for connecting your network to the Internet.

IP Address Details – Address Blocks

If your organization is approved, the RIR will assign your company a block of addresses¹¹. The block will be one of three sizes, small, medium and large, or tall, grande, or venti for you caffeine aficionados and addicts. The IANA block sizes are called Class A, which is large, Class B, which is medium, and Class C, which is small. Note that Class A, B, and C addresses are now often referred to as legacy address classes, as the current approach to IP addressing is based on a different scheme called Classless Inter-Domain

	Start	End	Subnet Mask	# of Networks	# of Hosts
Class A	1.0.0.0	127.0.0.0	255.0.0.0	126	16,777,214
Class B	128.0.0.0	191.255.0.0	255.255.0.0	16,382	65,534
Class C	192.0.0.0	223.255.255.0	255.255.255.0	2,097,150	254

Table showing Class A, B, and C network blocks, and the associated network addresses.

Routing (CIDR).

If you are assigned a Class A block of addresses, the network number will be between 1.0.0.0 and 127.0.0.0. For example, any network that starts with 3.0.0.0 or 67.0.0.0 would be part of a Class A block of IP addresses. Notice that there are only 126 Class A blocks.

The organization that's assigned the Class A block can subdivide it using netmasks any way they wish. And they certainly will use subnets, because if they use the default subnet mask of 255.0.0.0, they will have a single network segment with ~16.7 million devices on the same segment and in the same collision domain. As an example, say you are assigned 67.0.0.0. You can use the subnet mask 255.255.0.0 and hand out network numbers 67.0.0.0 through 67.255.0.0 to the LAN administrators in your organization. This would allow each LAN administrator to either build a single network segment with 65534 hosts, or the LAN administrator could also use netmasks to further subdivide the network address. For example, say you've been assigned the network number 67.143.0.0. You could hand out network numbers 67.143.0.0 through 67.143.255.0. In other words, you could build 255 networks with 254 hosts on each network, which is much more practical solution.

If you are assigned a Class B block of addresses, the network number will be between 128.0.0.0 and 191.255.0.0. For example, any network that starts with 132.86.0.0 or 177.1.0.0 would be part of a Class B block of IP addresses. Notice that there are 16,382 Class B blocks, and each block can hold a maximum of 65534 hosts. Like Class A blocks, Class B blocks are typically subdivided into multiple network

¹¹ <https://www.meridianoutpost.com/resources/articles/IP-classes.php>

numbers. For example, if you used a netmask of 255.255.255.0 with a Class B block of address, it would result in 255 network numbers, with a max of 254 host per network segment, which is once again a much more practical use of the block addresses.

If you are assigned a Class C block of addresses, the network number will be between 192.0.0.0 and 223.255.255.0. For example, any network that starts with 197.0.14.0 or 22.12.10.0 would be part of a Class C block of IP addresses. Notice that there are 2,097,150 Class C blocks, and each block can hold a maximum of 254 hosts. You can also subdivide a Class C block using netmasks, but you'll have to use advanced netmasks which you'll learn about below.

You may have noticed that there are some network addresses that aren't assigned in any of the blocks. There are certain network numbers in each block that are called private network numbers or non-routable network numbers. These network numbers are called private or non-routable because if a router is sent a packet with one of these addresses it will drop it. This means the private addresses can be used on a local network segment, but they'll never be able to reach another network segment or the Internet. This might seem a little weird, but it's actually a great feature that's made it possible to extend the life of IPv4. You'll learn all about this in a minute as it's a key feature of modern-day networking.

Here are the three blocks of private IP addresses in IPv4:

10.0.0.0 - 10.255.255.255

172.16.0.0 - 172.31.255.255

192.168.0.0 - 192.168.255.255

The IP address 127.0.0.1 is also not assigned as part of a block, because as you learned this is known as the loopback address, and it is reserved for the internal loopback function of a computer's network interface.

In the olden days, when the Internet was first starting, it was relatively easy to get a Class B or Class C network number. When I worked at Hanford in the 1980s and 1990s, I was able to get two Class B networks, and PNNL also had two class B networks. Since then, Hanford and PNNL returned at least one of the Class B network numbers each, and they may have returned more. You'll learn how they were able to do this and still connect computers to the Internet in a minute when you learn about Private IP Addresses.

Today, it's nearly impossible to get a block of addresses from IANA and RIRs, especially a Class A or Class B block. But even though you can't go directly to IANA or an RIR and get a block of addresses, you can get an IP address or set of addresses. You'll just get them from your ISP or upstream network provider.

IP Address Details – Non-Routable IP Addresses

When the Internet first started the designers planned for growth and thought that having over 4 billion IP addresses would be more than enough. But, as the popularity of the Internet surged, and more and more devices needed an IP address it seemed like 4 billion wasn't going to be nearly enough. One of the problems was the sheer number of devices that wanted to connect and needed IP addresses, and another problem was that since IP addresses are handed out in large blocks there were many addresses

that ended up being “wasted”. That is, if an organization received a Class B block of 65535 addresses, but only needed 30,000, the remaining 35,000 would not be used, which was a huge waste of numbers.

The number of available addresses started to shrink at the same time that the demand was rapidly growing, forcing the IETF to look for solutions. They came up with two, the first of which was create a new version of IP and increase the size of the IP address, from 4 numbers to 6 numbers. This new address version was introduced in 1998 and is called IPv6¹². While some systems use IPv6, at this point in time IPv4 is still more commonly used. Even though IPv6 was released decades ago there’s been a huge lag in adoption. The reason for the delay is that another process, the use of private or non-routable IP addresses, freed up a huge percentage of the IPv4 address pool.

Here are a couple of last notes about IPv6 before we start on non-routable addresses. The first thing is that at this point if someone is talking about IP addresses and doesn’t specify a version, they’re probably talking about IPv4 because it is still by far the most commonly used. If someone is talking about IPv6, they will spell it out, but if they just say IP address, they are most likely referring to IPv4. The second is that IPv4 was designed in the 1970s¹³, and while the design has proven to be great, by the 1990s it was possible to see areas that could be improved. So, in addition to creating a larger address pool, IPv6 makes some other technical improvements¹⁴ to increase the overall performance. You’ll learn more about IPv6 later in the class, but for now we want to concentrate on non-routable IP addresses as they are used everywhere, including your home network.

The first thing we need to do is define what a non-routable IP address is. As the name states, a non-routable address is an IP address that will not pass through a router. When a router sees a network packet where the destination IP address is a non-routable IP address, it will drop the packet. Non-routable IP addresses can be used to send packets to other devices on the same network segment, including the default gateway/router, but they’ll never be passed outside of the local network segment. It’s important to note that you can still send packets to the default gateway using non-routable IP addresses. You’ll see why this is important in a minute.

Non-routable IP addresses are used for building private networks. By private network, we mean a network that may be connected to the Internet, but the devices on the network can’t be seen by any devices outside the network. And by outside the network we mean any devices on the non-LAN side of the router. For your home network and most small networks, outside the network means any device on the Internet.

You can think of a private network as something like a secret military base. The base has streets and buildings, and the buildings have street addresses like normal buildings. People inside the base can send mail to each other by using the street addresses, but no one outside the base knows anything about the base layout. They don’t know the street names or building addresses, they don’t even know what streets or buildings exist.

¹² <http://www.steves-internet-guide.com/ipv6-guide/>

¹³ <https://www.geeksforgeeks.org/history-of-tcp-ip/>

¹⁴ https://www.tutorialspoint.com/ipv6/ipv6_address_types.htm

This is what happens with most home networks, and many intranets for larger businesses and organizations. If you look at the IP addresses assigned to the devices on your home network or almost any home network, you'll see that they almost certainly start with 192.168. You'll also see IP addresses that start with 192.168 on almost all the computers at CBC, and IP addresses that start with 172.16 or 192.168 on most computers and devices connected to a network that's connected to the Internet.

But how is this possible? After all, isn't it a requirement for every computer and device connected to the Internet to have a unique IP address? Here's how the non-routable addresses have been implemented and how most network connections are made today. This is made possible through a system called Network Address Translation or NAT, which runs on the LAN routers or your home router.

The NAT system acts like the mailroom workers at the secret military base. Anytime someone wants to send mail to the outside world, they first send it to the mailroom. The workers in the mailroom write down some data about your message, mainly your address as the original sender, and the address you're sending the message to. They then take the message out the envelope you used and place it in a new envelope. When they address the new envelope they use the original recipient address, but they use the address for the base mail room as the sender's address.

They send your message in their new envelope to your original recipient. The recipient receives your message, and when they address their reply, they look at the envelope they got the message in and see that it came from the base mail room, so they send their message back to the base mail room address. In fact, they have no idea what your real address is, since it's nowhere on the envelope the message came in. They send their reply, and since it's sent to the base mail room when it arrives it's delivered to the mail room. The mail room inspects their list and sees that it's a response to a message you sent. They take the message out its envelope and place it in a new envelope with your address on the secret base as the recipient address. The process that the mail room uses, keeping lists of outbound mail, changing envelopes and addresses, and using their lists to forward incoming mail is known as address translation.

This system has a couple of benefits. The first is increased privacy, since any mail that is sent out from the base looks like it's coming from a single address and no one on the outside world can gather any new information about the buildings on the base by looking at the sender's address. The second is that addresses inside the base can be exact duplicates of addresses on other bases as these private addresses are never seen by anyone outside of the base.

On a network the address translation works the same way as the mail room process, with the LAN router doing the translation. Let's use a home network for a walk through, but the process will be the same for a LAN run by any organization. When you connect your home network to your ISP one of the first things that happens is your router sends a signal to the ISP saying it's ready to go to work by sending out a DHCP request. Your ISP will then assign an IP address to your router from its pool of available addresses. The following figure shows the IP configuration for a wireless router which is connected to the ISP CenturyLink, and CenturyLink assigned the wireless router the address 71.223.54.169.

3. IPv4 Addressing


PPP Parameter	Status
Modem IPv4 Address:	71.223.54.169
Modem IPv4 Subnet Mask:	255.255.255.255
DNS Address #1:	205.171.3.65
DNS Address #2:	205.171.2.65
Remote Gateway Address:	71.32.112.11

If we look at the devices connected to the wireless router, we can see that they've all been assigned addresses that start with 192.168.0 which is the Class C non-routable network number. Almost all home routers also act as DHCP servers, which is how each device on this network got its IP address. That is, when the device connected to the wireless network, the device also sent a DHCP request, and the router, acting as a DHCP server assigned the device their IP address, netmask, default gateway address, and DNS server addresses.

DEVICE NAME	IP ADDRESS	MAC ADDRESS	CONNECTION TYPE
Unknown	192.168.0.5	00:f3:61:a7:64:f8	5 GHz CenturyLink3272
amazon-9412c0d28	192.168.0.86	44:d5:cc:cd:f1:0b	5 GHz CenturyLink3272
amazon-987f40404	192.168.0.87	dc:91:bf:d0:09:17	5 GHz CenturyLink3272
TONY-s-S21-Ultra	192.168.0.104	d6:e4:a8:d8:80:85	5 GHz CenturyLink3272
CVD-Q2827	192.168.0.103	84:5c:f3:ca:82:b7	5 GHz CenturyLink3272
YSPC	192.168.0.105	14:18:c3:73:e0:79	5 GHz CenturyLink3272

Whenever a device on the home LAN sends a network packet to the Internet, the packet will be grabbed by the router. For example, let's say the computer with the IP address 192.168.0.103 wants to view a web page at google.com. The computer builds the network packet asking for the web page, using Google's IP address as the destination IP and 192.168.0.103 as the sender's IP address.

The computer then uses the netmask and determines that google.com is on a different network segment, so it builds an ethernet frame, places the IP packet in the data section of the ethernet frame, and addresses the frame to the default gateway's MAC address, which in this case the wireless router. The home router unpacks the ethernet frame, sees the network packet, and knows it needs to route it upstream to get it to google.com. But before the router forwards the network packet it will run the NAT process and record some data about your network packet. The router will record your device's IP address as the original sender, and the IP address of the device the packets are being sent to. The router will then build a new network packet sending it to the original recipient IP address but changing the sender's address to the "outside" address used by the router, in this case 71.223.54.169. The following figure shows an example NAT table, used by the router to keep track of outgoing network packets.



NAT Status

Protocol	Timeout	Source IP	Source Port	Destination IP	Destination Port
tcp	0	192.168.0.1	443	192.168.0.103	62568
tcp	0	192.168.0.1	443	192.168.0.103	62569
tcp	0	71.223.54.169	44958	184.72.149.114	443
tcp	0	127.0.0.1	47732	127.0.0.1	12345
tcp	0	127.0.0.1	12345	127.0.0.1	47732
tcp	0	192.168.0.1	443	192.168.0.103	62572
tcp	0	192.168.0.1	443	192.168.0.103	62571
tcp	0	192.168.0.1	443	192.168.0.103	62562
tcp	0	192.168.0.1	443	192.168.0.103	62570
udp	0	71.223.54.169	50191	34.211.237.22	6363
tcp	7370	192.168.0.104	51942	35.82.251.20	443
tcp	7308	192.168.0.104	41044	172.217.14.67	443
tcp	7307	192.168.0.104	56810	142.250.141.188	5228

When Google gets the network packet, it will build a response and send it back to 71.223.54.169 which is the IP address of your router. When router receives any network packets from its external interface, it inspects its NAT list to see which device made the original request. In this case the router sees that this packet is a response to the request you sent to Google. To get the network packets back to your computer, when the router receives the packet it extracts the data, then builds a new network packet, placing the data from Google inside the new packet, and using your computer's IP address as the recipient address.

Since your IP address isn't in the IP packet sent to google, google.com and any other routers that forwarded the network packets won't know that the packet actually came from your computer or your computer's IP address. This keeps the IP addresses of all the devices on your network private and hidden from the outside world.

Besides keeping your internal IP addresses private, using non-routable addresses has another huge benefit. Most devices on local network segments can use private addresses, and the only device on a network segment that needs a “real” IP address is the default gateway. This frees up almost all of the IPv4 addresses that were previously being used by devices on local network segments, which has greatly extended the life and usefulness of IPv4.

IP Address Details – Advanced Subnetting & Netmasks

Now let’s turn our attention to netmasks and advanced subnetting. Let’s start by doing a quick review of why a network administrator would choose to sub-divide a network number. This is done to make network numbers more flexible, make more efficient use of the IP addresses available for a given network number, and at the same time reduce the number of devices in a collision domain, especially for organizations that received large blocks of IP addresses. That is, say you were an organization and receive a Class B block of addresses that use the network number 131.27.0.0. This means that you have all the IP addresses between 131.27.0.1 and 131.27.255.254, or 65534 IP addresses you can assign to devices on your network. But you certainly don’t want all 65534 devices to be connected to the same LAN because this would require that all the devices be connected to the same network segment or plugged into the same router because this would be physically impossible. And all the devices would be in the same ethernet collision domain, which means there would be a lot of contention for the network media, lots of collisions, which would cause lots of delay for network traffic.

So instead of having one large network, you’ll probably want to subdivide your network using the netmask. There are three types of netmasks, that I’m going to call simple netmasks, binary netmasks, and complex binary netmasks.

The simple netmasks are the ones that you’ve already learned about, where the bits in an octet are either set to all 1s or all 0s. These consist of 255.0.0.0, 255.255.0.0 and 255.255.255.0. These netmasks make it really easy find the network portion of an IP address because it will consist of an entire octet.

The binary netmasks and complex binary netmasks are netmasks that require using binary numbers to determine the network and host portions of each address. They’re used like super simple netmasks, where the 1 bits in the netmask designate the network, but the numbers in the octets won’t be 255. These types of netmasks aren’t as easy for humans to use because you have to write the netmask in binary, but they provided added versatility and more efficient use of IP addresses. For example, even a class C subnet mask of 255.255.255.0 is not very efficient as it only allows 1 network with 254 hosts, and this is a large number of devices to connect to a single network segment and a large number of devices in a single collision domain. With a binary netmask you could subdivide this into two networks, or 4 networks, and with a complex binary netmask you could subdivide it into 5 networks.

Here are some examples of binary netmasks that will illustrate what they look like and how they’re used.

Binary Netmask - Example 1

Let’s start with an example of subnetting the class C address 212.13.177.0, dividing it into two subnets. To do this we’ll use a netmask of 255.255.255.128. The first step to understanding what’s going on with this netmask is to write it out in binary. In this case the binary netmask is:

11111111.11111111.11111111.10000000

To use this to divide an IP address into the network portion and the host portion, find all the bits that are set to 1. The bits that are set to 1 represent the network portion, and the bits that are set to 0 represent the host portion.

Now let's look at using this netmask to decide if two IP addresses, 212.13.177.14 and 212.13.177.134, are on the same network segment or not. The first step is to write each address in binary, then we'll compare the network portion of the addresses to determine if they're on the same subnet or not.

212.13.177.14 = 11010100.00001101.10110001.00001110

212.13.177.134 = 11010100.00001101.10110001.10000110

When we apply the netmask to each number, we can skip the first three octets because they're obviously the same. I'm going to write the last octet of binary numbers for the netmask and both IP addresses on subsequent lines to make it easy to see the network bits, and if they match or not.

Netmask:	10000000
IP 1:	00001110
IP 2:	10000110

In this case the two IP addresses are on different networks as IP address 1 has a 0 in the network portion and IP address 2 has a 1.

Now let's check the two IP addresses 212.13.177.14 and 212.13.177.68. These numbers in binary are:

212.13.177.14 = 11010100.00001101.10110001.00001110

212.13.177.68 = 11010100.00001101.10110001.01000100

You might be able to see the answer already, but let's write out the last octet of the netmask and the two IP addresses to make it crystal clear.

Netmask:	10000000
IP 1:	00001110
IP 2:	00000100

In this case the two IP addresses are on the same network as IP address 1 has a 0 in the network portion and IP address 2 also has a 0.

The next thing to see is how many networks and how many hosts per network we can achieve with this netmask. In this case, the network bit can be set to 0 or 1, which means that we can only have two different (sub) networks. For the number of hosts per network, we have 7 bits to work with. This means that we can have 2^7 or 128 total addresses per network. But we don't use 0 because this represents the network, and we don't use all 1's, or 127 because this is the broadcast address. Subtracting these two possible host values leaves $128-2$ or 126 possible hosts.

Binary Netmask - Example 2

Let's do another example of subnetting the class C address 212.13.177.0, this time dividing it into four subnets. To do this we'll use a netmask of 255.255.255.192. In this case the binary netmask is:

```
11111111.11111111.11111111.11000000
```

Notice that this is just like the 255.255.255.168 netmask, it just has one more 1 bit. Once again, to use this to divide an IP address into the network portion and the host portion, find all the bits that are set to 1. The bits that are set to 1 represent the network portion, and the bits that are set to 0 represent the host portion.

Now let's look at using this netmask to decide if two IP addresses, 212.13.177.10 and 212.13.177.234, are on the same network segment or not. The first step is to write each address in binary, then we'll compare the network portion of the addresses to determine if they're on the same subnet or not.

```
212.13.177.10 = 11010100.00001101.10110001.00001010
```

```
212.13.177.134 = 11010100.00001101.10110001.11101010
```

Once again, once you write the IP addresses in binary the answer is probably clear, but let's complete the process by writing the last octet of the netmask and the two IP addresses to check. I'm going to skip the first three octets because they're still the same, and only write the last octet of binary numbers for the netmask and both IP addresses on subsequent lines to make it easy to compare the network bits.

```
Netmask:    10000000  
IP 1:      00001010  
IP 2:      11101010
```

In this case the two IP addresses are on different networks as IP address 1 has a 00 in the network portion while the network bits in IP address 2 are 11.

Now let's check the two IP addresses 212.13.177.14 and 212.13.177.33. These numbers in binary are:

```
212.13.177.14 = 11010100.00001101.10110001.00001110
```

```
212.13.177.68 = 11010100.00001101.10110001.00100001
```

You might be able to see the answer already, but let's write out the last octet of the netmask and the two IP addresses to make it crystal clear.

```
Netmask:    10000000  
IP 1:      00001110  
IP 2:      00100001
```

In this case the two IP addresses are on the same network as both IP addresses have a 00 in the network portion.

The next thing to see is how many networks and how many hosts per network we can achieve with this netmask. In this case, the network bit can be set from 00 to 11, which means that we can only have four different (sub) networks. For the number of hosts per network, we have 6 bits to work with. This

means that we can have 2^6 or 64 total addresses per network. But we don't use 0 because this represents the network, and we don't use all 1's, or 63 because this is the broadcast address. Subtracting these two possible host values leaves $64-2$ or 62 possible hosts.

As you can see, each time we add a bit to the network portion of the netmask, we subtract a bit from the host portion. Adding bits to the network portion allows for more possible networks, but fewer hosts on each network. The following table shows the possible values for the binary netmasks for subnetting a Class C network, along with the number of network addresses and the possible number of hosts and number of usable hosts per network. Note that you'd never use the last two netmasks as they don't have any usable host numbers.

Subnet Mask	Binary Netmask (Last Octet)	Networks	Hosts	Usable Hosts	Broadcast Address	CIDR
255.255.255.0	00000000	1	256	254	0.0.0.255	/24
255.255.255.128	10000000	2	128	126	0.0.0.127	/25
255.255.255.192	11000000	4	64	62	0.0.0.63	/26
255.255.255.224	11100000	8	32	30	0.0.0.31	/27
255.255.255.240	11110000	16	16	14	0.0.0.15	/28
255.255.255.248	11111000	32	8	6	0.0.0.7	/29
255.255.255.252	11111100	64	4	2	0.0.0.3	/30
255.255.255.254	11111110	128	2	0	0.0.0.1	/31
255.255.255.255	11111111	256	0	0	0.0.0.0	/32

You could also use this process to subnet a Class B or Class A network using a binary netmask. I'm not going to cover the process here as the idea is impractical. That is, there are good reasons to make subnet to create more networks, even if they have fewer hosts per network. But there's not a good reason to take a Class B or Class A network and create subnets with 500 or more hosts per network. We'll go through one example of using a binary netmask on a Class B address, but we're not going to look at every possible netmask. The only time I suggest you go through the exercise of using binary subnet masks on Class B or Class A networks is if you're planning on taking a certification exam on networking like the Cisco CCNA or CompTIA Network+.

Binary Netmask - Example 3

Here's an example of subnetting the class B address 155.13.0.0 using a netmask of 255.255.254.0. In this case the binary netmask is:

```
11111111.11111111.11111110.00000000
```

Notice that this is much like the 255.255.255.0 netmask, it just has one less 1 bit. Like always, to use this to divide an IP address into the network portion and the host portion, we need to first find all the bits

that are set to 1 as they will represent the network portion, while the bits that are set to 0 represent the host portion.

Now let's apply this netmask to decide if two IP addresses, 155.13.8.10 and 155.13.177.199, are on the same network segment or not. The first step is to write each address in binary, then we'll compare the network portion of the addresses to determine if they're on the same subnet or not.

```
155.13.8.10      = 10011011.00001101.00001000.00001010
155.13.177.199 = 10011011.00001101.10110001.11000111
```

Once again, once you write the IP addresses in binary the answer is probably clear, but let's complete the process by writing the third octet of the netmask and the two IP addresses to check. I'm going to skip the first two octets because they're obviously the same for both IP addresses, and also skip the last octet since it's all host bits. This means we'll only write the third octet of binary numbers for the netmask and both IP addresses on subsequent lines to make it easy to compare the network bits.

```
Netmask:      11111110
IP 1:         00001000
IP 2:         10110001
```

In this case the two IP addresses are on different networks as the network portions of the IP addresses are different.

The next thing to see is how many networks and how many hosts per network we can achieve with this netmask. In this case, the network bit can be set with 7 bits which means that we have 2^7 or 128 (sub) networks. For the number of hosts per network, we now have 9 bits to work with, 1 from the third octet plus the 8 from the fourth octet. This means that we can have 2^9 or 512 total addresses per network. But we don't use 0 because this represents the network, and we don't use all 1's because this is the broadcast address. Subtracting these two possible host values leaves $512-2$ or 510 possible hosts.

Complex Binary Subnet Masks

If you're liking counting in binary, then you're now in for a treat as we look at complex subnet masks. If you're not digging binary, you might not like this part so much. But keep in mind you won't need to know this unless you work as a network administrator or want to try a certification exam. So far, the netmasks we've looked at, both the simple netmasks and the binary netmasks have had a solid row of either 255's or 1's if we're counting in binary. In technical terms we describe the bits in these numbers as being contiguous, which means all the 1's are adjacent without any 0's between any of the 1's. I'm not sure who invented the term contiguous, and sometimes wonder if it came to be when some tech writer misspelled continuous. In any case, a complex subnet mask is one that has bits set to 1 that are not contiguous. A complex subnet mask can be used to create a subnet that is not a power of two. For example, a complex subnet mask of 255.255.255.248 has 29 network bits and 3 host bits. This type of netmask is extremely versatile, and it can allow a network administrator to make the optimal use of their pool of IP addresses.

Binary Netmask - Example 4

Here's an example of subnetting the class C address 212.13.177.0 using the netmask 255.255.255.217. The first step to understanding what's going on with this netmask is to write it out in binary. In this case the binary netmask is:

```
11111111.11111111.11111111.11011001
```

To use this to divide an IP address into the network portion and the host portion, find all the bits that are set to 1. The bits that are set to 1 represent the network portion, and the bits that are set to 0 represent the host portion.

Now let's look at using this netmask to decide if two IP addresses, 212.13.177.14 and 212.13.177.134, are on the same network segment or not. The first step is to write each address in binary, then we'll compare the network portion of the addresses to determine if they're on the same subnet or not.

```
212.13.177.14 = 11010100.00001101.10110001.00001110
```

```
212.13.177.134 = 11010100.00001101.10110001.10000110
```

When we apply the netmask to each number, we can skip the first three octets because they're obviously the same. I'm going to write the last octet of binary numbers for the netmask and both IP addresses on subsequent lines to make it easy to see the network bits, and if they match or not.

```
Netmask:    11011001
IP 1:       00001110
IP 2:       10000110
```

In this case the two IP addresses are on different networks as the network bits in the IP addresses are different.

Now let's check the two IP addresses 212.13.177.78 and 212.13.177.108. These numbers in binary are:

```
212.13.177.14 = 11010100.00001101.10110001.01001110
```

```
212.13.177.68 = 11010100.00001101.10110001.01101100
```

You might be able to see the answer already, but let's write out the last octet of the netmask and the two IP addresses to make it crystal clear.

```
Netmask:    11011001
IP 1:       01001110
IP 2:       01101100
```

In this case the two IP addresses are on the same network as they have identical bits in their respective network portions.

DHCP Details

In this section you'll learn the details of DHCP including how the client can request, release or renew its IP information, the protocol used to exchange information between the client and server, and how to configure a DHCP server.

History

Dynamic Host Configuration Protocol (DHCP) is a protocol used to assign IP addresses and other network configuration parameters, such as subnet masks, default gateways, and DNS servers, to network devices dynamically. This allows network administrators to manage IP address allocation and network configuration easily, without having to manually assign IP addresses to each device.

DHCP evolved from a protocol known as the Bootstrap Protocol (BOOTP) which was introduced in the early 1980s during an effort by several UNIX vendors to provide something called diskless workstations. As the name implies, diskless workstations were computers that didn't have a local hard drive, but instead stored everything on a network drive. The push behind this was that at the time hard drives were relatively expensive and networks were getting faster (although they were ridiculously slow by today's standards). Plus, it would make administration easier as both operating system software and files, and application programs, could be administered from a central location instead of having to install and manage the OS and applications on each separate computer. It turned out to be a good idea, but it was ahead of its time, as at the time networks weren't fast enough to make it practical. But even though the diskless workstations weren't adopted in their entirety, the protocol they used to boot, bootp, was found to be useful and eventually evolved into the DHCP we know and love today.

In 1993, the Internet Engineering Task Force (IETF) released the first version of DHCP as an extension to BOOTP, which added features such as automatic IP address allocation, lease management, and support for multiple vendor-specific options. The current version of DHCP, DHCPv6, was released in 2003 and provides support for IPv6 networks. The specifications for DHCP are documented in a series of IETF RFCs (Request for Comments), including RFC 2131 for DHCPv4 and RFC 3315 for DHCPv6. These documents outline the message formats, options, and procedures used by DHCP clients and servers to communicate and exchange configuration information.

DHCP has become a widely adopted protocol and is supported by most operating systems and network devices. DHCP servers are commonly deployed in enterprise networks, internet service providers, and home networks to manage IP address allocation. DHCP is built into Windows, Linux, macOS, iPhones, Android phones, and almost every operating system or device that require network connectivity. It works so well and so transparently that most people have no idea that their devices are using DHCP.

Overview

Before we jump into the details, let's start with an overview of the DHCP process which will show the main components and terminology. Remember that DHCP (Dynamic Host Configuration Protocol) is a network protocol used to automate the process of assigning IP addresses and configuring network parameters for devices on a network. Its purpose is to simplify the management of IP addresses and network settings by dynamically allocating and renewing them as needed. DHCP eliminates the manual configuration of IP addresses, subnet masks, default gateways, DNS server addresses, and other network

parameters on individual devices. By centralizing the IP address management, DHCP enables efficient utilization of available IP addresses and facilitates easy addition, removal, and movement of devices within a network.

Here are the main components of DHCP:

1. **DHCP Server:** A DHCP server is a network device or software application responsible for assigning IP addresses and other configuration parameters to DHCP clients on a network. It manages a pool of available IP addresses and leases them to clients upon request.
2. **DHCP Client:** A DHCP client is a device, such as a computer, smartphone, or network printer, that requests network configuration information from a DHCP server. It is responsible for obtaining an IP address, subnet mask, default gateway, DNS server addresses, and other DHCP options.
3. **DHCP Relay Agent:** A DHCP relay agent is a network device that forwards DHCP messages between DHCP clients and DHCP servers that are on different subnets or networks. It relays the DHCP requests and responses to ensure communication between the clients and servers across network boundaries.

The DHCP client, server, and Relay Agents communicate using the following types of messages:

1. **DHCP Discover Message:** The DHCP Discover message is a broadcast message sent by a DHCP client to discover available DHCP servers on the network. It is used to initiate the IP address assignment process. The Discover message contains information about the client's hardware address, network segment, and DHCP options it supports.
2. **DHCP Offer Message:** When a DHCP server receives a DHCP Discover message, it responds with a DHCP Offer message. This message is sent as a unicast directly to the client and includes an IP address lease offer, subnet mask, default gateway, DNS server addresses, and other configuration options.
3. **DHCP Request Message:** After receiving one or more DHCP Offer messages, the DHCP client selects a DHCP server and sends a DHCP Request message. This message confirms the acceptance of the offered IP address and configuration parameters from the chosen DHCP server.
4. **DHCP Acknowledgment Message:** Upon receiving the DHCP Request message, the selected DHCP server sends a DHCP Acknowledgment message to the client. This message acknowledges the IP address lease and provides the client with the confirmed network configuration details. It also specifies the lease duration, which determines how long the client can use the assigned IP address before it needs to renew.
5. **DHCP Lease Release:** DHCP lease release occurs when a DHCP client voluntarily relinquishes its assigned IP address before the lease expiration. The client sends a DHCP Release message to the

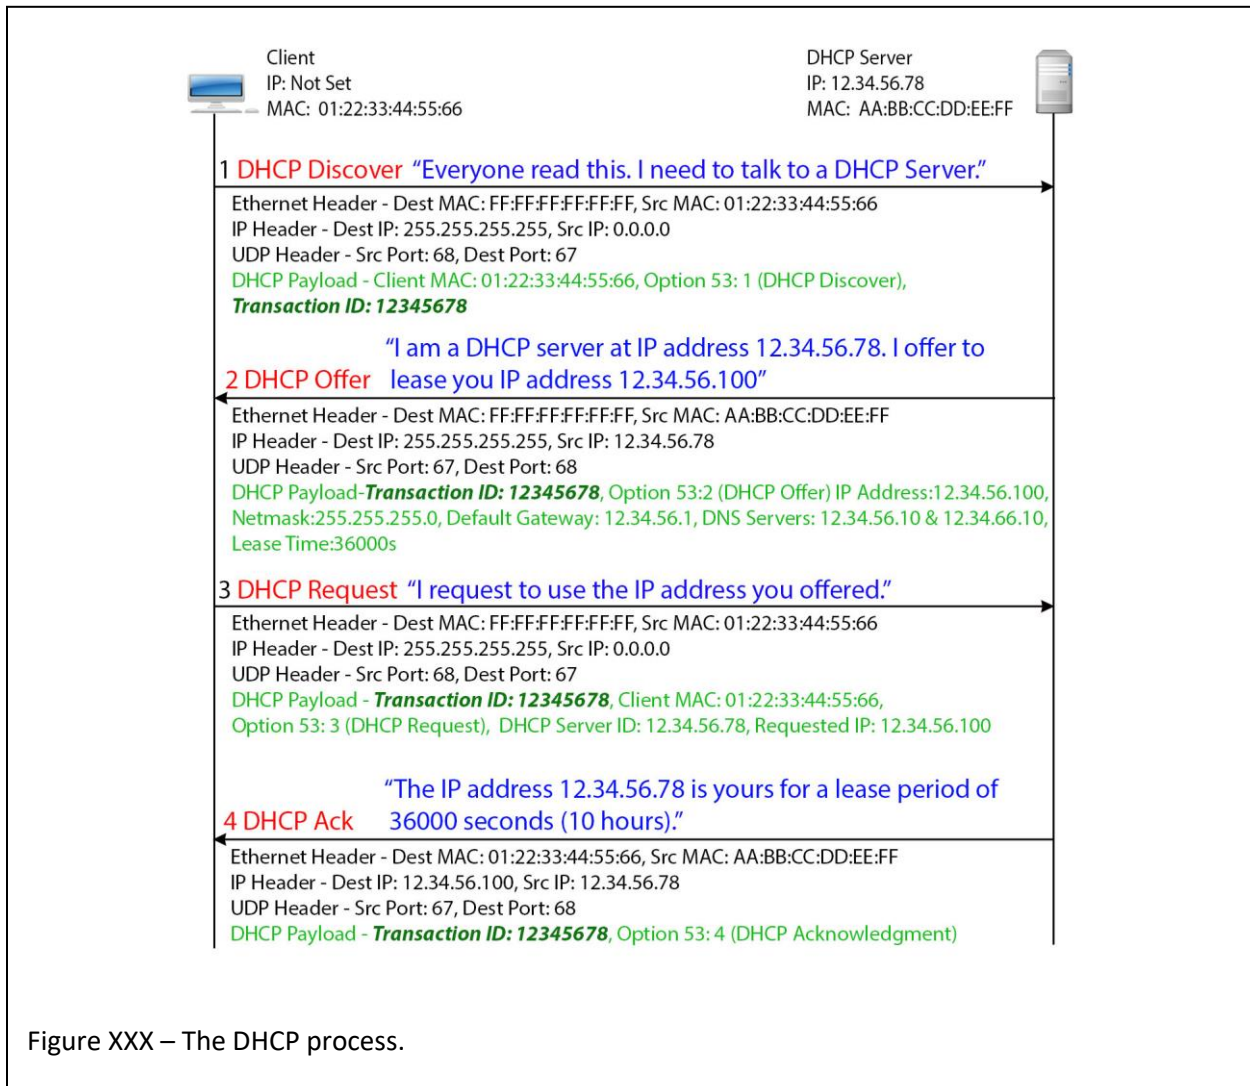
DHCP server, indicating that it no longer requires the IP address. The server then makes the IP address available for reassignment to other clients.

Here are some terms that are used with DHCP leases:

1. **Lease Duration:** The lease duration refers to the period of time that a DHCP client is allowed to use the assigned IP address. It is specified by the DHCP server in the DHCP Acknowledgment message. When the lease duration approaches expiration, the client must renew the lease to continue using the IP address. The lease values are specified in seconds. For example, a lease value of 3600 would be 1 hour.
2. **DHCP Lease Renewal:** DHCP lease renewal is the process by which a DHCP client extends the lease duration for its assigned IP address. Before the lease expires, the client sends a DHCP Request message to the DHCP server, requesting a renewal of the lease. If the server approves the renewal, it sends a DHCP Acknowledgment message with an updated lease duration.
3. **Binding:** When a server provides an IP address to a client, it keeps track of the client's MAC address and the leased IP address by binding the two addresses together. This is actually done in a database, which the DHCP server uses to keep track of the IP addresses it has available, and the IP addresses that are currently on lease to a client.

DHCP Process

The main steps in the DHCP process involves a client first finding a DHCP server using a process called Discovery, then the server replying with an offer of an IP address and other networking information. Here are the details of this process and a look at the IP packets and ethernet frames used. You're shown the UDP and IP packets and ethernet frames, and the DHCP data or payload that's transmitted in the network packets, so you can see how the process is accomplished even when the client has no IP configuration.



Note that for a complete understanding of this process you will need to know about UDP (User Datagram Protocol) and port numbers, which haven't been covered yet. You'll get a quick explanation of UDP packets and port numbers here, and the complete explanations in later chapters. UDP packets are similar to IP packets and ethernet frames, in that they are used to transport network data. The difference is that ethernet frames are used to move data from point-to-point, IP packets are used to move data from end-to-end, and UDP packets are used to get data to the correct application once the data has reached its final destination. That is, a computer can be running several applications that are all using the network simultaneously, so as data comes up the network stack, there needs to be some way to decide which program gets the data. This routing of network data as it comes up the stack is done using something called port numbers, which are just numbers that are assigned to specific applications. Popular server applications such as web servers, email servers, and DHCP servers are assigned what are

called well known port numbers. In the case of DHCP, the server port number is 67, which means that any incoming network data bound for port 67 will be given to the DHCP server.

With that background in place, let's look at the details of the DHCP process:

1. **DHCP Discover:** When a DHCP client connects to a network, it broadcasts a DHCP Discover message to discover available DHCP servers. The purpose of this step is for the client to find a DHCP server or servers and initiate the IP address assignment process. At this point the only network configuration data the client knows is its own MAC address, and it has no idea what the MAC address or IP address of the DHCP server may be. To get the DHCP Discover message to the DHCP server the client sends this message as an ethernet broadcast, FF:FF:FF:FF:FF:FF, which means it will be seen and read by every device on the local network segment.

The DHCP Discover message is built using the DHCP protocol, which defines several message fields and the possible values to place in each field. You'll learn the details of the DHCP protocol below, but for now we'll just look at two fields called the Options field and the Transaction ID field.

The DHCP Options field contains many possible options that are identified by a number. In the DHCP process the value of option #53 is used to specify the type of message. For a Discover message option 53 is set to 1, for an Offer message option 53 is set to 2, for a Request message it is set to 3, and for an Acknowledgment it is set to 4.

The Transaction ID field is used to distinguish one DHCP Discovery session from other sessions. The value in the field is set to a random number by the client, and then used in the subsequent Offer, Request, and Acknowledgment messages. This is done in case multiple clients are asking for IP addresses at the same time. Normally something like this would be handled using the client's IP address, but since the client doesn't have an IP address yet the Transaction ID field is used.

In the example in the figure, the DHCP Discover message is placed inside what is referred to as the DHCP payload. The term payload comes from the airline industry, where the passengers or freight loaded into a plane are called the payload because they're the things that pay for the cost of the flight. With network data the network protocols like IP and ethernet are delivery devices, like the plane, while the data loaded into the network packets are called the payload. In this example the DHCP payload has Option 53 set to the value 1, which means it's a Discover message, and the Transaction ID is set to 12345678.

The next thing that happens on the network stack on the client computer is the DHCP payload is loaded into a UDP packet, which is called a datagram. In this case the main thing to configure in the UDP datagram will be the source and destination port numbers. The destination port number will be set to the well-known number for DHCP, which is 67. The source port number will be set to 68 which is the port assigned to a DHCP client.

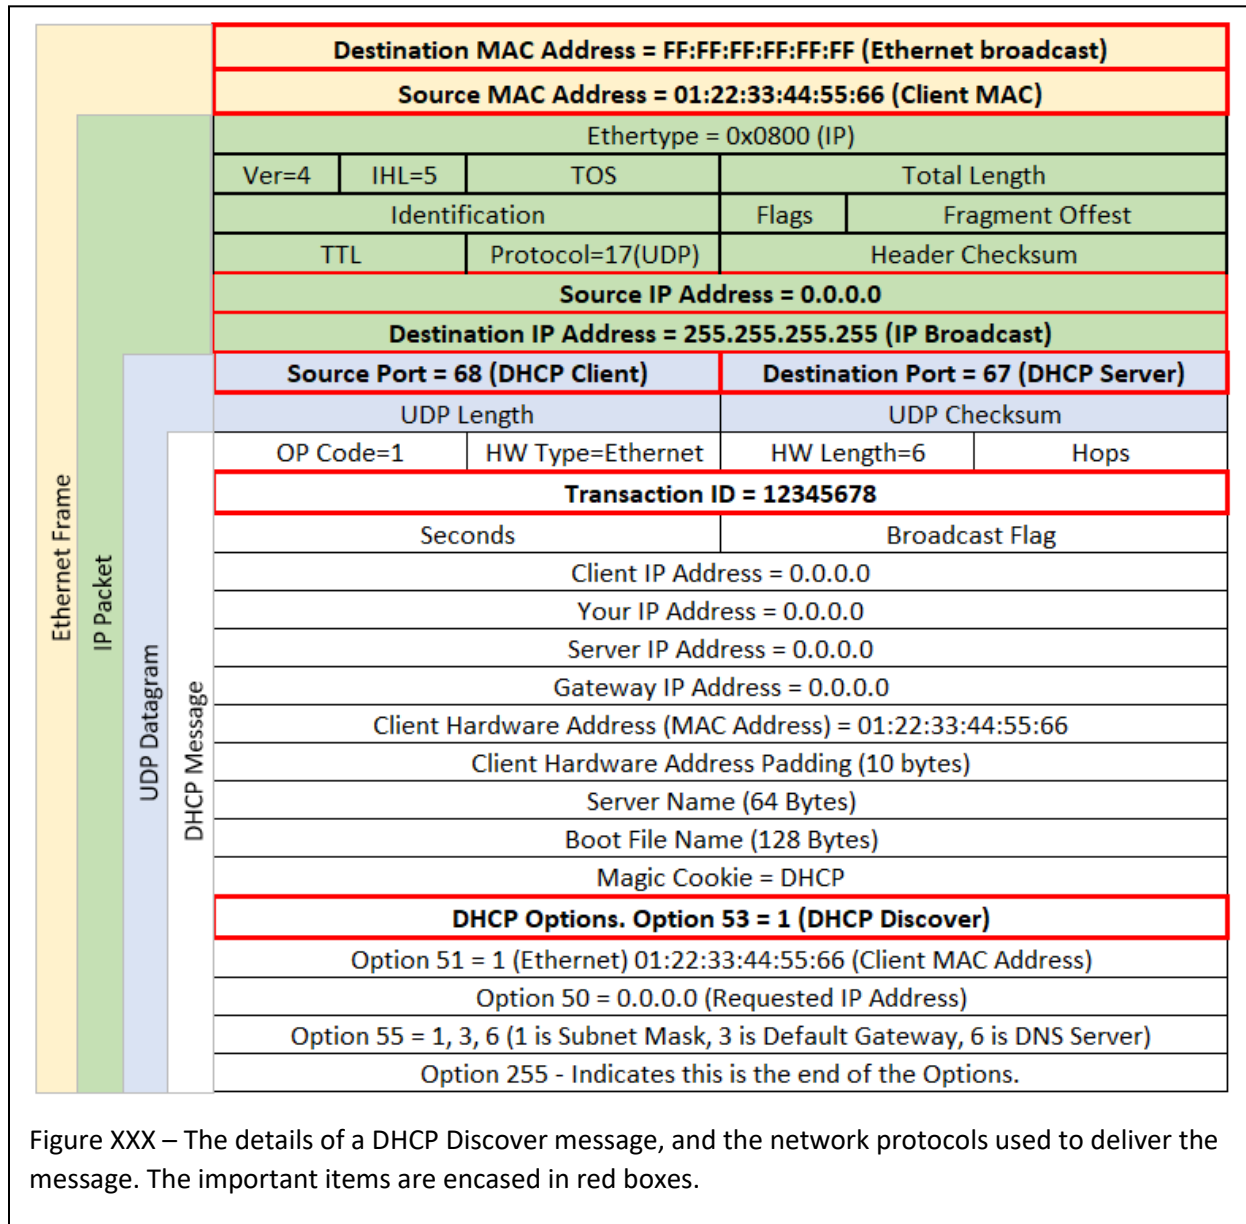


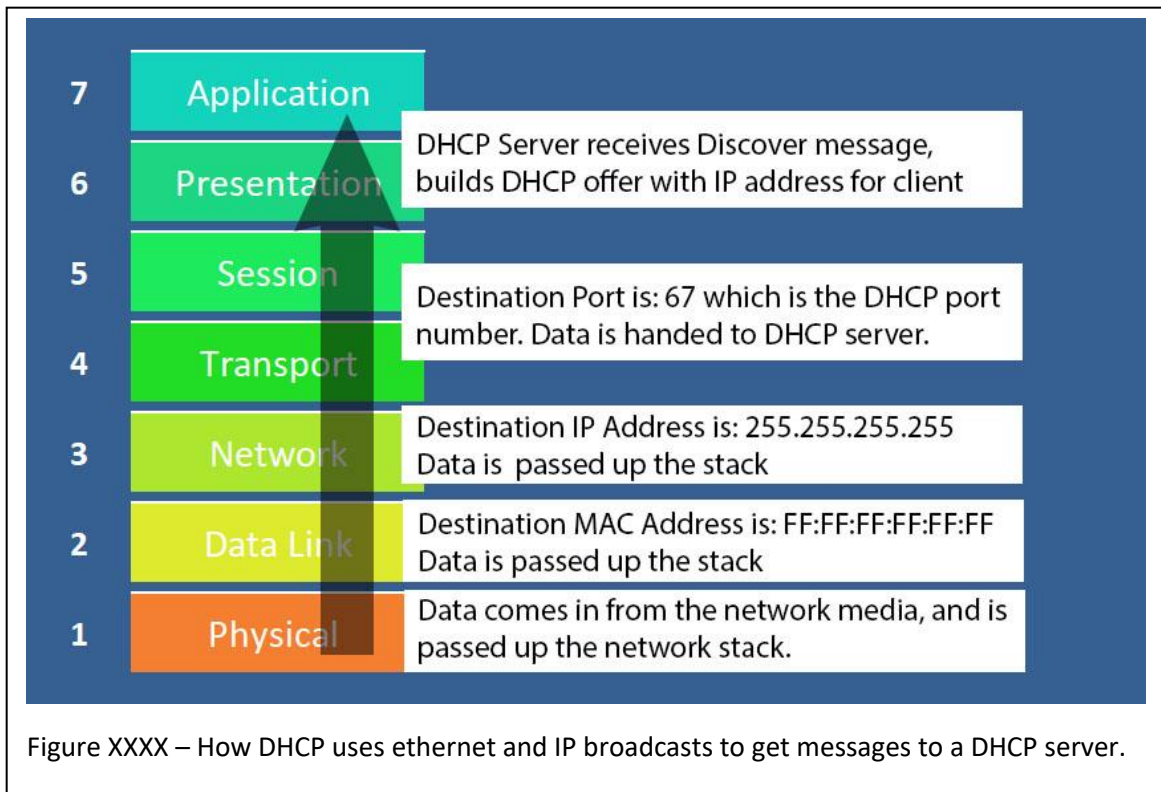
Figure XXX – The details of a DHCP Discover message, and the network protocols used to deliver the message. The important items are encased in red boxes.

The UDP datagram is then placed inside an IP packet. At this point the client doesn't have an IP address, which is why it's running DHCP in the first place, so the Source IP Address is set to 0.0.0.0. The destination IP address is set to 255.255.255.255 which is the broadcast address for IP and instructs all devices on the network segment to look at the IP packet, and hand it up the network stack to the next layer for further processing.

The IP packet is then placed inside an ethernet frame. In the ethernet frame the Source MAC Address will be the client's MAC address, which is known because it's burned into the firmware on the computer's NIC. For this example we'll use a fictional MAC address of 01:22:33:44:55:66.

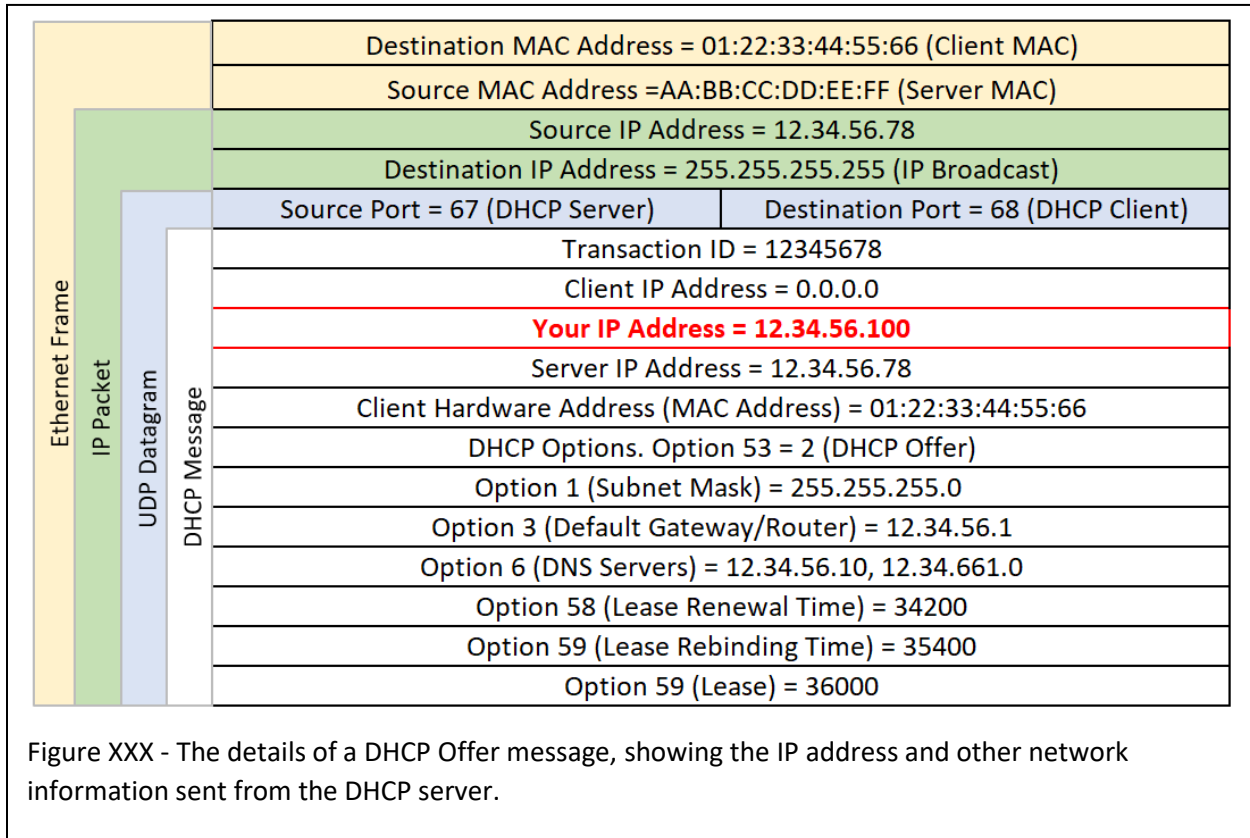
The Destination MAC Address should be the MAC address of the DHCP server, but the client doesn't know what this is, and can't do an ARP yet, because it would need to know the IP address of the DHCP server for ARP to work. So in this case the Destination MAC Address is set to FF:FF:FF:FF:FF:FF which is the broadcast MAC address and instructs all devices on the network segment to look at the frame and pass it up the network stack to the IP layer.

The result of all this is that the DHCP Discover message will be sent out as an ethernet broadcast. Because it's a broadcast all the devices on the network will read the ethernet frame, then pass the data up the network stack to the IP layer. The IP layer will see the IP broadcast address and try to pass the data up the network stack to the application listening to port number 67, which is the DHCP port. If a computer isn't running a DHCP server, the UDP packet will be dropped. But if the computer is running a DHCP server, the UDP packet and the DHCP Discover message will be handed to the DHCP server.



- DHCP Offer:** Upon receiving the DHCP Discover message, a DHCP server will respond by building a DHCP Offer message. The DHCP server will select one IP address from its pool of available IP addresses and add this to the Offer message. Each DHCP Offer message will also contain a subnet mask, default gateway IP address, DNS server addresses, and other configuration options. After the DHCP server sends the Offer message it reserves the offered IP address for the client, ensuring that it won't be assigned to any other client during this process. Note that

the reservation is only temporary at this point. The DHCP server won't make it official until it gets a response from the client.



The DHCP Offer message will also include three options related to the lease time. These are Option 51, the Lease Time, Option 58, the Lease Renewal Time, and Option 59, the Lease Rebinding Time. All these values are in seconds, so in the example the Lease Time is 10 hours, the Lease Renewal Time is 9 hours and 30 minutes, and the Lease Rebinding Time is 9 hours and 50 minutes. You'll learn the details about these three options below in the section on lease times.

Once the DHCP message is built, it's encapsulated in a UDP datagram which is sent to port 68, the DHCP client port. The UDP datagram is then placed inside an IP packet. The thing to note about the IP packet is that Destination IP Address is the IP broadcast address, 255.255.255.255. This is still being sent as a broadcast because the client hasn't yet received this message, so the client is still working without an IP address. Using the broadcast IP address ensures that the client computer's network stack will pass the data up the stack and that it will get to the DHCP client.

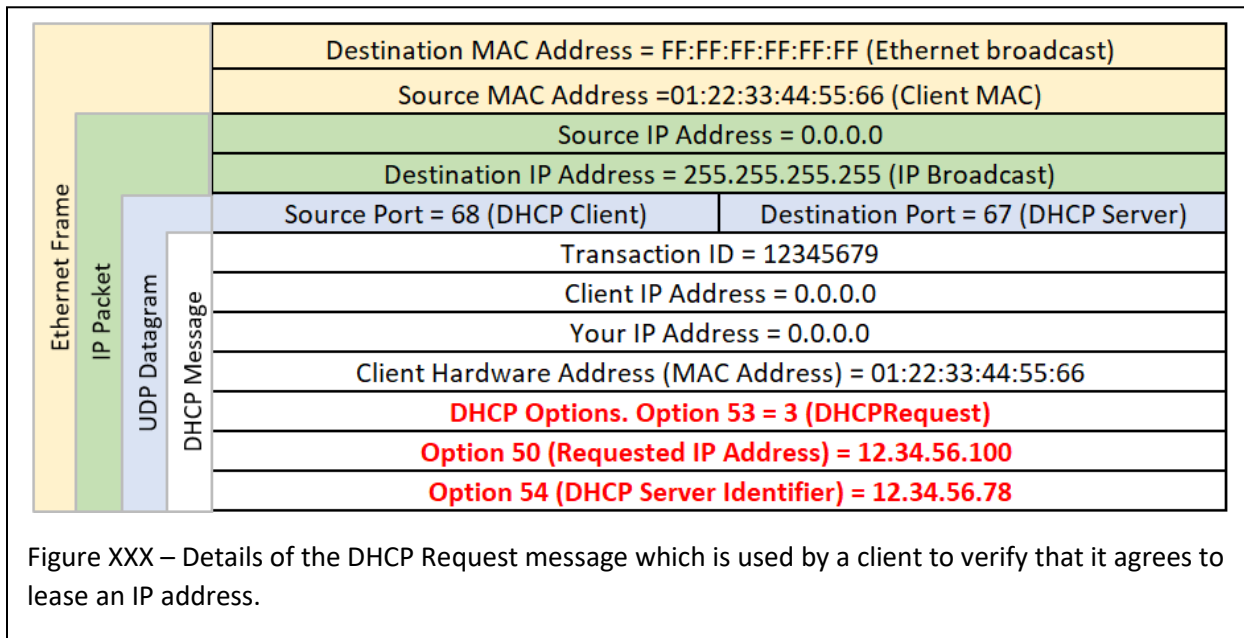
The IP packet is placed inside an ethernet frame which is addressed to the client's MAC address. The DHCP server knows the client's MAC address because it was sent in the ethernet frame containing the DHCP Discovery message.

The ethernet frame is then transmitted across the network and when the client receives it, the client unpacks the IP packet and hands it up to the network layer. The network layer sees the IP broadcast address, so it unpacks and reads the UDP datagram. The client's network stack then checks the destination port, and since it's set to 68, hands the DHCP message to the DHCP client.

The DHCP client reads this message, and now has the IP address sent by the DHCP server. The client also has a subnet mask value, IP address for the default gateway/router, and IP address for a DNS server. In the case of this example, the DHCP server offers the client a lease for the IP address 12.34.56.100, using a subnet mask of 255.255.255.0, the default gateway/router IP address of 12.34.56.1, and IP addresses of 12.34.56.10 and 12.34.66.10 for the Primary and Secondary DNS servers.

The DHCP client will also check the Transaction ID in the DHCP Offer to ensure that it matches the Transaction ID it sent in the DHCP Request. If it doesn't match, then it will ignore the offer. It's possible for an organization to run more than one DHCP server, so it's possible that multiple servers will send out DHCP Offer messages. If this happens, the DHCP client will take the IP address from the first offer it receives.

3. **DHCP Request:** After the DHCP client receives an Offer it sends a DHCP Request message which confirms that the client agrees to lease the IP address. If the client receives multiple DHCP Offers it selects one, typically the first it received, but only sends a single DHCP Request message. The Request message confirms the acceptance of the offered IP address and configuration parameters. The Request message is also sent as a broadcast, which will inform any and all DHCP servers that the client has chosen a specific DHCP server to lease an IP address from.



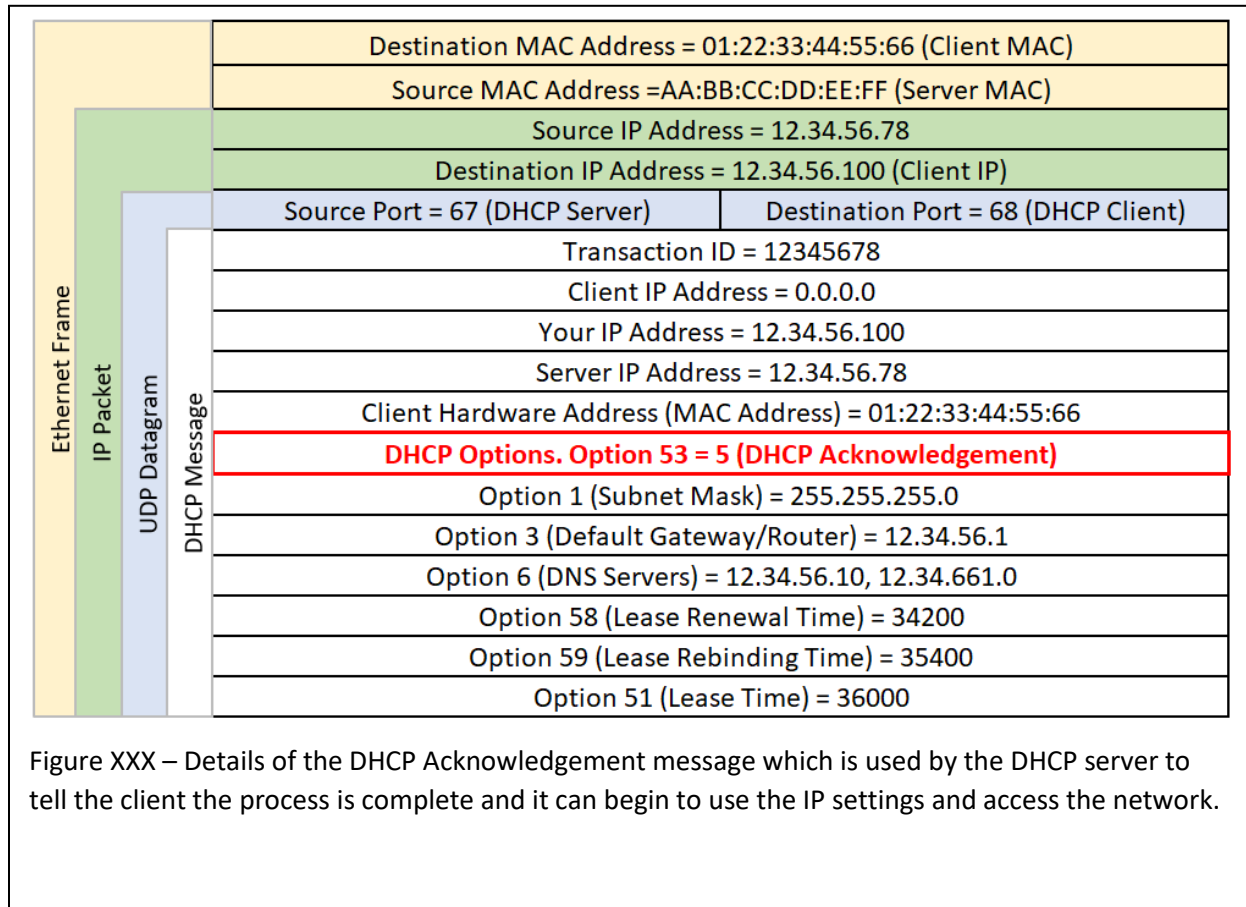
The key parts of the DHCP message are Options 50, 53, and 54. Option 53 is set to 3 which indicates that this is a Request message. Option 50 is used to tell the DHCP server which IP address the client is agreeing to lease, and Option 54 is set to the DHCP server's IP address and is used to identify the DHCP server. Option 54 is only needed if there's more than one DHCP server that may have made an offer, but it's always set to handle this possibility. If a DHCP server receives a DHCP Request but Option 54 contains the IP address for a different DHCP server, it will ignore the request.

The IP and ethernet destination addresses are also set to the broadcast addresses to handle the possibility of multiple DHCP servers. If there are multiple DHCP servers, they will all pass the data up the network stack to the DHCP server. Once the data hits the DHCP server it will only be processed by the server with the IP address set in Option 54 of the DHCP message. The other DHCP servers will know that their lease offers have not been accepted and they can return the IP address they offered back to the pool of available addresses.

Notice that the DHCP client still uses the IP address 0.0.0.0 in the IP packet instead of the IP address it's agreeing to lease. This is because the lease isn't official until after the next DHCP message, the DHCP Acknowledgment message, has been received.

4. **DHCP Acknowledgment:** Upon receiving the DHCP Request message, the DHCP server will know that the client has accepted the lease. At this point the DHCP server will finish the lease process on its end by marking the offered IP address as "assigned" or "in-use" in its lease database and associating the IP address with the client's MAC address. The association between the leased IP address and the client's MAC address is referred to as binding, and it's done to help ensure that the same IP address is assigned to the same client upon subsequent lease renewals.

Once the DHCP server has updated its lease database it builds an Acknowledgement message to let the client know the DHCP process is complete, and the client can begin using the IP address and other settings. The Acknowledgement message repeats the IP settings it previously sent the client, but Option 53 is set to a 5 which indicates that this time the message is an Acknowledgment.



When the client receives the Acknowledgement message, it knows that the process is also complete on its side, and it can begin to use the IP address and other settings to access the network.

Lease, Renewal, and Rebinding Times

During the DHCP process there are several values exchanged that deal with lease times, the Lease Time, Lease Rebinding Time, and Lease Renewal Time. In this section you'll learn what the meaning of these values and how they're used. Let's start with the Lease Time. When a client leases an IP address from a DHCP server it doesn't get to use it forever, as the lease will expire at some point, which is specified by the Lease Time. This time is used by the DHCP server and starts when the server sends the DHCP Acknowledgement. At this time the leased IP address and client MAC address are bound in the DHCP server's database, and the clock starts ticking down. If the client doesn't renew the address before the lease runs out, the server removes the binding from its database and returns the IP address to the pool of available addresses.

If a client wants to continue using a leased IP address, it needs to send a DHCP Renewal message to the DHCP server before the Lease Time expires. Renewing a lease is a simpler process than creating the

original lease, so renewing a lease is preferable to stopping all network traffic on a client and getting a new IP address. Plus, with a lease the client can keep using the same IP address, as opposed to possibly getting a completely different IP address, which might happen if the DHCP lease process is started from the beginning. The amount of time the client waits before sending the DHCP Renewal message is specified by the Lease Renewal Time. The Lease Renewal Time should obviously be less than the Lease Time.

The Lease Rebinding Time is used by the client in case the DHCP server doesn't respond to a renewal message. For example, if the DHCP server that issued the original lease was replaced or taken offline, the client won't be able to renew the lease. If the DHCP server doesn't respond and the Lease Rebinding Time is reached, the client will broadcast a Lease Rebinding message. This is sent as a broadcast so that any DHCP server will respond. The Rebinding message also includes the client's current IP address, in the hope that a new DHCP server can be found, and that the new DHCP server will be able to issue a new lease for the same IP address. The Lease Rebinding Time should be greater than the Lease Renewal Time, but less than the Lease Time.

If you're a network administrator and need to set the lease times, you need to take a couple of things into consideration. The first is how static or dynamic will the devices connecting to the network be. That is, will devices be needing addresses for days or weeks at a time, like in a business office, or will the devices only need an address for an hour or a few hours, like in a coffee shop or airport, or will be something in between like in your home where some devices are permanent, but others may come and go. The second factor is how many IP addresses you have to work with and how many devices will need an address. In most cases you'll be working with non-routable IP addresses, so you'll typically be working with ~250 IP addresses from the 192.168.1.0 network. But, if you become a network admin, you may have situations where you have smaller or larger pools of addresses to work with.

Here are a couple of scenarios and suggestions for values for the lease times. In the first scenario assume you're setting up a wireless network in a location like a coffee shop where the expectation is that there will be devices connecting and disconnecting from the network frequently, with the average connection lasting an hour or less. The coffee shop has 15 tables and a few other places to sit and is licensed for a maximum occupancy of 100 people. In this case shorter times are probably better because most customers aren't going to stay for an extended period, and even if they do, they can renew an IP address with little interruption. You also want to reclaim unused IP addresses frequently, because when a customer leaves the shop they won't need the address, and you may need the IP address for the constant stream of new customers. That is, if you set the lease time to something like 24 hours, your DHCP server will most likely run out of addresses because it won't reclaim unused addresses quickly enough. For this scenario a good lease time would be an hour or 3600 seconds. For the renewal time and rebinding time the general rule of thumb is that the renewal time should be 50% of the lease time, or 30 minutes (1800 seconds) in this case, and the rebinding time should be 87.5% of the lease time or 52.5 minutes (3150 seconds). Keep in mind that these percentages are just guidelines and not a strict rule, but it does give you some place to start when you're selecting times.

If you're setting up DHCP on a network in an office, or some location where the connected devices will remain fixed for weeks and months at a time, you can afford to make the lease time longer. For example, at the college we have several computer labs where the IP addresses and IP configuration settings for the computers in the labs are obtained through DHCP. In this case the lease time is set to 8

days or 691,200 seconds. This might seem like a long time, but the computers in these labs are rarely changed, and in fact ran with static IP addresses for many years. The renewal times are 7.5 days or 561,600 seconds, and the rebinding time is set to 6 days and 22 hours or 684,000 seconds. The renewal and rebinding times aren't set to the typical 50% and 82.5%, but they still leave plenty of time for the devices to renew, or rebind if necessary.

The last scenario to think about is in your home where your home router is most likely acting as your DHCP server. This situation falls somewhere between the coffee shop scenario and the computer lab. That is, some of your home devices will be static, while others such as your laptops and phones may be connected and disconnected frequently. Plus, you may have friends and family members that drop by and want to connect their devices and need IP addresses on a more dynamic basis. But even if allow your friends and family members connect to your network you most likely won't be concerned about running out of addresses like you would be with the coffee shop scenario. Home wireless routers will come set with a default DHCP lease time, with most being set to either 24 hours or 72 hours (3 days) depending on the router brand and model.

Router Brand	Default DHCP Lease Time
TP-Link	24 hours
Netgear	24 or 72 hours
Linksys	24 hours or 7 days
ASUS	24 hours
D-Link	24 or 72 hours

Table XXX – Typical DHCP Lease times for major home router manufacturers.

The amount of control you have over the DHCP lease times also varies by manufacturer and model. Some won't have any way to change the default values, some will let you change all three values, and others will allow you to set the DHCP Lease Time but not the renewal and rebinding times. In the last case, the renewal and rebinding times will typically be set to the 50% and 82.5% of the lease time.

If you want to check or control the settings on your home router you'll have to login to the router as the administrator, and then find the configuration settings. The process for logging in as administrator and finding the DHCP settings varies by router brand and model, so you'll have to search the Internet for instructions for your router. Note that this is just something you can do if you're curious. I don't suggest mucking around with these values, or any of the settings on your home router, unless you're having

Use Router as DHCP Server

Starting IP Address: 192 . 168 . 1 . 2

Ending IP Address: 192 . 168 . 1 . 254

DHCP Lease Time (range: 1-24 hours): 24 Hours

Figure XXX – DHCP configuration on a Netgear Nighthawk router.

specific issues that you're trying to correct. Your home router is definitely a case where the saying of "If it ain't broke, don't fix it" applies.

The following figure shows the DHCP configuration settings for a Netgear Nighthawk router. As you can see, the only configurable setting is for the DHCP Lease Time, and the default value is 24 hours.

You can also use Wireshark to view the DHCP settings by starting Wireshark, setting the filter to only view DHCP packets, then connecting a new device to your home network. You can then drill down into the DHCP Offer and DHCP Acknowledge messages to see the lease, renewal, and rebinding values.

Labsim Section 4 Notes

4.1.14 Practice Questions

Several of the questions have to do with advanced subnet masking. This is a subject you'll want to know if you become a Network Administrator, or if you want to pass a certification exam. But, if there are any subnetting questions on the exam they will only cover basic subnetting. So, my point is that you can spend some time working with advanced subnetting if you really want to understand it at this point, but you will not be tested on advanced subnetting.

Also, make sure that you're familiar with CIDR notation for subnet masks.

At least one of the questions requires converting between decimal and binary. It would be good if you can do this yourself. But, if not you can always use an online convertor. Just search for something like "convert decimal to binary" or "convert binary to decimal".

Section 4.2 and 4.3 in wrong order

Labsim covers APIPA before they cover DHCP which is a mistake in my opinion. Until you know what DHCP does, the APIPA info won't make much sense. So, I strongly suggest that you do Section 4.3 before Section 4.2, and read or view the extra material on DHCP before you look at the APIPA material.

Although if you do section 4.3 first, one of the DHCP quiz questions will be hard to answer as it requires knowledge of APIPA. So ... I suggest going through the material in 4.3 and 4.2 before trying the Practice Questions in either section.

4.3.11 Practice Questions

One of the questions requires knowledge from the DHCP Relay section, which is the next section.

4.4 DHCP Relay

This section is optional. While it's crucial that you understand the basics of DHCP and setting up a DHCP server, it's not important that you understand how to configure a DHCP relay at this point. This is rather

advanced information that you may use if you're a network administrator setting up DHCP, but I think it might just muddy the waters at this point in your education. So, you can do this section if you want, but it's optional.

4.5.9 Configure DNS Caching on Linux

This video is optional as it is obviously meant for Linux systems and users, which means there's a whole boat load of background information you will need to understand what this is talking about. You may choose if you already know how to work with Linux as an administrator, but as I said it's optional and you will not see any questions about this on exams in this class. I also did not see any questions about this in the Practice Questions for this section, but that could change. That being said, Testout included this information in this Labsim because you may see a question about this if you take a certification exam.

4.6 NTP

This is an application protocol, so it operates further up the network stack than most of what you're learning about now.

4.6.3 Configure NTP on Linux

This video is optional as it is meant for Linux systems and users, which means there's a whole boat load of background information you will need to understand what this is talking about. You may choose if you already know how to work with Linux as an administrator, but as I said it's optional and you will not see any questions about this on exams in this class. I also did not see any questions about this in the Practice Questions for this section, but that could change. That being said, Testout included this information in this Labsim because you may see a question about this if you take a certification exam.

4.7 IPv6

This section is optional. While it's good to know about IPv6, the use of non-routable IP addresses has extended the life of IPv4. And, most of networking is still built around IPv4, so it would be better to learn as much as you can about IPv4 first. Once you have a firm grasp of how IPv4 works you'll be able to quickly pick up IPv6.

4.8 3 Practice Questions

There's one question about IPv6 addresses that may be confusing if you skipped the IPv6 section. In any case, make your best guess, and then repeat the quiz if you didn't get the correct answer.

4.9.3 Use the ip Command

This video is optional as it is meant for Linux systems and users, which means there's a whole boat load of background information you will need to understand what this is talking about. You may choose if you already know how to work with Linux as an administrator, but as I said it's optional and you will not see any questions about this on exams in this class. I also did not see any questions about this in the Practice Questions for this section, but that could change. That being said, Testout included this information in this Labsim because you may see a question about this if you take a certification exam.

4.9.9 Practice Questions

These are great questions for checking your total understanding of how TCP/IP, DHCP, and DNS work together. Try and take your time to understand the questions and answers.

4.10.7 Practice Questions

The first question has output from the ip a command, which is a Linux command. Don't worry if you don't recognize this.

4.11.5 Use dig

This video is optional as it is meant for Linux systems and users, which means there's a whole boat load of background information you will need to understand what this is talking about. You may choose if you already know how to work with Linux as an administrator, but as I said it's optional and you will not see any questions about this on exams in this class. I also did not see any questions about this in the Practice Questions for this section, but that could change. That being said, Testout included this information in this Labsim because you may see a question about this if you take a certification exam.

4.11.7 Practice Questions

An inordinate number of the questions have to do with dig, which is a Linux utility. If you skipped the dig material these questions won't make much sense. But, since you probably want the points, you can note the correct answers, take the quiz again, and supply the correct answers the second time around.

Ways to Check Your Comprehension

The material in this section will be included in Test 2 in Labsim, which isn't due for a few weeks. If you want a couple of other ways to check your comprehension before Test 2 you can use the Practice Questions in Labsim, which you have to complete as part of your homework, or you can use the Practice Test Questions that I've set up for you.

The Practice Test Questions are completely optional, but they will provide several more hands-on labs and questions that will be similar to those on the actual test. This is a great way to get more hands-on experience, as well as a good way to check your comprehension and prepare for the real test. The Practice Test is completely optional, and you can take it as few or as many times as you want. If you do

decide to take the Practice Test you should note that it may have questions over items that were not assigned in class or questions that have to do with Linux systems.