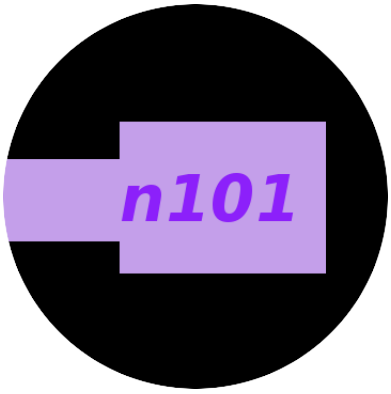

Networking 101
An Introduction to Networking

Don Colton
Brigham Young University–Hawai‘i

December 28, 2016

Preface

Note (April 2012): This book is mostly complete. A year ago it was in active development. Content is still being revised but things are mostly stable. Suggestions are welcome: email doncolton2@gmail.com

Networking 101: This is a "first things" book. What are the first things we need to know about networking? Answering that question and teaching those things to you is our goal.

Networking 101 is generally targeted towards students studying toward a bachelor's degree in Computer Science, Information Systems, or Information Technologies. It also applies to persons that are simply interested in a general introduction to data communications and networking. The depth of coverage is suitable to an introductory one-semester course that meets for about forty hours.

We focus on target concepts, topics, and skills that are normally expected of students who have passed such a course. We also cover the basic skills and concepts that support these target objectives. And we do some preparation for more advanced courses in networking.

Target Skills

So, what is typically expected of students who have passed an introductory course in networking? What would our friends and family expect of us?

(a) Help me set up my home network. I don't understand all these settings.

(b) Help me set up my wireless home network. Does it matter if I let my neighbors share it?

(c) Help me with all these firewalls and things. It sounds like a good thing,

but it just seems to get in the way, especially when I am gaming.

(d) Help me share a printer with my roommates.

(e) Help me share documents and files with my friends and family.

These are the target skills, which we could also call exit skills. They are the skills that enable students to overcome their own networking problems and to solve networking problems for others.

Skill: Internet: Students should know how the Internet works. This includes topics such as connecting to the Internet, using the domain name system and using dynamic host configuration. These basics are covered starting with chapter 1 (page 10).

Skill: Home Network: Students should be able to properly set up a home network. Home Networking is covered starting with chapter 12 (page 75).

Skill: Wi-Fi: Students should be able to properly set up wireless networking. This includes channel selection, WEP/WPA, SSID, and antenna considerations. Wireless Networking is covered starting with chapter 19 (page 117).

Skill: Security: Students should know what security they have and whether they need more. Includes password selection, firewalls, and issues with opening up ports for gaming or whatever. Does sharing our Wi-Fi put us in any danger? This is covered starting in chapter 23 (page 145).

Skill: Servers: Networks often involve the sharing of printer and files. This can be done by adding network-ready printer or storage. But often it is done by sharing parts of existing computer systems, such as their printer or hard drive. How this is done depends a lot on the operating system of the computer that will be doing the sharing.

We will address these tasks in the context of Microsoft Windows. Specifically we will look at printer sharing, file sharing, and configuring ad hoc wireless networks. This is covered starting with chapter 17 (page 105).

Basic Skills

The target skills, mentioned above, are things you probably already intend to learn.

Some of the target skills can be memorized, but many cannot. Instead, they must be understood. Basic concepts come up like domain names, SSIDs, IP addresses, network masks, and ports.

The basic skills and concepts are those other things we have to learn first before we can be truly proficient in the target skills. They are things that perhaps we didn't intend to learn. They may not be as glamorous. We may not know they even exist. But we don't score baskets in basketball (the target skill) unless we are also good at dribbling (the basic skill).

Skill: Addressing: Students should understand the following concepts: IPv4 address, subnet mask, port, address class (A, B, C), MAC address, collision domain, broadcast domain, what's a LAN, what's a router, what's a switch, broadcast addresses. This is covered mostly in chapter 27 (page 179).

Skill: Theory: Understand the following basic concepts: OSI 7-layer stack, protocol data units (packets, frames, etc.), udp, tcp, arp, ports (21, 22, 25, 80, 443). This is covered starting with chapter 6 (page 36).

Skill: Power Tools: Students should be able to properly use these tools: ping, traceroute, ipconfig, dig, nmap, ssh, telnet, ftp, and Wireshark. These are the tools of the trade often used by networking professionals. Students should be either skilled or familiar with them. They are covered in chapters 31 (page 216), 32 (page 233), and 33 (page 239).

Notation

The following notation is used in this book to help identify important types of content.

Skill: This is a skill that should be acquired.
Mem: Memorize this. It is a "Must Know" item, something we must know upon completion of the course.
FYI: For Your Information. It is a "Should Know" item, something we should probably know, but we are not expected to memorize it at this time.

Q: This is a typical question that uses the skill.
A: This is a correct answer for the question just given.
A? This is an explanation of how the answer could be derived.

Test Bank

As material is covered in the book, exam questions are inserted to show what the student should be learning. These exam questions appear throughout the book, together with acceptable answers. At the back of the book, Appendix A (page 320) is a Test Bank. It repeats these same questions that appeared throughout the book, but without their answers.

The Test Bank is a way for students to test themselves by reviewing the questions and making sure they know at least one acceptable answer.

(The acceptable answers provided are often just the most simple answer I have found that covers the question adequately. There are often much more complete and accurate answers that go far beyond the minimal answer shown as "acceptable." Please forgive me for that.)

The Test Bank is also a way for teachers to be reminded of specific things that I think students should be able to answer.

In many cases the questions and answers summarize material that is presented nearby in substantially greater detail.

In other cases the questions and answers are the actual presentation of that material. This is especially true when the specific material is something simple like vocabulary, and repetition would be tiresome and redundant.

Following is the format in which questions and answers are presented.

Exam Question 1 (p.320)**:** What are target skills?
Acceptable Answer: Target skills are the skills we intend to gain.

Exam Question 2 (p.320)**:** What are basic skills?
Acceptable Answer: Basic skills are the skills we must learn before we can be good at the target skills.

The questions are hot-linked to make it easy for the student to jump back and forth between the test bank and the content chapters.

Contents

Unit I

Networking Basics

Chapter 1

Exploring The Web

Contents

For most people, the first exposure to the **Internet** comes in the form of a **web** browser. It is a living newspaper. It is a source of entertainment. It is a window into the library of the world. For retrieving content, it uses a (somewhat cryptic) system called the **URL**.

The information itself lives on servers. A server is normally another computer, somewhere else in the world, that provides services to people like ourselves. The URL is the browser's way of finding that server and requesting the content that we desire.

What is the Web? The web, or world wide web, is a common name for the Internet. But it is not the whole Internet. It refers specifically to that part of the Internet where web pages live.

What is the Internet? The Internet includes the whole collection of all web sites and other services (of which there are many) that are connected together in a world-wide network of resources and components. It is much bigger than the web, but the web is its most familiar face.

Exam Question 3 (p.320)**:** What's the difference between the Web and

the Internet?
Acceptable Answer: The Web refers to web sites. The Internet is bigger. It includes web sites and other things like network time synchronization, email, gaming, Skype (VoIP), and Google Earth.

1.1 The URL

Exam Question 4 (p.320)**:** What does URL stand for?
Acceptable Answer: uniform resource locator

URL is usually pronounced as three separate letters, like "you are ell," and not like "earl" or "ural."

Exam Question 5 (p.320)**:** What does URI stand for?
Acceptable Answer: uniform resource identifier

URI is actually more technically correct, but somehow **URL** seems to roll off the lips easier and has become the de facto name for web addresses.

`http://en.wikipedia.org/wiki/Uniform_Resource_Locator` has more on URLs.

A URL is more than just a domain name. It consists of several parts. Let's look at a simple URL.

```
http://doncolton.com/networking/2011/book.pdf
```

Exam Question 6 (p.320)**:** List in any orderthe three most common parts of a URL.
Acceptable Answer: protocol, domain, path

`http:` is the **scheme** (or **protocol**).

`doncolton.com` is the network location (or domain name).

`/networking/2011/book.pdf` is the path.

In the next several sections we will talk about each of these pieces in greater depth. URLs can also include `:` and `?` and `+` and `#` and more special "syntax" to indicate other special things. We address this in chapter 2 (page 17).

1.2 Protocols

Exam Question 7 (p.320)**:** In networking, what is a protocol?

Acceptable Answer: A protocol is a standardized method (set of rules) for communicating.

Between people of varying cultures, we follow protocol to ensure we are properly understood. This often implies respectful interaction. People can also use common sense to recognize the intentions of those with whom they communicate.

Similarly in networking a protocol is a standardized method of communicating. The standards tend to be very specific because, in contrast to human protocols, we normally cannot rely on common sense to figure out what was meant.

http is probably the most common **protocol** on the Internet.

Exam Question 8 (p.320)**:** What does HTTP stand for?
Acceptable Answer: hyper text transfer protocol

http is the language with which the browser will talk to the server.

Exam Question 9 (p.321)**:** List in any order five popular protocols.
Acceptable Answer: http: hyper text transfer protocol.
file: a file right on the user's local computer.
https: hyper text transfer protocol with security.
mailto: email address (and possibly more).
ftp: file transfer protocol.

1.3 Domain Names

`doncolton.com` is an example of a domain name. It is intended to identify a specific computer that contains or has access to the resources we want.

There is a whole system of domain names. The most popular set is called the **dot com** domain names. These are used by nearly every business that has a presence on the Web.

Other common sets are **dot edu** for schools, **dot org** for organizations, **dot net** for networking entities, and **dot us** and its brethren, the country codes for political units around the world.

Exam Question 10 (p.321)**:** Does capitalization matter with domain names?

Required Answer: no

Exam Question 11 (p.321)**:** Does capitalization matter with URLs?
Required Answer: yes

In some parts of the URL, such as the path or file name, capitalization matters. In other parts, such as the domain name, it does not.

Because capitalization does not matter with domain names, it is all the same to the network if we can say "doncolton.com" or "DonColton.COM".

This can confuse people because they see the URL written with different capitalization in the domain name and they assume they can use different capitalization everywhere in the URL. Not true.

Exam Question 12 (p.321)**:** What is the structure of a domain name?
Acceptable Answer: little-endian

Domain names are hierarchical. The most specific, least substantial part, or little end, comes first. The later parts own the earlier parts.

The terms **little-endian** and **big-endian** are discussed in section 1.5 (page 14).

The domains x.abc.com and y.abc.com are related to each other. Both are sub-domains of (owned by) abc.com.

The domains abc.com and abc.org might not be related to each other. Or they might be. We cannot tell from the domain name.

Chapter 4 (page 25) looks at the domain name system in greater detail.

1.4 Paths

`/networking/2011/book.pdf` is an example of a path.

Exam Question 13 (p.321)**:** Does capitalization matter in the path portion of a URL?
Required Answer: yes

In contrast to domain names, frequently capitalization DOES matter with the path. "book.pdf" is usually not the same as "Book.pdf" or "Book.PDF". It could be, but often it is not. This can be confusing.

Exam Question 14 (p.321)**:** What is the normal structure of the path?
Acceptable Answer: big-endian

Paths are normally hierarchical. The most general, most substantial part, or big end, comes first. The later parts are subdivisions of the earlier parts.

By design, the path is intended to be hierarchical. Strictly speaking, the path is not actually required to have any special meaning. It is just the name of the object we wish to retrieve.

`/networking/2011/` would usually be the directory or folder on the server.

`book.pdf` would usually be the filename on the server. If we decide to save a copy of it on our local computer, this would probably be name suggested by the browser.

`.pdf` would usually indicate the type of file at that location.

Because path names were historically built out of folder names and file names, and often still are, we can sometimes trim off the last piece and retrieve again. This may get us additional information. For example, if "/networking/2011/book.pdf" is the path we are given, we could try to retrieve just "/networking/2011/" or "/networking/" or "/". We might find a directory containing several items, including "book.pdf" as well as "answerkey.pdf" and "quiz1.pdf". Or maybe not. But it may not hurt to look.

1.5 Endians

In computing and networking it is common to see the words big-endian and little-endian to describe hierarchical relationships.

Exam Question 15 (p.321)**:** What does big-endian mean? Give an example.

Acceptable Answer: Big-endian indicates (1) a hierarchical relationship among parts, (2) with the most broad, inclusive, substantial, general end coming first. The later parts are subdivisions of, owned by, or contained in the earlier parts. Example: year-month-day.

In human communication, it is common that the most important or significant things are stated first. Big-endian means that the general end comes first (because the big end is more important or significant). Little-endian means that the precise end comes first (because the little end is more important or significant). For this reason, you should not rely on words like important or significant in explaining endians.

Also, avoid words like "big" or "little" because they just create circular definitions.

Numbers are big-endian. In the number 123, the 1 represents 100s, the 2 represents 10s, and the 3 represents 1s.

Telephone numbers are big-endian. In the number 808-675-3478, the 808 is an area code. Within that area code, 675 is a telephone exchange. Within that exchange, 3478 is an individual telephone number.

Times are big-endian. 12:47:13 means the hour 12, and within that, the minute 47, and within that, the second 13.

USA postal codes (zip codes) are big-endian. 96762 is within the 90000 zone, meaning the western USA. 967 is the state of Hawaii, except Honolulu. 96762 is the town of Laie, within non-Honolulu Hawaii, within the western USA.

/courses/networking/2011/winter/textbook is an example of a big-endian file name. File names are generally written in a big-endian notation. "courses" would be the outer folder, and it includes another folder, "networking," that includes another folder, "2011," that includes another folder, "winter," that includes a file, "textbook."

Scientific dates are big-endian, as in 2011-03-15, where 2011 is the year, and within that 03 is the month, and within that, 15 is the day. This is handy for sorting.

Exam Question 16 (p.321)**:** What does little-endian mean? Give an example.

Acceptable Answer: Little-endian indicates (1) a hierarchical relationship among parts, (2) with the most precise, specific, narrow end coming first. The later parts are increasingly more broad, inclusive, and general. Later parts own or contain the earlier parts. Example: day-month-year.

Geographical locations are often little-endian. In Laie, Hawaii, USA, the smallest part, Laie, is listed first. It exists within Hawaii, the second part, which exists within USA, the third part.

European dates are little-endian, as in 15 March 2011, or 15/3/2011, where 2011 is the year, and within that 03 or March is the month, and within that, 15 is the day.

The advantage of little-endian notation is that it is the most easy notation to shorten when it is relative to some obvious base. Time (basis is now) and

location (basis is here) are good examples. If it is currently March 5, and I want to refer to March 15, I can say "on the 15th" and the current month and year will be understood to be the base of reference.

Exam Question 17 (p.321)**:** When is little-endian better?

Acceptable Answer: Little-endian is better when (a) the item is relative to some obvious base of reference. Then only the difference from the basis is needed and (b) the obvious part is often omitted.

Exam Question 18 (p.321)**:** What does mixed-endian mean? Give an example.

Acceptable Answer: Mixed-endian indicates a hierarchical relationship among parts, with a cultural expectation about the ordering of the parts, but being neither big-endian nor little-endian. Example: month-day-year.

USA dates are mixed-endian, as in March 15, 2011, or 3/15/2011, where 2011 is the year, and within that 03 or March is the month, and within that, 15 is the day.

Addresses can be mixed-endian, as in 531 Main Street, Apt 3-B. In this example, Main Street is the largest item. It contains 531, a building on Main Street. 531 contains apartments, of which 3-B is one of them.

Chapter 2

Parts of the URL

As we mentioned above, URLs can also include `:` and `?` and `+` and `#` and more to indicate special things. These act as signposts in finding the meaning of the URL.

http://en.wikipedia.org/wiki/URI_scheme gives more information about the various parts of the URL.

Exam Question 19 (p.321)**:** What is syntax?
Acceptable Answer: how we say something

Syntax is the form or format taken by language. It is the "how."

Communication can be divided into (a) what we mean and (b) how we say it. What we mean, the meaning, is called **semantics**. How we say it, the format or form, is called **syntax**. Between form and substance, form is syntax and substance is semantics.

The URL is divided into parts. These parts can be identified by special characters that act as markers, introducing or separating various pieces. These markers are part of the syntax of URLs.

Exam Question 20 (p.321)**:** In a URL, where does @ (at) go?
Acceptable Answer: after username, before domain name

The @ is a syntax marker that comes after the username (and password) and before the domain name. It appears this way in mailto URLs. It also appears with certain protected web sites.

Exam Question 21 (p.321)**:** In a URL, where does : (colon) go?
Acceptable Answer: three places: (a) after scheme, (b) between username

and password, (c) between domain name and port

The : is a syntax marker that can appear in three places. It appears after the scheme (protocol). It appears between username and password, if these are present. It appears after the domain name, if a port is specified.

Normally the port number is not included in a URL because it can be guessed from the protocol.

We talk more about ports in chapter 30 (page 208).

Exam Question 22 (p.321)**:** In a URL, where does ? (question mark) go?
Acceptable Answer: after path, before query

The ? is a syntax marker that comes after the path and introduces the query, which consists of name=value pairs.

The **query** can specify names and values to be sent to the server. This is most often seen with search engines like Google.

Exam Question 23 (p.321)**:** In a URL, where does = (equals) go?
Acceptable Answer: in query, after name and before value

The = is a syntax marker that connects names and values in the query section of the URL, like name=value, or like a=5.

Exam Question 24 (p.321)**:** In a URL, where does & (ampersand) go?
Acceptable Answer: in query, between name=value pairs

The & is a syntax marker that appears between the name=value pairs of a query, like name=value&a=5.

Exam Question 25 (p.321)**:** In a URL, where does ; (semi-colon) go?
Acceptable Answer: in query, between name=value pairs

The ; is a syntax marker that appears between the name=value pairs of a query, like name=value;a=5. It serves the same role as the &.

Exam Question 26 (p.321)**:** In a URL, where does + (plus) go?
Acceptable Answer: in place of space

The + is a syntax marker that replaces the space character. Spaces are not allowed in URLs.

Exam Question 27 (p.321)**:** In a URL, where does % (percent) go?
Acceptable Answer: (a) introduce a percent code. (b) percent codes replace characters that have special meaning

The % is a syntax marker that introduces a hexadecimal percent code. (b)

These codes are used in place of characters that would otherwise have special meaning, like syntax markers.

Many characters, including syntax markers themselves, have special meaning in a URL and therefore are not permitted to appear directly as part of a name or value. These characters are transmitted as percent codes.

When **percent codes** reach the web server, they are then converted back to their original form.

Exam Question 28 (p.321)**:** In a URL, where does # (hash) go?
Acceptable Answer: at the end of the URL, right before the fragment id

The # is a syntax marker that introduces the fragment ID. It tells the browser what part of a web page to jump to.

The **fragment** is useful in large web pages to jump right to the section of interest, even if it is in the middle of the page. Specifically it matches an "id" in the HTML of the web page.

If we say `#abc` at the end of the URL, and `id='abc'` someplace in the web page, the browser will try to put that id at the top when the page is displayed.

A really complicated URL could look something like this.

```
http://u:p@example.com:8080/path?query=123&x=5#abc
```

In this example, `http` is the protocol or scheme, `u` is the username, `p` is the password, `example.com` is the network location or domain name, `8080` is the port number, `/path` is the path, `query=123` is the first keyword/value pair of the query, `x=5` is the second keyword/value pair of the query, and `abc` is the fragment id.

Chapter 3

How The Internet Works

Contents

This chapter provides an overview, from a networking point of view, of how the Internet works.

Computers belong to local area networks. Local area networks belong to the Internet.

Networking means that information is being passed from one computer to another. The information is sent in packets, each consisting of about 1500 bytes (characters) of data and another 20 or so bytes of control information.

3.1 The Meaning of Internet

The "inter" part of **Internet** means "among" or "between." The "net" part of Internet means "network." Basically the Internet is a very large collection of local area networks.

We sometimes discriminate between big-I Internet and little-i internet. With a big I, it means "the" Internet, meaning the world-wide network of networks

that connects nearly all computers. With a little i, it means "an" internet: any other network of networks. A little-i internet might be operated by a large organization such as a military or a corporation.

Often a little-i internet is called an **intranet**. The "intra" part of intranet means "within."

3.2 Local Area Networks

A typical computer belongs to a single local area network.

Each local area network, or **LAN**, can have its own rules for communication. Today (2011) most LANs use the **Ethernet** protocol for communication. In the past other protocols have been popular, including Apple's AppleTalk and Novell's IPX.

Exam Question 29 (p.322)**:** What does LAN stand for?
Required Answer: local area network

Exam Question 30 (p.322)**:** What protocol do most LANs use for communication?
Required Answer: ethernet

TCP and IP are used by the Internet, at layer 3. LANs work at layer 2.

http://en.wikipedia.org/wiki/Ethernet has more on Ethernet.

Chapter 6 (page 36) has more on layers.

When a computer wants to communicate with another computer, the first question to resolve is whether they are on the same LAN or not. If they are on the same LAN, then the rules for that LAN are used for delivery of the message.

If the sender and receiver are not on the same LAN, the message gets passed to a router.

3.3 Routing Between LANs

A computer that belongs to more than one local area network can act as a router, and can pass messages between those LANs.

The moving of a message from one LAN to an adjacent LAN is called a hop.

It is normal for messages to make ten or more hops before they reach their destination.

The router must send the packet on its next hop toward its ultimate destination. To do this, each router must belong to two or more LANs.

Routers talk to their neighboring routers. They learn what networks each one can reach. This enables them to pick the best direction for the next hop.

Chapter 36 (page 258) goes deeper into routing.

3.4 IP Addresses and Ports

Each computer on the Internet has an **IP** address.

Exam Question 31 (p.322)**:** In networking, what does IP stand for?
Required Answer: internet protocol

Each program on that computer, if it wants to receive messages, has a **port** number.

When we talk to a computer, we don't just talk to a computer. We talk to a computer program inside that computer. We specify the program by specifying a port number. Port numbers range from 0 to 65535.

Exam Question 32 (p.322)**:** What is a software port?
Acceptable Answer: number telling which program should get the message

A software port is a number that indicates which computer program should receive the message.

We talk more about ports in chapter 30 (page 208).

Today (2011) the prevailing IP address structure is called **IPv4**, for version 4. There is a new structure called **IPv6** that is expected to become popular, but the pace of change has been very slow. In this book we will mostly ignore IPv6 until chapter 43 (page 296) and focus on IPv4.

`http://en.wikipedia.org/wiki/Ipv6` introduces IPv6.

When we refer to an IP address without saying whether it is IPv4 or IPv6, we mean IPv4.

The IP address consists of four numbers connected by dots. Each number

can be between zero and 255. While 255 might seem like a strange limit, it is the largest number that can be expressed in binary using eight bits.

Based on the IP address of the sender and the receiver, it can be told whether they are on the same local area network or not.

If they are on the same LAN, the packet can be delivered immediately. If not, the packet gets sent to a router.

3.5 Fundamental Principles

The design of the Internet is quite remarkable. Following are some fundamental principles on which that design is based. These principles help explain why things are as they are.

Probably the number one consideration is survivability. We must survive hackers and terrorists. We must survive the future improvements of technology.

Survivability. By this we mean that enemies cannot take it down by destroying one or two well-placed pieces at the core. Removing a leader may destroy a small organization, but removing any item from the Internet must not destroy it.

Exam Question 33 (p.322)**:** What is Survivability?
Acceptable Answer: ability to still function but with reduced performance when the network is partly destroyed

Survivability is the ability to continue functioning, perhaps with reduced performance, if part of the network has been destroyed.

Avoiding Centralization. By this we mean that although there are definite economies to be had by consolidating power in at the top, this flies in the face of Survivability. If there is a "top" then we can destroy it and the rest will die. That must not be allowed.

Exam Question 34 (p.322)**:** Why does the Internet avoid centralization?
Acceptable Answer: improve survivability

Centralization is a threat to survivability.

Local Autonomy. So far as it is possible, each piece of the Internet must be able to grow at its own pace, and with local direction. Central registration and central decisions should not be necessary, except rarely.

Layered Approach. It should be possible to change the way something is done without messing up how everything else is done. Wired and wireless may be great for yesterday, but if new connection technologies like 4G or satellite or fiber-to-the-home are invented, it should be easy to slot them in without rewriting the rest of the Internet.

Scalability. In 2000, Businessweek wrote that Internet traffic is doubling every three months. (Oct. 9, 2000).

http://newsroom.cisco.com/dlls/2008/ekits/Cisco_Visual_Networking_Index_061608.pdf is an analysis by networking giant CISCO Systems. As of 2008, they estimate traffic will double every two years for the next five years.

Wildly exponential growth cannot go on forever. Growth seems to be slowing, but is still very impressive. We are through the early adoption phase in many places. The technologies are starting to mature. Eventually the world will be saturated by networked devices, as every person on the planet has their own computer of the future. But those days are still in front of us.

The Internet must be built to withstand massive growth for many years into the future.

Self-healing. This is another form of survivability. But this time we mean at the edges. Networks should be self-healing. Add a new piece. Take something away. The network should notice and adjust. It should not be necessary for a skilled human to fix things.

Chapter 4

Domain Names and DNS

Contents

Domain names are used extensively with the Internet. It is important to understand domain names and how they work. The most important reason is to avoid being scammed.

`doncolton.com` is a domain name.

DNS is the Domain Name System.

Domain names are normally intended to identify a specific computer that contains or has access to the resources we want. For larger situations, they identify a group of machines, any one of which can do the same thing.

There is a whole system of domain names. The most popular set is called the **dot com** domain names. These are used by nearly every business that has a presence on the Web.

Exam Question 35 (p.322)**:** What does DNS stand for?

Required Answer: domain name system

Exam Question 36 (p.322)**:** What service does DNS provide?
Acceptable Answer: converts domain names into IP addresses

Exam Question 37 (p.322)**:** Is there any special meaning to the order of the parts in a domain name? If so, what?
Acceptable Answer: yes. they are little-endian.

By this we mean they are expressed with the most specific part first, and the most general part last. The parts are separated by dots.

Example: fourth.third.second.first.root

By fourth, we mean the fourth level. First is the top level domain. Root is the basis that leads to all the other levels. Every browser knows in advance where the root level is, but for all the other levels, it consults the next higher level to find it.

It is similar to the way a geographical location might be named:

Laie, Hawaii, USA

Laie is part of Hawaii. Hawaii is part of USA. USA is a primary name, part of the geographical universe.

n101.doncolton.com

n101 is part of doncolton. doncolton is part of com. com is a primary name, part of the DNS universe.

4.1 The Root Domain

The root level of the domain name system consists of 13* special servers located around the world. They each contain a small amount of information, about 1/5 of a megabyte, or 4% of the size of a typical photograph. But that information leads to all the top level domains in the world. When a new top level domain is created or moved, which is very rare, the root servers must be updated.

Having 13 servers helps ensure that the loss of any one of them will not disable the whole system. It supports the fundamental principle of survivability.

* http://en.wikipedia.org/wiki/DNS_root has more information about

the root zone of the dns system. It points out that the original 13 servers have been joined by many more that also serve as roots.

4.2 Top Level Domains

The dot com piece of the domain name is called the Top Level Domain, or **TLD**. There are only a few of these. New ones are strictly controlled and should be rare.

Other common TLDs are **edu** for schools, **org** for organizations, **net** for networking entities, and the two-letter country codes such as **us** for the United States.

http://en.wikipedia.org/wiki/List_of_Internet_top-level_domains has a complete list of top level domains, including country domains.

Top level domains cannot be purchased. Setting them up can be a very long and involved political process.

Skill: Know the important terminology relating to networking.

Exam Question 38 (p.322)**:** What does TLD stand for?
Required Answer: top level domain

4.3 Second Level Domains

Below each top level domain, it is generally possible to purchase a second level domain from the appropriate domain registrar.

We say "generally" because some TLDs are more restrictive. .edu, for example, will only provide domain names for educational institutions.

Example: doncolton.com is a second level domain purchased from a registrar. The com part is the top level domain. The doncolton part is the second level domain.

Second level domains are typically owned by organizations, and represent an organization or similar entity.

There is often very little restriction on who can purchase second level domains, and they could be owned by anyone including people pretending to be someone else.

Example: microsoft.com is owned by Microsoft Corporation. micro.soft.com looks similar but is owned by Software Research, Inc. There is no deception here because it is assumed people will know the dots are very important in defining who owns the web site.

Trademark can come into play. If someone owns a trademark in a certain name and someone else registers it for their own domain, there are administrative procedures by which the trademark owner can gain control of it. This helps to ensure that domains like microsoft.com are really owned by Microsoft Corporation, and not by an imposter.

Skill: Know the important terminology relating to networking.

Exam Question 39 (p.322)**:** What is Cyber Squatting?
Acceptable Answer: It is the practice of registering domain names that people would probably assume are really owned by someone else.

Deception can still be a problem. It seems common for organizations to contract with other service providers for parts of their web presence.

Example: `boh.com` is the domain name for the Bank of Hawaii website, where I do some of my online banking. `cibng.ibanking-services.com` claims to also be an authorized Bank of Hawaii website. Should I believe that? On the face of it, no. It could be a fake. But it actually is an authorized BOH website. Through authorized contractors, that kind of naming often happens. I believe it creates a dangerous situation because it helps to condition people to trust anything that looks legitimate.

4.4 Domain Registrars

Domain registrars provide the service of selling available domain names to those qualified to buy them.

Exam Question 40 (p.322)**:** Can someone own a domain name?
Required Answer: no

Domain names are leased for one or more years at a time.

For now, the correct answer is to say that domain names cannot be owned.

Actually, the courts are wrestling with this in some places. It boils down to this: you own an exclusive license to use that domain name. But it is virtual property, not physical property. This creates some odd difficulties

at law.

It is common to pay around $10 per year for a domain name, depending on the TLD and the registrar.

http://en.wikipedia.org/wiki/Domain_registrars has more on registrars, including links to lists of registrars.

One of the important services provided by the registrar is the associating of your domain name with the DNS server that you are authorizing and requesting to give out your IP address when someone wants it.

Otherwise there is danger of identity theft, or domain hijacking, where customers to your web site can be directed to a copy-cat site.

Because of such dangers, registrars take great care with DNS issues.

4.5 Effective Top Level Domains

Some second level domains (and lower) are effectively the same as top level domains. For example, in the United Kingdom, the domain name ".co.uk" is a second level domain name that is treated a lot like a top level domain name. It is the UK version of the ".com" top level domain.

Such second (and lower) domain names are called effective top level domains because new domains can be registered under them.

An **effective top level domain** is also called a **public suffix**.

This starts to be important because of web **cookies**, a method used to share information between related web sites. Recognizing public suffixes is an important way to avoid unauthorized sharing of cookie information between unrelated web sites.

4.6 Sub Domains

After you acquire a domain at any level, you are generally permitted to create any number of sub domains. You can use letters and digits and a few special characters such as the dash (hyphen).

Example: soft.com is a second level domain. micro.soft.com is a third level domain. It belongs to the soft.com domain.

There are generally no controls over sub domains except that you must first control the higher level domain.

Example: If I own doncolton.com, I could create pentagon.doncolton.com, even if I have no association with the Pentagon. Similarly, I could create citibank.doncolton.com even if I have no association with Citibank Corporation.

Exam Question 41 (p.322)**:** Is byuh.doncolton.com controlled by BYUH? Why or why not?

Acceptable Answer: No. Sub domains are controlled by the next higher domain.

Sub domains may have no relationship to more famous domains or brands that they look like.

Each sub domain is still a domain, and it can have further sub domains of its own. So doncolton.com can have the sub domain n101.doncolton.com, which can have the sub domain testbank.n101.doncolton.com, which can have a further sub domain.

4.7 The www Sub Domain

Often domain names start with **www**, which stands for World Wide Web. It is not a requirement, but it is probably the most familiar way to introduce a URL, instead of saying "http://".

Example: Instead of saying http://disney.com/, the Walt Disney Company may find it more desirable to simply say www.disney.com. It is a marketing thing, and it is very common.

Strictly speaking, www.somedomain.com does not have to lead to the same web pages as somedomain.com without the www, but practically speaking they almost always lead the same place.

4.8 DNS Resolution

DNS, the domain name system, is a service that converts domain names into IP addresses.

Annual fees paid by second level domain holders are used, among other

purposes, to fund the central part of the domain name system. The central part consists of the root servers.

http://en.wikipedia.org/wiki/Root_servers has more information on the root servers.

Chapter 33.2 (page 240) discusses the Domain Information Groper, a command that resolves domain names into IP addresses and tells how the resolution was verified.

Basically the resolution works like this:

The root servers convert the last chunk of a domain name, that is, the top level domain, into the identity of a machine.

That machine converts the next chunk into the identity of another machine.

The process is recursive. By that we mean it continues until the last chunk is converted, giving the identity of the final machine.

For example, n101.doncolton.com is resolved by checking the root server. It returns the IP address for the .com DNS server. The .com server returns the IP address for doncolton.com, which returns the IP address for n101.doncolton.com.

We will talk much more about IP addresses, starting in chapter 27 (page 179) below.

Chapter 5

DHCP: Host Configuration

Contents

Before a computer can use the Internet, it needs a few things. It needs to have an IP address. It needs to know the **gateway** (router) that can forward messages beyond the local area network. It needs to know the location of certain services, such as DNS. **DHCP** is generally used by computers to find out these configuration details when they first boot up.

The requesting computer will make a DHCP request by way of a general broadcast asking for configuration information.

Exam Question 42 (p.322)**:** What does DHCP stand for?
Required Answer: dynamic host configuration protocol

5.1 Host

Exam Question 43 (p.322)**:** What is a host?
Acceptable Answer: any computing device on the network

Host is just another name for computing device. It might be a server, a desktop computer, a laptop, a tablet, a smart phone, or even a router.

5.2 Configuration

Getting information into the right place is called **configuration**.

For a computer to make requests and receive replies over the Internet, it must have an IP address.

Exam Question 44 (p.322)**:** What is configuration?
Acceptable Answer: settings that control how something operates

Configuration adjusts the settings that control something, such as a computer.

Driving example: Programming is knowing how to drive the car. Configuration is telling it where to go.

Basketball example: Programming is knowing how to shoot a basket. Configuration is knowing which basket to shoot at.

Getting this particular information into the right place on the host computer is called **host** configuration.

5.3 Dynamic v Static

Configuration by hand is called **static** configuration, since it tends to be stable across time, not changing until human intervention happens.

Automatic configuration is called **dynamic** configuration, since it can change without anyone noticing.

Dynamic (automatic) does not require the users to understand what is going on or to alter settings or make decisions.

Exam Question 45 (p.323)**:** What does dynamic mean?
Acceptable Answer: can change settings on its own

Dynamic means something that can change its settings automatically, with no outside intervention. It is the opposite of static.

Exam Question 46 (p.323)**:** What does static mean?
Acceptable Answer: does not change on its own

Static means something that does not change unless there is outside intervention. It is the opposite of dynamic.

5.4 Typical Scenarios

In the early days, computers were few and big and did not move. It was practical to assign these addresses by hand.

With the advent of laptop computers and other portable devices, and further driven by large collections of end-user computer workstations, configuration by hand has become very unpopular and most configuration is automated.

For servers, this configuration information is often **hard coded**, by which we mean that someone actually typed it in and saved it on that computer. When the computer first turns on, it checks for the information. If it finds it, it is happy and life moves forward.

For most other computers, and especially for portable computers like laptops and tablets, this information changes depending on what networks are nearby. The computer has to ask for the information. This asking is done by way of a **DHCP** request.

In effect, the computer asks "Please tell me my name and where I can find the front door." Or, more accurately, please tell me my IP address, net mask, local gateway, and location of the DNS servers I should use.

Exam Question 47 (p.323)**:** What does DHCP provide? Include a specific example.

Acceptable Answer: (a) DHCP provides parameters necessary to use the network. (b) Specific examples include IP address, net mask, gateway, and DNS servers.

Exam Question 48 (p.323)**:** How does a typical laptop computer discover its own IP address?

Acceptable Answer: dhcp

Exam Question 49 (p.323)**:** How does a typical server computer discover its own IP address?

Acceptable Answer: Server IP addresses are normally static. The IP address is stored within the server. It does not have to ask for it.

Servers have static IP addresses to prevent accidental changes. If the IP address of a server is changed, then the DNS must also be updated to properly direct requests for its services.

DHCP can be configured to always give the same IP address to certain machines, based on their MAC addresses.

Unit II

OSI Model

Chapter 6

The OSI Model

Contents

The **OSI** networking model is the standard way of thinking about and talking about the different layers that make up the Internet.

Each layer has specific duties. By keeping things in separate layers, changes and upgrades are more easily made. Better local area networking can be implemented without redesigning the whole Internet.

This is a good thing.

Students should be able to apply the OSI model to networking hardware and software. Mostly we are concerned with the bottom four layers.

`http://en.wikipedia.org/wiki/OSI_model` has more.

6.1 Open Standards

OSI stands for Open Systems Interconnection. It provides an open standard for connecting networks together. This is important.

Vendors (makers and sellers) want your business. Once they get you to buy something, they want to own you. They want all of your business. They don't want you buying anything from their competitors.

One way they do this is by making it so competing products just won't work. They can do this in many different ways. Mostly they do it by having a **proprietary** (private or secret) interface or by digitally signed, certified components. Proprietary means they own it and nobody else can use it. I try to avoid proprietary things whenever I can.

Imagine buying a car, and then finding out you must buy gasoline and tires and everything else for your car from that same source. Life is simple. All your future choices have been made for you. And there is no competition so they can charge whatever they want.

In the early days of computing, this was common. Diskettes (the precursor to USB drives) could only be used on the machine where they were recorded. Maybe they could be used on other machines by the same manufacturer. But usually they could not be used on machines by other manufacturers.

To get a file from an IBM Personal Computer to an Apple Personal Computer was very difficult. Nearly impossible. Sharing information was difficult.

As networks started to be available, there was a serious effort to agree on common rules so that every computer could connect to and communicate with every other computer.

6.2 Seven Layers of the Stack

The OSI networking model consists of seven layers: application, presentation, session, transport, network, data link, and physical. They are numbered from 7 down to 1.

http://en.wikipedia.org/wiki/OSI_model has more on the OSI model.

These layers are also called the network **stack**. Specifically, the **network stack** is the collection of software that makes it possible to connect with

other computers on the Internet.

Skill: Know, by number and by name, the seven layers of the OSI model.

Application Presentation Session Transport Network Datalink Physical

There are phrases (called mnemonics) that can be used to remember the order of the layers. You can make up your own. Or pick one of these phrases and use the first letter of each word to remind yourself of the OSI layer that starts with the same letter. Here is the most popular phrase I have encountered.

a p s t n d p :: All People Seem To Need Data Processing

http://en.wikipedia.org/wiki/List_of_mnemonics has several listed under Computing.

Skill: Know which common networking items match other layers.

6.3 OSI Layer 7: The Application Layer

Exam Question 50 (p.323)**:** What is layer 7 of the OSI model?
Required Answer: application

Exam Question 51 (p.323)**:** What layer number is the Application layer of the OSI model?
Required Answer: 7

Layer 7, the top layer, is the application layer. We will mostly ignore it.

6.4 OSI Layer 6: The Presentation Layer

Exam Question 52 (p.323)**:** What is layer 6 of the OSI model?
Required Answer: presentation

Exam Question 53 (p.323)**:** What layer number is the Presentation layer of the OSI model?
Required Answer: 6

Layer 6 is the presentation layer. As data flows down the stack on the sending end, this layer converts from the character set of the local computer into a general network character set. It also provides for any **compression**

or **encryption** that might be taking place. As data flows up the stack on the receiving end, this layer converts from the general network character set into the character set of the receiving computer, which may be different from the character set of the sending computer.

Exam Question 54 (p.323)**:** At which OSI layer is encryption / decryption?
Required Answer: 6 or presentation

Exam Question 55 (p.323)**:** At which OSI layer is data compression?
Required Answer: 6 or presentation

6.5 OSI Layer 5: The Session Layer

Exam Question 56 (p.323)**:** What is layer 5 of the OSI model?
Required Answer: session

Exam Question 57 (p.323)**:** What layer number is the Session layer of the OSI model?
Required Answer: 5

Layer 5 is the session layer. We will mostly ignore it.

As originally designed, this layer divides the data up into segments. Pieces are sent across the network. Checkpoints occur from time to time so if the connection fails, a restart is possible. As data flows up the stack, this layer collects the segments that have arrived and places them in order. When all the segments have been received, it passes the assembled data chunk up the stack.

In reality, TCP at layer 4 ends up handling this.

6.6 OSI Layer 4: The Transport Layer

Exam Question 58 (p.323)**:** What is layer 4 of the OSI model?
Required Answer: transport

Exam Question 59 (p.323)**:** What layer number is the Transport layer of the OSI model?
Required Answer: 4

Layer 4 is the transport layer.

As data flows down the stack, this layer divides the data up into segments. The maximum size of a segment is normally 1500 characters. It is called the MTU, or Maximum Transmission Unit. The segments are numbered and sent on down the stack, one by one. As data flows up the stack, this layer collects the segments that have arrived and places them in order. When all the segments have been received, it passes the assembled data chunk up the stack to layer 5.

Exam Question 60 (p.323)**:** What does MTU stand for?
Acceptable Answer: maximum transmission unit

MTU is the number of bytes of data that can be sent in a single packet. Normally this number is 1500. If you try to send more at once, TCP will break it up into smaller pieces to stay within the MTU requirements.

Exam Question 61 (p.323)**:** What is the typical value for MTU (in bytes)?
Acceptable Answer: 1500

TCP and UDP also define the rules for ports on your computer. Ports help to separate the network traffic that is coming to your computer. This is important because computers are required to be able to carry on several conversations at the same time. Ports help to keep the conversations separate.

Exam Question 62 (p.323)**:** At which OSI layer are software ports?
Required Answer: 4 or transport

Exam Question 63 (p.324)**:** What does TCP stand for?
Required Answer: transmission control protocol

Exam Question 64 (p.324)**:** Which protocol provides for guaranteed delivery of information?
Required Answer: tcp

TCP, the Transmission Control Protocol, provides guaranteed delivery through retransmission as needed.

Exam Question 65 (p.324)**:** At which OSI layer is TCP?
Required Answer: 4 or transport

TCP is suitable for file transmission where accuracy is more important than avoiding delays.

`http://en.wikipedia.org/wiki/Transmission_Control_Protocol` has

more on TCP, the Transmission Control Protocol.

Exam Question 66 (p.324)**:** Which protocol provides for fast (but not guaranteed) delivery of information?
Required Answer: udp

UDP, the User Datagram Protocol, provides fast delivery but does not guarantee delivery.

Exam Question 67 (p.324)**:** What does UDP stand for?
Required Answer: user datagram protocol

UDP is suitable for live voice, live sound, and live video, where avoiding delays is more important than absolute accuracy, and dropped data can usually be ignored.

http://en.wikipedia.org/wiki/User_Datagram_Protocol has more on UDP, the User Datagram Protocol.

Exam Question 68 (p.324)**:** At which OSI layer is UDP?
Required Answer: 4 or transport

Exam Question 69 (p.324)**:** At which OSI layer do we find segments?
Required Answer: 4 or transport

Exam Question 70 (p.324)**:** What is the Protocol Data Unit at OSI layer 4?
Required Answer: segment

6.7 OSI Layer 3: The Network Layer

Layer 3 is the network layer. It defines the rules for wide area networks. Usually this is the Internet, or some piece of it involving many computers. For the Internet, this always relies on IP addresses. Routers and gateways operate at this level.

Layer 3 is sometimes called the **Internet layer**.

Skill: Know which common networking items match layer 3: Network.

Exam Question 71 (p.324)**:** What is layer 3 of the OSI model?
Required Answer: network

Exam Question 72 (p.324)**:** What layer number is the Network layer of the OSI model?

Required Answer: 3

Exam Question 73 (p.324)**:** At which OSI layer is the Internet?
Required Answer: 3 or network

Exam Question 74 (p.324)**:** At which OSI layer are Wide Area Networks?
Required Answer: 3 or network

Exam Question 75 (p.324)**:** What does WAN stand for?
Required Answer: wide area network

Exam Question 76 (p.324)**:** At which OSI layer is IP Addressing?
Required Answer: 3 or network

Exam Question 77 (p.324)**:** In networking, what does IP stand for?
Required Answer: internet protocol

Exam Question 78 (p.324)**:** At which OSI layer is Logical Addressing?
Required Answer: 3 or network

Exam Question 79 (p.324)**:** At which OSI layer does a router operate?
Required Answer: 3 or network

Exam Question 80 (p.324)**:** At which OSI layer does a gateway operate?
Required Answer: 3 or network

Exam Question 81 (p.324)**:** At which OSI layer is Network Address Translation?
Required Answer: 3 or network

Exam Question 82 (p.324)**:** At which OSI layer is Port Address Translation?
Required Answer: 3 or network

Exam Question 83 (p.324)**:** At which OSI layer do we find packets?
Required Answer: 3 or network

Exam Question 84 (p.324)**:** What is the Protocol Data Unit at OSI layer 3?
Required Answer: packet

6.8 OSI Layer 2: The Data Link Layer

Layer 2 is the data link layer. It defines the rules for local area networks. Typically this is Ethernet and relies on MAC addresses. Switches and

bridges operate at this level.

Skill: Know which common networking items match layer 2: Data Link.

Exam Question 85 (p.324)**:** What is layer 2 of the OSI model?
Required Answer: datalink

Exam Question 86 (p.325)**:** What layer number is the Data Link layer of the OSI model?
Required Answer: 2

Exam Question 87 (p.325)**:** What does MAC stand for?
Required Answer: media access control

Exam Question 88 (p.325)**:** How does a typical laptop computer discover its own MAC address?
Acceptable Answer: it is burned in on the network card

The network card is called a NIC, or Network Interface Controller.

Exam Question 89 (p.325)**:** What does NIC stand for?
Acceptable Answer: network interface controller

Exam Question 90 (p.325)**:** At which OSI layer is the Local Area Network?
Required Answer: 2 or datalink

Exam Question 91 (p.325)**:** At which OSI layer is MAC Addressing?
Required Answer: 2 or datalink

Exam Question 92 (p.325)**:** At which OSI layer is Physical Addressing?
Required Answer: 2 or datalink

Exam Question 93 (p.325)**:** At which OSI layer is Ethernet?
Required Answer: 2 or datalink

Exam Question 94 (p.325)**:** At which OSI layer does a switch operate?
Required Answer: 2 or datalink

Exam Question 95 (p.325)**:** At which OSI layer does a bridge operate?
Required Answer: 2 or datalink

Exam Question 96 (p.325)**:** What is a multi-port bridge called?
Required Answer: switch

Exam Question 97 (p.325)**:** What is a two-port switch called?

Required Answer: bridge

Exam Question 98 (p.325)**:** At which OSI layer do we find frames?
Required Answer: 2 or datalink

Exam Question 99 (p.325)**:** What is the Protocol Data Unit at OSI layer 2?
Required Answer: frame

6.9 OSI Layer 1: The Physical Layer

Exam Question 100 (p.325)**:** What is layer 1 of the OSI model?
Required Answer: physical

Exam Question 101 (p.325)**:** What layer number is the Physical layer of the OSI model?
Required Answer: 1

Layer, 1, the bottom layer, is the physical layer. It defines the rules for the physical connection between the computer and the Internet. Typically this is Wi-Fi or Cat 5 cable or Fiber-optic cable. Hubs and repeaters operate at this level.

Exam Question 102 (p.325)**:** At which OSI layer is wireless signal?
Required Answer: 1 or physical

Exam Question 103 (p.325)**:** At which OSI layer is coaxial cable?
Required Answer: 1 or physical

Exam Question 104 (p.325)**:** At which OSI layer is cat5 cable?
Required Answer: 1 or physical

Exam Question 105 (p.325)**:** At which OSI layer is fiber-optic cable?
Required Answer: 1 or physical

Exam Question 106 (p.325)**:** What is a multi-port repeater called?
Required Answer: hub

Exam Question 107 (p.325)**:** What is a two-port hub called?
Required Answer: repeater

Exam Question 108 (p.325)**:** At which OSI layer does a hub operate?
Required Answer: 1 or physical

Exam Question 109 (p.326)**:** At which OSI layer does a signal repeater

operate?
Required Answer: 1 or physical

Exam Question 110 (p.326)**:** At which OSI layer do we find bits?
Required Answer: 1 or physical

Exam Question 111 (p.326)**:** What is the Protocol Data Unit at OSI layer 1?
Required Answer: bit

Chapter 7

IP Addressing Preview

Contents

This chapter is a simplified introduction to IP addresses and ports.

IP stands for Internet Protocol. It is frequently seen in the acronym TCP/IP.

An IP address is an Internet address.

There are two major versions of IP addressing. IPv4 (version 4) was new in 1983. IPv6 (version 6) is intended to replace it but adoption has been slow. As of 2016 IPv6 is making serious headway but IPv4 still rules.

We cover IPv4 addressing in much more depth in chapter 27 (page 179). We cover ports in much more depth in chapter 30 (page 208).

We cover IPv6 addressing in chapter 43 (page 296).

7.1 Routers Use IP Addresses

The purpose of a network, such as the Internet, is to send information from one device (computer) to another.

The movement of information is done using routers. Each **router** is connected to another in an unbroken chain from sender to receiver.

Actually, once you get very far from home, each router is connected to many other routers. The router has the job of deciding where to send the message for its next **hop** across the Internet.

Routers do this by looking at the IP address. IP addresses are somewhat geographic. Based on the IP address, and the tables of addresses that it has been able to build, the Router picks the next Router and forwards the message.

7.2 What Does an IP Address Look Like?

What does an IP address look like?

The current IP address standard is called IPv4, meaning version 4. It is still in common use.

The future standard is IPv6. We will not discuss it in this book, except to say that IPv4 is running out of addresses, and IPv6 addresses are much bigger than IPv4 addresses.

Each IPv4 address is composed of four numbers, like this: 192.168.0.100.

The numbers are separated by dots.

Each number can be, at a minimum, zero, and at a maximum, 255.

The reason for the strange maximum is that 255 is the largest number you can represent using eight bits.

For comparison, in the USA, a ZIP code always starts with five digits. Each digit can be between zero and nine. The five-digit zip code identifies a specific post office someplace in the country.

ZIP codes have five digits that range from zero to nine.

IP addresses have four numbers that range from zero to 255.

7.3 Networks and Hosts

Each IP address can be broken into two pieces. The first piece is the network and the second piece is the host.

In the case of 192.168.0.100, the network is 192.168.0 and the host is 100.

The break does not always come right before the last number. The rules can be complicated. We talk about them in chapter 27 (page 179).

All machines in the same local area network have the same network number, but each machine has a different host number.

7.4 Local Area Networks

All the computers that are close to you are considered to be in your Local Area Network. We use the abbreviation **LAN** for Local Area Network.

Local computers are the ones that are close enough that you can talk to them directly.

Other computers are far away enough that you cannot talk to them directly. In that case, you have to send your message by way of a Router. The Router passes along the message, and when a response comes, it returns the response to you.

All computers in the same LAN have IP addresses that look very similar.

If your IP address is 192.168.0.100, then the other computers in your LAN will have 192.168.0.x for their IP address, where x can be any other number between 1 and 254.

Why not zero to 255? Good question. Essentially it is because the first and last address are reserved for special purposes. Zero is the network itself. 255 is called the (default) broadcast address. If you want to send a message to everyone on your local area network, you could send it to 192.168.0.255.

As we said above, we cover IP addresses in much more depth in chapter 27 (page 179).

7.5 Ports Lead to Programs

Besides having a destination IP address, messages also have a destination Port number.

When you talk to a computer, you don't just talk to a computer. You talk to a computer program inside that computer. You specify the program by specifying a port number. Port numbers range from 0 to 65535.

In effect, when the computer starts up, it also starts up a bunch of programs. Some of those are willing to receive messages. Those programs tell the computer, hey, if you get a message for port 123, send it to me.

Each program on that computer, if it wants to receive messages, has a **port** number.

Exam Question 112 (p.326)**:** What is a software port?
Acceptable Answer: A port is a number that indicates which computer program should receive the message.

Chapter 8

Converting Between Bases

Contents

It is important to be able to do **number base conversion** between the different numbering systems that are commonly used in networking.

Base two, also called binary, is the natural basis for network communications, due to the most simple communication being on/off or yes/no or true/false.

Easy: Bases that are powers of two are extremely easy to convert back and forth with base two. A power of two is any number you can reach by repeatedly multiplying 2. The powers of two are 1, 2, 4, 8, 16, 32, 64, 128, 256, and so on.

Base ten, also called decimal, is the number base most commonly used by humans, probably due to us having ten fingers.

Hard: Converting back and forth between binary number and decimal numbers is more difficult.

8.1 2→8: Convert Binary to Octal

To **convert binary to octal**, first divide the number into groups of three bits, starting at the back.

1011001010001 becomes 1 011 001 010 001.

Next convert each group into an octal digit. Leading zeroes (zeroes at the front of a number) do not matter, except for deciding where to divide into groups. The last digit is worth 1. The digit before it is worth 2. The digit before it is worth 4. Each digit is worth twice as much as the one that comes next.

1 011 001 010 001 becomes 13121.

binary	3-bit binary	octal
0	000	0
1	001	1
10	010	2
11	011	3
100	100	4
101	101	5
110	110	6
111	111	7

Exam Question 113 (p.326)**:** Convert binary 11110100010000 to octal.
Acceptable Answer: 36420

8.2 2→16: Convert Binary to Hex

To **convert binary to hex**, first divide the number into groups of four bits, starting at the back.

1011001011001 becomes 1 0110 0101 1001.

Next convert each group into a hex digit. Zeroes at the front of a number are called **leading zeroes** and do not affect the value of the number. The last digit is worth 1. The digit before it is worth 2. The digit before it is worth 4. The digit before it is worth 8. Each digit is worth twice as much as the one that comes next.

1 0110 0101 1001 becomes 1659.

binary	4-bit binary	hex	decimal
0	0000	0	0
1	0001	1	1
10	0010	2	2
11	0011	3	3
100	0100	4	4
101	0101	5	5
110	0110	6	6
111	0111	7	7
1000	1000	8	8
1001	1001	9	9
1010	1010	A	10
1011	1011	B	11
1100	1100	C	12
1101	1101	D	13
1110	1110	E	14
1111	1111	F	15

Exam Question 114 (p.326)**:** Convert binary 11111111011001011001 to hex.
Acceptable Answer: FF659

8.3 8→2: Convert Octal to Binary

To **convert octal to binary**, just convert each digit, one by one. You can start at the front or the back. It does not matter.

Make sure you use three bits per octal digit.

13121 becomes 001 011 001 010 001.

Finally, remove the leading zeroes and close up the spaces.

001 011 001 010 001 becomes 1011001010001.

Exam Question 115 (p.326)**:** Convert octal 16471 to binary.
Acceptable Answer: 1110100111001

8.4 16→2: Convert Hex to Binary

To **convert hex to binary**, just convert each digit, one by one. You can start at the front or the back. It does not matter.

Make sure you use four bits per hex digit.

1659 becomes 0001 0110 0101 1001.

Finally, remove the leading zeroes and close up the spaces.

0001 0110 0101 1001 becomes 1011001011001.

Exam Question 116 (p.326)**:** Convert hex 64209 to binary.
Acceptable Answer: 1100100001000001001

8.5 10→2: Convert Decimal to Binary

Exam Question 117 (p.326)**:** Convert decimal 162 to binary.
Acceptable Answer: 10100010

To **convert decimal to binary**, there are several ways.

My Way: Here is a method that I like.

Start with the decimal number on the left, then a comma, then a blank answer (binary number) on the right. Repeat this process: If the decimal number is even, put a 0 on the front of the answer. If the decimal number is odd, put a 1 on the front of the answer. Subtract the 0 or 1 from the base 10 number, and divide the rest by two.

162, is our starting point.
81,0 : 0 because 162 is even, half of 162 is 81.
40,10 : 1 because 81 is odd, half of 80 is 40.
20,010 : 0 because 40 is even, half of 40 is 20.
10,0010
5,00010
2,100010
1,0100010
,10100010 is our final answer.

We can short-circuit the process and end early when we know the rest of the conversion.

When we get to 5, and we know 5 is 101,

5,00010
,10100010 is our final answer.

The Other Way: Here is the most popular other method.

Make a list of powers of two until you surpass the number you are converting. 162 is our starting point, so the powers would be these:

1, 2, 4, 8, 16, 32, 64, 128, 256.

We can stop with 256 because it is greater than 162.

Reverse the list (or write the list in this order to begin with):

256, 128, 64, 32, 16, 8, 4, 2, 1.

Now, subtract the largest power of two possible from 162. The largest one is 128. 162 - 128 = 34.

Repeat the process: 34 - 32 = 2. 2 - 2 = 0.

For each number you subtracted, you get a 1. For each number you skipped, you get a zero.

256=0, 128=1, 64=0, 32=1, 16=0, 8=0, 4=0, 2=1, 1=0.

Collect the ones and zeroes:

010100010

Throw away the leading zero:

10100010 is our final answer.

Both methods are guaranteed to work.

8.6 2→10: Convert Binary to Decimal

Exam Question 118 (p.326)**:** Convert binary 10001000 to decimal.
Acceptable Answer: 136

To **convert binary to decimal**, there are several ways.

My Way: Here is a method that I like.

Start with an answer (base 10) of zero. Start with the first digit in the binary number. Repeat this process: double the (base 10) answer and add the next digit from the binary number.

We will write the base 10 number to the left of the comma, and the remaining base 2 number to the right of the comma. We start with nothing. At each step, we double the base 10 number and add in the next binary digit. (Special thanks to my student Zhiwei Hou for demonstrating this notation, which I see as an improvement over the notation I was using before.)

0,10001000 is our starting point.
1,0001000 : double 0, add 1, = 1.
2,001000 : double 1, add 0, = 2.
4,01000
8,1000 : we will double 8 and add 1 giving 17.
17,000 : we will double 17 and add 0 giving 34.
34,00
68,0
136, we have arrived at our final answer.

For a short-cut, you can immediately convert as many digits from the front as you happen to have memorized. Say you know that 1000 is 8. Then you can remove the 1000 and start with 8.

8,1000 is our starting position.
17,000 : we will double 17 and add 0 giving 34
34,00
68,0
136, we have arrived at our final answer.

The Other Way: Here is the most popular other method.

10001000 is our starting point.

Make a list of powers of two.

256, 128, 64, 32, 16, 8, 4, 2, 1.

Assign the bits to the powers, starting at the little end.

256, 128=1, 64=0, 32=0, 16=0, 8=1, 4=0, 2=0, 1=0.

Keep the numbers that got ones assigned to them.

128=1, 8=1.

Add them up. 128 + 8 = 136.

We have arrived at our final answer.

Chapter 9

Anatomy of a Hop

Contents

With your knowledge of the OSI model and of IP addresses, we can explain in greater detail how information moves across the Internet.

Looking at all this, it may seem like a miracle that it ever worked. I am sure many people felt that way when they first got it going.

You will probably notice that some layers did a lot of work, while others seem to do little or nothing. Mostly that is because we have simplified it greatly, just keeping the parts we think you need to know.

Each layer has a fairly straightforward set of responsibilities. By breaking them up using the OSI model, the program for each layer was kept as simple as possible, thus making it easier to write and debug. In general, this is a good approach to software development.

In the following sections, we will look in detail at how data moves through the network. Another interesting view of this is given in section 25.6 (page 169), where we examine a hacking legend, the **Ping of death**, that relies on the stack to do its dirty work.

9.1 Down the Stack

Information starts out as a big blob of whatever, created by some program, and destined for delivery to another program running on (probably) another computer, someplace across the Internet.

The application program hands it to layer 7, the top layer of the stack, together with the IP address and port to which it should be sent.

Layer 7 is the application layer. It verifies that everything is in order with the request, and it hands the data down the stack to layer 6.

Layer 6 is the presentation layer. It converts the data into network encoding. It also does any encryption required. It also does any data compression required. Then it hands the data down the stack to layer 5.

Layer 5 is the session layer. It pretty much just hands the data down the stack to layer 4.

Layer 4 is the transport layer. It has to do some serious work. First, it chops the entire block of data into segments of (typically) 1500 bytes, based on the **MSS** (maximum segment size) or **MTU** (maximum transmission unit). Each is called a segment. A header is added to each segment, which is then handed down, one by one, to layer 3.

The segment header is typically 20 bytes long and includes the source port number, the destination port number, an offset within the whole message, and a sequence number as well as a checksum.

Layer 3 is the network layer. It puts a packet header on the segment and calls the result a packet. Then it hands it to layer 2.

The packet header is generally 20 bytes, (the same size as the segment header,) and includes the source IP address, the destination IP address, time to live, and a header checksum.

Layer 2 is the data link layer. It puts a frame header on the packet and calls the result a frame. Then it hands it to layer 1.

The frame header is generally about 20 bytes, and includes the source MAC address, the destination MAC address. After the data there is a **CRC** (cyclic redundancy check) which is essentially like a checksum, and covers the whole frame.

Layer 1 is the physical layer. It sends the bits out one by one across the physical medium of the local area network.

9.2 Encapsulation

By the time we reach layer 1, we have the following basic data structure.

frame header (20 bytes, including MAC addresses)
packet header (20 bytes, including IP addresses)
segment header (20 bytes, including ports)
segment data (up to 1500 bytes).

9.3 Up the Stack

If the ultimate destination is in the same local area network, we are ready to push it back up the stack. If not, we have routers and hops in front of us. But let's skip that for the moment and pretend we have reached the ultimate destination.

Layer 1 is the physical layer. It receives the bits one by one from across the local area network. When it has a complete frame, it hands it up the stack to layer 2.

Layer 2 is the data link layer. It verifies the header on the frame and removes it, revealing the packet. Then it hands it to layer 3.

Layer 3 is the network layer. It verifies the header on the packet and removes it, revealing a segment. Then it hands it to layer 4.

Layer 4 is the transport layer. Again, it has to do some serious work. It looks at the offset and length of each segment. It removes the segment headers and reassembles the segment data in a buffer. With TCP it replies to the sender so the sender will know which segments are still missing. When the data is complete, it hands it up the stack to layer 5.

Layer 5 is the session layer. It pretty much just hands the data up the stack to layer 6.

Layer 6 is the presentation layer. It converts the data from network encoding into the local coding on the receiving computer. It also reverses any encryption that was done. It also reverses any data compression that was done. Then it hands the data up the stack to layer 7.

Layer 7 is the application layer. It places the data where the application program can get it and notifies the application program that the data is available.

9.4 Routers and Hops

The simple case we have examined above passes the data down the stack, across the local area network, and back up the stack, from one computer to another.

In many cases the computers do not share the same local area network. We must enlist the aid of **router**s that will help the data along its way.

Each router stands at the intersection of two (or more) local area networks. It belongs to both networks.

Chapter 36 (page 258) goes deeper into routing.

Local area networks exist at OSI layer 2. Routers exist at OSI layer 3. If the router is not the ultimate destination for the data, it performs a **hop**, moving the data from one local area network to the next one.

Data will pass through the following layers: 1, 2, 3, 2, 1.

The data passes up through layers 1 and 2 as always.

The activity at layer 3 is different.

Layer 3 is the network layer. The router extracts the destination IP address. It consults its routing tables to pick a path that will take it closer to its destination. Then it passes it back down to layer 2 with a new intermediate destination.

The data passes back down through layers 2 and 1 as always.

This 1, 2, 3, 2, 1 action is called a **hop**.

Exam Question 119 (p.326)**:** What is a hop?
Acceptable Answer: moving a packet from one lan to the next lan

A hop is the activity performed by a router when it receives a packet on one local area network, and passes it across to a different local area network, one step closer to its final destination.

9.5 To Infinity and Beyond

(Yeah, Buzz Lightyear says that in Toy Story.)

Early on, it was realized that routing mistakes could happen, either due to programming errors or temporary conditions.

Consider an example in driving a car from point A to point B. We can set out with good directions, and then somewhere along the way, the road is closed. So we get new directions and head out again.

Suppose the new directions were flawed. The road is closed in that direction too. We get new directions and head out again. Only it turns out that these directions put us back on our original trail.

In a case like this, we will repeat endlessly, first going one way and then another, bouncing back and forth between the places the road is closed.

The same thing can happen in the Internet. Usually it is due to temporary conditions, like the physical unplugging of a wire from one router and plugging it into another. Or it can happen automatically when a wire is cut due to construction. Or it can happen automatically when electrical power is lost at one router, and it suddenly disappears from the network.

If the packet never dies, it can continue bouncing back and forth forever. To prevent this, packets are given a **TTL**, time to live. It is part of the packet header. Each time a hop occurs, the TTL counts down by one. When it reaches zero, the packet dies.

Exam Question 120 (p.326)**:** What does TTL stand for?
Required Answer: time to live

TTL (Time To Live) is the number of hops that a packet is allowed to take before the packet must reach its destination or die.

Exam Question 121 (p.326)**:** What is the purpose of TTL?
Acceptable Answer: It limits looping.

It does not prevent looping, but it stops it from becoming excessive.

Mitigate means limit the damage that something could cause. Network loops happen for various reasons. If a packet falls into a network loop, it will just go around and around and never be delivered. As more and more packets fall into that loop, a kind of gridlock happens, just like rush-hour traffic in a crowded city.

So, TTL does not prevent looping. It just limits it. It prevents infinite looping.

For normal packets, TTL should never reach zero. For packets that are being improperly routed in a loop, TTL prevents infinite travel around those loops.

When a packet dies, or **times out**, the packet is dropped. To be helpful, the

router that dropped the packet sends a death notice back to the originator. The death notice simply says that the packet died, and where it died.

If a death-notice packet dies, it just dies. Nobody gets told about it.

Chapter 36 (page 258) goes deeper into routing.

9.6 Trace Route

By manipulating the TTL values in packets, we can find the route that they are taking through the Internet. This is the approach used by the **traceroute** command, also called **tracert**.

Trace route sends the first packet with a TTL of 1. That packet reaches the first router where the TTL is reduced by one and becomes zero. The packet dies. The first router sends the death notice back. Trace route reports the location of death. It is the first hop in the route.

Trace route sends another packet. This one has a TTL of 2. That packet reaches the first router where the TTL is reduced by one and becomes 1. It continues forward to the second router where the TTL becomes zero. The packet dies. The second router sends the death notice back. Trace route reports the location of death. It is the second hop in the route.

By gradually running up the numbers, trace route can feel out the entire route that the packet will take. It can report all this information to the technician that is watching.

See section 32.1 (page 233) for more on trace route.

Chapter 10

Address Sharing (NAT)

Contents

The IPv4 address space is running out. Within organizations and in homes, the number of outward facing computers with routable IP addresses is limited.

http://en.wikipedia.org/wiki/Network_address_translation has more on NAT.

Exam Question 122 (p.327)**:** What does having a routable address mean?
Acceptable Answer: packets can be sent to your IP address across the Internet

Exam Question 123 (p.327)**:** What does having a non-routable address mean?
Acceptable Answer: packets can be sent to your IP address across your LAN, but the Internet requires NAT.

Routable refers to an IP address. Those addresses that are officially usable

as sources and destinations must start with a number between 1 and 223, with certain addresses reserved (not routable).

If you have a non-routable address, you can receive packets across your local area network, or perhaps across your corporate network, but not across the full Internet.

If you have a non-routable address, you share the routable address of someone else that will receive the packet on your behalf. They use NAT to send the packet along to you.

Exam Question 124 (p.327)**:** List in any order the five non-routable IP address blocks
Required Answer: 10.0.0.0/8, 127.0.0.0/8, 169.254.0.0/16, 172.16.0.0/12, 192.168.0.0/16

Network Address Translation, **NAT**, has been used since the 1990s as a partial solution to this problem. Because it is in place, it is also being used to solve other problems.

Exam Question 125 (p.327)**:** What does NAT stand for?
Required Answer: network address translation

Exam Question 126 (p.327)**:** What does PAT stand for?
Required Answer: port address translation

PAT is another name for Network Address Translation. NAT is the generally used terminology.

For comparison, look at regular postal mail. In the USA, postal mail is often delivered directly to a home or to a post office box. Although it is possible for each person living at that home to have their own, personal mail box, it is far more common to have one mail box that covers everyone. The address is shared.

In networking, for requests and replies to be transmitted across the Internet, both the sender and the receiver must have a valid IP address.

Since the addresses are becoming scarce, ISPs charge extra for each additional IP address. It has become common to share IP addresses within an organization or a home.

10.1 Network Address Translation Model

In the general model of address sharing we have the following components.

(a) An interior computer with a non-routable address.

(b) A router with a routable address.

(c) A destination with a routable address.

Sending: The interior computer sends a message as normal. The router receives it as normal. **The router modifies the message to identify itself as the sender.** The message then traverses the Internet until it reaches the destination.

Receiving: The destination computer sends a reply. The reply traverses the Internet until it reaches the router. **The router modifies the message to identify the interior computer as the new receiver.** The router delivers the message to the interior computer.

10.2 Man In The Middle

This is almost identical to a hacking technique called the **Man in the Middle** attack.

http://en.wikipedia.org/wiki/Man-in-the-middle_attack has more.

Exam Question 127 (p.327)**:** What does MITM stand for?
Acceptable Answer: man in the middle

The network is composed of many local area networks that connect with each other. At the point where two networks touch, there is a router. There are many such routers in the path followed by a typical packet as it goes from the sender to the receiver.

Sometimes one of the steps in the path goes through a device controlled by a hacker, who can intercept the messages going each way and either record them or modify them or both. When this happens, we call it a Man in the Middle Attack.

Exam Question 128 (p.327)**:** Explain Man In The Middle.
Acceptable Answer: (a) hacking attack. (b) two devices think they are talking privately. (c) instead both are talking to a spy.

The third computer is the man in the middle. The man in the middle is

not just listening in. It is actively passing the messages along. It could choose to drop messages, change them, or create false messages.

NAT changes messages in well-accepted ways but does not drop them or create false messages. NAT is legitimate. The router is in the middle, but everyone is happy. NAT is not a man in the middle attack.

10.3 Faking Out the Interior Computer

For this to work the interior computer has to send its messages to the router. No problem there. This will happen anyway.

Because we do not have enough routable IP addresses to go around, the router fakes out the interior computer by giving it a non-routable address, like 192.168.0.100. Other interior computers (there could be many) also have their own non-routable addresses.

The router hands out these non-routable addresses as part of its task of providing DHCP services to the local area network.

10.4 Faking Out the Destination Computer

When an out-bound packet arrives at the router, he plays his man in the middle role and updates the packet to make it appear that it is coming from him instead of the interior computer.

This is called network address translation.

The message is then forwarded to the destination computer which processes it as though the router were the true sender.

A reply is then sent back to the router.

The router then modifies the reply so it can continue on to its true destination, the interior computer.

10.5 Remembering the Lies

The trick about telling lies is that you have to remember what you told people. In this case, the router has to remember which interior computer is

being represented by the modified message. It needs to know this so when the reply comes back it can send it to the correct computer.

The original message has a source IP address and a source Port number. The IP address belongs to the interior computer (as assigned by DHCP). The Port number is made up by the interior computer.

The modified message will have a different IP address and Port number. The IP address belongs to the router. The Port number is made up by the router.

The router does this by having a **NAT address pool**. The key to each address pool entry is a router port number. The data part of the address pool entry includes the internal computer's IP address and port number.

When a reply comes back, the router looks up the port number and finds the original computer IP address and port number.

Example:
- True source is: 192.168.0.100
- True source uses port: 12345
- Router receives: 192.168.0.100:12345
- Router picks 60123 as the new port number
- **Router remembers: 60123 stands for 192.168.0.100:12345**
- Router IP is: 1.2.3.4
- Router sends as: 1.2.3.4:60123
- Destination gets message and responds
- Router receives as: 1.2.3.4:60123
- **Router converts 60123 into 192.168.0.100:12345**
- Router sends as: 192.168.0.100:12345
- True source receives the reply

Exam Question 129 (p.327)**:** In NAT, how does the router remember the original sender?

Acceptable Answer: The router keeps a table that matches outside port numbers with inside IP address and port number pairs.

The router maintains a table called the NAT address pool, which is based on router outside port number. The outside port number can be translated back to the original sender IP address and port number.

10.6 Time Out

The router does not have infinite memory. Eventually the address pool table will fill up. Something will have to be thrown away.

The router does this by using a ten-minute timeout. (The exact time limit may vary from router to router.)

Each time an entry is used, its count-down timer is reset to ten minutes.

When a timer reaches zero, that entry gets deleted. The interior computer is not active, or the browser has been closed. The space can be freed up for other computers.

This is called **garbage collection**.

Exam Question 130 (p.327)**:** What is garbage collection?
Acceptable Answer: reclaiming resources that are no longer in use

Notice that you are not reclaiming "data." You are, instead, reclaiming the storage space that was being used by the data.

After collection, the resources can be recycled and used for something else.

If an entry is needed but the address pool is full, the entry closest to timing out can be deleted.

Exam Question 131 (p.327)**:** How long do NAT address pool entries last?
Acceptable Answer: (a) ten minutes of inactivity, or (b) if table is full the longest inactive expires

Time-out expiration and garbage collection can cause problems.

True Example: An interior computer can be connected to a server somewhere. Maybe the connection is by way of SSH. If the end user walks away from their computer, the connection becomes quiet. Eventually the connection is dropped from the NAT table. But the connection is still alive so far as the server is concerned.

What to watch for: If you are experiencing timeouts, and if it seems to happen exactly ten minutes after you went idle, then maybe this is the problem.

How to avoid problems: You can often arrange to send a keep-alive message every five minutes or so. This keeps the connection from timing out and

being dropped.

Exam Question 132 (p.327)**:** What is a keep-alive?
Acceptable Answer: A message sent between two computers just to keep the connection from timing out and being dropped.

Keep-alive messages reset the timeout value. This prevents having the connection closed before you are ready.

Each keep-alive also reassures each machine that the other machine is still up and running and interested in continuing the conversation.

10.7 Be The Router

It is also possible to do Internet connection sharing from one computer to another.

Let's say you do not have a router, but you have two computers. Computer A is connected to the Internet. Computer B is not.

It is often possible to connect computer B to computer A, possibly through a USB cable, or possibly through Wi-Fi with computer A acting as the hot spot.

Chapter 11

Peer to Peer with NAT

Contents

Network Address Translation, **NAT**, is discussed in chapter 10 (page 62). In the current chapter, we see how NAT affects a few of the peer-to-peer activities we may want to do between computers.

11.1 What is a Client?

When you sit down at a computer and start up a browser, your computer is acting as a client. It is making requests to other computers that act as servers.

We use the word **client** to describe a computer that mostly makes requests.

We use the word **server** to describe a computer that mostly responds to requests.

Strictly speaking, every computer does both. Every client computer also has moments when it must act as a server. Otherwise it will be dropped from the network. It must answer certain kinds of requests, like "do you exist?" No answer means someone else may be assigned your IP address since you do not seem to be using it anymore.

11.2 Peer to Peer

Peer-to-peer communication is simply communication between two computers of similar rank. Normally we mean two clients.

Peer-to-peer communication is nice because direct communication is more efficient and less subject to problems elsewhere in the network. If a server goes down somewhere, peer-to-peer traffic is not affected.

In contrast, most traffic on the Internet is between a client and a server. Normally it is traffic between a web browser (running on a client computer) and a web server.

Peer-to-peer activity commonly includes things like these:

- skype (telephone-like communication).
- netmeeting (multi-person telephone-like communication).
- gaming (locally hosted).
- file sharing (legal or illegal).

11.3 Hidden Computers

For any two computers to talk to each other across the Internet, each must have a routable address.

Let us invent some terms we can use in this discussion.

"visible" means a computer that has a routable address.

"hidden" means a computer that does not have a routable address. Instead it has a non-routable address such as 192.168.x.x.

"NAT proxy" means a computer that gives hidden computers (its clients) access to the Internet. The NAT proxy shares its IPv4 address with its clients by providing Network Address Translation. It keeps track of things

by using a network address translation table.

Hidden computers cannot receive messages unless the NAT proxy will forward them. Hidden computers cannot be attacked by hackers because the NAT proxy stops all packets except responses to ongoing conversations.

The hidden computer must initiate contact with the outside world. It does this by sending a packet through the NAT proxy. The NAT proxy makes a note in its NAT table so when a reply comes back, the reply can be given to the hidden computer. The NAT table entry expires after about ten minutes.

By initiating contact, the hidden computer is opening a small hole in the firewall provided by the NAT proxy. Packets that use that hole in the firewall are considered to have been invited. They are part of an ongoing conversation.

Unless there is an entry in the NAT table, no message can reach the hidden computer.

11.4 Hidden Peers

There is no problem for a hidden computer to talk to a visible computer. The NAT proxy handles things.

There is a big problem for a hidden computer to talk to another hidden computer. Whoever talks first creates a firewall hole that allows replies to make it back. But the peer cannot receive that first message because its firewall is still in place.

11.5 Port Forwarding and DMZ

The first solution to un-hiding a peer is to have its NAT proxy poke a permanent hole in the firewall. This is done by identifying a specific port number at the firewall and having all traffic for that port automatically forwarded to the hidden computer.

This only has to be done by one client, which then becomes the “server” for all the other clients.

This works but there are several problems with it.

(a) The hidden computer operator must get the NAT proxy operator to

open up the hole in the firewall. This is beyond the technical skills of many home users.

(b) Only one hidden computer can own a given specific port. That may work okay when the connection is for hosting a game and nobody else in the family is likely to want to host at the same time, but it's not so great if the connection is for video chat.

Setting up DMZ is easier than setting up port forwarding, but has similar problems.

DMZ stands for demilitarized zone. Demilitarized zone suggests that the place has no military to protect it. In our case, it is not protected by the router or its firewall.

Exam Question 133 (p.327)**:** What does DMZ stand for?
Acceptable Answer: demilitarized zone

Port forwarding can be set up with triggers, but it's still tricky.

These methods are really a last resort. They work but they are difficult to set up and use.

11.6 Hosted Connections

Instead of talking directly to each other, we can introduce a permanent man in the middle (**MITM**). This MITM talks to each of the hidden computers, passing along the information they provide.

The big plus is that this is easy to set up. And for the MITM the other big plus is that they control the conversation and can charge money for keeping it going.

Another big plus is the conversation can easily have more than two clients.

The big minus is the MITM.

If the MITM gets bogged down with lots of traffic, it becomes a bottle-neck that slows everything else down.

If the MITM gets greedy, the clients become unhappy.

11.7 Brokered Connections

Both peers can talk to a broker. The broker is a visible computer whose purpose is to get peers talking to each other.

Both peers must initiate a conversation with the broker. This creates the small holes in the firewalls. The broker can then share those holes with the peers, and the peers can reach each other directly after that.

So far as setting up goes, this is just like the MITM scenario provided above. But the MITM does not host the connection after the conversation gets started, so the load on the MITM is much smaller. This makes it easier for low-budget and open source projects to provide brokers.

This kind of connection is used for things like skype.

Unit III

Home Networking

Chapter 12

Home Network Components

Contents

One of our target objectives is that you be able to correctly set up a home network. Let's start by introducing (or reviewing) about the components you will need.

Skill: Know the components of a home Internet system. Know their names. For each name, be able to describe what it is and what it does.

The components of a home Internet system are generally items like these:

12.1 ISP

There is a cable or phone line coming in from the ISP (Internet Service Provider) that connects you to the Internet. You typically pay a monthly fee (in 2010 in the USA it is often around $50) in return for this service. Typically this is the slowest part of your network. Without it, you are not

on the Internet.

Exam Question 134 (p.327)**:** What does ISP stand for?
Required Answer: internet service provider

http://netindex.com/ has speed statistics.

http://www.netindex.com/ is a good source of statistics for upload and download speeds throughout the world.

Typical download speeds range from around 12 Mb/s to 20 Mb/s.

Mb/s means megabits per second. It is also written **Mbps**. The "b" is small to indicate bits. If we said MB/s that would mean megabytes per second.

Exam Question 135 (p.327)**:** What is a typical broadband download speed in megabits per second (2013, Worldwide)?
Acceptable Answer: 15

Typical upload speeds range from around 4 Mb/s to 7 Mb/s.

Exam Question 136 (p.328)**:** What is a typical broadband upload speed in megabits per second (2013, Worldwide)?
Acceptable Answer: 5

Exam Question 137 (p.328)**:** What is bandwidth?
Acceptable Answer: bits transmitted per second

Bandwidth measures how many bits per second can be transmitted. Bandwidth is important to downloading files and viewing web pages.

Exam Question 138 (p.328)**:** What is throughput?
Acceptable Answer: bits transmitted per second

Throughput measures how many bits per second can be transmitted. Bandwidth and **throughput** mean the same thing. Bandwidth is the more commonly used term.

Exam Question 139 (p.328)**:** List in either order the two measures of network speed.
Acceptable Answer: bandwidth, latency

Exam Question 140 (p.328)**:** List in any order the three measures of network speed.
Acceptable Answer: upload bandwidth, download bandwidth, latency

For consumers, download bandwidth and upload bandwidth are often dif-

ferent. It can take a long time to upload a picture, for instance, as part of sending an email, but a much shorter time to download it when receiving that email.

Exam Question 141 (p.328)**:** What is latency?
Acceptable Answer: time between sending and receiving a packet

Latency is also called **ping** time or **lag**. Latency is especially important to gamers.

In the home market, typically an ISP is a cable TV company or a telephone company. Satellite companies also provide Internet.

12.2 DeMarc

Exam Question 142 (p.328)**:** What does demarc stand for?
Acceptable Answer: demarcation

Exam Question 143 (p.328)**:** Why is the demarc important?
Acceptable Answer: ISP is responsible outside the demarc. You are responsible inside the demarc.

The demarcation point, or demarc, is a fancy word for the point where the ISP stops being responsible. They are responsible for everything on the outside of the demarc and you are responsible for everything on the inside.

12.3 Modem

The word **modem** is an abbreviation that stands for modulator / demodulator.

Exam Question 144 (p.328)**:** What does modem stand for?
Acceptable Answer: modulator demodulator

The **modem** is the connection between your Local Area Network and your ISP. Typically the modem is right at the demarc. Typically the modem is a box that has three connectors: (a) ISP or WAN, (b) LAN or Ethernet, (c) Power.

What is modulation?

Modulation and demodulation are just the acts of changing how the content is represented. With modulation it is converted into a form that is best for

transmission. With demodulation it is converted into a form that is best for local use.

To understand those words, let's talk briefly about radio. Commonly we refer to FM radio and AM radio. FM is Frequency Modulated. AM is Amplitude Modulated.

When content (music or voice) is sent by AM radio, a basic radio frequency is selected, and then the amplitude (height) is modified a bit. Those modifications exactly match the music or voice being transmitted.

When content (music or voice) is sent by FM radio, a basic radio frequency is selected, and then the frequency itself is modified a bit. Those modifications exactly match the music or voice being transmitted.

At the radio tower, the original content is "modulated" with radio waves. The result is then transmitted.

At your personal radio, the radio waves are received and "demodulated" to retrieve the original content, which is then available to be played back to you through your speakers.

Now let's consider your home network.

On the ISP side of the modem, there may be cable (coaxial), phone line (DSL), fiber, wireless, satellite, or even power lines. In any case, we say that the data on the ISP (WAN) side is transmitted as waves because it is more suitable for long-distance transmission.

Exam Question 145 (p.328)**:** What does WAN stand for?
Acceptable Answer: wide area network

On the LAN side of the modem, there is typically a Cat 5 cable slot. We say that the data on the LAN side is transmitted as bits because it is more suitable for direct use by computers. The bits are represented as voltages.

The modulator takes voltages and turns them into frequencies (waves or photons) on the physical media. The demodulator converts frequencies back into voltages.

Exam Question 146 (p.328)**:** What does a modem do?
Acceptable Answer: convert between voltages and frequencies

Voltages: Within a local area network using cat5 cable (or similar) the information is in the form of voltages on the wires. Voltages are a better way to send information short distances.

Frequencies: Within the wide area network provided by DSL or cable, the information is in the form of frequencies on the wires. Frequencies are related to modulation, and are a better way to send information long distances.

Modulation: Ordinary radios use frequency modulation to send sounds great distances. AM radio uses Amplitude Modulation. FM radio uses Frequency Modulation.

12.4 UTP Cable or Wire

There are two kinds of cable that we have mentioned. One is on the ISP side. It is typically coaxial cable and is used to connect your cable modem to the Internet.

The other is on the LAN side. It is typically **Ethernet** cable which is usually **UTP** cable. UTP stands for Unshielded Twisted Pair. It is also referred to as Cat 5 cable (or similar words). It is the cable that you use to connect computers to routers and switches and gateways inside your house. UTP cable uses eight-bladed modular connectors, properly called **8P8C**, and commonly called **RJ45**, that clip into physical ports on computers and routers.

Exam Question 147 (p.328)**:** What does 8P8C stand for?
Acceptable Answer: eight position eight contact

Exam Question 148 (p.328)**:** In RJ45, what does RJ stand for?
Required Answer: registered jack

The RJ series of jacks is standardized by the US FCC (Federal Communications Commission).

Comparison between 8P8C (RJ45) (Ethernet) and 6P2C (RJ11) (telephone) connectors.

http://en.wikipedia.org/wiki/Ethernet_over_twisted_pair has more, including specifications for which wires go in which slots.

http://en.wikipedia.org/wiki/Ethernet_crossover_cable has simi-

lar information on crossover cables. In earlier years, **crossover** cables were required between similar devices, but since 1998 **automatic crossover** has been increasingly used.

Automatic crossover is called **Auto-MDIX**. **MDI** stands for Medium Dependent Interface. **MDIX** stands for MDI crossover.

Exam Question 149 (p.328)**:** What is the technical term for a connection that can use either straight-through or crossover cables?
Required Answer: auto-mdix

When a connection is auto-MDIX, it means you can use a straight-through cable or a crossover cable. Either one will work fine.

http://en.wikipedia.org/wiki/MDIX has more.

Exam Question 150 (p.328)**:** What does UTP stand for?
Required Answer: unshielded twisted pair

Exam Question 151 (p.328)**:** What does UTP do?
Acceptable Answer: It is a physical cable for carrying bits between devices within a local area network.

UTP wire and **UTP** cable are the same thing.

Exam Question 152 (p.328)**:** What does Cat 5 stand for?
Required Answer: category five

Cat 5 cable is a UTP cable with a quality rating of 5. It is better than cat 3 cable, and not as good as cat 6 cable.

Exam Question 153 (p.328)**:** Which is better quality, cat5 or cat6?
Required Answer: cat6

12.5 Computer

Although it is less common, it is possible to connect one computer directly to the Modem. This may be required during **troubleshooting**. Normally we go through a router instead.

Exam Question 154 (p.328)**:** What is troubleshooting?
Acceptable Answer: When you have trouble, you try to figure out what the cause is and how to fix it.

However, if there is network trouble that you cannot fix, and you need

to bring in someone from the ISP, they will want your computer to be directly connected to their modem. That removes issues caused by your network. If they can determine that their connection works well to your single computer, without router or anything else, then they will conclude it is not their problem.

Before calling them, you should make that test yourself. Connect your computer directly to the modem. Can you get Internet? If not, then call the ISP. But if you can get Internet, then you need to debug your home network.

Exam Question 155 (p.328)**:** What two times should you connect your computer directly to a modem?

Acceptable Answer: (a) when it is the only computer on your home network. (b) when you are troubleshooting the network.

Even if you only have one computer, it is still worth while to have a router for its firewall benefits.

Chapter 13

Home Router

Contents

13.1 Router

It is more common to connect the Modem to a Router.

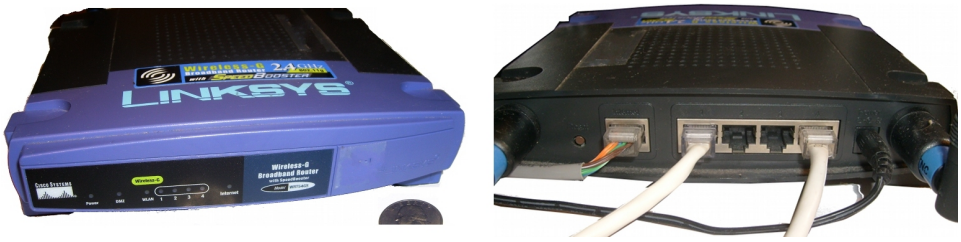

First, the Router acts as a **firewall** to prevent hackers from reaching your computers.

Second, it also acts as a gateway between the Internet and the LAN (local area network) inside your home or business.

Third, it runs a DHCP server to provide internal IP addresses within your LAN.

Fourth, it provides NAT services so your interior computers can talk to the Internet.

Fifth, it often provides Wi-Fi for laptops and other devices.

Exam Question 156 (p.329)**:** List in any order the five services a typical home router provides.
Acceptable Answer: firewall, gateway, dhcp, nat, wifi

Exam Question 157 (p.329)**:** What does a firewall do?
Acceptable Answer: It stops new conversations from reaching you.

Every computer can act as both a client and a server. Clients initiate conversations and servers respond to them.

Firewalls stop outsiders from creating new conversations with you. Instead, they can only respond to conversations that you started.

Exam Question 158 (p.329)**:** What does a gateway do?
Acceptable Answer: provide access to other networks

It provides a path from your computer to the Internet.

Exam Question 159 (p.329)**:** What does DHCP do?
Acceptable Answer: provide network configuration details so you can use the network

Exam Question 160 (p.329)**:** What does NAT do?
Acceptable Answer: replace one ip address with another

NAT can translate any kind of IP address into any kind, but it is most useful converting between routable and non-routable. Normally it translates a private, non-routable IP addresses that is being used within the LAN into a (shared) routable address that can be used on the Internet.

13.2 Wi-Fi / WAP

Exam Question 161 (p.329)**:** What does Wi-Fi do?
Acceptable Answer: provide a wireless connection between a device and an access point

Home Routers normally provide Wi-Fi signals within the home. It is called

a Wireless Access Point, or WAP. The WAP allows computers to attach to the Internet without using physical cables. Wi-Fi is a shared connection that is slower than a wired connection, but usually still faster than your ISP. It is subject to disruption from cordless phones and microwave ovens.

Exam Question 162 (p.329)**:** What does WAP stand for?
Required Answer: wireless access point

Exam Question 163 (p.329)**:** What does WLAN stand for?
Required Answer: wireless local area network

Exam Question 164 (p.329)**:** How fast is 802.11b Wi-Fi in Mb/s (theoretical max)?
Required Answer: 11

Exam Question 165 (p.329)**:** How fast is 802.11g Wi-Fi in Mb/s (theoretical max)?
Required Answer: 54

Exam Question 166 (p.329)**:** How fast is 802.11n Wi-Fi in Mb/s (theoretical max per channel)?
Required Answer: 150

- standard max typical
- 802.11ac 800
- 802.11ad 7000 na

Wi-Fi within a local area network is generally much faster than the broadband connection between your gateway and your ISP, as long as obstacles do not slow it down very much. The ISP connection will normally be the slowest part of the network.

http://en.wikipedia.org/wiki/IEEE_802.11 has an article about 802.11 standards and speeds.

802.11a Wi-Fi tops out at 54 Mb/s. It was not popular with consumers. It uses the 3.7 and 5 GHz bands.

802.11b Wi-Fi tops out at 11 Mb/s. This was the first popular standard, making its debut in 2000. It uses the 2.4 GHz band.

802.11g Wi-Fi tops out at 54 Mb/s. Typical best speeds are around 25 Mb/s. This replaced 802.11b in about 2003. This was the second popular standard. It uses the 2.4 GHz band. Because the Wi-Fi is usually backwards-compatible with the 11b standard, it is often advertised as 802.11b/g.

802.11n Wi-Fi tops out at 150 Mb/s per stream, with up to four stream for a total of 600 Mb/s. It was adopted as a standard in 2009, and is currently (2013) the popular standard. It uses both the 2.4 GHz and 5 GHz bands, and multiple data streams. Its range is about double that of a/b/g. Because the Wi-Fi is usually backwards-compatible with the 11b and 11g standards, it is often advertised as 802.11b/g/n.

802.11ac This is the newly emerging standard. Because the Wi-Fi is usually backwards-compatible with the 11b and 11g standards, it is often advertised as 802.11b/g/n/ac. But because b/g/n/ac is getting pretty long, you may see it as just 802.11ac and you will find the rest listed in the fine print.

Wi-Fi bandwidth is shared among all users.

Exam Question 167 (p.329)**:** How many connections can a Wi-Fi access point handle?
Acceptable Answer: 20+

The theoretical limit is hundreds or even thousands. The usable number is probably more around 20. The practical limit depends on how chatty each device will be. It is a bandwidth problem instead of a number of connections problem.

Some places configure their WAPs to limit the number of connections, apparently in hopes of keeping the bandwidth from being spread too thin. This can be deceptive, however. Someone wandering through your area may snatch up one of your connection slots and it may remain allocated to them even after they have left the area, thus denying service to a new person entering the area.

Exam Question 168 (p.329)**:** How fast (in Mb/s) is a wired connection?
Acceptable Answer: 100 is common

10 Mb/s is the old standard. It is called Ethernet. It normally requires cat3 cable or better.

100 Mb/s is very common. It is called Fast Ethernet. It normally requires cat5 cable or better.

1000 Mb/s is becoming common. It is called gigabit Ethernet. It normally requires cat6 cable or better.

Exam Question 169 (p.329)**:** What is duplex (in general)?
Acceptable Answer: Duplex tells whether you have two paths, in and

out, that can operate at the same time.

Exam Question 170 (p.329)**:** What is half duplex?
Acceptable Answer: can send and receive but NOT at the same time

Half duplex is like a one-lane driveway to a house. Cars can come in or out, but only one way at a time.

Wi-Fi is typically half duplex. Old-style networks that use hubs instead of switches are half duplex.

Exam Question 171 (p.329)**:** What is full duplex?
Acceptable Answer: can send and receive at the same time

Wired connections are typically full duplex. Modern networks use switches instead of hubs to achieve full duplex communication.

If your router or network card has an option for selecting half, full, or auto, you should select auto.

13.3 Switch

Home Routers normally support about four local (wired) connections, plus a large number of Wi-Fi (wireless) connections. Each connection can handle one device, such as a computer, a printer, a Blu-Ray player, a gaming console, or a security system. If you need more than four wired connections, you need to get another Router or a Switch.

Exam Question 172 (p.329)**:** What benefits does a switch provide?
Acceptable Answer: (a) more wired ports. (b) full-duplex communication.

13.4 Computers / Hosts

In this category, we include computers, printers, Blu-Ray players, gaming consoles, security cameras, and any other device that connects to the LAN or the Internet.

13.5 Wiring

As mentioned above, typically wiring is an **Ethernet Cat 5** (or better quality) cable with eight-blade connectors, called **8P8C** (**RJ45**) modular connectors. You can buy these in various lengths, or you can make your own. Wire is generally much faster than Wi-Fi, but both may be faster than the ISP connection.

Exam Question 173 (p.329)**:** Name five components of a typical home Internet system.

Acceptable Answer: isp, modem, router, wap, computer

Chapter 14

Selecting the Pieces

Contents

14.1 Selecting the Pieces

There are several important decisions involved in setting up a home network. Some are obvious like how much money you can afford to spend. You should know how much each piece costs.

Let's run through the others.

14.1.1 Pick Your ISP

ISPs are categorized by the method they use to provide signal to you: cable, telephone, wireless, satellite, or power line.

Exam Question 174 (p.329)**:** List in any order the six categories of ISP.
Acceptable Answer: cable, dsl, fiber, wireless, satellite, powerline

Wireless is not the same as Wi-Fi. Wi-Fi is small-radius, local, and not provided by an ISP. Wireless would be like 3G or 4G, as used by cell phones, or WiMax or "super Wi-Fi".

Often there are several ISPs available, including the cable TV company and the telephone company. **WiMax** or **4G** wireless is starting to be a real option, but often has capacity limits. Electric companies seem poised to get into this market as well. And there is always satellite, which may be your only option in remote locations or on a boat.

Read through the contracts carefully.

Ideally the ISP provides a single **routable** IP address for you to use. (More is better, but more is also unusual.)

Ideally the ISP specifies how much bandwidth you will receive. Normally this is specified in Mb/s up and down, meaning megabits per second for uploads to the Internet and downloads from the Internet.

Ideally the ISP provides a money-back guarantee if you are not satisfied within the first 14 days (or some other reasonable period of time).

Ideally the ISP does not restrict the number of computers that you will eventually attach. If this matters to you, make sure you check carefully. In the early days of home Internet, it was common to charge extra for each computer. Some ISPs may still do that. Currently it seems common to just restrict the total bandwidth.

Ideally the ISP does not charge extra for going over the amount of traffic you are allowed on your data plan. For wireless Internet, typically there is a limit. For wired Internet, generally there is no limit, but you need to read carefully to find out. If there is a data plan, and you reach your maximum, what happens? Do they limit your usage, or do they bill you extra? Find out.

Maybe they provide a modem as part of the package, and maybe you have to buy your own. Find out.

Call them to find out how long it would be until they installed your connection. It may be a couple of days. It may be a month. Find out before you make your final decision.

14.1.2 Make a Floor Plan

Sketch out the locations of all the devices that you expect to be on the network. You may want to over-estimate and then cut back if you cannot afford to have them all.

Include the walls. Include the major furniture. Those things can affect the signal strength of the Wi-Fi. If you do not plan to use Wi-Fi, then you will still need to run cables somehow.

14.1.3 Wireless Devices

You should probably plan on having Wi-Fi. It really is the best way to provide connections for portable devices like laptops, tablets (like the iPad), and handhelds (like the iPod Touch).

Figure out where the wireless devices will be used. Normally they are used in living spaces, including bedrooms, kitchen tables, and patio areas.

Normally wireless is used by visitors including friends and relatives that may drop by or stay overnight when they are in your area.

Do you intend to provide wireless access to neighbors or renters? Do you want the option to do that later?

14.1.4 Wired Devices

How many extra wireless access points? How many computers? How many printers? How many security systems? How many other things?

Each wireless access point can normally provide three or four ports for nearby wired devices.

Where will you run the wires? If you do not own the house, your options are more limited. Often you can put them under throw rugs, behind furniture, and across the tops of cupboards. Sometimes you can run them out one window and back in another.

It is also possible to set up a wireless bridge. This means that you have an access point in reverse, more in the role of a receiver than a transmitter. It supports wired devices, but it communicates wirelessly with the central access point.

If you do own the house, you have a few more options. Sometimes you can run wires under the floor, in the crawl space under the house, or in the attic space over the house. Sometimes you can run them outside the house, poking out through a wall in one place and back in through a wall someplace else.

14.1.5 Central Equipment

You may be able to have the cable company or the telephone company locate the modem exactly where you want. But good luck if you want to move it later. Pick carefully.

If you have many visitors to your home, for example, friends of your children, you may want a location that you can secure, that is, lock down so people don't mess with it. But you want to be able to get to it easily yourself.

You also need electric power to that location, probably including a **UPS** (uninterruptible power supply, also called a battery backup).

Exam Question 175 (p.330)**:** What does UPS stand for?
Acceptable Answer: uninterruptible power supply

The router will probably be located very close to the cable modem. It is not necessary, but it is convenient that way. You need to secure it too. You might as well secure them together.

You want the hotspots to be reasonably close to all the wireless devices you intend to support. You want it reasonably close to the living room, for instance.

Exam Question 176 (p.330)**:** What is a hotspot?
Acceptable Answer: wireless access point

A hotspot, also called a Wi-Fi hotspot, it is just another name for a Wireless Access Point.

You may find that you need more than one hotspot, maybe because the house is large, or maybe because the walls are thick or contain metal studs. Wi-Fi repeaters may help.

The wireless devices (laptops, iPads, etc.,) need to be reasonably close to the hotspot.

If they will be too far away, you need to create another hotspot, or you need to provide a wired connection.

14.2 Installing the Modem

It is good to have a plan before you start buying things. Once you are ready to move forward, call your chosen ISP and arrange to get the cable or DSL line and modem installed.

This part can take several weeks, depending on the backlog for installations, and other factors like your distance from population centers.

When the installer comes, make sure you have a computer ready to connect with the modem, either a laptop or a desktop with a sufficiently long cable. Also make sure you have electricity and maybe a **UPS** at the spot they will install your connection.

During installation, there are several things you want to do.

1. Make sure the equipment gets located where you want it.

2. Make sure your computer can access the Internet directly through the modem. Make sure you can pull up web pages.

3. Make sure your upload and download speeds are in the range you were expecting. Use something like http://speedtest.net/ to check it out while the installer is present. If the speed is surprising (too high or too low), ask about it.

14.3 Adding the LAN

After the installer is gone, disconnect your computer from the modem. This will cut you off from the Internet.

14.3.1 Configure the Router

You can actually configure the router before the installer comes or after the installer leaves.

Connect your computer to your main router. Do not attach the router to the modem yet. You are still off the Internet.

Use the MAC clone feature, if available, to copy your laptop MAC address into the router. This will make it possible for you to plug your laptop directly back into the modem for **troubleshooting** later. MAC clone means that your router will pretend to have the same MAC address as your laptop (or computer).

Set an administrative password on the router. Nobody should know this but you and very trusted people. Chapter 21 (page 133) talks more about passwords.

Configure the DHCP on the router. Normally this includes NAT, and NAT provides a firewall for you. Pick an address range. The default is probably okay.

Exam Question 177 (p.330)**:** What does default mean?
Acceptable Answer: the setting that is in effect unless you change it

Default is the setting or choice or action that will be in effect or will happen if you do not specify something else.

During configuration, normally each option will have a default value that is acceptable to keep just as it is.

Exam Question 178 (p.330)**:** List in either order the two router configuration default values you should not keep.
Required Answer: passwords, ssid

If you care about security, you should change the passwords and your SSID.

Configure the Wi-Fi on the router. Use WPA2 for your encryption. Pick a Wi-Fi password. It should not be the same as the administrative password. Everyone that shares the Wi-Fi will know this password.

Exam Question 179 (p.330)**:** List in any order the two passwords a home router normally has.
Required Answer: admin, wi-fi

Save all your settings.

14.3.2 Connect the Router

Connect the router to the modem. Your computer should still be connected to the router.

At this point, you should be able to see the Internet again. Make sure you can.

You should also test the Wi-Fi to make sure it is working. Walk around inside and outside the house and see what the signal strength is in various places.

If you have more routers, switches, wireless bridges, wireless access points, or wired devices, add them now, one by one.

As you add each item, test it to make sure it is functioning properly.

Then have a party. You will deserve it.

Chapter 15

Making Your Own Cat5 Patch Cable

Contents

Patch cable is used for wired connections, to connect between routers, switches, and computers. It uses stranded-core copper wire for flexibility.

It is distinguished from infrastructure wiring that is generally hidden within walls and terminates at a punch-down block. It generally uses cheaper solid-core copper wire.

You can buy patch cable ready-made, and that is probably what most people would do. However, sometimes it is better to make your own patch cable. In this section, we teach you how.

15.1 Equipment

You will need the following things: (a) **cable**, (b) **ice cubes**, (c) **crimper**, and (d) **tester**.

The **patch cable** wire comes on a spool or in a box. Or we can repurpose

an old patch cable that is no longer working. For patch cable, we want the stranded core.

The **ice cubes** are **8P8C** modular connectors (male). They are transparent so you can see what you are doing. They have an opening on one end where you will insert the properly-trimmed cable, and they have eight gold-plated contacts on the other end, where you will connect to your infrastructure wiring. There is a spur that locks the connector into its socket. Ice cubes can be purchased in bulk for under five cents each. (In 2012 I purchased 100 for $1 plus $2 for shipping.)

The **crimper** is used to squeeze the ice cube, causing the eight contacts to be driven into the wires of the cable. It also has several sharp blades for preparing the cable by cutting either the whole cable, or the outer sheath of the cable. I think I spent about $25 for a crimper the last time I bought one.

The **cable tester** is used to verify that your cable is working properly. It can be very hard to tell whether it is working. Testers can be purchased on the Internet for around $10, which is an amazingly low price, but you can spend much more if you want.

15.2 Steps

Use the crimper to cut off one end of your cable, to a nice, square cut.

Use the crimper to cut carefully into the outer sheath of the cable, one half inch from the end. Your goal is to be able to remove the outer sheath without

damaging any of the eight wires that are inside it, or their individual plastic coatings.

Push the outer sheath down, exposing more of the wire that is inside.

Untwist the wires. You will find there are four twisted pairs. They are numbered. Pair number 1 is blue. Pair number 2 is orange. Pair number 3 is green. Pair number 4 is brown. Within each pair, one wire is solid colored, and one wire has a white stripe. Or maybe it is mostly white with a colored stripe.

This is an example of wires that are too long. The sheath is not pinched inside the ice cube.

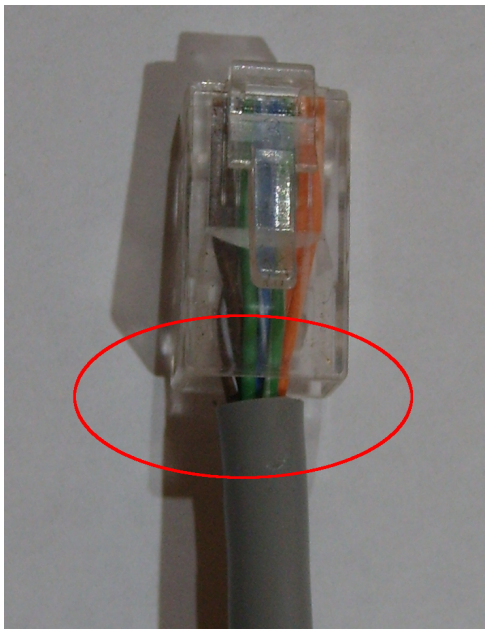

This is an example of wires that are okay. The sheath fits well into the ice cube.

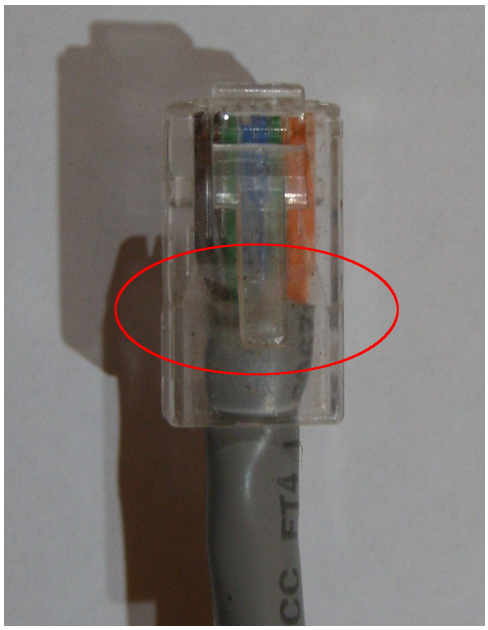

15.3 Cable Length

For Cat5 cable, the official maximum length is 100 meters. To go farther than that, you need a repeater or switch.

Exam Question 180 (p.330)**:** What is the maximum length (in meters) for Cat5 cabling?
Acceptable Answer: 100

15.4 Cable Termination

We terminate our cable on each end by putting the eight individual wires of the cable into the eight individual channels of an **8P8C** modular connector (the ice cube).

Lazy Termination: The wires must be in the same order on each end of the cable. For short cable runs, maybe five or ten feet, that is what really matters.

Professional Termination: For longer cable runs, we also need to worry about cross talk. And if someone else will inspect our work, we should follow

professional standards. These standards are designed to reduce cross talk. Consistent wiring rules let us wire each end without needing to look at the other end.

Cross Talk: For long runs the individual wires within the cable can affect each other to produce a situation known as **cross talk**. That is where the signal traveling on one wire is picked up on an adjacent wire. It is like the sending wire is a radio station, and the receiving wire is a radio receiver. The twists inside the cable are carefully calculated to cancel out as much cross talk as possible.

The wiring order standard is called "TIA/EIA-568-B.1-2001." **TIA** stands for Telecommunications Industry Association. **EIA** stands for Electronic Industries Alliance. The T568 standard specifies two accepted ways to do the wiring: **T568A** and **T568B**.

http://en.wikipedia.org/wiki/Category_5_cable has details.

Numbering: With the spur down, and the gold-plated contacts up, and the opening towards us, the contacts are numbered from left (1) to right (8).

pin	T568A	T568B
1	white/green	white/orange
2	green	orange
3	white/orange	white/green
4	blue	blue
5	white/blue	white/blue
6	orange	green
7	white/brown	white/brown
8	brown	brown

Exam Question 181 (p.330)**:** With T568 wiring, are the striped wires odd or even?
Required Answer: odd

The striped wires go in the odd-numbered channels: 1, 3, 5, and 7.

Exam Question 182 (p.330)**:** With T568 wiring, are the solid-color wires odd or even?
Required Answer: even

The solid-colored wires go in the even-numbered channels: 2, 4, 6, and 8.

Exam Question 183 (p.330)**:** With T568A wiring, what color goes in slots

1 and 2?
Required Answer: green

Exam Question 184 (p.330)**:** With T568B wiring, what color goes in slots 1 and 2?
Required Answer: orange

For T568A, the green stripe goes in slot 1 and the green solid goes in slot 2. For T568B, the orange stripe goes in slot 1 and the orange solid goes in slot 2.

Exam Question 185 (p.330)**:** With T568A wiring, what color goes in slots 3 and 6?
Required Answer: orange

Exam Question 186 (p.330)**:** With T568B wiring, what color goes in slots 3 and 6?
Required Answer: green

For T568A, the orange stripe goes in slot 3 and the orange solid goes in slot 6. For T568B, the green stripe goes in slot 3 and the green solid goes in slot 6.

Exam Question 187 (p.330)**:** With T568A wiring, what color goes in slots 4 and 5?
Required Answer: blue

Exam Question 188 (p.330)**:** With T568B wiring, what color goes in slots 4 and 5?
Required Answer: blue

For both T568A and T568A, the blue solid goes in slot 4 and the blue stripe goes in slot 5. These are the center two slots.

Exam Question 189 (p.330)**:** With T568A wiring, what color goes in slots 7 and 8?
Required Answer: brown

Exam Question 190 (p.330)**:** With T568B wiring, what color goes in slots 7 and 8?
Required Answer: brown

For both T568A and T568B, the brown stripe goes in slot 7 and the brown solid goes in slot 8.

The only difference between T568A and T568B is that green and orange

swap positions. Blue and brown stay the same with both methods.

A cable that is wired T568A at both ends, or T568B at both ends, is called a **straight-through cable**.

A cable that is wired T568A at one end, and T568B at the other end, is called a **crossover** cable. Pin 1 goes to 3, pin 2 goes to 6, and vice versa.

Chapter 16

Network Speed

There are two very different kinds of numbers that are commonly associated with speed: **latency** and **bandwidth**.

Latency measures the time it takes to transmit a small amount of information from your computer to another computer, and to receive a response. Gamers often refer to this as **lag**.

Download Speed measures the time it takes to download a large file from the Internet to your local computer. Mostly it measures bandwidth.

Upload Speed measures the time it takes to upload a large file from your local computer to the Internet. Mostly it measures bandwidth.

Example: Telephone lines are very fast in terms of latency but slow in terms of upload and download bandwidth. A satellite link is typically slow in latency but fast in upload and download bandwidth.

http://en.wikipedia.org/wiki/Bandwidth_(computing) has a bandwidth table.

Skill: Critically compare the bandwidth characteristics of several types of physical communication media.

Skill: Explain how bandwidth and latency impact throughput in a data communications channel.

http://speedtest.net/ provides a combined measure of network speed, including download file transfer speed, upload file transfer speed, and latency.

http://www.pingtest.net/ provides a more carefully calculated average ping time (latency) for your network connection.

http://en.wikipedia.org/wiki/Latency_(engineering) provides additional discussion of **latency**.

http://en.wikipedia.org/wiki/Lag provides additional discussion of **lag**.

Chapter 17

Servers

Contents

Networks often involve the sharing of printer and files. In this chapter we show how this can be done by sharing parts of existing computer systems.

The exact how-to depends a lot on the operating system of the host computer. We have chosen to address these tasks in the context of Microsoft Windows. We will look at printer sharing, file sharing, and configuring ad hoc wireless networks.

17.1 Printer Sharing

Skill: Print Server: MS Windows provides the ability for your computer to act as a local print server. (Print servers are also commonly done as separate interior computers.)

todo: add more

17.2 File Sharing

Skill: File Server: MS Windows provides the ability for your computer to

act as a local file server. Normally the protocol is SMB (Server Message Block). This provides file-sharing capability between interior computers (and possibly exterior).

todo: add more

Chapter 18

Troubleshooting the Network

Contents

One common request is to get the network back up and running.

Generally the network has stopped working for someone, and they have come to you, the expert, for assistance. (If you know more than they do, that makes you the expert.)

The person making the request is probably just overwhelmed by the complexity of the network. They don't know what to check. Is the problem with their machine only? Did someone knock out a utility pole that is carrying the Internet?

The best way to solve the problem is by narrowing down the possibilities. We do this by ruling out things that are still working.

Mostly we use **ping**.

18.1 When Trouble Strikes

If you are setting things up, or making a change to the system, that's a lot different from having things suddenly stop working.

For setting things up, skip the rest of this section.

If things suddenly stopped working, usually one thing failed. Try to determine where the problem might be.

Use your computer to run some commands that test the network.

Try the `ping 127.0.0.1` command. 127.0.0.1 is the special address for localhost. That's you. If the ping fails, there is something wrong with your computer. Try turning it off and then on.

We talk more about ping in section 31.2 (page 221).

18.1.1 Global Broadcast Ping

A global broadcast ping can quickly tell you a lot about yourself and the local area networks you are currently part of.

Exam Question 191 (p.331)**:** What is a Global Broadcast Ping?
Acceptable Answer: It is a ping to every device that you can reach.

Normally with a global broadcast ping you can only reach devices that are part of your own current local area network.

That is because devices beyond your local area network can only be reached by going through a router, and routers normally do not pass along global broadcast ping requests.

When they do not pass along a request, they are said to filter it out.

If routers did pass along that request, you would get replies from every device in the entire Internet, so it should be pretty obvious why they filter it out nowadays.

Exam Question 192 (p.331)**:** What is the command to do a Global Broadcast Ping?
Required Answer: ping 255.255.255.255

255.255.255.255 is a special reserved address for global broadcasts. You

would use it when you do not know your own IPv4 address yet.

Exam Question 193 (p.331)**:** With Global Broadcast Ping, who is the first responder?

Acceptable Answer: The device sending the ping is normally the first to respond.

You are your own first responder because you are the closest device to yourself on the network.

You might not be your own first responder if your own computer intentionally ignores such requests.

Exam Question 194 (p.331)**:** With Global Broadcast Ping, who is the second responder?

Acceptable Answer: The router (gateway) connecting you to the rest of the Internet is normally the second device to respond.

Your router or gateway is almost always the second responder.

Try the `ping 255.255.255.255` command. This is called the global broadcast ping. If it fails, there is something wrong with your computer. Try turning it off and then on. If it works, all machines on your local area network should respond to you.

Some local devices may not respond to your ping request. The most common reasons are because (a) they are currently turned off, (b) they are turned on but their wiring is disconnected or broken, (c) they have been configured to ignore broadcast ping requests, (d) they have been configured to ignore all ping requests.

Here is a sample global broadcast ping request and responses.

```
> ping 255.255.255.255
PING 255.255.255.255 (255.255.255.255): 56 data bytes
64 bytes from 192.168.1.100: icmp_seq=0 ttl=64 time=0.153 ms
64 bytes from 192.168.1.1: icmp_seq=0 ttl=64 time=1.667 ms (DUP!)
64 bytes from 192.168.1.101: icmp_seq=0 ttl=64 time=49.637 ms (DUP!)
64 bytes from 192.168.1.103: icmp_seq=0 ttl=64 time=63.141 ms (DUP!)
64 bytes from 192.168.1.100: icmp_seq=1 ttl=64 time=0.143 ms
```

The output has been shortened. On a unix system (like OS X or Linux), the output would continue until you press **control-c**, which is the special key-stroke combination for cancelling the current command. On a Microsoft

system, generally you will get the first four sets of responses and then it will automatically stop.

The first line, ping 255.255.255.255, is the command we typed in.

The second line is the ping command telling us what it will do.

The next set of lines say `icmp_seq=0`. They are the first set of responses that we received. Each one tells the IP address of the device that is responding.

After a brief delay, we get another set of responses from the same devices. These will say `icmp_seq=1`.

After a brief delay, we get another set of responses from the same devices. These will say `icmp_seq=2`.

And so it continues until we break out of it, usually by hitting **control-c** to cancel the current operation.

From these sets of lines we learn that there are four responsive devices in our local area network. There may be other devices that are not responding but usually every device that can respond will respond.

The last three of those responses are marked (DUP!). That is because ping normally expects one response and anything beyond that is considered to be a duplicate response.

The response that is not marked (DUP!) is the first response that we received. Normally that response is from ourselves. We can deduce that in this case our IPv4 address is almost certainly 192.168.1.100.

The next response is normally from our router because it is normally the closest device to us on the network. In this case, the router is almost certainly 192.168.1.1.

The other responses are from other devices in our LAN. In this case they are 192.168.1.101 and 192.168.1.103. They are probably other computers.

18.1.2 ipconfig (or ifconfig)

Try the `ipconfig` command. It should tell you your Internet address. It will be something like 192.168.0.100.

Microsoft calls the command **ipconfig** and Unix-based systems like Macintosh OS X and Linux call the command **ifconfig**.

18.1.3 169.254.x.x: Self-Assigned Addresses

If your Internet address starts with 169.254, you were not assigned an IPv4 address within a reasonable amount of time, so your computer simply made one up. To try again for an assigned address, turn off your machine and then turn it back on. The 169.254.x.x series of addresses are used when your machine cannot find a DHCP server. Normally your Router is a DHCP server. If you do not have a Router, then your ISP provides the DHCP server.

If you still get a 169.254 address, the problem is with your wiring or your Router (if you have one), or your ISP (if you do not have a Router).

18.1.4 Other Things to Ping

If your Internet address is 192.168.0.100, you would:

`ping 192.168.0.100` If it fails, there is something wrong with your network interface card. This would be rare.

Figure out your gateway address. Normally it is the same as your Internet address, except the last number is 1.

If your Internet address is 192.168.0.100, you would:

`ping 192.168.0.1` If it fails, there is probably something wrong with your Router. Try turning it off and then on.

You can do all of the above steps without involving anyone else. If they all work, but you still cannot reach the Internet, then start involving other people.

See how many devices are affected. Are other computers in your house able to reach the Internet? For any device that is working, you can rule out problems with all the pieces between it and the Internet.

If most devices are working, typically we would focus our attention on the devices that are not working.

If no devices are working, typically we would focus our attention on the Modem.

18.2 Have a Laptop Available

We will assume that you have a laptop that you can use for **troubleshooting**. If it is not a laptop, we need to have enough cable to reach from your computer to the Modem and to the Router.

In the instructions that follow, we will refer to the computer you are using as "your laptop." We will assume it has the **ipconfig** command available. If you are using Mac OSX or Linux, the command may be **ifconfig** instead.

Make sure that Wi-Fi is turned OFF on your laptop. Otherwise, you might accidentally pick up a signal from some neighbor. You could think you have Internet when really it is your neighbor that has Internet.

18.3 Check the Modem

Before Trouble Strikes:

You should know what the flashing lights on your Modem normally look like.

During Initial Setup:

Generally during initial setup, you will have a list of instructions from the ISP. You should follow those.

After Trouble Strikes:

If any device can communicate with the Internet, then the Modem is okay.

Devices with their own personal connection to the Internet do not count. For example, a cell phone may have Internet access. It does not use your Modem, so it does not prove that your Modem is working.

If no device can communicate with the Internet, we begin by unplugging the power from the Modem. This will remove it from the Internet. To be safe, we should leave it unplugged for 60 seconds or more. This gives the ISP time to recognize that the Modem is no longer connected. The ISP will release the line. You want that to happen. A shorter disconnect may be ignored by the ISP.

After 60 seconds, apply power to the Modem. Also connect your laptop to the network port of the Modem.

It is typical for a Modem to go through another 60 seconds of setup and

internal testing while it establishes a connection with the ISP.

Normally there is a characteristic pattern of flashing lights on the Modem. The exact pattern differs from Modem to Modem. If you have documentation, you can consult it. Otherwise wait until the lights seem to be flashing normally.

Next, try to use your laptop to reach the Internet.

If you are successful, then you know that your Modem is working okay and the cable is okay. Use ipconfig or ifconfig to discover your WAN IP address. Make a note of it. Then move on to the next item.

If you are not successful, try a different cable. Make sure you try your best cable.

If you are still not successful, you might try calling a nearby friend to see if their Internet is working. If theirs is down too, it makes your call to the ISP much faster because you can report that both you and your nearby friend lost Internet.

The ISP:

If you are still not successful, call your ISP. Let them know that you cannot reach the Internet. If they have heard complaints recently, they may just believe you and say that help is on the way.

If you are the first caller, normally they will assume the problem is with your equipment. (Normally they would be right.) They may give you steps to follow. Probably they are the same steps you just followed. You can (pretend to) follow them again, and report the results.

Do not bother trying to explain to the ISP what you have tried already. You are probably talking to a level-1 technician. They are the new kid on the block where they work. They have a script they are required to follow with you.

The tech may start by checking things from their end of the connection. They can use ping or other tools to see if your modem is online. They may find that your whole neighborhood is offline due to a power outage or equipment failure. Eventually they will either solve the problem, or they will move you up to a level-2 technician.

If they can reach your modem, it will almost certainly mean that the problem is with your own equipment.

18.4 Check the Router

Before Trouble Strikes:

You should know the internal URL for your Router. Most Routers use a web interface for configuration and information reporting. Normally the URL for the Router is one of these:

http://192.168.0.1/

http://192.168.1.1/

http://192.168.10.1/

Use ipconfig on your laptop to find your own IP address on the internal network. Whatever the first three numbers are, like 192.168.10, those numbers will be the same for the Router.

After Trouble Strikes:

If you know the Modem is working, turn off your Router for a minute. Then insert the Router into the system and turn it back on.

Connect your laptop to the router using a wired port.

If your laptop can access the Internet, then your Router is working.

If your laptop can access the Router web interface (mentioned above), check it to see what WAN address the Router is using. Normally it will be the same as the IP address you got when you were directly connected.

18.5 Check the Wiring

Wiring is usually stable.

Problems can come from getting your feet tangled in that wire behind your desk.

Problems can come from construction that may disturb the wiring.

Problems can come from harsh environments, for example, living near the ocean or having really bad weather. This can result in corrosion.

Sometimes wiring problems can be solved by just unplugging and replugging the **8P8C** connectors several times. Sometimes it helps to clean the connectors with ordinary alcohol and a cotton swab.

Sometimes wiring problems are due to a loose connection. Wiggle the cable near each connection. If you discover that it makes a difference, you have found a problem.

Unit IV

Wireless Networking

Chapter 19

Wi-Fi Configuration

Contents

The two major categories of networks are wired and wireless.

When you set up **Wi-Fi**, you are setting up a **wireless access point**.

`http://en.wikipedia.org/wiki/Wi-Fi` has more on Wi-Fi.

An access point, particularly one that you intend to share with the public, is also called a **hotspot**.

Wiring your house for Internet can be a real pain. That's why most people don't do it. Instead they go with a wireless solution called Wi-Fi.

Essentially, Wi-Fi lets computers and other devices communicate using radio signals just like those used by cordless telephones (which are not the same as cell phones).

The upside of Wi-Fi is that it is easy to install. Very easy. Most home routers you might buy these days come with Wi-Fi already installed. Most laptops have Wi-Fi installed. Desktops are even starting to come with Wi-Fi already installed.

The downsides of Wi-Fi are that it requires some configuration, its range is limited, it suffers from interference, it is less secure, and you may pick up free loaders.

Wi-Fi is generally considered to mean the same thing as **WLAN**.

19.1 Wi-Fi Channels

Skill: Students should understand Wi-Fi channels.

http://en.wikipedia.org/wiki/IEEE_802.11 has more on wireless.

Exam Question 195 (p.331)**:** What is 802.11?
Acceptable Answer: international standard for wireless networking

There are three major frequency ranges for **WLAN**s (Wireless Local Area Networks). They are 2.4 GHz (gigahertz), 3.6 GHz, and 5 GHz. Each range is divided into a number of channels.

Exam Question 196 (p.331)**:** What does GHz stand for?
Required Answer: gigahertz

Exam Question 197 (p.331)**:** What are the two main Wi-Fi frequency ranges?
Acceptable Answer: 2.4 GHz and 5 GHz

802.11a is the first standard to be worked on. It uses the 5 GHz frequency band.

802.11b is the second standard to be worked on, but the first standard to be adopted. Being "first to market" it became very popular. It uses the 2.4 GHz frequency band.

The old-school Wi-Fi band is the 802.11b band. Improvements were made. The first major improvement was called 802.11g. The second was called 802.11n. This is often written 802.11b/g/n.

In the USA, the 2.4 GHz range has 11 channels, numbered 1 through 11. Some countries reserve a bigger range around 2.4 GHz which allows more

channels to exist. Japan has 14 channels.

Exam Question 198 (p.331)**:** What 802.11b Wi-Fi channels exist (in the USA)?
Acceptable Answer: 1 - 11

The 3.6 GHz range is used by 802.11y. It requires a special license to use it in the USA. It uses channels 131 through 138.

Exam Question 199 (p.331)**:** What is Channel 196?
Acceptable Answer: wi-fi channel in 5 GHz range

http://en.wikipedia.org/wiki/List_of_WLAN_channels has good details on the various channels that exist.

The 5 GHz range is used by 802.11a/h/j/n. It uses a variety of channel numbers, going as high as 196.

802.11n uses both the 2.4 GHz and the 5 GHz frequency bands for maximum speed.

19.2 2.4 versus 5.0

Much of the newer networking equipment can use either frequency band for communication. If you have the opportunity to choose, which should you choose?

19.2.1 2.4 GHz Advantages

The 2.4 band is more widely supported. Everything you buy that is Wi-Fi capable can work with the 2.4 band.

The 2.4 band has a longer range than the 5.0 band. The 2.4 signal degrades more slowly as you get farther from the base station. In open air you can get about 100 meters away and still make contact. For 5.0 you do not get as much distance.

Exam Question 200 (p.331)**:** Which Wi-Fi band gets better distance, 2.4 or 5.0?
Acceptable Answer: 2.4

19.2.2 5.0 GHz Advantages

The 5.0 band is less crowded. With 2.4 you are competing with your neighbors in two ways: (a) some hardware can only use 2.4 (but over time this may change), and (b) 2.4 signals simply travel farther. With 5.0 your neighbors are less likely to be using it, and even if they are they are less likely to have their signals reach you.

Exam Question 201 (p.331)**:** Which Wi-Fi band has less competition, 2.4 or 5.0?
Acceptable Answer: 5.0

The 5.0 band has more channels. With 2.4 you are pretty much limited to three channels (1, 6, and 11) because of frequency overlap. With 5.0 there are somewhere between 9 and 25 channels available, depending on how you count, and none of them overlap. So 5.0 has **many more** usable channels: eight times as many.

Exam Question 202 (p.331)**:** Which Wi-Fi band has more usable channels, 2.4 or 5.0?
Acceptable Answer: 5.0

Bottom line: 5.0 gives lots more channels with lots fewer users (less interference), but somewhat less distance and fewer devices that are capable of operating.

The popular wisdom seems to be that if your equipment can use the 5.0 band, you will get better results by using it.

19.3 Wi-Fi Channel Selection

When you buy a wireless router, odds are good that the channel is pre-set to 6 (the middle). Then again, so is everybody else's.

Newer wireless access points may be set to "auto." They automatically pick a channel based on doing their own site survey.

You may want to pick a channel. The good choices are 1, 6, and 11. To make a good choice among those, you need to do a site survey.

Exam Question 203 (p.331)**:** What is a site survey?
Acceptable Answer: You check for Wi-Fi signals that already exist.

A **site survey** includes running a computer program or using a special

device that will **sniff** the air for Wi-Fi signals. It will tell you what your competition is. If nobody is near you, then it does not matter which channel you pick. If lots of networks show up in your survey, you should pick the channel that will create the smallest amount of competition.

Exam Question 204 (p.331)**:** In networking, what does sniff mean?
Acceptable Answer: listen to the network traffic and try to learn from it.

Why 1, 6, and 11? It turns out that using commonly available equipment, there is too much bleed-over between channels. A study by CISCO, currently the largest and most influential networking company, showed that placing active channels any closer than 1, 6, and 11, resulted in less successful traffic. Specifically, they showed that 1,4,8,11 gets lower throughput than 1,6,11. It is like painting more lines on the highway without making the road any wider. Sure you can narrow the lanes by 20% and cars will still fit, and you may get an extra lane out of the deal, but the cars will have to drive slower to stay in their lanes. Overall you do not win.

Exam Question 205 (p.332)**:** What 802.11b channels are commonly usable (in the USA)?
Acceptable Answer: 1, 6, 11

Exam Question 206 (p.332)**:** Why are many Wi-Fi channels not used?
Acceptable Answer: signal bleed

There is too much signal bleed between adjacent Wi-Fi channels.

19.4 Wi-Fi SSID

Skill: Students should understand Wi-Fi SSID names.

When you set up your wireless access point, you have to give it a name. The name is called an SSID, for Service Set Identifier.

Exam Question 207 (p.332)**:** What does SSID stand for?
Required Answer: service set identifier

Exam Question 208 (p.332)**:** What is the purpose of the SSID?
Acceptable Answer: it names the wireless access point

When you connect to a Wireless Access Point, you have to know the SSID. Normally it broadcasts its SSID so you can know it.

Exam Question 209 (p.332)**:** How many characters long can an SSID be?

Acceptable Answer: 32

If you don't set the SSID, it will probably be something well-known like "LinkSys" or "D-Link" or some other manufacturer name. That would be bad because people would know that you did not bother to configure your access point. They could use the fact that you appear to be lazy to identify you as an easy target. They may try to break into your router and do bad things.

It is possible (but uncommon) to use any character in an SSID. It does not have to be English or European or even printable. Technically, it is just a string of bits. It has no meaning other than to advertise that your access point exists, and to match when someone tries to connect.

19.4.1 Hidden SSIDs

Exam Question 210 (p.332)**:** Why are some SSIDs hidden?
Acceptable Answer: Hiding it supposedly provides additional security.

Exam Question 211 (p.332)**:** Does hiding your SSID improve security?
Required Answer: no

Google search "why is hidden ssid a bad idea" for lots of commentary.

Actually, it does not provide much security, and it has negative aspects that actually reduce security. But at one time it was believed to make security stronger. It actually does not. See "problems" in the next section for more information.

Before anyone can connect to your wireless network, they must know your SSID.

Normally you would broadcast the SSID. Anyone within range can see your SSID. Based on the SSID they may be able to tell whose signal it is.

If you do not broadcast it, people have to guess it or be told. Or they can just **sniff** your traffic until someone else connects and learn it that way. Having a hidden SSID was intended to provide some security, but it only provides a little, and it actually reduces the security of the client. Hiding the SSID is called **security by obscurity.**

Normally people set the SSID to something like their family name, and they broadcast it. If you did a site survey, you will know the SSID names of the nearby networks, except the hidden ones.

19.4.2 Hiding SSIDs Causes Problems

If your WAP's SSID is hidden, then the client must broadcast it when it attempts to connect. It does this by broadcasting its preferred network list. This seems like a bad idea because it violates the privacy of the client.

19.5 Wi-Fi Security

Skill: Students should know about Wi-Fi security.

Exam Question 212 (p.332)**:** List in any order the three Wi-Fi security methods that are commonly used.
Acceptable Answer: wep, wpa, wpa2

By default, your Wi-Fi connection does not use a password. It has no security. This means that anyone near you can connect and use your Internet. You might be okay with that. Or you might be upset. When neighbors walk near your house at night, do you care that your windows cast some light on the sidewalk? Or when neighbors water their lawn with your hose, do you care?

While it is nice to share, there are some risks involved.

Normally you should put a password on your Wi-Fi signal. Then only those who know the password can get access to the Internet through your network.

Chapter 21 (page 133) has more information on passwords.

There are three common protocols for security.

The oldest is called **WEP**, which stands for Wired Equivalent Privacy. Sadly it is easily broken. It is not secure.

http://en.wikipedia.org/wiki/Wired_Equivalent_Privacy has more on WEP.

Exam Question 213 (p.332)**:** What does WEP stand for?
Required Answer: wired equivalent privacy

Next we have WPA (and WPA2) which stands for Wi-Fi Protected Access. These have largely replaced WEP.

http://en.wikipedia.org/wiki/Wi-Fi_Protected_Access has more on

WPA and WPA2.

Exam Question 214 (p.332)**:** What does WPA stand for?
Acceptable Answer: wi-fi protected access

You should use WPA2 if possible. If not, use WPA. If not, use WEP.

Exam Question 215 (p.332)**:** Which is better: WEP or WPA?
Required Answer: WPA

Chapter 20

Wi-Fi Antennas and Signal Strength

Contents

Antennas are a great mystery. If you point them in the right direction, you get better signal reception. If you point them wrong, you get poor reception. When you drive through a tunnel, your reception is typically bad. If you live in the mountains, your reception is typically bad.

Wi-Fi suffers from two major signaling problems.

First, by law in the USA, Wi-Fi transmitters are limited to 100 milliwatts of broadcast power. That is to prevent your transmitter from jamming the neighbors. Everybody has to stay under 100 mW.

Exam Question 216 (p.332)**:** What is the legal maximum Wi-Fi signal (in milliwatts) in the USA?
Acceptable Answer: 100

Second, the strength of the signal gets smaller as you get farther from the access point. When it gets weak enough, you can no longer transmit and receive useful information.

This results in a coverage area, the places in your house or nearby where you can get a good enough signal.

There are some ways to improve the situation.

One solution is to put in more access points, with enough overlap in the coverage areas that you can always find a good signal in your house.

Another solution is to use one or more directional antennas. These can be used to push your 100 milliwatts in a certain direction rather than scattering it uniformly. It is like the difference between having a curved mirror behind your light bulb, and just a bare light bulb. Both glow just as brightly, but the mirror pushes the light in a specific direction.

What you need to know: The dispersion of signal is at right angles (90 degrees) to the surface of the antenna. A folded dipole antenna creates a 5 dB gain in a pancake shape. A cantenna creates a 10 dB gain in a single direction. A parabolic antenna creates a 15 dB gain in a single direction.

20.1 Antenna Shapes

Antenna design is a science and an art. It is not easy to do correctly. When looking for an **antenna**, unless we are studying electrical engineering ourselves, we just buy something. These are some options.

http://fpv-forum.com/index.php?topic=1311.0 is a helpful forum post that gives a summary of antenna shapes.

The strength of the signal is called its **gain**, and is measured in dB as compared to a non-directional antenna.

Whip / Dipole: The most common shape for a Wi-Fi antenna is just a plain wire, with one end attached to the transmitter and the other end just poking out into space. This is like the radio antenna you might see on a car. This is a directional antenna. The signal radiates in a pancake shape from the middle of the antenna. It does **not** shoot out from the ends. The gain can be in the 2 dB (dipole) to 6 dB (folded dipole) range.

Dish / Parabolic: This is often seen with satellite antennas. It is highly directional, and points in a single concentrated direction. The gain can be maybe 15 dB.

Pringles Can: This is a famous example of using an everyday object to create an antenna. It basically creates a parabolic antenna, highly direc-

tional, pointing in a single direction. Do a Google search on "pringles can antenna" or "**cantenna**" for some enjoyable reading.

20.2 Signal Strength

Absolute signal strength is measured in **dBm**, which means decibels on a 1-milliwatt scale.

Exam Question 217 (p.332)**:** What does dBm stand for?
Acceptable Answer: decibels milliwatt (decibels on a milliwatt scale)

http://en.wikipedia.org/wiki/DBm has more information.

Warning: Math is coming at you. More numbers. Fortunately, these are easy to work with.

Signal strength is like measuring earthquakes. A 5.0 earthquake is 10 times as powerful as a 4.0 earthquake. This is often called a logarithmic scale, or an exponential scale. A 6.0 earthquake is 10 times as powerful as a 5.0. A 6.0 is two orders of magnitude, or 100 times, larger than a 4.0.

With Wi-Fi power the orders of magnitude are measured in **dB** (without the m). dBm is an absolute measure of signal strength. dB is a relative measure that compares two strengths.

Exam Question 218 (p.332)**:** What is the difference between dBm and dB?
Acceptable Answer: dB is a relative measure. dBm is an absolute measure.

Skill: Know a bit about wireless transmission power measurement.
FYI: **dBm** = **decibel**s based on a 1-milliwatt scale.
FYI: 20 dBm = 100 mW = 100 milliwatts of transmission power.
FYI: 10 dBm = 10 mW = 10 milliwatts of transmission power.
FYI: 0 dBm = 1 mW = 1 milliwatt of transmission power.
FYI: -10 dBm = 0.1 mW = 1/10 milliwatt of transmission power.
FYI: -20 dBm = 0.01 mW = 1/100 milliwatt of transmission power.

Skill: Know how much power is typical in Wi-Fi situations.
FYI: 30 dBm: microwave oven leakage (noise when operating).
FYI: 20 dBm: strongest legal (outdoor) Wi-Fi signal strength (USA)
FYI: 15 dBm: typical laptop Wi-Fi signal strength.

FYI: -70 dBm: Weakest Wi-Fi signal that can be usefully received.

Well, that's all fine and good, but what does that have to do with my Wi-Fi reception, really?

The important thing is called **SNR**, the signal to noise ratio. If the noise gets too high, it drowns out the signal. It is like trying to whisper to each other at the beach, with the ocean in the background. If you get too far apart, you will not be able to tell what the whisper was supposed to mean.

Exam Question 219 (p.332)**:** What does SNR stand for?
Required Answer: signal to noise ratio

Exam Question 220 (p.332)**:** In what units is SNR measured?
Required Answer: decibels (dB)

Take microwave ovens. As stated in the list above, a microwave leaks 30 dBm of noise when it is operating. The strongest Wi-Fi signal is 20 dBm. That gives us a SNR of 10 dB. Oops, make that minus 10 dB. Signal (20) minus noise (30) is 20 minus 30 = -10. That's way more noise than signal. So, start the microwave and lose your Wi-Fi, if you are close enough.

http://xkcd.com/654/ has a cute comic about microwave ovens and Wi-Fi.

When the SNR is zero or below, it is impossible to detect what was sent.

When the SNR is positive, you can detect something. The higher it is, the faster you can recognize it. It is kind of like trying to see in a dark room. The darker the room, the longer it takes to recognize something. The brighter the room, the faster you can recognize it.

For a practical limit, unless your SNR is 20 dB, it will be too slow to be useful.

Exam Question 221 (p.332)**:** What is the minimum SNR (in dB) needed for a usable connection?
Acceptable Answer: 20

20.3 Signal Loss

Skill: Students should understand Wi-Fi signal loss (dB).

The farther you get from a bright light, the less bright it seems. The same

holds true for Wi-Fi signal. Every meter farther you go, the signal drops by 0.5 dB.

This is based on the 802.11 frequencies that are commonly used for Wi-Fi.

Beyond regular Wi-Fi, there are other frequencies that get better range. One example where this comes into play is the so-called Super Wi-Fi, also known as White Spaces Wi-Fi.

Different frequencies are blocked by different materials. Lead is used to block x-rays. Regular walls will block visible light but not x-rays.

http://en.wikipedia.org/wiki/White_spaces_(radio) has more.

Back to regular Wi-Fi.

In a perfect world, with -90 dBm background noise, and with 20 dBm initial signal strength, you can lose 90 dB of signal before it gets down to -70 dBm and becomes too weak to be useful.

That's 180 meters of distance. In a perfect world.

In a typical world, you get more like 100 meters. You can only lose about 50 dB before the signal is too weak to be useful.

Exam Question 222 (p.332)**:** What is the typical range (in meters) for Wi-Fi signals?
Acceptable Answer: 100

Exam Question 223 (p.333)**:** For typical Wi-Fi, how much signal (in dB) can be used up before the SNR is too low for useful communication?
Acceptable Answer: 50

Skill: Students should be able to answer questions about Wi-Fi, including estimating the range of a Wi-Fi signal given the distance and obstructions between the base station and the receiver.

Skill: Given a floor plan, tell how much signal will be present at various places.

Here are some signal loss numbers. They are good numbers, but not all walls are the same. Not all floors are the same.

Exam Question 224 (p.333)**:** For typical Wi-Fi, how much signal (in dB) is lost per ten meters of open air?

Acceptable Answer: 5

Exam Question 225 (p.333)**:** For typical Wi-Fi, how much signal (in dB) is lost per interior wall (plaster-board, wooden studs)?
Acceptable Answer: 5

Exam Question 226 (p.333)**:** For typical Wi-Fi, how much signal (in dB) is lost per exterior wall (wood, brick, cement block, metal studs)?
Acceptable Answer: 10

Exam Question 227 (p.333)**:** For typical Wi-Fi, how much signal (in dB) is lost per floor (thick plywood, support beams)?
Acceptable Answer: 15

20.4 Signal Loss Example

If you put your base station in one corner of the house, and you want your laptop to be able to communicate in the other corner of the house, will it work?

We need numbers. How far apart are the corners? Let's say 20 meters. So we lose 10 dB right there. We had about 50 dB to work with. Now we are down to 40.

How many floors? Good news. We are on the same floor. Nothing lost there.

How many exterior walls? Good news again. There are none. We are inside the house the whole way. But, wait a second. There is a "wet wall" when we go through a bathroom. That's a wall with pipes in it. Those walls are thicker and have more interference. We should count that like an outside wall. -10 dB leaves us with 30 dB.

How many interior walls? Hmm. Draw a floor plan. Looks like maybe five interior walls. Each one costs us 5 dB. That's 25 dB lost. We are now down to 5 dB available SNR.

Will that be enough? Maybe. But the signal will be weak. We will not get a good fast connection. We are talking maybe one bar instead of four bars.

Maybe we should think about another place to put the base station. Or maybe we should run a wire to the other corner of the house and put up another base station.

Is this analysis accurate?

It is close. But you would need to actually do a site survey with actual signal measurements from your laptop or some other tool. Then you would know how much signal you are getting. Maybe it is fine. Maybe not. And who knows if there is another internal wall that is really an external wall in disguise.

Oh, and I almost forgot. Walls and floors are not the only problems. Furniture is a problem too. Kitchen cupboards are kind of like walls. Refrigerators and stoves are kind of like walls. Couches and beds are kind of like walls. They reduce the amount of signal that gets through.

Exam Question 228 (p.333)**:** Name at least three typical indoor obstacles that affect Wi-Fi signal strength.

Acceptable Answer: Open air, interior walls, exterior walls, floors, furniture, cupboards, appliances, people.

Exam Question 229 (p.333)**:** Name at least three typical outdoor obstacles that affect Wi-Fi signal.

Acceptable Answer: Open air, people, trees, rain, buildings.

Exam Question 230 (p.333)**:** What conflict happens with 802.11b networks?

Acceptable Answer: microwave cooking

Cordless phones are also a problem, but since cell phones have gotten so popular, cordless phones are not seen much anymore.

Bluetooth: To a small extent, bluetooth devices can conflict with 802.11b networks, but bluetooth tends to be weak and should not be a problem. On the other hand, everything else can conflict with bluetooth. Bluetooth tends to have short effective distances.

Unit V

Security

Chapter 21

Passwords

Contents

Passwords are supposed to make sure that only authorized people get access to your assets, be that data, the ability to change data, or some other resource that is yours.

The basic concept is called **shared secret**. A shared secret is something that is known only to you and the other party to the activity. On the Internet you cannot "see" who you are dealing with. The normal way they prove their identity is by telling you something that nobody else would know.

Passwords are simply shared secrets of this type.

You should have passwords (or something better) to protect your assets.

Exam Question 231 (p.333)**:** Why are weak passwords a significant problem in networks?
Acceptable Answer: hackers get in and cause trouble

Weak passwords are those that can be discovered easily.

Generally anyone that knows the password can get into the account. They will then have the same access as the true owner.

Often we say this was done by a **hacker**, and breaking into an account can be called hacking in.

The hacker may see information that they should not see. For example, customer credit card numbers, student grades, or patient's medical information.

The hacker may change information that they should not change. For example, changing student grades, or transferring funds.

21.1 How Hackers Hack

To avoid becoming the victim, it is helpful to know how hackers hack. Badguy hackers have several methods for breaking into your accounts, including email, banking, social media, and just plain old shopping websites.

Online Dictionary Attack: First, they run a dictionary attack online against a very large number of accounts in hopes of finding a username and password that will let them in. If the username is varied, it is difficult to stop the attack. Dictionary attacks take a long time because they must go through the network, and the network is slow. A dictionary is a list of common passwords, which probably includes words from an actual dictionary, but also common passwords from other sources. They can use a botnet to carry out this attack to make it faster, but it is still slow.

Social Engineering: Alternately, they try to convince someone to let them in. They call tech support and claim to be your secretary that is in a hurry and in a lot of trouble and desperately needs to get into your account. Or you are stranded at the airport. Or at a police station. It's all about telling a convincing lie. There are pre-written scripts out there.

Escalation: Second, once they find some username and password that gets them in, they try to do a privilege escalation attack to gain more rights than that user would normally have. The goal is to get root access. This kind of attack depends on they operating system for that website, and whether it is up to date on its security patches. With bad enough security, they could skip the first step and start here.

Download: Third, once they have enough access privileges, they download the password file. These passwords are encrypted, or, more technically, "hashed," to hide their original values.

High-Speed Offline Cracking: Fourth, once they have downloaded the password file, they can use high-speed offline cracking technology to try lots of possible passwords. This is incredibly faster than the dictionary attacks. Short passwords or those based on simple patterns are broken almost immediately.

Reused: Fifth, once they have your username and password pair, they will try it on other websites. If you are using the same password in more than one place, they might get into those other places.

21.2 Measures of Password Quality

Commonly people will say that, among other things, (a) a password should be long, (b) a password should contain special characters, and (c) a password should be changed frequently.

But why?

One goal is to defeat hackers that might harm our assets. Hackers have three main approaches: dictionary attacks, brute-force attacks, and social engineering.

Dictionary Attack: When we say dictionary attack, what we really mean is trying everything on a list of commonly used passwords. These lists can be very long.

Uncommon: Your best defense is to avoid commonly used passwords, which includes most normal words plus maybe a few digits added to the end.

Brute Force: When we say brute force, we mean trying every possible password, starting with the shortest ones and working our way up, one letter at a time.

Long: Your best defense is a long password that has a variety of characters so it will not be discovered very soon.

Reused Passwords: If a hacker breaks into website A, where you have an account, and they steal all the passwords, over time they will crack many of them. Maybe they crack yours because it is too short. They can then try the same password at other websites where you might have accounts.

They might get lucky. People often use the same password on more than one website.

Unique: Your best defense is to use a different password for every website that matters. For websites that do not matter, anything is fine.

Social Engineering: When we say social engineering, we mean they call somebody pretending to be somebody, and they fool them into resetting your password. Hi, I'm Bob's boss, and he is out sick today, and we desperately need a file that is in his email. Can you reset his password? Never mind that you are not Bob's boss, and Bob is not even out sick today.

Of course, if you are a very important target, maybe they will just kidnap you and whack you with a tire iron until you give them the password. In a high-stakes game, you never know what will happen.

Crimes of Opportunity: Another goal is to avoid crimes of opportunity, where someone accidentally learns your password and then uses it. Maybe they found a scrap of paper where you had written it down. Maybe they were watching as you keyed it in. (This is called shoulder surfing.)

Memorable: For writing, your best defense is to memorize your password. If you write it down, keep the written copy someplace safe.

Complex: For shoulder surfing, your best defense is to have something complex enough that there is no fast and easy way to remember it. Something like Aloha123 is easy to remember once you have seen it.

Exam Question 232 (p.333)**:** List in any order the four measures of password quality.

Acceptable Answer: easy for you to remember, not easy for others to remember, not easy to guess, not used elsewhere

There are several good measures of password quality. (a) Memorable: How easy is it for you to remember? (b) Complex: How difficult is it for anyone else to guess? (c) How difficult is it for anyone else to remember, should they happen to accidentally see it? (d) Unique: If it is discovered, will it work on other websites?

The common recommendations appear to be designed to meet the first three goals. Or at least (b) and (c). They don't help much with (a). In fact, they defeat (a), which leads to people writing down their passwords, which leads to defeating (b) since you can find it written somewhere convenient, like on a yellow sticky on the wall or under the telephone or in the desk drawer.

Do not simply believe the common recommendations. Think about them.

The first real question is whether you need to remember your password or not.

The best strategy for passwords that you do not need to remember is simply to use a random password generator, and make the longest password you are allowed to have. This satisfies (b) and (c) but not (a). However, if you have an automated way to store the password, then (a) should not be an issue.

Examples of automated storage include building it into a program that does some task in your behalf. Database access is a typical example of this. Another example is by using special "key ring" programs that generate and save passwords for you based on the web site you are visiting or some other criteria.

For the rest of this section, we will focus on passwords you must remember.

21.3 How To Pick A Bad Password

Hackers and others that wish to guess your password have several typical approaches. (a) If they know you, they can try combinations of personal information such as your telephone number or the name of your spouse or significant other or pet. (b) Whether they know you or not, they can try lists of common passwords. This is called a brute force attack, or a dictionary attack.

Here is a list of the 13 most common passwords found on **Gawker** when hackers broke in during December of 2010: 123456, password, 12345678, lifehack, qwerty, abc123, 111111, monkey, consumer, 12345, 0, letmein, trustno1.

Would you use any of those? Apparently many did. We can attribute it to not thinking, or maybe to not caring. After all, if I have an account on Gawker, do I really care if someone else knows the password?

It would be much more interesting to look at a collection of passwords for online banking, where presumably more people would care.

But why stop at 13? Hackers have lists of thousands of common passwords. They can try each of those in an attempt to break into your account. This is called a **dictionary attack**. If you care, you need to pick something they

will not find. If you don't care, see "password reuse" below.

Do a Google search on "common passwords" for lots more. It makes for very interesting reading.

Exam Question 233 (p.333)**:** What is the problem with short passwords?
Acceptable Answer: too easy to guess

Yes, stay away from short passwords. Hackers will also try a brute force attack with all passwords, starting with the blank password, then going through the 26 letters one by one, then the digits and special characters. Then all possible two-character passwords. Then all possible three-character passwords. Depending on their connection, they can get up to five or six characters pretty fast.

For every character longer that you make your password, assuming it is unpredictable, you increase the cracking time by a factor of maybe 50. If an eight-character password takes a minute to crack, a nine-character password will take an hour, and a ten-character password will take two days.

Exam Question 234 (p.333)**:** What is the problem with long passwords?
Acceptable Answer: too hard to remember

Exam Question 235 (p.333)**:** What is a dictionary attack?
Acceptable Answer: using common passwords in hopes of finding one that works

Exam Question 236 (p.333)**:** What is the problem with dictionary passwords?
Acceptable Answer: too easy to guess

By common, I mean something that is on those dictionary lists used by hackers.

21.4 How To Pick A Good Password

The best strategy that I have found for passwords you must remember, and especially one you will share with others (like a Wi-Fi password), is to select a moderately long but memorable phrase and reduce it to the first letters of each word. Then mess with the letters.

Memorable Phrase

For example, **Lincoln**'s **Gettysburg Address** starts with the words: "Four

score and seven years ago our fathers brought forth on this continent a new nation, ..."

These words are familiar to many school children in the USA. By themselves, they may satisfy (a) and (b), but not (c) because they would be immediately recognized if seen. Plus they take a long time to type.

Going with the initials, we have "Fsasyaofbfotcann". It now satisfies (a), (b), and (c).

If there is very little chance that anyone else will see your password, you can just type in the whole phrase and not worry about (c).

This also makes for a very nice shared password, because the "insiders" can be told the secret for remembering it.

Exam Question 237 (p.334)**:** What do I recommend for a password?
Acceptable Answer: initials of a memorable phrase

Mess with the Letters

Beyond this, it could be further modified by replacing the "F" with a "4" and maybe the "s" with a "7". Many other replacements might be considered, such as using digits or special characters that are shaped similarly to the letters they replace. "A" might be replaced with "4". "s" might be replaced with "5". "O" (oh) might be replaced with "0" (zero). The password is reduced to gibberish that nobody would guess or remember if seen, but still you could create it as needed.

After messing we might have "4sa7yaofb40tc1nn".

21.5 Reuse: How Many Passwords Do I Need?

Like me, you probably have lots of accounts on lots of web sites. But remembering lots of passwords is just a pain. What to do? Write them all down? Or have one password (or a few) that you use over and over again?

My own strategy is to have one or two passwords that I use every place I don't care about. I call them my junk passwords. I don't use them on my banking accounts. I don't use them on my email. But if there is a web site that could do me little or no harm if someone else stole my identity there, they get a junk password. Ho hum.

For high-value targets, like bank accounts, where identity theft could cause

me serious grief, I use a better password, something harder to guess.

Why worry?

Every web site where you create an account has to remember your account name and password. Maybe they encrypt it. Maybe they don't bother. Maybe they are trustworthy. Maybe they are evil.

Once they have your username and password, can it fall into the hands of the bad guys? Well, up above we read about Gawker and the fact that a file of passwords was stolen. I imagine that the usernames were stolen at the same time. So it could happen.

Or maybe somebody sets up a junk web site just to harvest passwords.

http://xkcd.com/792/ has a cute comic about this.

And once they have a huge file, even if they are encrypted, they can make a massive effort to decrypt every password in the file.

And then they go visiting all the web sites they know of to see if their new list of usernames and passwords can get them in the door.

For those doors that pop open, they can decide whether a human should take over and step through the door. They can exploit their ability to get into your account on a totally different web site than the one that was compromised.

Don't let a low-value junk web site have the username and password that you use on a high-value web site.

Exam Question 238 (p.334)**:** What is a high-value password?
Acceptable Answer: password for a high-value asset

It's not about the password itself, or the quality of that password. It's about the thing that it protects.

High-value examples include private records, email accounts, bank accounts, Facebook, blogs, eBay, anything that could cause you serious pain if you lost it.

Exam Question 239 (p.334)**:** What is a low-value password?
Acceptable Answer: password for a low-value asset

Low-value examples include web site memberships that we do not really care about. If you could easily walk away from it and feel very little pain, then

it is a low-value asset.

Exam Question 240 (p.334)**:** Does it matter if a low-value password is easy to guess?
Required Answer: no

Hackers will use passwords harvested from one website to try to log into other websites. This works because the login name is often the same (typically an email address), and people hate to remember very many passwords.

Exam Question 241 (p.334)**:** If several high-value passwords are the same is that okay?
Required Answer: no

Exam Question 242 (p.334)**:** If several low-value passwords are the same is that okay?
Required Answer: yes

Password managers exist. In my opinion one of the best is **LastPass**. Another is **RoboForm**.

`https://lastpass.com/` is free to download and use on computers and laptops. They have a premium version for cell phones.

Exam Question 243 (p.334)**:** Are password managers a good thing?
Acceptable Answer: yes

With a password manager, you just have to learn one master password, which better be good. Then all the other passwords can be randomly generated and be different from one another. The password manager fills in the blanks for you when you are asked to login to a website.

`https://askleo.com/are_password_managers_safe/` A friend of mine has written a nice article about using password managers in general, and LastPass in specific. His recommendation and others convinced me to change over to using a password manager several years ago. I did so very cautiously, but have been very satisfied with the results.

21.6 Change Passwords How Often?

Some authors suggest that passwords should be changed every three months. Others suggest every year. Some suggest never.

Exam Question 244 (p.334)**:** List up to three problems with changing passwords frequently.
Acceptable Answer: written down, forgotten, similar to priors

When passwords must be changed frequently, they are often written down or similar to past passwords. Or they are forgotten. Some people I know rely on the "forgot my password" feature a lot, and simply pick passwords that they never intend to remember.

Exam Question 245 (p.334)**:** What is the problem with changing passwords rarely?
Acceptable Answer: more time for hackers to guess them

Short passwords can be discovered quickly. Long passwords can take years or centuries to be discovered by brute-force search. With a long enough password, there is no reason to change it unless it may have been discovered.

21.7 What If You Die?

Sharing passwords is bad. It opens you up to having other people use your accounts. But what if you die? Do you want your email account to die with you? Or your Facebook account? Or your bank account? Who owns your music in the cloud or your other digital assets?

This area of law seems to be still developing (2015), but in some cases you can contact the account provider with a certificate of death and evidence that you were appointed to manage the estate, and they will reset the password for you. That could take a long time, and in the mean time will you be hurt by being locked out?

I have read about automated systems where you can save an email to be sent to someone special if you fail to check in every month or so. Such an email could contain your passwords as well as other final thoughts, wishes, and apologies for being such a jerk in life.

Or maybe you just trust your spouse or child and you give them your passwords.

Or maybe you put the passwords in a sealed envelope and hide it in a safe. (That could make it hard to change your passwords very often.)

Chapter 22

Security Protocols

Contents

There are several protocols that are mostly automatic and let you, the user, benefit from a secure connection without the headache of setting it up. You still have to do a few small things, but they are small.

22.1 VPN: Virtual Private Networks

Virtual Private Networks are becoming common.

http://en.wikipedia.org/wiki/Vpn has more.

These provide a "tunnel" from wherever you are, at home, in a hotel room, at an airport, in a coffee shop, from there to someplace inside your network at your job, as though you were actually located in your office.

22.2 TLS (SSL)

When using web pages, **HTTPS** is the secure protocol. It utilizes **SSL**, the **Secure Sockets Layer**, which is now called **TLS**, **Transport Layer Security**, to provide encryption.

Chapter 23

Authentication

Contents

Security has two main objectives. (a) To keep secrets. (b) To prevent unauthorized changes.

Authorized people should be able to see the secrets. Unauthorized people should not be able to see the secrets.

Authorized people should be able to make alterations, additions, corrections, or any other kinds of changes. Unauthorized people should not.

23.1 Hackers and Identity Theft

For a bad guy to break in, the easiest way is to look like an authorized person.

In the largest sense, identity theft is pretending to be someone else and then having others believe you. In your pretended role as someone else, you could

do good or bad things, but normally we worry about you doing selfish things that damage the true owner of the identity.

Normally identity can be proven in four ways. Be: it can be based on something unique that we are, such as DNA, finger prints, or retina scans. Have: it can be based on something unique that we possess, such as an ID card, a key, or a document. Know: it can be based on something unique that we know, such as a password or a shared secret. Do: it can be based on the ability to physically perform certain acts, or perhaps respond to certain "challenges."

Exam Question 246 (p.334)**:** List in any order the four types of things used to prove identity (four single words).
Required Answer: know, have, are, do

Know: Things you know prove who you are. Passwords are the prime example.

Have: Things you have prove who you are. Physical keys are the prime example.

Are: Things about you personally prove who you are. Fingerprints are the prime example. Biometrics is the generic name.

Problem: Biometrics don't change. One problem with biometrics is that you cannot change them. If someone figures out how to hack your biometrics, then they can use it on all websites. On the other hand, things you know, have, or can do could be different for each website. But with biometrics this is a problem. How do you change an iris scan? (Minority Report?)

Problem: Biometrics do change. Another problem with biometrics is that we ourselves may change. How good is my voiceprint if I have a cold? How good is my fingerprint if I cut my finger while working on a project?

Do: Things you can do prove who you are. This is often done in a challenge / response configuration, where you are given some pattern and you must complete it.

In real life, identity is confirmed by having certain documents, such as a passport or a driver's license. We present these documents to prove our identity. Someone checks the documents against our physical person to see if they believe we are who we claim to be.

In network life, identity is difficult to confirm in that way. Instead we

normally rely on passwords.

If the password is lost or forgotten, we fall back to more general security questions, such as mother's maiden name, or city of your first memory. All of these are examples of shared secrets.

Another option is to have a trusted source (such as an email account) confirm the identity. This simplifies the job at one stage, but simply moves the real job to another stage.

Exam Question 247 (p.334)**:** What is multi-factor authentication?
Acceptable Answer: Several UNRELATED factors must be presented.

Multi-factor means you must have several (multi) unrelated things (factors) that only an authorized person would be likely to have.

Exam Question 248 (p.334)**:** Is it multi-factor authentication if you have both a password and a security question?
Required Answer: no

A security question is something you know. A password is something you know. A security question is just an alternate password. If a hacker can get your password, they can probably get your security question. It is not good enough to be considered multi-factor.

2FA means two-factor authentication.

Exam Question 249 (p.334)**:** What does 2FA stand for?
Acceptable Answer: two factor authentication

Exam Question 250 (p.334)**:** What is single sign-on?
Acceptable Answer: You authenticate once and then receive credentials (such as a cookie) that are (a) used on related websites (b) in place of authenticating again.

Normal sign-on gives you access to a single resource or website. Within the same website, we do not call it single sign-on.

Authenticating (signing on) can be a pain, so after we do it we may want to avoid having to do it again.

The authentication process may give your computer a special cookie (authentication credential) that proves who you are, and that you already signed on. If the cookie is present, you are not required to log in again for related websites.

Example: I may have a gmail account and also a Google AdSense account.

Both are supported by Google, but they are otherwise unrelated. They could have the same username and password. If logging into one gets me automatically logged into the other, then single sign-on is happening. If I have to log into each one independently, then single sign-on is not happening there.

Hacker Defined

In general, a **hacker** is someone who goes beyond the routine ways of using a tool. They invent creative new ways to use it, ways that were not intended by the inventors. This can be good or bad.

Exam Question 251 (p.334)**:** What is a hacker?
Acceptable Answer: Someone that goes beyond the routine ways of using a tool.

Another phrase for this is "thinking outside the box." The "box" would be the routine ways things are done. Outside the box would be creative, unusual ways of doing things.

Exam Question 252 (p.334)**:** Is hacking bad?
Required Answer: no

In computing, hackers are often divided into two categories: **white hat** hackers (the good guys) and **black hat** hackers (the criminals). White hat hackers are also referred to as **ethical hackers**.

Exam Question 253 (p.334)**:** What does black hat mean?
Acceptable Answer: It means a bad-guy hacker.

By common usage, the word **hacker** by itself often indicates one of the **black hat** variety.

Exam Question 254 (p.335)**:** What does white hat mean?
Acceptable Answer: It means a good-guy (ethical) hacker.

White hat hackers are often employed to do Pen Testing, meaning penetration testing. They are hired by companies to try to break in to their own company. The purpose is to find security weaknesses so they can be fixed before black hats exploit them.

Exam Question 255 (p.335)**:** What is pen testing?
Required Answer: penetration testing

http://en.wikipedia.org/wiki/Hacker has more.

23.2 Man in the Middle

One problem with passwords is the **Man in the Middle** attack, previously mentioned in section 10.2 (page 64).

http://en.wikipedia.org/wiki/Man-in-the-middle_attack has more.

Let's say Alice and Bob want to communicate. They cannot see each other, so they rely on a shared secret, a password, to verify each other's identity.

Alice (to Bob): What's the password?
Bob (to Alice): Frog lips.
Bob (to Alice): What's the password?
Alice (to Bob): Spaghetti.

Each knows the correct password. Life is good. They communicate, knowing who they are talking too.

But now we add Charlie. He is in the middle. To Alice, Charlie pretends to be Bob. To Bob, Charlie pretends to be Alice.

Alice (to Charlie): What's the password?
Charlie (to Bob): What's the password?
Bob (to Charlie): Frog lips.
Charlie (to Alice): Frog lips.
Bob (to Charlie): What's the password?
Charlie (to Alice): What's the password?
Alice (to Charlie): Spaghetti.
Charlie (to Bob): Spaghetti.

Charlie knows everything that is going on. Alice and Bob have no way to verify that they are speaking directly to each other. The secret is out. It is no longer a secret.

Instead of talking in the clear (using **clear text**), what if we encrypt all our messages?

23.3 Replay Attack

Now that Charlie knows the passwords, he can do business with either Alice or Bob, pretending to be the other. That's because the password does not change.

An interesting alternative is to use a formula as the shared secret. This is called challenge / response.

Let's say the formula is $2x + 1$. The challenge would be a number, like 5. The response would be the result from the formula, 11. (2x5+1=11.) This has the advantage of defeating a replay attack. Of course, in real life the formula might be much more complicated.

Challenge/Response does not prevent Man in the Middle, but it can limit it to times that both Alice and Bob are present.

23.4 Secure Connections

By using a network, you are necessarily moving information around. You may want that information to remain private between yourself and the person to whom you send it. Does the net have the ability to keep a secret?

As your first order of business when communicating, you can choose some keys by which your data will be encrypted. There is a clever way by which this is often done, but you don't really need to understand it. The important thing is that ALL communications can be observed and copied, but encrypted ones will remain obscure and secret.

When using web pages, **HTTPS** is the secure protocol. It utilizes **SSL**, the **Secure Sockets Layer**, which is now called **TLS**, **Transport Layer Security**, to provide encryption.

Exam Question 256 (p.335)**:** Is http considered to be secure? Why?
Acceptable Answer: No. Traffic (data) is not encrypted.

Exam Question 257 (p.335)**:** Is https considered to be secure? Why?
Acceptable Answer: Yes. Traffic (data) is encrypted.

http://en.wikipedia.org/wiki/Transport_Layer_Security has more.

Exam Question 258 (p.335)**:** What does SSL stand for?

Acceptable Answer: secure sockets layer

Exam Question 259 (p.335)**:** What does TLS stand for?
Acceptable Answer: transport layer security

SSL and TLS are the same thing. TLS is the new name for SSL. SSL (the old name) is still better known (as of 2011).

RSA is not typically used for SSL or TLS. A public-key system is used initially and just to generate a shared secret, but after that less expensive forms of encryption are used for the ongoing traffic.

Exam Question 260 (p.335)**:** How does SSL protect confidentiality of a TCP connection?
Acceptable Answer: Traffic (data) is (a) encrypted to (b) hide its meaning.

As mentioned previously, there are two kinds of encryption: the normal kind that hides the meaning, and **signing** that proves authorship. See section 24.6 (page 157) for more.

When using Wi-Fi, **WEP** was the first major security protocol, but it was found to be very insecure. WPA and later **WPA2** have replaced WEP as the standard for security in wireless communication.

When using a shell (command line terminal window), **ssh** is the standard security protocol. It has almost entirely replaced **telnet** which was used for command line access.

23.5 Encryption

There are two kinds of encryption.

Regular **encryption** is when we scramble the message so that only the intended recipient can decode and understand it. This is normally what we mean when we say encryption.

The other encryption, called **signing**, is when we scramble the message in a way that only the author could have done, thus proving authorship.

We use a key to encrypt a message. Really we are just changing one set of bits (called the **clear text**) into another set of bits (called the **cipher text**).

If we can reverse the encryption, we must have a procedure that exactly

undoes the scrambling caused by the original procedure.

If we can find another procedure and key that converts the **cipher text** into **clear text**, then the two procedures are symmetric.

Often we consider the procedure and the key together to simply be the key.

Exam Question 261 (p.335)**:** What are symmetric keys?
Acceptable Answer: Encryption keys that cancel each other out are symmetric.

Sometimes we have one-way encryption. We convert the **clear text** into **cipher text**, but there is no efficient way to convert **cipher text** back into **clear text**.

One-way encryption is good for storing encrypted passwords. Take the **clear text** password. Encrypt it. Compare the results to your stored copy. If they match, the password is good.

On the other hand, if the stored copy is revealed, this does not really help the hacker because they need the original password, and the encryption cannot be reversed.

Rot13 Encryption

We use a procedure, or algorithm, to encrypt a message. A very simple encryption algorithm from ancient times is to rotate the letters of the message. Rotate them one backward and IBM becomes HAL. Rotate them one forward and HAL becomes IBM. (Dave Bowman: Open the pod bay doors, HAL. HAL: I'm sorry, Dave. I'm afraid I can't do that.)

Rotation is an algorithm. The number of places rotated is a key. We can rotate by two, or three, or ten places.

Rotate 13 (rot13) is a simple encryption scheme that replaces each letter by the one 13 places away.

Exam Question 262 (p.335)**:** What does rot13 stand for?
Required Answer: rotate thirteen

Exam Question 263 (p.335)**:** How does rot13 work?
Acceptable Answer: Each letter is replaced by the one 13 places away.

A/N B/O C/P D/Q E/R F/S G/T H/U I/V J/W K/X L/Y M/Z

The characters “Hello” would become “Uryyb”.

The characters “Uryyb” would become “Hello”.

Rot13 is a very simple encryption method. **Rot13** is self-symmetric.

In rot13, we might say that the key is the number 13, and the algorithm is rot, or rotation.

http://en.wikipedia.org/wiki/Rot13 has more.

Clearly rot13 is a horrible way to protect a real secret. But it introduces our terminology, **key** and **symmetric keys**.

23.6 The Man is Not Impressed

Charlie, our man in the middle, is not impressed. If Alice and Bob start to communicate using rot13, it is weak encryption, easily broken.

What if Alice and Bob decide to use something stronger? Charlie might not be able to figure it out. Especially if they agreed in advance, before Charlie got into the middle of things.

Skill: Explain how secret-key algorithms protect data confidentiality during transport over a network.

If Alice and Bob have agreed on a way of encrypting information, then Alice can encrypt it before sending it to Bob. Bob can decrypt it. Since Charlie does not have the secret key, Charlie cannot tell what they are saying.

But what if Alice and Bob do not know each other? That changes everything. How can Alice and Bob set up a secure channel of communication if they have never exchanged their secret keys?

How can Alice and Bob create a true shared secret?

With Charlie in the middle, listening to everything, it would be very difficult.

Chapter 24

Public Key Systems

Contents

24.1 Public and Private Keys

Skill: Explain how public-key algorithms work to authenticate and how they negotiate secret keys.

We will use our knowledge of symmetric keys to solve this problem.

Imagine that Alice and Bob each create a pair of keys. Alice will have a public key and a private key. They are symmetric. Bob will have a public key and a private key. They are also symmetric.

Alice and Bob publish their public keys. Charlie intercepts them. Now everybody knows Alice's public key and Bob's public key.

But only Alice knows her own private key. And only Bob knows his own private key.

Exam Question 264 (p.335)**:** Who knows Alice's public key?
Acceptable Answer: everybody

Exam Question 265 (p.335)**:** Who knows Alice's private key?
Acceptable Answer: Alice

And the keys are symmetric. The public reverses the private, and the private reverses the public.

24.2 Private Messages to Bob

Alice can use Bob's public key to encrypt a message to him that nobody else can read. When Charlie intercepts the message, Charlie cannot read it. When Bob finally receives the message, Bob uses his private key to read it.

Nobody but Bob can read a message that is encrypted using Bob's public key. The message is truly private.

Exam Question 266 (p.335)**:** Whose key, and which key do you use to send a private message to Bob?
Acceptable Answer: bob public

Exam Question 267 (p.335)**:** What is the purpose of encrypting a message?
Acceptable Answer: Prevent others from understanding it.

Encrypting does not prevent others from intercepting or seeing or capturing the message. It does not prevent them from transmitting it again, maybe many times. But it does prevent them from understanding it.

But is it authentic? Can we tell that it came from Alice? Or is it possible that Charlie made it up, pretending to be Alice?

Because Charlie has Bob's public key, Charlie could make up the message and send it to Bob, pretending to be Alice.

How can Bob tell who sent the message?

24.3 Authenticated Messages from Alice

Alice can use her own private key to encrypt a message to Bob that nobody else could have written. When Charlie intercepts the message, Charlie can read it too. When Bob finally receives the message, Bob uses Alice's public key to read it.

Nobody but Alice could have sent a message that is encrypted using Alice's private key. The message is truly authentic.

But is it private? No. Charlie can read it. Anyone can read it.

Using the private key to encrypt the message is called signing the message. It proves authorship.

Exam Question 268 (p.335)**:** What is the purpose of signing a message?
Acceptable Answer: Prove authorship.

Signing does not prevent others from reading it, copying it, or sending it again. It does not prove who sent it or when it was sent.

Exam Question 269 (p.335)**:** Whose key, and which key do you use to sign a message?
Acceptable Answer: your private

Anyone can decrypt it with your public key.

Exam Question 270 (p.335)**:** How does signing prove authorship?
Acceptable Answer: private key is required and nobody else has it

If the author is the only person with the private key, then the author is the only person that could encrypt with that key.

Anything the public key can decrypt must be from the author.

24.4 Private Authenticated Messages from Alice to Bob

Alice can use her private key (the first layer of encryption) followed by Bob's public key (the second layer of encryption) to encrypt a message to Bob.

Notice that the last encryption is using the other person's public key.

When Charlie intercepts the message, he cannot open it because that would require Bob's private key.

When Bob receives the message, he can open the second layer of encryption with his own private key. Then he can use Alice's public key to open the first layer of encryption, revealing the original message.

Because only Alice could have created it, the message is authentic. It is from Alice.

Because only Bob could have read it, the message is private. Only the original author and Bob can know its contents.

Exam Question 271 (p.335)**:** How can Bob send a private, authenticated message to Alice?

Acceptable Answer: First, encrypt it with Bob's private key to prove authorship. Second, encrypt the result with Alice's public key to provide privacy.

We must use Bob's private key to prove authorship.

We must use Alice's public key to provide privacy.

If we do Bob's first, then Alice's, nobody can open the message. That is the best sequence.

If we do Alice's first, then Bob's, everybody can open the message. They just cannot read it. From a practical point of view it probably does not matter really which encryption is first and which is second, so long as Alice can guess the right order to decrypt things.

24.5 Man in the Middle Defeated

Using two symmetric keys, Alice and Bob are able to set up communication with each other. They can pass a message that is authenticated and private. That message can become their shared secret.

After that, Alice and Bob can use simpler forms of encryption based on their shared secret. Or they can go back to the two symmetric keys to create a new shared secret.

24.6 RSA Creates Symmetric Keys

RSA was invented by Rivest, Shamir, and Adleman back in about 1970, as a method for creating a secure channel between two parties, Alice and Bob,

that were previously unknown to each other. The difficulty in such a case is to establish a shared secret that can be used later for authentication.

http://en.wikipedia.org/wiki/RSA tells more about it.

In section 24.1 (page 154) we discussed public and private keys, and how symmetric keys make it possible to defeat a man in the middle attack.

Before RSA, symmetric key systems were common, but public key systems were not. RSA established the idea that a public key could be widely shared while a private key was kept confidential.

Before RSA, it was difficult to create a public key that could not be easily broken to find out the matching private key.

The major contribution of **RSA** was the creation of a system whereby these high-quality public and private keys could be easily created.

Exam Question 272 (p.335)**:** What do public-key systems make possible?
Acceptable Answer: strangers can create shared secrets

Exam Question 273 (p.335)**:** Why is RSA special?
Acceptable Answer: public keys are easy to make

What does an RSA public key look like? What does an RSA private key look like? We can use the commonly-available **ssh-keygen** command to create a key pair.

Here is a typical private key. The information is binary but it is stored using base64 encoding.

```
-----BEGIN RSA PRIVATE KEY-----
MIIEogIBAAKCAQEAvGf4pUNrh0irqFbR3+tKhcMebpm4GDC7e7JHNQtVhTdUIL
CXYp22s55dZcO8Z6IsAmRviwmTgtWoL2t/fiN3GN3ssj4/mszwqfm1bN7KdJTS
XMHNzK8GIbKzTbdB6nhhDSS1unvQhh6uLufIV7ps3Owy/gRbtamrJqZncvfadz
wqnw21J4bMguhSXRmcLI8xYXcrm86qol0h7A0iUuqkgd/PTwvD3CyW9Rj19DWd
RxEgseNMysCsNAyUqE8BZTBmBhxNVcI/pHNEwLY9ZsSmuDLT6P/ZxiWPdsVqyo
eg3G/aj+iy+MEhiK6V4hr7e4ciqK+m5sdLlWsXYIPkAc/KdwIDAQABAoIBAFXz
cMGkMkfzZSxU/0YSB1timl2Gjy+p+74ApO+UDh3KEf0f0tBQdusWkjE8zJTSrR
UiUKizY7JKWoHXnsAi2dBkXfkq4MV0PU6t9pCGSwJIaneaipKYPkndRmNr+QxI
gq4GU0mR9p1ZHcmUQBeVmthmgRTVVClaEBR/0a3nWKkP5szsFS0HK+uzkrlYcb
/WwOctk/S2HG829s9OPUgLXx+iuQRCxNAxJomSTAIoU0xzfcmiFbgCUo7OkpKA
Vq33KZxAN/GjuLV9FvU7LnPerNios6MrkyeHzb/OEG8KQejyl3CWhHpAS0lrd6
SV5UMpvcqsNwuz5ebJx/0TXyIvwNECgYEA7GOjZLK78y7chJUQq3cF0MpBSg/S
aXnIZnMzcho0F6OKtNipFor4VMYCj9BLCRBsAHab7Xj7HQuZQcKs/DggClF7wn
```

```
c8BGZOwoinSanEeKlhPwguk69xPGzXbmNMbPpSqlhfVvkgvj4lUqB50zElK5Vf
FnmksvwpJp0trzeLg7sCgYEAzAlIfeVreOhefwzmtyIAWs2T8DhqUtHvJsqc2x
DKFGul+oiBY88W6tc/LUyEPsmwkDOboeIQ0XkZWt41w3/dNbXg8o/JdfQhbCAH
7wC4UVdCxkjboa4mEvX2rvAaxvEGh19ekoAgBgxvU5fSs3Sao+443aBDEXo2fp
EhzWBtYnUCgYA8sNAweFsc8nkXaRYgj6xkkKj2lN1WyzeSZh1dPDEHqKtiBwZS
jVd4nGXn+nhpWVgBPnKGI7uFF4c/hPXYw6gmPgliQi6cio/KDj+0/+s7d3FKN5
WNwUt3UhJ1evUSPaqfpZTExa2ManQ2pJDikgmTCQSixpRsh6UhDtW92KrvJQKB
gBF0zd3vjeVXRBZsnSSX5R3frsODoB0b1vEjro3TfuaGWzn3CdLnOde3uLEAde
QFT3TW5X7RH408Zh3vGvxx6RQaTneSM7NCwsVEbElXb22OIKlkPF41zw4aOxO2
eSQCJQPZ9fZkvdimobOuLpmp5T0XOYhEmNha82SauCChGlqtAoGAb9af52myn4
aZu0HHy4prtDEzqK4GxndTpebCLYPaqyK7COlXhaTrDG7uhaKssgLyl7QGvTQI
6SVuXeXx1JBB9rXppcN58bbOxQRG9zTfxNOgzXQ0fbDw04zjwUW7xauR3lHQ1I
1sv6Wibgnv4HX8mLcnfY2xaoE8i2GoeGp2eVc=
-----END RSA PRIVATE KEY-----
```

Here is the matching public key. It is all one line, but I have broken it up onto several lines for ease of display. The user@computer at the end is replaced with the name of the key, which can be pretty much anything.

```
ssh-rsa AAAAB3NzaC1yc2EAAAADAQABAAABAQC8Z/ilQ2uHSKuoVtHf60qFwx
5umbgYMLt7skc1C1WFN1QgsJdinbaznl1lw7xnoiwCZG+LCZOC1agva39+I3cY
3eyyPj+azPCp+bVs3sp0lNJcwc3MrwYhsrNNt0HqeGENJLW6e9CGHq4u58hXum
zc7DL+BFu1qasmpmdy99p3PCqfDbUnhsyC6FJdGZwsjzFhdyubzqqiXSHsDSJS
6qSB389PC8PcLJb1GPX0NZ1HESCx40zKwKw0DJSoTwFlMGYGHE1Vwj+kc0TAtj
1mxKa4MtPo/9nGJY92xWrKh6Dcb9qP6LL4wSGIrpXiGvt7hyKor6bmx0uVaxdg
g+QBz8p3 user@computer
```

So, what does it all mean? The mathematics is a bit tricky. We will present a simplified version that has the major elements you should grasp.

The mathematics involves using really big prime numbers, called p and q. The private key consists of these two numbers. We are talking like 100 digits long, each. It turns out that really big prime numbers are actually quite easy to find. Then you multiply them together to get a 200-digit number, called n, that is the public key.

(It is a bit more complicated, but only a bit. For our purposes, we will pretend that p and q are the private key, and that n is the public key. It is close enough to the truth.)

Exam Question 274 (p.336)**:** What is a prime number?

Acceptable Answer: A number with no proper factors.

A proper factor is a smaller whole number that divides exactly into the original big number. It must be greater than 1 and less than the original number.

7 is a prime. 6 is not a prime because 2 and 3 are proper factors of 6. 27 is not a prime because 3 and 9 are proper factors of 27. 29 is a prime.

Using these three numbers, p, q, and n, messages can be transformed from **clear text** to **cipher text** and back.

RSA could be broken if anyone ever discovered a fast way to convert the public key back into the two parts of the private key. So far, more than forty years later, nobody has revealed such a way. (If they have it, they are not talking about it.)

Quantum Computing is sometimes mentioned as a technology that would enable **RSA** keys to be broken, but Quantum Computing is still very much in the theoretical realm. Practical Quantum Computers do not now exist, and may never exist. And even if they did, they may not be able to factor large public keys into their prime components.

Exam Question 275 (p.336)**:** Why are prime numbers used in encryption?
Acceptable Answer: easy to multiply but hard to find the original numbers

Exam Question 276 (p.336)**:** What does the RSA private key consist of?
Acceptable Answer: two large prime numbers

Exam Question 277 (p.336)**:** What does the RSA public key consist of?
Acceptable Answer: private key numbers are multiplied to create public key

Starting with the RSA private key, it is very easy to construct the public key.

Starting with the RSA public key, it is very difficult to construct the private key.

Exam Question 278 (p.336)**:** If RSA is so great, why are other things used?
Acceptable Answer: once you have a shared secret, other things are faster

Once you have a shared secret, other methods can be used that are much faster. But until you have a shared secret, they cannot be used.

RSA is slow, but it is the best known method for creating the shared secret on which faster methods of communication can then be based.

Chapter 25

Firewalls

Contents

In real life, a firewall is a wall in a building. It is built out of special materials that do not burn, or at least resist for a long time.

In networking life, a firewall is a device in a network. It is built out of special software that prevents bad things from getting through. Those bad things are messages, typically requests.

http://en.wikipedia.org/wiki/Firewall_(computing) has more.

How exactly does a firewall work? It filters out the bad packets from outside threats so they do not reach your computer.

Exam Question 279 (p.336)**:** What is an Outside Threat?
Acceptable Answer: threat from outside your lan

Outside threats can come from anywhere in the world. There are always people trying to hack into unprotected computers.

Why would someone hack into my computer?

The major reason is to recruit your machine into their **bot net**. A botnet is a network of computers that are, to some extent, robots under the command of their hacker overlord. Some botnets have over a million computers involved.

Exam Question 280 (p.336)**:** What is a botnet?
Acceptable Answer: A network of computers controlled by a hacker, usually without the knowledge of their real owners.

http://en.wikipedia.org/wiki/Botnet has more.

Exam Question 281 (p.336)**:** For what three things are botnets commonly used?
Acceptable Answer: Sending spam. Mining bitcoins. Doing attacks.

Mining bitcoins is just one example. What we really mean is doing something that takes lots of compute cycles. Another example is cracking password hashes.

Attacks are typically DDOS, Distributed Denial of Service attacks. Many widely separated computers try to talk to the same server at once. The server could become overwhelmed and confused and stop working. Or the server could just become so busy that legitimate traffic is not handled promptly.

The "distributed" part is important because an attack from a single computer can be easily fire-walled. When you see a flood of packets from the same place, you can have an automatic response to block that single source. But when the packets are coming from a large number of different places it is very hard to block them without blocking the traffic you still want.

Exam Question 282 (p.336)**:** What does DDOS stand for?
Acceptable Answer: distributed denial of service

Attacks could also be automated attempts to break into other computers, recruiting more zombies for the botnet, or finding good candidates.

If a hacker takes over your computer, you might not even know it. They want you to keep the machine turned on so they can use it. Your computer is referred to as a **zombie**. The hacker **own**s your computer. The hacker **elite speak** (or **leet speak** or **1337 speak**) term for **owned** is **pwned**, with the o turned into a p, as in "you have been pwned."

pwn is often pronounced "pown".

http://en.wikipedia.org/wiki/Zombie_computer has more.

http://en.wikipedia.org/wiki/Pwn has more.

Exam Question 283 (p.336)**:** What is a zombie?
Acceptable Answer: A computer that is part of a botnet.

Exam Question 284 (p.336)**:** What does PWN stand for?
Acceptable Answer: own

Pwn means you have been owned by a hacker. They control your computer.

Exam Question 285 (p.336)**:** What is an Inside Threat?
Acceptable Answer: threat from inside your lan

Normally computers inside your local area network are trusted. Computers you do not trust should be kept outside your LAN.

When someone joins your Wi-Fi hotspot, they may be inside your local area network. If so, they have bypassed one of your firewalls.

25.1 Client / Server

Network communication is based on two major roles, client and server. In the networking world, a client (possibly a web browser) wants some information. The client crafts a packet to request that information and sends it to a server. The server responds with the desired information.

What's tricky about this is the fact that the server does not know in advance who will try to communicate with it, or when it might happen. The server just sits there, ready to talk to (nearly) anybody.

To complicate matters even more, the server (machine) may have lots of servers (software) running at the same time. I mean, after all, any given server (software) may have nothing to do at the moment. We would hate for the server (machine) to go to waste. So we combine lots of servers (software) so the server (machine) can keep busy.

From this you should have noticed that we are using the word **server** in two distinct but related ways. The "machine" server is a physical device, typically a computer. The "software" server is a computer program, designed to respond to requests.

Normally this does not cause any confusion, but at the moment, here in this chapter, it could. So we are calling it out to be very specific about it.

Exam Question 286 (p.336)**:** What two things does server mean?

Acceptable Answer: software that answers requests, hardware where such programs run

When the request comes in from the client, the server (hardware) decides which server (software) should get the request. It does this by using port numbers.

25.2 Filtering Traffic

In its typical form, a firewall looks at the network traffic that is passing through. The firewall has a list of rules. The rules base decisions on things like protocol, port number, and source IP address. For example:

* If the packet uses protocol ICMP, send it through.

* If the packet is from xyz, send it through.

* If the packet is from abc, drop it.

* If the packet is for port hij, send it through.

* For any other packet, drop it.

Crafting these rules can be a bit tricky, and we are not really going to address it in this book, any more than to say that this is basically how it works.

25.3 Client as Server

There is a secret. It is the fact that every normal client (machine) is also set up to act as a server. You may not think of your desktop computer or your laptop as a server, but it is. Every client has certain ports open and ready to respond to messages from the Internet.

The simple and classic example of this is the ICMP (Internet Control Message Protocol) echo request, known as **ping**.

According to the rules of the Internet (RFC 1122), if your computer is part of the Internet it must accept and reply to echo requests.

This forces your computer to be a server, at least a little bit.

There are other examples of your computer acting as a server. It may, for example, be running **Remote Desktop** software that allows others to connect to your machine and use it. **Remote Desktop** is intended to be a

way for others to help you, but unintended visitors could exercise the same privileges.

25.4 Router as Firewall

Who can start a conversation? Normally your computer is just a client and starts all the conversations that it is a party to. But maybe you are hosting a game or something. Then your computer is acting as a server and other clients will start conversations with you.

When you start a conversation, NAT records your IP address and port number in its table. This allows other computers to respond to your requests.

Unless you do port forwarding or create a DMZ, outsiders cannot start conversations with you because the NAT table does not give them a way to reach you.

With NAT, your router acts as a firewall to all interior computers. Because the interior computers can only receive packets when they go to a port listed in the NAT address pool, new conversations (service requests) are simply dropped.

Exam Question 287 (p.336)**:** How can firewalls defend against network attacks on clients?

Acceptable Answer: outsiders cannot start conversations with clients

If the client starts a conversation with someone outside of the LAN, then a record is made in the NAT table and responses can be accepted. If there is no entry in the NAT table outside messages cannot get past the firewall into the LAN.

Firewalls can also prevent out-bound communication with dangerous locations. They can stop you from starting a conversation with a known bad guy, or someone that is administratively restricted, such as a porn website.

Exam Question 288 (p.336)**:** How can firewalls defend against network attacks on servers?

Acceptable Answer: outsiders can start a few conversations of specific types with servers

Servers, by their very nature, must handle conversations started by outsiders. Each conversation takes a certain amount of resources, and if too many conversations are requested at the same time, those resources can be

overwhelmed. In some cases, the server can actually crash (or "melt down") under that load.

Firewalls can filter out dangerous packets based on type of packet or place of origin.

If an outsider tries to start more than a few conversations with a server, this is seen as an attack, and the other conversations are dropped.

Only specific types of conversations may be allowed, meaning conversations directed at specific ports.

Exam Question 289 (p.336)**:** How does DDOS defeat firewall protection for servers?
Acceptable Answer: many outsiders can each start a conversation

Because each conversation is coming from a different client, the firewall cannot limit them based on IP address.

Firewalls provide a defense between yourself and the threats you are trying to avoid.

Outside Threats: Your home router probably provides the best overall defense against outside threats. Typically the router does network address translation (NAT) as part of sharing your Internet connection among several computers at the same time. This makes each of those computers effectively invisible to everyone outside of your local area network. However, your home router has no capability to protect you against inside threats.

Inside Threats: Your operating system typically provides the only defense against inside threats. The protection varies from system to system, and specific observations are beyond the scope of this document and this course. However, it is common to allow files or printers to be shared within your LAN. This sharing often requires special steps by the owner of the resource (the files or the printer).

DMZ and Port Forwarding are settings in your router. They can open up one machine, or parts of several machines, to communication (and therefore to possible abuse) from things outside your local area network. However, they are necessary if you want your machine to act as a server. This typically happens (a) for gaming, if you want to host a game, (b) for running a personal web site, and (c) for serving video, such as making surveillance cameras at home visible to you at work.

Exam Question 290 (p.336)**:** What does DMZ stand for?

Acceptable Answer: demilitarized zone

Demilitarized zone suggests that the place is not protected, at least by the router or its firewall.

Exam Question 291 (p.336)**:** What service does DMZ provide?
Acceptable Answer: It directs new conversation requests on all ports to one designated machine.

Old Conversations: Normally there are several machines inside the LAN. These machines are actively communicating with the Internet outside. Responses have to get back to the machine that made the request. This is expected traffic, ongoing conversations. It is recognized by having a match in the **NAT address pool** table.

New Conversations: This is traffic being sent to a server.

With DMZ, there is only one machine that receives the new conversation traffic for all different ports. That machine could be a combination of web server, mail server, ftp server, and more.

http://en.wikipedia.org/wiki/DMZ_(computing) has more.

Exam Question 292 (p.336)**:** What service does port forwarding provide?
Acceptable Answer: It directs new conversation traffic on a few ports to a designated server.

Port forwarding always goes to a server. Traffic is from outside clients who are initiating the conversation.

With port forwarding, there can be several machines, each of which receives the traffic for different ports (different services). One could be a web server. Another could be a mail server. Another could be an ftp server. Another could be a game server.

Providing wireless access to strangers can easily give them access to the inside of your network. At that point, your router is no longer providing a firewall between the stranger and your equipment and data. Generally it is a good idea to have WPA2 or better security on your wireless connections.

Exam Question 293 (p.336)**:** How can sharing your Wi-Fi be dangerous?
Acceptable Answer: Bad people might get directly into your LAN. This bypasses your main firewall.

25.5 Firewalls and Gaming

Gaming requires a tighter interaction than does using a web browser. With the browser, time delays can be lengthy and it will merely be annoying. In gaming, delays can get you kicked out of the game.

To improve performance, gaming computers often become servers to each other.

The difficulty with NAT is getting past the router firewall to the desired computer. The computer behind that firewall cannot be reached except through a port number. The port number must have been registered in the **NAT address pool** table.

But there is a way around it. In the router, under the security settings, there is generally a set of options for things like port forwarding.

Say your game requires you to have port 666 open. You can register your computer with the router so that any traffic requesting port 666 will be forwarded to you. That way the port is always open.

If you need more than one port, like say 666 through 671, the router can usually do that for you also.

There is one obvious drawback with port forwarding. Only one computer can be the receiver of packets labeled 666. If two people want to host a game at the same time, you cannot do it. Or even if one person has 666 tied up, and another person wants to host a game, you cannot do it. You are limited to one host per router.

It is not just gaming though. This also affects things like Skype.

Software developers are aware of these problems. At one time they mostly assumed each participant would have a personal and fully routable IP address. Now they are aware that many participants are behind NAT firewalls.

Over time many of these work-arounds may cease to be needed. But for now, port forwarding is the primary key to success in getting through the firewall.

25.6 The Ping of Death

The **Ping of death** provides one interesting example.

http://en.wikipedia.org/wiki/Ping_of_death has more.

Ping is a command that is used to debug the network. Its primary purpose is to find out whether computer A can talk to computer B. To achieve this, computer A can send a **ping** packet to computer B. Computer B receives the packet, which it was not expecting, and simply replies to it. It is called an **echo**.

Computer A acts as a client. Computer B acts as a server.

But ping is not the important part. Every packet goes through the same process, whether ping or response to a browser request. It is all the same.

Information physically arrives as a bits (layer 1) and is reassembled into a frame. The frame is handed up the stack to layer 2, where it is verified and the headers are discarded. Then the packet within the frame is handed up the stack to layer 3. The packet headers are examined and discarded and the segment within the packet is placed into a buffer where it can eventually be handed up the stack to layer 4.

At this point the mischief begins. With the **Ping of death**, the packet claims to be a fragment of the original packet. Fragments can happen. It is no big deal. They just get reassembled on the receiving end.

The target computer will take that data and begin to reassemble the original data. It has a pool of buffers available for this reassembly process. Each buffer is 65536 bytes big. 65536 is two to the 16th power, and it is a hard upper limit to the size of a packet.

If the new data starts near the end of that buffer, it is possible that it will go past the end of the buffer. This is called a **buffer overflow**. Anything nearby that was stored in memory beyond the real end of the buffer could be destroyed.

Once that memory had been destroyed (overwritten by the bogus packet fragment), any programs that relied on the destroyed data could fail. Specifically, computers who received and processed such a packet could crash immediately due to programming errors that did not anticipate such a bogus packet.

There are several solutions. Some are mentioned in the Wikipedia article. For our purposes the thing to notice is that computers behind NAT firewalls will probably not receive the packet in the first place.

Ping of death is just one example of malicious “letter bomb” packets that

cause trouble when they are opened by the receiving computer. And **buffer overflow** is just one way in which the malice can be carried out.

Another opportunity for trouble can come from having **Remote Desktop** activated. This tool is designed to let someone far away see your desktop and help you when you are stuck. But untrusted people could also see and modify your desktop to "help themselves."

Unit VI

IPv4 Addressing

Chapter 26

Number Bases

Contents

In chapter 7 (page 46) we introduced IPv4 addressing. Here we carry forward.

The main thrust of this chapter is an exploration of the IP address, more properly called the IPv4 address. "v4" means version 4. There is a new standard emerging, IPv6. It will eventually replace IPv4, but people are slow to change, and IPv4 is still the dominant way to provide addresses on the Internet.

As a basic skill, students must be able to manipulate IPv4 addresses and to answer various questions about them.

The most important question is whether two IP addresses are in the same LAN (local area network) or not.

Let's dive right in.

26.1 What is an IP Address?

The actual way that computers are identified on the Internet is by way of something called an **IP address**. These are typically written as four numbers connected by dots. For example, `216.228.254.20` is the IP address of the computer that sits on my desk at the university where I teach.

Numbers are a fine way for computers to find each other, but humans like words instead. We see that in the numerous clever attempts to convert telephone numbers into words. Imagine 1-555-SAVE-NOW. Compare that to 1-555-728-3669. Which one do you think is easier to remember? For me it is easier to dial the digits, but easier to remember the words.

Chapter 4 (page 25) reviews the issues that surround domain names.

A whole industry has grown up around the providing of domain names. They are big business. And it makes it pretty easy to identify the business or organization that you are visiting on the web.

Domain names are just a front door way to get to the IP address. The IP address is the real way by which your computer communicates with other computers on the Internet.

26.2 Vocabulary

Skill: Know the names of various notations used in IPv4 addressing.

Exam Question 294 (p.337)**:** What does a dotted quad number look like?
Acceptable Answer: It is a format where four numbers, usually in base 10, are connected by dots. For example, "123.123.123.123".

Besides dotted quad, there is another notation, used for net masks, called CIDR or "slash" notation. We will discuss it in section 28.1 (page 191). CIDR stands for Classless Inter-Domain Routing.

Exam Question 295 (p.337)**:** When an IPv4 address is written in x.x.x.x format, the possible value for x range from 0 to what?
Acceptable Answer: 255

26.3 Numbering Systems

Normally we use **base 10** when we do numeric calculations. Base 10 is the system of numbering that uses ten digits (zero through nine) to form numbers.

Base 10 is the number system we most commonly use in everyday life. But it is not the only one we commonly use.

The second most common system is **base 60**. We use it for keeping time. The minute after 10:59 is 11:00. The second after 2:59:59 is 3:00:00. We commonly divide time into hours, minutes, and seconds. Minutes run from zero to 59, and then they roll over, starting a new hour, with minutes returning to zero.

We also use **base 60** for geometric things like navigation: Degrees, minutes, and seconds. Just like with time, a minute is 1/60.

Very new, but coming up in third place is probably **base 16**, also called **hexadecimal**, **hex**, or a hex code. It shows up in colors on web pages, where `#66ff99` indicates a color with a certain amount of red, green, and blue. In base 16, the number after 39 is 3a (or 3A), and the number after 3f (or 3F) is 40. With hex, capitalization does not matter.

It is believed that humans use base 10 because that's how many fingers most people have. That makes base 10 somewhat natural for humans.

Computers use **base 2** for engineering reasons. It makes a lot of things easier. But numbers in base 2 can get long to write. They have slightly more than three times as many digits as the same number would have in base 10.

To get around that wordiness, base 2 is often converted into some other power of two for convenience.

We use powers of a base often for grouping.

In base 10 we group the 10s in groups of three, and separate them by commas. For example, the number after 999 is 1,000. The number after 999,999 is 1,000,000. It is even more obvious when we try to say the numbers. We say 999 thousand, 999. Or 999 million, 999 thousand, 999.

Sidenote:

Interestingly, some Asian countries group their powers of 10 in groups of 4. 10000 ("mahn") is the base for large numbers, instead of the 1000 that is

commonly used in western countries.

And some years ago the British usage was to group large numbers by six digits. A **million** was a one followed by six zeroes. A **billion** had 12 zeroes. A **trillion** had 18 zeroes.

http://en.wikipedia.org/wiki/Billion talks more about the short billion (nine zeroes) and the largely-obsolete long billion (twelve zeroes)

Current usage has a million at 6, a billion at 9, and a trillion at 12. It makes less sense, but it is still the current usage.

End of Sidenote.

26.4 Base 2 Groupings

Each digit in a base 2 number is called a **bit**, which is short for **binary digit**.

Starting with numbers in base 2, we can group by threes to get **base 8** (2x2x2, **octal**). We can group by fours to get **base 16** (2x2x2x2, **hex**). And we can group by eights to get **base 256**.

Octal lets us express numbers in a computer-friendly way with roughly the same number of digits as we would use in base 10. The down side is that grouping by three is awkward since there are usually an even number of total bits to be represented.

Hex (hexadecimal) lets us express numbers in a computer-friendly way with fewer digits than base 10. Grouping by four is wonderful. The down side is you need 16 digits, so beyond 9 we use the letters A through F, which is awkward.

Base 256 is a cross between base 10 and base 2. We group bits into groups of eight, also called octets. Then we translate each **octet** into base 10. (This is exactly what we do with minutes and seconds, except they are base 60.) Base 256, also called **dotted quad**, is the notation used for IP addresses, net masks, and many related concepts in networking.

Skill: Be familiar with the notation for powers, especially **powers of two**.
Mem: 2^5 means 2 to the fifth power, and means you multiply 2 by itself 5 times. 2x2x2x2x2 = 32. It is also written as `2^5`.
Mem: 2^n means 2 to the nth power, and means you multiply 2 by itself n

times. It is also written as `2^n`.
Mem: 10^n means 10 to the nth power, and means you multiply 10 by itself n times. It is also written as `10^n`.

Skill: Be familiar with the number bases used in networking.

Exam Question 296 (p.337)**:** Describe base 2.
Acceptable Answer: (a) It is a numbering system. (b) The digits range from 0 to 1. (c) Each digit represents 1 bit. (d) It is also called binary.

Its current popularity is because of its usefulness with computing.

Exam Question 297 (p.337)**:** Describe base 8.
Acceptable Answer: (a) It is a numbering system. (b) The digits range from 0 to 7. (c) Each digit represents 3 bits. (d) It is also called octal.

Its current popularity is because of its usefulness with computing. The advantage over base 2 is that numbers are shorter (only 1/3 as many digits). The advantage over base 16 is that only normal digits are used, not A-F.

Exam Question 298 (p.337)**:** Describe base 10.
Acceptable Answer: (a) It is a numbering system. (b) The digits range from 0 to 9. (c) Each digit represents about 3 1/3 bits. (d) It is also called decimal.

Because each digit uses about 3 1/3 bits, three digits uses about 10 bits. (Actually each digit uses log(10)/log(2) = 3.322 bits, but probably only an engineer or mathematician would care.)

Its use dates back at least to Egypt, 3000 BC, and is common in nearly all cultures today.

http://en.wikipedia.org/wiki/Base_10 has more.

Exam Question 299 (p.337)**:** Describe base 16.
Acceptable Answer: (a) It is a numbering system. (b) The digits range from 0 to 9 and A to F. (c) Each digit represents 4 bits. (d) It is also called hex or hexadecimal.

A=10, B=11, C=12, D=13, E=14, F=15.

Its current popularity is because of its usefulness with computing. The advantage over base 2 is that numbers are shorter (only 1/4 as many digits). The advantage over base 8 is that groups of 4 bits work much better than groups of 3 bits.

Exam Question 300 (p.337)**:** Describe base 60.

Acceptable Answer: (a) It is a numbering system. (b) The digits range from 0 to 59. (c) It is used for time and angle measurement, including navigation. (d) Because the digits are multi-character, in writing they are normally separated by another character.

For time, the separator is the colon ":" and for angle measurement we put a ' after the minutes and a " after the seconds.

The origin seems to be Sumerian or Babylonian, and dates from maybe 2500 BC. The full name of a second is a 'second minute'. 1/60 of a 'second' is actually called a 'third'.

http://en.wikipedia.org/wiki/Base_60 has more.

Exam Question 301 (p.337)**:** Describe base 64.

Acceptable Answer: (a) It is a coding system (not a numbering system). (b) The characters include A-Z, a-z, 0-9, and two others. (c) Each character represents 6 bits. (d) It is used for transmitting data as though it were text.

http://en.wikipedia.org/wiki/Base64 has more.

It is important for transmitting arbitrary text through email relays that may accidentally interpret the traffic (data) to be in-band commands.

http://en.wikipedia.org/wiki/In-band has more on in-band and out-of-band signalling.

Exam Question 302 (p.337)**:** Describe base 256.

Acceptable Answer: (a) It is a numbering system. (b) It uses multi-character 'digits' that range from 0 to 255. (c) Each digit represents 8 bits. (d) In writing the digits are normally separated by dots.

Chapter 27

IPv4 Addresses: Advanced

Contents

Exam Question 303 (p.337)**:** What is an octet?
Acceptable Answer: 8 bits

Bits and bytes are related, but they are not the same thing. A **bit** is one **binary digit**.

Exam Question 304 (p.337)**:** How many bits in a byte?
Required Answer: 8

In networking, a **byte** is almost always eight bits.

Exam Question 305 (p.337)**:** How many bits in a nybble?
Required Answer: 4

In eating, a nibble is a small bite. So half of a byte (in computing) is called a **nybble**. Some consider this to be cute. (I admit, I do too.)

Dotted binary looks like 10110101.10001001.11010010.01101001. It simply translates each part of dotted quad into base 2.

27.1 Writing Numbers

For brevity, instead of requiring people to explicitly say the word "octal" or "hex", a traditional notation has grown up that is now widely accepted in the computing world.

Octal numbers are traditionally written with a zero in front. Thus, octal 755 would be written as 0755.

Exam Question 306 (p.337)**:** The number 0755 is assumed to be in what number base?
Required Answer: 8

This tradition is to distinguish base-8 (octal) numbers from base-10 (decimal) numbers. Base-10 numbers are written with a non-zero digit at the front (except for zero itself, which has the same value in octal or decimal).

There is a potential problem when writing two-digit month numbers, such as 08 for August and 09 for September. Fortunately the meaning is obvious because 8 and 9 are not octal digits.

Exam Question 307 (p.337)**:** The number 755 is assumed to be in what number base?
Required Answer: 10

Hexadecimal numbers are traditionally written with a "**0x**" in front. Thus, hex 755 would be written as 0x755. It does not matter whether it is a lowercase x or an uppercase X. Hex 755 could be written as 0X755, but lowercase is easier to read.

Exam Question 308 (p.337)**:** The number 0x755 is assumed to be in what number base?
Required Answer: 16

There seems to be a developing tradition to express binary numbers with a prefix of "**0b**" or "0B". This is not universally accepted, but is gaining momentum.

The key notational concept seems to be that a leading zero means "this is not base 10" and the next character tells what base it really is. If 0-7, it is octal. If x or X, hex. If b or B, binary.

27.2 What's a Kilo?

Just as there is a short billion and a long billion (see 26.3, page 175 above), there is a "short" **kilo** and a "long" kilo. (I just made those names up. I don't know if anyone else calls them that. I am not sure what their standard name is.)

Skill: Understand that kilo and kilo may mean different things.

Exam Question 309 (p.337)**:** What are the two meanings of kilo?
Acceptable Answer: Ten to the third power, and two to the tenth power.

A marketing kilo is ten to the third, a thousand: 1000. Humans typically use base-10 numbers. Ten to the third power means three tens multiplied together. The result is 1000.

We usually write ten to the third as 10^3 or as `10^3`.

A computer engineering kilo is two to the tenth: 1024. Computers are designed to use base-2 numbers. Two to the 10th power means ten twos multiplied together. The result is 1024.

We usually write two to the tenth as 2^{10} or as `2^10`.

Because 1000 and 1024 are so close to the same number, they are both called a kilo.

Computer engineering is more closely tied to the underlying operation of the equipment. Since computational equipment including networking equipment uses base 2 as its native numbering system, 1024 makes good sense.

Marketing tends to like basing things on 1000 because it is easily defensible and it makes computing equipment look somewhat bigger than what the computer engineers would say.

Exam Question 310 (p.338)**:** What are the two meanings of meg?
Acceptable Answer: Ten to the sixth power. Two to the 20th power.

A marketing meg is ten to the sixth, a million: 1000 * 1000.

A computer engineering meg is two to the twentieth: 1024 * 1024.

Exam Question 311 (p.338)**:** What are the two meanings of gig?
Acceptable Answer: Ten to the ninth power. Two to the 30th power.

A marketing gig is ten to the ninth, a billion: 1000 * 1000 * 1000.

A computer engineering gig is two to the thirtieth: 1024 * 1024 * 1024.

Example: I bought a 2T external hard drive recently. By 2T, marketing meant 2,000,000,000,000 bytes. Computer engineering would have meant 2 x 1024 x 1024 x 1024 x 1024. So in computer engineering terms, my 2T drive is about 91% of 2T. In any case, I'm happy. It is a lot of storage. Just be aware that there are two languages in use here.

Skill: Be familiar with units of measure for information quantity.
Mem: **bit** means one binary digit, either a zero or a one.
Mem: **nybble** means four bits.
Mem: **byte** normally means eight bits.
Mem: **kilo** means $2^{10} = 1024$ for computing and 1000 for marketing.
Mem: **mega** means $2^{20} = 1024^2$ for computing and 1000^2 for marketing.
Mem: **giga** means $2^{30} = 1024^3$ for computing and 1000^3 for marketing.
Mem: **tera** means $2^{40} = 1024^4$ for computing and 1000^4 for marketing.
FYI: **peta** means 2^{50}. 1 petabyte is about 1000 terabytes.
FYI: **exa** means 2^{60}. 1 exabyte is about 1000 petabytes.
FYI: **zetta** means 2^{70}. 1 zettabyte is about 1000 exabytes.
FYI: **yotta** means 2^{80}. 1 yottabyte is about 1000 zettabytes.
Mem: **KB** means kilobyte, one thousand bytes.
Mem: **Kb** means kilobit, one thousand bits.
Mem: **Kbps** means kilobits per second.
Mem: **MB** means megabyte.
Mem: **Mb** means megabit.
Mem: **Mbps** means megabits per second.
Mem: **Mb/s** means megabits per second.
Mem: **GB** means gigabyte, one "gig".
Mem: **Gb** means gigabit.
Mem: **Gbps** means gigabits per second.

27.3 Popular Numbers

In Chapter 8 (page 50) we explained how to convert between the common number bases used in networking: base 2, 8, 10, and 16.

Not all numbers are equally popular. In day-to-day living, the number 100 comes up a lot more often than the number 97. Similarly, 5 is more popular than 3 or 4 or 6 or 7.

Numbers like 10, 100, and 1000 are popular because they are "round," meaning that they end in a lot of zeroes.

In networking as well some numbers are much more popular than others. And it is because they are round in some sense. For networking, round means how they look in base 2, binary. And they come up a lot more often than the non-round numbers.

Skill: Quickly recognize and use the common IPv4 numbers: Powers of 2. These are the only numbers that appear in (dotted quad) subnet block sizes. (Spaces have been added for clarity, but normally they are left out.)
Mem: binary 00000000 is 0
Mem: binary 0000000 1 is $2^0 = 1$ (multiply no 2s)
Mem: binary 000000 1 0 is $2^1 = 2$ (multiply one 2)
Mem: binary 00000 1 00 is $2^2 = 4$ (multiply two 2s)
Mem: binary 0000 1 000 is $2^3 = 8$ (multiply three 2s)
Mem: binary 000 1 0000 is $2^4 = 16$ (multiply four 2s)
Mem: binary 00 1 00000 is $2^5 = 32$ (multiply five 2s)
Mem: binary 0 1 000000 is $2^6 = 64$ (multiply six 2s)
Mem: binary 1 0000000 is $2^7 = 128$ (multiply seven 2s)

Skill: Quickly recognize and use the common IPv4 numbers: Negative Powers of 2. These are the only numbers that appear in (dotted quad) net masks. They are also the boundaries between the address classes. (Spaces have been added for clarity, but normally they are left out.)
Mem: binary 11111111 is $256 - 2^0 = 255$
Mem: binary 1111111 0 is $256 - 2^1 = 254$
Mem: binary 111111 00 is $256 - 2^2 = 252$
Mem: binary 11111 000 is $256 - 2^3 = 248$
Mem: binary 1111 0000 is $256 - 2^4 = 240$ (also start of class E)
Mem: binary 111 00000 is $256 - 2^5 = 224$ (also start of class D)
Mem: binary 11 000000 is $256 - 2^6 = 192$ (also start of class C)
Mem: binary 1 0000000 is $256 - 2^7 = 128$ (also start of class B)
Mem: binary 00000000 is $256 - 2^8 = 0$ (also start of class A)

We will talk about classes and network masks shortly.

27.4 Pre-1981 Network.Host Addressing

Before 1981, the Internet, then called the ARPANet, used 8-bit network addresses and 24-bit host addresses. This is before the introduction of classes A, B, and C.

IP addresses were originally conceived in a network.host naming format. A **host** is an individual computer.

The idea was that within a network, hosts could communicate directly to one another. Whatever happens in Las Vegas, stays in Las Vegas, so to speak. The Internet did not want to know or be bothered by the details of how each local area network was organized. The Internet only cared about communication between networks.

If computer 20.1234 wanted to talk to computer 20.5678, they could do so without involving the Internet. Both were on network 20. They could just talk by whatever rules were used on network 20.

If computer 20.1234 wanted to talk to computer 30.5678, they would need to involve the Internet. In the simple case, there would be one machine that had two wires. One wire would connect it to network 20 and the other wire would connect it to network 30. That machine would act as a gateway, or go-between, for networks 20 and 30. Often the gateway got a special number, like .1. It could be 20.1 on the 20 network, and 30.1 on the 30 network.

Typically the connecting wires were telephone lines.

Of course, for 256 networks to communicate, it is not necessary for each to have a private line to the other 255. If six degrees of separation works for humans, why not computer networks as well?

http://en.wikipedia.org/wiki/Six_degrees_of_separation has more on this six degrees concept. It makes for interesting reading. Someone even made a movie.

Long story short, by careful programming we could have a message from 20.1234 go to machine 99.1111 through a number of intermediaries. It might go first to 20.1, alias 30.1, which could pass it along to 30.77, alias 77.1, and so on until it reached the 99 network for final delivery.

Routing involves having a small number of core machines that know where everything is. Pass a message to one of these core machines and they can pass it on to the proper branch of the network. The core routers were operating based on hand-crafted routing tables. Each time a new network came online, the tables would be adjusted. This did not happen often at first, so it was no big problem.

Chapter 36 (page 258) goes deeper into routing.

27.5 Classful Addressing

http://en.wikipedia.org/wiki/Classful_network gives good background on the origin of classful networks.

There were to be 256 networks, and each network could have like 16 million hosts. But that's not how things turned out. As prices for computers came down, it turned out that there was a large demand for many more networks, and the networks were much smaller than 16 million hosts. So the numbering got rearranged.

Classful addressing started about 1981.

The original networks 1 through 126 remained as before. They are called Class A networks. Each one has 16 million hosts.

The 64 networks from 128 to 191 became Class B. Instead of just being those 64 networks, each was divided 256 ways, creating 16 thousand smaller networks. You don't get something for nothing though. Each network had only 1/256 as many hosts, which is $2^{16} = 65536$ hosts. (We are being a bit sloppy here. It is actually only 65534. We will clean up our sloppiness later.)

The 32 networks from 192 to 223 became Class C. Instead of just being those 32 networks, each was divided 256 ways, and then 256 ways again, creating around 2 million networks, each having $2^8 = 256$ (really 254) hosts.

And life was good for a while. This way of dividing things became known as **classful addressing** and it prevailed for many years. Eventually the addresses started to run out. Then these address classes were further divided, not by the powers that be, but by the end users, the companies that were using them.

They started to create sub-networks, called subnets, and the addressing was called **classless addressing**. We will get to that in section 28.2 (page 192), but first we need to concentrate on the world of Classful Addressing.

Skill: Given an IPv4 address, tell what the class is.
Mem: 1-126.0.0.0 is class A. (starts with 0xxxxxxx)
Mem: 128-191.0.0.0 is class B. (starts with 10xxxxxx)
Mem: 192-223.0.0.0 is class C. (starts with 110xxxxx)
Mem: 224-239.0.0.0 is class D. Multi Cast. (1110xxxx)

Mem: 240-255.0.0.0 is class E. Experimental. (1111xxxx)

Exam Question 312 (p.338)**:** For 19.19.19.19, what class is it?
Acceptable Answer: Class A

Exam Question 313 (p.338)**:** For 199.199.199.199, what class is it?
Acceptable Answer: Class C

27.6 Network Masks

Since the network.host split had become variable, instead of always being an 8.24 split, programs had to be updated. Programmers are lazy and hate to update anything twice if they can spend the extra time to get it updated correctly the first time.

To handle this variation in splits, a lot of programming was done using the knowledge that class A was always an 8.24 split, and class B was always a 16.16 split, and class C was always a 24.8 split.

net mask: Given a computer address, like 20.1234, programmers needed to separate that address into its two parts. They did this mathematically by using something called a mask. The mask consisted of ones for the things they wanted to keep and zeroes for the things they wanted to get rid of. The logic behind this relates to AND and OR, with 1 standing for TRUE and 0 standing for FALSE. Computers can do this calculation really fast.

For class A, the net mask is 11111111.00000000000000000000000. It consists of eight 1s followed by 24 0s.

For class B, the net mask is 1111111111111111.0000000000000000. It consists of 16 1s followed by 16 0s.

For class C, the net mask is 111111111111111111111111.00000000. It consists of 24 1s followed by eight 0s.

Programmers, being human and not machines (a statement that some might dispute), found two improvements to this notation. The first was to separate the 32 bits into four groups of 8. Thus:

For class A, the net mask is 11111111.00000000.00000000.00000000.

For class B, the net mask is 11111111.11111111.00000000.00000000.

For class C, the net mask is 11111111.11111111.11111111.00000000.

The second was to convert each set of 8 into base 10 numbering. Maybe this was for the sake of their managers. In any case, the result was shorter.

For class A, the net mask is 255.0.0.0.

For class B, the net mask is 255.255.0.0.

For class C, the net mask is 255.255.255.0.

That is because 11111111 (eight 1s) is the base 2 (binary) number that is equivalent to 255 in base 10 (decimal). And 00000000 in any base is always the same: 0.

Skill: Understand net masks.

Exam Question 314 (p.338)**:** What does a net mask look like?
Acceptable Answer: 32 bits, with all the 1s first, then all the 0s.

For IPv4, the total is 32 bits. It is normally written in base 10 notation, consisting of four numbers, each of which represents 8 bits. Here is an example written in base 2.

```
11111111.11111111.11110000.00000000
```

It can also be written in CIDR notation, with a slash followed by the number of binary 1s. (CIDR stands for Classless Inter-Domain Routing. We explain it in section 28.1, page 191.) Here is that same example written in CIDR notation:

```
/20
```

Exam Question 315 (p.338)**:** In a net mask, what do the 1s mean?
Acceptable Answer: The 1s indicate the part of the address that is the network number.

Exam Question 316 (p.338)**:** In a net mask, what do the 0s mean?
Acceptable Answer: The 0s indicate the part of the address that is the host number.

27.7 Special Addresses

Some IPv4 addresses were given special significance. They were not just any old address. Here is the list.

Skill: Tell the IPv4 special address ranges.

Mem: 0.0.0.0 is the local network.
Mem: 10.x.x.x is the class A private address range.
Mem: 127.0.0.0 is the localhost address range.
Mem: 127.0.0.1 is the local host.
Mem: 169.254.x.x is the **link local** self-assigned address range.
Mem: link local is also called **APIPA**.
Mem: 172.16-31.x.x is the class B private address range.
Mem: 192.168.x.x is the class C private address range.
Mem: 255.255.255.255 is the IPv4 global broadcast address.

Exam Question 317 (p.338)**:** What is the IPv4 special address range for the local network?
Required Answer: 0.0.0.0/8

Exam Question 318 (p.338)**:** What is the Class A Private Address Range?
Required Answer: 10.0.0.0/8

Exam Question 319 (p.338)**:** In the Class A Private Address Range, how many (classful) networks are there?
Required Answer: 1

Exam Question 320 (p.338)**:** In the Class A Private Address Range, what is the first IP address?
Required Answer: 10.0.0.0

Exam Question 321 (p.338)**:** In the Class A Private Address Range, what is the first usable host address?
Required Answer: 10.0.0.1

Exam Question 322 (p.338)**:** In the Class A Private Address Range, what is the last IP address?
Required Answer: 10.255.255.255

Exam Question 323 (p.338)**:** In the Class A Private Address Range, what is the last usable host address?
Required Answer: 10.255.255.254

Exam Question 324 (p.338)**:** What is the local host Private Address Range?
Required Answer: 127.0.0.0/8

Exam Question 325 (p.338)**:** What is the Link Local (APIPA) Private Address Range?

Required Answer: 169.254.0.0/16

Exam Question 326 (p.338)**:** How are Link Local (APIPA) addresses assigned?
Acceptable Answer: each computer picks its own

They are not "assigned." They are self-chosen or self-assigned. Because they are self-assigned, there is some small risk of duplication within a LAN. Two computers could pick the same IP address. Because there are $2^{16} = 65536$ APIPA addresses, the risk is small.

Exam Question 327 (p.338)**:** What is the Class B Private Address Range?
Required Answer: 172.16.0.0/12

The /12 means the first 12 bits constitute the fixed part of the address. The other 20 bits can be almost anything.

Exam Question 328 (p.338)**:** In the Class B Private Address Range, how many (classful) networks are there?
Required Answer: 16

Exam Question 329 (p.338)**:** In the Class B Private Address Range, what is the first IP address?
Required Answer: 172.16.0.0

Exam Question 330 (p.339)**:** In the Class B Private Address Range, what is the first usable host address?
Required Answer: 172.16.0.1

Exam Question 331 (p.339)**:** In the Class B Private Address Range, what is the last IP address?
Required Answer: 172.31.255.255

Exam Question 332 (p.339)**:** In the Class B Private Address Range, what is the last usable host address?
Required Answer: 172.31.255.254

Exam Question 333 (p.339)**:** What is the Class C Private Address Range?
Required Answer: 192.168.0.0/16

Exam Question 334 (p.339)**:** In the Class C Private Address Range, how many (classful) networks are there?
Required Answer: 256

Exam Question 335 (p.339)**:** In the Class C Private Address Range, what is the first IP address?

Required Answer: 192.168.0.0

Exam Question 336 (p.339)**:** In the Class C Private Address Range, what is the first usable host address?
Required Answer: 192.168.0.1

Exam Question 337 (p.339)**:** In the Class C Private Address Range, what is the last IP address?
Required Answer: 192.168.255.255

Exam Question 338 (p.339)**:** In the Class C Private Address Range, what is the last usable host address?
Required Answer: 192.168.255.254

Exam Question 339 (p.339)**:** What is the IPv4 global broadcast address?
Required Answer: 255.255.255.255

RFC 1918 at http://datatracker.ietf.org/doc/rfc1918/ covers private address ranges: 10, 172, 192.

http://datatracker.ietf.org/doc/rfc3927/ covers link local addresses, also known as **APIPA**.

Chapter 28

IPv4 Addresses: Classless

Contents

28.1 CIDR Routing

Instead of having 256 networks, Classful Addressing gave us millions. This had an impact on the routing tables in those core routers. Nobody wants millions of lines of routing information, especially when 256 lines was working just fine, thank you very much.

CIDR started about 1993 as a way to consolidate groups of networks into a single entry in the routing table. It also opened the door to classless addressing.

Chapter 36 (page 258) goes deeper into routing.

Exam Question 340 (p.339)**:** What does CIDR stand for?

Acceptable Answer: classless inter-domain routing

Exam Question 341 (p.339)**:** What does CIDR notation look like?
Acceptable Answer: It is written as a slash followed by a number between 0 and 32; for example, /21

CIDR is also called **slash notation**. It uses a /16 to indicate that the first 16 bits of the IP address are the network bits. If we are consolidating 4 class B addresses, and they each have the same first 14 bits in their network number, we can call it a /14 super-net.

Exam Question 342 (p.339)**:** What is the other name for slash notation?
Acceptable Answer: cidr

Exam Question 343 (p.339)**:** What is the other name for CIDR notation?
Acceptable Answer: slash

Skill: Explain CIDR "slash" notation.
Mem: /8 = 11111111.00000000.00000000.00000000 = 255.0.0.0
Mem: /16 = 11111111.11111111.00000000.00000000 = 255.255.0.0
Mem: /24 = 11111111.11111111.11111111.00000000 = 255.255.255.0
Mem: /xx means the first xx bits of the net mask are 1s. The rest are 0s. There are 32 total bits.

Skill: Given an IPv4 address, tell the default Net Mask in CIDR and dotted quad notation.
Mem: For class A, CIDR is /8, net mask is 255.0.0.0
Mem: For class B, CIDR is /16, net mask is 255.255.0.0
Mem: For class C, CIDR is /24, net mask is 255.255.255.0

Exam Question 344 (p.339)**:** For 199.199.199.199, what is the default Net Mask in CIDR and dotted quad notation?
Acceptable Answer: /24 and 255.255.255.0.

28.2 Classless Addressing

As mentioned above, organizations started to create sub-networks, called subnets, and the addressing was called **classless addressing**. We are now ready to talk about it.

Making some original addresses into class B and class C addresses was simple and elegant. 16 million hosts per network was a bit much. As computer

prices came down and networking prices came down, there was great pressure to further subdivide the B and C addresses (and the A addresses too).

It was too late to redesign the overall plan for Classful Addressing. Too many people were using it. Too many programs depended on it.

But CIDR and the net mask concept provided an opportunity to go the other way. Instead of aggregating things into super-nets, they could be further divided into subnets.

To keep clear the division between network bits and host bits, it would be necessary to utilize an explicit subnet mask telling exactly how many bits were network. There was already an implicit mask based on the first number in the IP address. Now things were going to be expressed and not just assumed.

28.3 Subnet Block Size

Calculating **subnet block size** is useful.

Subnetting is the process of dividing a class A, B, or C address block into smaller blocks, and associating those blocks as separate networks.

Since base 2 is very convenient for networking, things are most efficient when done by doubling or splitting in half. That keeps it simple for the network programmers and fast for the networking equipment.

The cost of this subnetting was that net masks got more complicated. Instead of just consisting of 255s and 0s, we got a bunch of other numbers that you met back in section 27.3 (page 182).

Skill: Given a CIDR subnet, tell what the subnet mask is in dotted quad notation. (Be able to do this for /8 through /30.)

Q: For /21, what is the subnet mask?
A: /20 = 11111111.11111111.11110000.00000000 = 255.255.240.0.
A: /21 = 11111111.11111111.11111000.00000000 = 255.255.248.0.
A: /22 = 11111111.11111111.11111100.00000000 = 255.255.252.0.
A: /23 = 11111111.11111111.11111110.00000000 = 255.255.254.0.
A: /24 = 11111111.11111111.11111111.00000000 = 255.255.255.0.
A: /25 = 11111111.11111111.11111111.10000000 = 255.255.255.128.
A: /26 = 11111111.11111111.11111111.11000000 = 255.255.255.192.
A: /27 = 11111111.11111111.11111111.11100000 = 255.255.255.224.

A: /28 = 11111111.11111111.11111111.11110000 = 255.255.255.240.
A: /29 = 11111111.11111111.11111111.11111000 = 255.255.255.248.

Subnet block size is a key number in calculating many other important numbers.

We will show you several ways to calculate the subnet block size. You will have to do this skillfully, so read through these methods and pick the one that makes the most sense to you. Then learn it well.

(p2) Powers of Two Method: The block size is just 2 raised to the power of however many bits are available for the host address. Remove as many 8s as you can before converting to base 10. Then add a .0 for each 8 you removed.

(spl) Split Block Method: Figure out which quad has the split within it. For /8, /16, and /24, the split falls between quads which makes everything very easy. For the rest, /9-15 are in quad 2, /17-23 are in quad 3, and /25-30 are in quad 4.

Within that quad, determine how many host bits there are, and take 2 raised to that power.

(sub) Subtraction Method: By careful design, the subnet mask plus the block size always equals zero. If you know the net mask, you can subtract it from zero to get the block size. It is very convenient to do this in base 256.

(add) Addition Method: Sometimes it is easier to add than subtract. Given the subnet mask, what do we need to add to reach zero? Remember that the numbers are in base 256.

(tbl) Table Method: You can construct the following table from memory. Then use it to find the answers to subnetting questions.

The first four columns are labeled Q1 (quad 1) through Q4 (quad 4). If the slash is between /1 and /8, then the important quad is quad 1. If the slash is between /17 and /24, then the important quad is quad 3.

The "pow" column has powers of two. The "neg" column has negative powers of two (in base 256). If you add the numbers in columns five and six, you always get 256.

Skill: Given a CIDR subnet, tell what the subnet block size is in dotted quad notation.

Exam Question 345 (p.340)**:** For /10, what is the subnet block size?

Subnet Table

Q1	Q2	Q3	Q4	pow	neg
/1	/9	/17	/25	128	128
/2	/10	/18	/26	64	192
/3	/11	/19	/27	32	224
/4	/12	/20	/28	16	240
/5	/13	/21	/29	8	248
/6	/14	/22	/30	4	252
/7	/15	/23	/31	2	254
/8	/16	/24	/32	1	255

Acceptable Answer: 0.64.0.0

(p2) 32 minus 10 = 22 host bits. $2^{22} = 2^6$ x 2^8 x 2^8 = 0.64.0.0.

(spl) 10 is in quad 2, with a 2/6 split. $2^6 = 64$. Put 64 in quad 2.

(sub) 0.0.0.0 minus 255.192.0.0 (/10 net mask) = 0.64.0.0.

(add) 255.192.0.0 (/10 net mask) + 0.64.0.0 = 0.0.0.0.

(tbl) /10 is in Q2 and has pow 64. Put a 64 in quad 2.

Exam Question 346 (p.340)**:** For /23, what is the subnet block size?
Acceptable Answer: 0.0.2.0

(p2) 32 minus 23 = 9 host bits. $2^9 = 2^1$ x 2^8 = 0.0.2.0.

(spl) 23 is in quad 3 with a 7/1 split. $2^1 = 2$. Put 2 in quad 3.

(sub) 0.0.0.0 minus 255.255.254.0 (/23 net mask) = 0.0.2.0.

(add) 255.255.254.0 (/23 net mask) + 0.0.2.0 = 0.0.0.0.

(tbl) /23 is in Q3 and has pow 2. Put a 2 in quad 3.

Exam Question 347 (p.340)**:** For /28, what is the subnet block size?
Acceptable Answer: 0.0.0.16

(p2) 32 minus 28 = 4 host bits. 2^4 = 0.0.0.16.

(spl) 28 is in quad 4 with a 4/4 split. 2^4 16. Put 16 in quad 4.

(sub) 0.0.0.0 minus 255.255.255.240 (/28 net mask) = 0.0.0.16.

(add) 255.255.255.240 (/28 net mask) + 0.0.0.16 = 0.0.0.0.

(tbl) /28 is in Q4 and has pow 16. Put a 16 in quad 4.

Skill: Given an IPv4 subnet mask, tell what the subnet block size is in dotted quad notation.

Exam Question 348 (p.340)**:** For 255.255.248.0, what is the subnet block size?
Acceptable Answer: 0.0.8.0

(sub) 0.0.0.0 minus 255.255.248.0 = 0.0.8.0

(add) 255.255.248.0 + 0.0.8.0 = 0.0.0.0

(tbl) 248 in the neg column, in quad 3, gives you /21.

28.4 Subnet Count

Calculating **subnet count** is useful.

There are two ways to count subnets. The old way, in effect until about 2005, is called **no subnet-zero**. The new way, in effect since 2005, is called **subnet-zero** or "with" subnet-zero.

The subnet-zero issue is what to do with the first and last usable subnets in the network. The old-school idea was that no subnet should use those slots. Any exam questions you might find online from that era probably assume this old-school approach.

The current idea is that every subnet slot can be used. Under this method, which is widely accepted, the number of subnets is just the power of two that matches the number of bits available for the subnet portion of the address.

With no subnet-zero, just take the subnet-zero count and subtract two from it.

Skill: Given an IPv4 address and a CIDR subnet, tell how many subnets there are, assuming subnet-zero.

Q: For 150.150.150.150/21, how many subnets are there (with subnet-zero)?
A: $2^5 = 32$.
A? 21-16 (class B) = 5 bits.

Skill: Given an IPv4 address and a CIDR subnet, tell how many subnets there are, assuming no subnet-zero.

Q: For 150.150.150.150/21, how many subnets are there (no subnet-zero)?

A: $2^5 - 2 = 30$.

28.5 First Usable Subnet

Calculating the **first usable subnet** is useful.

To determine the first usable subnet, you need to know the subnet block size and the (classful) network address.

With subnet-zero, the first usable subnet address is the network address.

With no subnet-zero, the first usable subnet address is the network address plus the block size.

Skill: Given an IPv4 address and a CIDR subnet, tell what the first usable subnet address is in dotted quad notation, assuming subnet-zero.

Q: For 150.150.150.150/21, what is the first usable subnet address (with subnet-zero)?
A: 150.150.0.0.
A? With subnet-zero, the first usable subnet is always the same as the network address.

Skill: Given an IPv4 address and a CIDR subnet, tell what the first usable subnet address is in dotted quad notation, assuming no subnet-zero.

Q: For 150.150.150.150/21, what is the first usable subnet address (no subnet-zero)?
A: 150.150.8.0.
A? With no subnet-zero, the first usable subnet is always the network address plus the block size. 150.150.0.0 + 0.0.8.0 = 150.150.8.0.

28.6 Last Usable Subnet

Calculating the **last usable subnet** is useful.

To determine the last usable subnet, you need to know the subnet block size and the (classful) network address of the NEXT network.

With subnet-zero, the last usable subnet address is the next network address minus the block size.

With no subnet-zero, the last usable subnet address is the next network address minus two block sizes.

In calculating the next network address, be careful to remember whether the network is class A, B, or C.

Skill: Given an IPv4 address and a CIDR subnet, tell what the last usable subnet address is in dotted quad notation, assuming subnet-zero.

Q: For 150.150.150.150/21, what is the last usable subnet address (with subnet-zero)?
A: 150.150.248.0.
A? It is always the next network address minus the block size. 150.151.0.0 minus 0.0.8.0 = 150.150.248.0.

Skill: Given an IPv4 address and a CIDR subnet, tell what the last usable subnet address is in dotted quad notation, assuming no subnet-zero.

Q: For 150.150.150.150/21, what is the last usable subnet address (no subnet-zero)?
A: 150.150.240.0.
A? It is always the next network address minus the block size x2. 150.151.0.0 minus 0.0.8.0 x2 = 150.150.240.0.

28.7 Current Subnet

Calculating the **current subnet** is useful.

There are four questions you should be able to answer about the current subnet. What is the subnet address, the first usable host address, the last usable host address, and the broadcast address.

To solve this, you need to know the subnet block size. You need to be able to count in multiples of that number.

If the block size is 1, 2, 4, or even 8, counting is pretty easy.

For 16, the multiples would be 0, 16, 32, 48, 64, 80, 96, 112, 128, 144, 160, 176, 192, 208, 224, and 240. That's pretty hard.

For 32, the multiples would be 0, 32, 64, 96, 128, 160, 192, and 224. That's pretty hard.

For 64, the multiples would be 0, 64, 128, and 192. That's pretty easy again.

For 128, the multiples would be 0 and 128. That's way easy.

So, watch out for the 16s and 32s.

You don't actually need all of those numbers. You just need the one that is before and the one that is after the current IP address.

Let's work through an example: For 1.2.3.150/28, what are the subnet address, first usable host, last usable host, and broadcast address?

Quad 4 is where the split occurs. The block size is 16. The multiples of 16 that surround 150 are 144 and 160.

Why quad 4?

Before /8 is quad 1. After /8 and before /16 is quad 2. After /16 and before /24 is quad 3. After /24 and before /32 is quad 4.

/28 is after /24 and before /32, so quad 4 is where the split occurs.

Why multiples of 16?

/32 is the end of the quad. From /28 to /32 we have four bits for host addressing. $32 - 28 = 4$

Four bits of host addressing gives us 16 addresses. $2^4 = 16$

The **subnet address** is 144 because that is the multiple of 16 that comes at or before 150.

The **first usable host** is always just one beyond the subnet address.

The **next subnet address** is 160 because that is the multiple of 16 that comes after 150.

The **broadcast address** is always just one before the next subnet address.

The **last usable host** is always just two before the next subnet address.

Skill: Given an IPv4 address and a CIDR subnet, tell what the current subnet address is, meaning the subnet in which that IPv4 address occurs.

Q: For 150.150.150.150/21, what is the current subnet address?
A: 150.150.144.0.

Skill: Given an IPv4 address and a CIDR subnet, tell what the first usable host address is in dotted quad notation.

Q: For 150.150.150.150/21, what is the first usable host address?
A: 150.150.144.1.

Skill: Given an IPv4 address and a CIDR subnet, tell what the last usable host address is in dotted quad notation.

Q: For 150.150.150.150/21, what is the last usable host address?
A: 150.150.151.254.

Skill: Given an IPv4 address and a CIDR subnet, tell what the default broadcast address is in dotted quad notation.

Q: For 150.150.150.150/21, what is the default broadcast address?
A: 150.150.151.255.

An important question that is often asked is whether two IP addresses are in the same subnet or not.

Skill: Given an IPv4 address and a CIDR subnet, tell whether another specified IPv4 address is in the same LAN.

Q: For 150.150.150.150/21, is 150.150.149.149 in the same LAN?
A: Yes.

Q: For 150.150.150.150/21, is 150.150.155.155 in the same LAN?
A: No.

28.8 IPv4 Summary

You probably thought this unit would never end. It just about didn't.

What's the take-away? What do you really need to know?

When you are configuring a home router or home wireless, you need to pick an IP address range.

Normally there is a default. Often it is 192.168.0.0/24. You can run into trouble if you have more than one router. They cannot both have the same address. Something has to change.

If you understand how the addressing works, you can make that change.

What else?

As you move into more advanced networking classes, you will be making networks with not just one or two routers, but with lots of them. You need to organize your addresses so they make sense and do not cause conflicts.

This leads you into subnets and block sizes and how many subnets you can

have and how many hosts can be in each subnet. It leads to easy-to-ask questions like, "what is the first usable subnet and the last usable subnet?" And questions like, "what is the first usable host in a subnet, and the last usable host, and the broadcast address?" These show up a lot on certification tests.

What else?

Sometimes the network around you is failing. Common causes are things like the servers dying, power outages, or wires being cut by construction equipment.

If you have networking skills, you can get around some of those problems. You can be up and running while everyone else is taking a break.

Maybe this is not such a smart idea after all ... .

Just kidding.

When you are on the road, and you want to get on the Internet, it can be very helpful to know what's going on and how to fix it.

Chapter 29

VLSM

Contents

Why do all the subnets have to be the same size? Good question. Originally subnets were the same size because the subnet mask was imposed at the network level by the routers. Eventually the routing protocols started carrying around the subnet masks so they could be different in each subnet.

Chapter 36 (page 258) goes deeper into routing.

Exam Question 349 (p.340)**:** What does VLSM stand for?
Acceptable Answer: variable length subnet mask

http://en.wikipedia.org/wiki/VLSM has more on **CIDR** and **VLSM**.

29.1 Explanation

VLSM is a method for dividing up a block of IP addresses into several subnets that are not all the same size.

In this section we show how this can be done. To keep things simple, we will use a very small address space, but the same principles apply to the full-sized IPv4 address space.

We will pretend that IP addresses are only three bits in size. In IPv4 they

are actually 32 bits in size.

Subnet sizes must always be a power of two, that is: 1, 2, 4, 8, 16, 32, etc. In IPv4 the smallest allowed is 4. We will pretend we can use all the way down to 1.

This grid shows how a block of IP addresses can be divided into different subnets.

The top row, outlined here in red, shows the four possible netmasks. A netmask of 110 means that the first two bits identify the network, and the last bit identifies the host within that network.

000	100	110	111
xxx	0xx	00x	000
xxx	0xx	00x	001
xxx	0xx	01x	010
xxx	0xx	01x	011
xxx	1xx	10x	100
xxx	1xx	10x	101
xxx	1xx	11x	110
xxx	1xx	11x	111

The first column represents the IP addresses that fall within the 000 subnet. It contains all eight of the addresses.

000	100	110	111
xxx	0xx	00x	000
xxx	0xx	00x	001
xxx	0xx	01x	010
xxx	0xx	01x	011
xxx	1xx	10x	100
xxx	1xx	10x	101
xxx	1xx	11x	110
xxx	1xx	11x	111

Using a netmask of 100, we have two subnets. One is the 0xx subnet. The other is the 1xx subnet. The xx part represents the host address.

000	100	110	111
xxx	0xx	00x	000
xxx	0xx	00x	001
xxx	0xx	01x	010
xxx	0xx	01x	011
xxx	1xx	10x	100
xxx	1xx	10x	101
xxx	1xx	11x	110
xxx	1xx	11x	111

Split: Using a netmask of 100, subnets have a size of four. But not just any four. They must be four that have the same network number. Here we have selected four IP addresses from the middle of the range. Two of them have 0xx as their IP address, and two of them have 1xx as their IP address. But according to the rules, all of them must have the same network number, and these do not. Two are in network 000 and two are in network 100. This does not work.

000	100	110	111
xxx	0xx	00x	000
xxx	0xx	00x	001
xxx	0xx	01x	010
xxx	0xx	01x	011
xxx	1xx	10x	100
xxx	1xx	10x	101
xxx	1xx	11x	110
xxx	1xx	11x	111

Using a netmask of 110, subnets have a size of two. There are four subnets. Each of the four possible subnets is outlined in red.

000	100	110	111
xxx	0xx	00x	000
xxx	0xx	00x	001
xxx	0xx	01x	010
xxx	0xx	01x	011
xxx	1xx	10x	100
xxx	1xx	10x	101
xxx	1xx	11x	110
xxx	1xx	11x	111

Overlap: The subnets we actually use are not allowed to overlap. In this example, we have selected two subnets. One is the 1xx subnet that includes four hosts: 100, 101, 110, and 111. The other is the 101 subnet including a single host. The problem here is that they overlap. Both subnets include 101. That is not allowed.

000	100	110	111
xxx	0xx	00x	000
xxx	0xx	00x	001
xxx	0xx	01x	010
xxx	0xx	01x	011
xxx	1xx	10x	100
xxx	1xx	10x	101
xxx	1xx	11x	110
xxx	1xx	11x	111

Here we show a way that the eight addresses can be divided into four subnets. One of the subnets has a netmask of 100 and includes four IP addresses. Another has a netmask of 110 and includes two IP addresses. The last two have a netmask of 111 and each includes only a single IP address.

000	100	110	111
xxx	0xx	00x	000
xxx	0xx	00x	001
xxx	0xx	01x	010
xxx	0xx	01x	011
xxx	1xx	10x	100
xxx	1xx	10x	101
xxx	1xx	11x	110
xxx	1xx	11x	111

29.2 Steps To Follow

A network can be divided into subnets so long as each subnet starts on the proper boundary. A subnet of size 0.0.0.4 (/30) must start on a multiple of 0.0.0.4. A subnet of size 0.0.64.0 (/18) must start on a multiple of 0.0.64.0.

To divide a network into subnets, first identify the LANs (broadcast domains) that will be involved. For each LAN determine how many hosts (sometimes called users) will be involved.

For those small router-to-router subnets there are only two hosts (the routers). In this case, you need four IP addresses, and the subnet size is 0.0.0.4.

Others subnets will have actual servers or end users. When you count hosts within a LAN, be sure to include any routers and other less-obvious things like servers, networked storage, photocopy machines, and printers. And in real life you would probably estimate on the high side to allow for a reasonable amount of future growth, but on an exam just go with what they tell you.

Exam Definition: A **user** and a **host** may be the same thing, and maybe not. If there is no other information provided, assume the words mean the same thing.

In any case find the number of hosts, add two (to allow for the first and last IP address in the block), and then round up (if necessary) to the next power of 2.

For example, if you have 47 hosts, you need 49 IP addresses and a subnet size of 0.0.0.64. If you have 63 hosts, you need 65 IP addresses and a subnet size of 0.0.0.128. The result will be your minimum subnet size for that LAN.

At the end of this process, you should know how many LANs you have, and how many IP addresses you need in each LAN.

Assign the locations for your biggest blocks first. If you get them lined up properly, everything else will fit snuggly.

Look at the range of IP addresses that are available to you. Starting with the largest LAN, assign it a location that starts with a multiple of the LAN size. For example, if the LAN size (rounded up to a power of 2) is 64, you must start it at x.x.x.0 or x.x.x.64 or x.x.x.128 or x.x.x.192.

Continue assigning LANs from largest to smallest. So long as you start each on a compatible boundary, any assignment is legal.

Common Errors: (a) make sure each LAN starts on the right boundary, which is a multiple of the LAN block size. (b) make sure the LANs do not overlap each other, or any restricted addresses.

Chapter 30

Ports

Contents

In this chapter we collect together and expand on things we have already said about ports. There is also a discussion of ports in the section on firewalls in chapter 25 (page 162).

Besides having a destination IP address, messages also have a destination Port number.

When you talk to a computer, you don't just talk to a computer. You talk to a computer program inside that computer. You specify the program by specifying a port number. Port numbers range from 0 to 65535.

In effect, when the computer starts up, it also starts up a bunch of programs. Some of those are willing to receive messages. Those programs tell the computer, hey, if you get a message for port 123, send it to me.

Each program on that computer, if it wants to receive messages, has a **port** number.

Exam Question 350 (p.340)**:** What is a software port?
Acceptable Answer: a number that tells which program should receive the message

We have mentioned ports several times. In this section we go over everything we know and try to bring it together.

The word **port** has two meanings. A port (physical) is typically a mostly square hole on the back of a router or switch. You plug in a network cable (**8P8C**, Cat 5 or similar) to connect it to another piece of equipment, such as a computer.

A port (software) is a number between 0 and 65535. You can think of it as a post office box within the computer. The number tells which server (software) should receive the message.

Exam Question 351 (p.340)**:** Software port numbers range from 0 up to what number?
Acceptable Answer: 65535

http://en.wikipedia.org/wiki/List_of_TCP_and_UDP_port_numbers has more information.

Port numbers have a sixteen-bit range, meaning that it covers all of the whole numbers that can be expressed in sixteen bits. 65535 is equal to $2^{16} - 1$.

http://en.wikipedia.org/wiki/Port_numbers has more on port numbers.

30.1 Important Port Numbers

Here are some of the more important port numbers. This is obviously not a complete list.

http://www.iana.org/assignments/port-numbers has an official, complete list of generally recognized port numbers.

30.1.1 80: HTTP

Port 80 is probably the best-known port in the computer world. It is the port normally used by web servers. When you use the **http** protocol in a web address, it tells the browser to talk to port 80 on the server.

Exam Question 352 (p.340)**:** What is port 80 normally used for?
Acceptable Answer: http

Exam Question 353 (p.340)**:** What port does http normally use?
Required Answer: 80

Exam Question 354 (p.340)**:** What does http stand for?
Required Answer: hyper text transfer protocol

http://en.wikipedia.org/wiki/Hypertext_Transfer_Protocol has more.

30.1.2 443: Secure HTTP

Exam Question 355 (p.340)**:** What is port 443 normally used for?
Acceptable Answer: https

Exam Question 356 (p.340)**:** What port does https normally use?
Required Answer: 443

Exam Question 357 (p.340)**:** What does https stand for?
Required Answer: hyper text transfer protocol secure

30.1.3 21: File Transfer Protocol

Exam Question 358 (p.340)**:** What is port 21 normally used for?
Acceptable Answer: ftp

Exam Question 359 (p.340)**:** What port does ftp normally use?
Required Answer: 21

Exam Question 360 (p.340)**:** What does ftp stand for?
Required Answer: file transfer protocol

http://en.wikipedia.org/wiki/File_Transfer_Protocol has more.

30.1.4 22: Secure Shell

Exam Question 361 (p.340)**:** What is port 22 normally used for?
Acceptable Answer: ssh

Exam Question 362 (p.341)**:** What port does ssh normally use?
Required Answer: 22

Exam Question 363 (p.341)**:** What does ssh stand for?
Required Answer: secure shell

30.1.5 23: Telnet

Exam Question 364 (p.341)**:** What is port 23 normally used for?
Acceptable Answer: telnet

Exam Question 365 (p.341)**:** What port does telnet normally use?
Required Answer: 23

30.1.6 25: Email

Exam Question 366 (p.341)**:** What is port 25 normally used for?
Acceptable Answer: email

Exam Question 367 (p.341)**:** What port does smtp normally use?
Required Answer: 25

Exam Question 368 (p.341)**:** What does smtp stand for?
Required Answer: simple mail transfer protocol

Email uses the SMTP protocol.

http://en.wikipedia.org/wiki/Simple_Mail_Transfer_Protocol has more.

30.1.7 53: DNS

Exam Question 369 (p.341)**:** What is port 53 normally used for?
Acceptable Answer: dns

Exam Question 370 (p.341)**:** What port does dns normally use?

Required Answer: 53

Exam Question 371 (p.341)**:** What does dns stand for?
Required Answer: domain name system

The domain name system converts human-readable domain names into numeric IP addresses.

30.1.8 123: Network Time Protocol

NTP is used to keep clocks synchronized across the Internet.

Exam Question 372 (p.341)**:** What is port 123 normally used for?
Acceptable Answer: ntp

Exam Question 373 (p.341)**:** What port does ntp normally use?
Required Answer: 123

Exam Question 374 (p.341)**:** What does ntp stand for?
Required Answer: network time protocol

Often we think of a port like a post office box, and an IP address like a post office.

Many people share the same post office. Many computer programs share the same computer. The most common example of a computer program is a web browser. Other examples include email programs and music downloading tools.

Normally each person (or family) that will receive mail has its own post office box. Normally each computer program that will receive messages has its own port number.

When a letter arrives for someone, the postal workers place the letter in that person's post office box, based on the box number written on the letter. When a network message arrives, it is delivered to the program that owns the port number in that message.

30.2 Making a Request

The purpose for being on the network is to be able to send and receive messages (information). Normally the process starts with a request for information.

The device making the request is called the **client**. The computer that will receive the request and respond to it is called the **server**.

Exam Question 375 (p.341)**:** What does a client do?
Acceptable Answer: make requests

Normally the client expects to receive a response from a server. (In rare cases, the client may not expect a response.)

Exam Question 376 (p.341)**:** What does a server do?
Acceptable Answer: answer requests

The requests come from clients. Normally the server replies to the client. It is said to "service" the requests.

Notice that receiving requests is not the same as answering requests. Hackers send requests to random places all the time. In many cases those requests go unanswered. But if the request gets an answer, then they know they are talking to a server, even if it is a very limited server.

Receiving requests is not the same as receiving messages. Requests start new conversations. The server is ready, willing, and able to service the requests, but it does not know when or if ever the requests will come.

Messages, on the other hand, if they are not requests, are responses to previous messages. Because they are expected, firewalls typically provide a fast way for such messages to get through.

Computers that are on the network but are not servers are called clients. Normally clients are placed behind a firewall to prevent new conversations from reaching them.

It is common for several client programs to be communicating at the same time. Each program will be expecting replies to its requests.

Because a reply is normally expected, the sending computer always includes a return address with the request. The return address includes a port number, assigned seemingly at random, where the reply can be delivered. The port number identifies the program that should be given the reply when it arrives.

30.3 Servicing a Request

Requests are handled by servers. The word **server** can refer to the entire machine on which the request is serviced, or it can refer to the specific computer program that will service the request.

Exam Question 377 (p.341)**:** What two things does server mean?
Acceptable Answer: software that answers requests, hardware where such programs run

Randomly assigned port numbers work fine for clients, but do not work well for servers. Instead, servers use well-known port numbers to receive their requests.

It is important to know port numbers because you may need to open holes in your firewall to allow those ports to be used, or you may want to specifically block them to prevent those services (like email).

http://en.wikipedia.org/wiki/Port_numbers has more on port numbers.

30.4 Being a Server

Normally a desktop or laptop computer acts as a client.

Sometimes it acts as a server. Common examples of this include hosting an online game, or participating in a peer-to-peer network, or engaging in an online phone call using something like Skype.

When you act as a server, you normally have to use a well-known port number for that service so others can connect to you.

30.5 Security

Security is often done on a port-by-port basis. Each well-known port is connected to a service. These are the **vectors** or avenues through which bad guys can try to attack you.

Unit VII

Power Tools

Chapter 31

Basic Power Tools

Contents

In this chapter we look at several tools commonly used by networking professionals for setting up and **troubleshooting** networks.

ipconfig is introduced in 31.1 (page 217).

ping is introduced in 31.2 (page 221).

traceroute is introduced in 32.1 (page 233).

ftp is introduced in 32.2 (page 236).

telnet is introduced in 32.3 (page 237).

Wide Examples

Many of the sections in the chapters on tools include printouts of the actual use of the tool. In many cases these printouts are wide, extending beyond the width of the page. Some material is lost.

We decided to allow these printouts extend as they will, and in some cases be cut off at the edge of the page, because (a) the material lost is of small value, (b) folding the lines makes things more confusing (and it is more work), and (c) shrinking the font size is harder on the eyes, especially because (a) the material lost is of small value.

31.1 ipconfig

What is your IP address? What is your networking configuration?

Most computers have a command that allows you to see the network configuration details, and sometimes to change them.

On Microsoft Windows, that command is **ipconfig**.

On Linux and Mac OS X systems, that command is **ifconfig**.

The details that are reported are somewhat cryptic. Examples are shown below. This command provides a dump of almost all the information that might be useful to a networking professional. Most details are beyond the scope of this book, but the IPv4 address is important and can be found here.

If the IPv4 address starts with 169.254, it was self-assigned. It means that your computer did not find a DHCP server within a reasonable amount of time, so it made up an IP address for itself. You will not be able to communicate to the Internet. If you should be able to reach the Internet, you may need to do something that will cause the machine to try again to get an IP address. Sometimes that means turn it off and back on.

31.1.1 ipconfig (Windows)

http://en.wikipedia.org/wiki/Ipconfig has more.

The **ipconfig** command will tell you your own IP address. This is important information. It will often give you your gateway address as well. If not, you can usually guess it based on your own IP address.

Following is a response to a Windows **ipconfig** request.

The machine is running Windows 7. It discovers an IPv4 address of 192.168.0.20 and a gateway of 192.168.0.1.

```
windows7> ipconfig

Windows IP Configuration

Ethernet adapter Local Area Connection:

   Connection-specific DNS Suffix  . : earthlink.com
   Link-local IPv6 Address . . . . . : fe80::8886:4c85:99f1:d903%11
   IPv4 Address. . . . . . . . . . . : 192.168.0.20
   Subnet Mask . . . . . . . . . . . : 255.255.255.0
   Default Gateway . . . . . . . . . : 192.168.0.1

Tunnel adapter isatap.earthlink.com:

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

Tunnel adapter Teredo Tunneling Pseudo-Interface:

   Connection-specific DNS Suffix  . :
   IPv6 Address. . . . . . . . . . . : 2001:0:4137:9e76:100e:3901:bf34:edc9
   Link-local IPv6 Address . . . . . : fe80::100e:3901:bf34:edc9%12
   Default Gateway . . . . . . . . . : ::
```

Here is ipconfig on the same machine using the /all option.

```
windows7> ipconfig /all

Windows IP Configuration

   Host Name . . . . . . . . . . . . : Nebula
   Primary Dns Suffix  . . . . . . . :
   Node Type . . . . . . . . . . . . : Hybrid
   IP Routing Enabled. . . . . . . . : No
   WINS Proxy Enabled. . . . . . . . : No
```

```
   DNS Suffix Search List. . . . . : earthlink.com

Ethernet adapter Local Area Connection:

   Connection-specific DNS Suffix  . : earthlink.com
   Description . . . . . . . . . . . : Realtek PCIe GBE Family Controller
   Physical Address. . . . . . . . . : 40-61-86-99-56-7A
   DHCP Enabled. . . . . . . . . . . : Yes
   Autoconfiguration Enabled . . . . : Yes
   Link-local IPv6 Address . . . . . : fe80::8886:4c85:99f1:d903%11(Preferred)
   IPv4 Address. . . . . . . . . . . : 192.168.0.20(Preferred)
   Subnet Mask . . . . . . . . . . . : 255.255.255.0
   Lease Obtained. . . . . . . . . . : Thursday, March 24, 2011 12:33:32 PM
   Lease Expires . . . . . . . . . . : Saturday, March 26, 2011 12:33:33 PM
   Default Gateway . . . . . . . . . : 192.168.0.1
   DHCP Server . . . . . . . . . . . : 192.168.0.1
   DHCPv6 IAID . . . . . . . . . . . : 239100294
   DHCPv6 Client DUID. . . . . . . . : 00-01-00-01-13-2A-FB-2F-40-61-86-99-56-7A
   DNS Servers . . . . . . . . . . . : 207.69.188.186
                                       207.69.188.187
   NetBIOS over Tcpip. . . . . . . . : Enabled

Tunnel adapter isatap.earthlink.com:

   Media State . . . . . . . . . . . : Media disconnected
   Connection-specific DNS Suffix  . : earthlink.com
   Description . . . . . . . . . . . : Microsoft ISATAP Adapter
   Physical Address. . . . . . . . . : 00-00-00-00-00-00-00-E0
   DHCP Enabled. . . . . . . . . . . : No
   Autoconfiguration Enabled . . . . : Yes

Tunnel adapter Teredo Tunneling Pseudo-Interface:

   Connection-specific DNS Suffix  . :
   Description . . . . . . . . . . . : Teredo Tunneling Pseudo-Interface
   Physical Address. . . . . . . . . : 00-00-00-00-00-00-00-E0
   DHCP Enabled. . . . . . . . . . . : No
   Autoconfiguration Enabled . . . . : Yes
   IPv6 Address. . . . . . . . . . . : 2001:0:4137:9e76:100e:3901:bf34:edc9(Preferred
   Link-local IPv6 Address . . . . . : fe80::100e:3901:bf34:edc9%12(Preferred)
```

```
   Default Gateway . . . . . . . . . : ::
   NetBIOS over Tcpip. . . . . . . . : Disabled
```

31.1.2 ifconfig (Linux, Mac OS X)

http://en.wikipedia.org/wiki/Ifconfig provides an overview of **ifconfig** (Unix).

Following are typical responses to **ifconfig** requests.

Here is an excerpt from ifconfig from a Linux Ubuntu install. It discovers an IPv4 (inet) address of 216.228.254.20. It does not reveal a gateway, but it does reveal a broadcast address of 216.228.254.255.

```
ubuntu> ifconfig

Link encap:Ethernet  HWaddr 00:21:9b:4c:6d:39
inet addr:216.228.254.20  Bcast:216.228.254.255  Mask:255.255.255.0
inet6 addr: fe80::221:9bff:fe4c:6d39/64 Scope:Link
UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
RX packets:1317297 errors:0 dropped:0 overruns:0 frame:0
TX packets:670804 errors:0 dropped:0 overruns:0 carrier:0
collisions:0 txqueuelen:100
RX bytes:403240486 (403.2 MB)  TX bytes:224166179 (224.1 MB)
Memory:febe0000-fec00000
```

Here is ifconfig from an Apple MacBook Pro running OS X. It discovers an IPv4 (inet) address of 192.168.1.126. It does not reveal a gateway, but it does reveal a broadcast address of 192.168.1.255.

```
macosx> ifconfig

lo0: flags=8049<UP,LOOPBACK,RUNNING,MULTICAST> mtu 16384
        inet6 ::1 prefixlen 128
        inet6 fe80::1%lo0 prefixlen 64 scopeid 0x1
        inet 127.0.0.1 netmask 0xff000000
gif0: flags=8010<POINTOPOINT,MULTICAST> mtu 1280
stf0: flags=0<> mtu 1280
fw0: flags=8863<UP,BROADCAST,SMART,RUNNING,SIMPLEX,MULTICAST> mtu 4078
        lladdr 00:1f:f3:ff:fe:0a:29:ca
```

```
        media: autoselect <full-duplex>
        status: inactive
en1: flags=8823<UP,BROADCAST,SMART,SIMPLEX,MULTICAST> mtu 1500
        ether 00:1e:c2:bf:90:51
        media: autoselect (<unknown type>)
        status: inactive
en0: flags=8863<UP,BROADCAST,SMART,RUNNING,SIMPLEX,MULTICAST> mtu 1500
        ether 00:1f:5b:eb:56:59
        inet6 fe80::21f:5bff:feeb:5659%en0 prefixlen 64 scopeid 0x6
        inet 192.168.1.126 netmask 0xffffff00 broadcast 192.168.1.255
        media: autoselect (100baseTX <full-duplex,flow-control>)
        status: active
vmnet8: flags=8863<UP,BROADCAST,SMART,RUNNING,SIMPLEX,MULTICAST> mtu 1500
        ether 00:50:56:c0:00:08
        inet 172.16.202.1 netmask 0xffffff00 broadcast 172.16.202.255
vmnet1: flags=8863<UP,BROADCAST,SMART,RUNNING,SIMPLEX,MULTICAST> mtu 1500
        ether 00:50:56:c0:00:01
        inet 172.16.20.1 netmask 0xffffff00 broadcast 172.16.20.255
```

31.2 ping

In case of network problems, **ping** is often the first tool that is used to understand the problem.

The **ping** command verifies connectivity and also measures **latency** (network delay).

If ping works, the network is connected and working between the two computers that were involved. Lots of things had to be working. Ping proves that they were all working.

If you let ping run for a while, you will see one of three things. (a) nothing is getting through. (b) everything is getting through. (c) sometimes things get through and sometimes they do not.

Do not worry about the first ping or two. Sometimes parts of the network go to sleep when they are not in use, and it can take a second or two to wake up. Specifically, they may have tables of information that has gotten stale, and they may need to refresh it.

If you have case (c), we call that an **intermittent** error. Intermittent means

that it happens sometimes but not other times.

Exam Question 378 (p.341)**:** What does intermittent mean?
Acceptable Answer: Sometimes it happens. Sometimes it does not happen.

Intermittent (IN-ter-MIT-unt) often refers to failure or success. Intermittent failure and intermittent success are the same thing. Sometimes it works. Sometimes it does not work.

Exam Question 379 (p.341)**:** What is another word for intermittent?
Acceptable Answer: flaky

Flaky is another word for intermittent.

If you have an intermittent failure on a ping, it often means the cable has a loose connection somewhere.

Exam Question 380 (p.342)**:** What does a successful ping prove?
Acceptable Answer: All the pieces involved were working AT THAT MOMENT.

Exam Question 381 (p.342)**:** What does a failed ping prove?
Acceptable Answer: ONE OR MORE pieces involved were not working at that moment.

If ping fails, there are many possible reasons. More testing is required. Maybe the cable is broken. Maybe a router is powered off. Maybe a firewall is stopping the echo request or the echo response. It is really hard to tell what the exact cause might be.

Doing the exact same test again may tell you if the problem is intermittent or constant.

Exam Question 382 (p.342)**:** What one statistic does ping report?
Required Answer: latency

Latency is the round-trip echo time between two computers.

It also reports the IP address of the computers involved (but that is not really a statistic).

ping is sometimes said to stand for Packet Internet Groper. Really it stands for the sounds made by submarines and ships using active sonar equipment. They send out a pulse of sound waves and listen for the echos that come back.

http://en.wikipedia.org/wiki/Ping has more.

The syntax is `ping 12.34.56.78` where 12.34.56.78 is replaced by the IPv4 address that you wish a response from. Generally you can use a domain name instead, and ping will use DNS to convert it into an IPv4 address for you.

Within the network, **ping** sends a small message to someone. By Internet rules they were required to respond. In today's world of heightened paranoia and security worries, these requests are sometimes filtered by firewalls.

The conventional wisdom is to begin troubleshooting by pinging these four things: (1) localhost, (2) yourself, (3) something else inside your LAN, typically your gateway, and (4) something beyond your LAN.

We mention the conventional wisdom because you may find it on other networking exams, so it's good to know.

We specifically call it "conventional" because it may not actually be the smartest way to do things. But on some tests you need to pretend that it is.

Exam Question 383 (p.342)**:** Using ping to troubleshoot, conventional wisdom says you should ping what first?
Acceptable Answer: yourself, using your localhost address: 127.0.0.1

This demonstrates that your own network stack (software) is working.

Exam Question 384 (p.342)**:** Using ping to troubleshoot, conventional wisdom says you should ping what second?
Acceptable Answer: yourself, using your LAN address

This demonstrates that your own network stack (software) is working, and your network interface card is working.

Exam Question 385 (p.342)**:** Using ping to troubleshoot, conventional wisdom says you should ping what third?
Acceptable Answer: something else inside your LAN, typically your router or gateway

This demonstrates that your own network stack (software) is working, and your network interface card is working, and your wiring is working, and something else on your LAN is working.

Exam Question 386 (p.342)**:** Using ping to troubleshoot, conventional wisdom says you should ping what fourth?
Acceptable Answer: something beyond your LAN

This demonstrates that your own network stack (software) is working, and your network interface card is working, and your wiring is working, and your router is working, and something beyond your LAN is working.

That is the conventional wisdom. You start small and expand, especially when you are having trouble.

In an ideal world, all four steps will be successful. In a paranoid world (like we actually live in), some of them will fail even though they should be successful.

Actually, because your local machine may have firewalls that interfere, you should go in this order: (x) ping 255.255.255.255, (3) ping your gateway, (4) ping something beyond your network, and if necessary (1) ping localhost, and (2) ping yourself.

Each ping can give you information about things that are working. Once you have figured out what is working, and narrowed down what is not working, then you can go about fixing it.

If (4) works for even one web site, then your Internet connection is up and running. None of the other pings matter.

If (4) fails, try other web sites. Perhaps the first one is just blocking ping.

If (4) fails for all web sites, but (3) works, then your personal computer is working and your gateway (router) is working, but your ISP link is probably down. Power off your modem. Wait a minute. Then apply power again. Wait a minute. Then try pinging it again. Power off and power on will reset the modem and often solves the networking problem. If the problem persists, you are ready to call your ISP and report that your Internet seems to be down.

If (4) fails and (3) fails but (2) works, then your personal computer is working and your gateway is down. Remove power to the gateway. Wait a minute. Then apply power again. Wait a minute. Then try pinging it again. Power off and power on will reset the gateway and often solves the networking problem.

31.2.1 ping localhost

The (conventional) first target for **ping** is yourself, using your inside address. This address is called the **localhost** address, and is the same for every

computer.

Exam Question 387 (p.342)**:** What is localhost?
Acceptable Answer: Localhost is the computer you are using.

On every computer, this address will be 127.0.0.1.

Exam Question 388 (p.342)**:** What is the IPv4 address of localhost?
Required Answer: 127.0.0.1

Following is a successful ping to the local interface. The machine is running Ubuntu Linux. Notice that the times are very short. After the first ping, subsequent pings are around 0.009 ms, or 9 millionths of a second. It is this fast because the local area network is not actually involved.

```
ubuntu> ping 127.0.0.1
PING 127.0.0.1 (127.0.0.1) 56(84) bytes of data.
64 bytes from 127.0.0.1: icmp_seq=1 ttl=64 time=0.026 ms
64 bytes from 127.0.0.1: icmp_seq=2 ttl=64 time=0.010 ms
64 bytes from 127.0.0.1: icmp_seq=3 ttl=64 time=0.008 ms
64 bytes from 127.0.0.1: icmp_seq=4 ttl=64 time=0.008 ms
```

Here is a failed ping to the localhost interface. The machine is running Mac OS X. In this case, ping does work for other things, as we find out later, but not for localhost. Because it works on other hosts, we presume it fails on localhost because of a firewall.

```
macosx> ping 127.0.0.1
PING 127.0.0.1 (127.0.0.1): 56 data bytes
Request timeout for icmp_seq 0
Request timeout for icmp_seq 1
Request timeout for icmp_seq 2
Request timeout for icmp_seq 3
```

Exam Question 389 (p.342)**:** What things does a successful ping to localhost prove?
Acceptable Answer: Your network stack (software) is working.

A failed ping could mean the stack is not working, or it could mean that ping response is firewalled. If any other pings are successful, then it is a firewall. Try other pings before you come to a conclusion.

If all pings fail, it could be that your internal network protocol stack has failed and you will need to reinstall your network drivers. Reinstalling is work, so try everything else first.

31.2.2 Step 2: ping Yourself

If you can ping localhost successfully, the (conventional) second target is to ping your own outside IPv4 address. By which we mean your Local Area Network address.

Mostly this works if you have a static IP address.

You can find this address through **ipconfig**.

Exam Question 390 (p.342)**:** How can you (the user) find your IP address?
Acceptable Answer: Use ipconfig (or ifconfig).

If you do not have ipconfig or ifconfig available, you could also ping 255.255.255.255 and see the responses to identify your own IPv4 address.

Your computer itself typically uses DHCP to discover its IP address, or it has a static IP address that is keyed in by a human.

Nowadays, many computers are mobile (laptops, etc.,) and get their IPv4 address through a DHCP process. But if the network is down, DHCP does not work. If DHCP does not work, your computer will use APIPA to self-assign an address.

On a typical home network, your address will be 192.168.0.100.

Following is a successful ping. The machine is running Ubuntu Linux. Notice that the times are very short. After the first ping, subsequent pings are around 0.009 ms, or 9 millionths of a second. It is this fast because the local area network is not actually involved.

```
ubuntu> ping 192.168.0.100
PING 192.168.0.100 (192.168.0.100) 56(84) bytes of data.
64 bytes from 192.168.0.100: icmp_seq=1 ttl=64 time=0.022 ms
64 bytes from 192.168.0.100: icmp_seq=2 ttl=64 time=0.009 ms
64 bytes from 192.168.0.100: icmp_seq=3 ttl=64 time=0.008 ms
64 bytes from 192.168.0.100: icmp_seq=4 ttl=64 time=0.009 ms
```

Following is an unsuccessful ping. The machine is running Mac OS X. Other pings work, so we conclude that this one is being blocked by a firewall.

```
macosx> ping 192.168.0.100
PING 192.168.0.100 (192.168.0.100): 56 data bytes
Request timeout for icmp_seq 0
Request timeout for icmp_seq 1
Request timeout for icmp_seq 2
Request timeout for icmp_seq 3
```

Exam Question 391 (p.342)**:** What things does a successful ping to your own LAN address prove?

Acceptable Answer: 1. Your network stack (software) is working and 2. your network interface hardware is working.

It does not prove that your network is working. Let's use an analogy. All you did is step outside, turn around, and knock on your own front door. You did not verify that the sidewalk or street exists, or that people can get around out there.

If ping fails to get a response from yourself, it could be that your machine has a firewall that prevents response. Try other pings before you come to a conclusion.

If all pings fail, it could be that your network interface card has failed and your computer needs a hardware repair. Repairs are expensive, especially needless ones, so try everything else first.

31.2.3 Step 3: ping Your Neighbor

By neighbor, we mean another device on your local area network.

Even though conventional wisdom identifies this as step 3, this is probably the thing you should try first. It is listed as step 3 because that is the conventional wisdom.

Ping something else (besides yourself) that is inside your local area network. Usually you should try your gateway.

You can usually find your gateway address through **ipconfig**. If not, you can often guess it.

To guess, change the last number of your own IPv4 address to 1 (the first legal host address) and try that. If not, change it to 254 (the last legal host address) and try that. If not, change it to 255 (the broadcast address) and

try that.

Exam Question 392 (p.342)**:** What things does a successful ping to your neighbor prove?

Acceptable Answer: 1. Your network stack (software) is working, 2. your network interface hardware is working, 3. your network connection (cable) is working, 4. your neighboring computer is working.

Using the same analogy as above, you step outside, walk down the street, and knock on your neighbor's door. If you get a response, then lots of things must be working.

If the gateway ping is successful, then you can conclude that your network stack is fine, and your network interface card is fine, even though they may have failed to ping properly.

If the gateway ping is successful but the ping to localhost failed, you can conclude that a firewall is in place, preventing the localhost pings.

Following is a successful ping to a gateway. The host machine is running Mac OS X. The gateway is a home networking router. Notice that the times are short, around 2 ms, or 2000 millionths of a second. It takes more time than localhost because the local area network is involved.

```
macosx> ping 192.168.0.1
PING 192.168.0.1 (192.168.0.1): 56 data bytes
64 bytes from 192.168.0.1: icmp_seq=0 ttl=63 time=3.479 ms
64 bytes from 192.168.0.1: icmp_seq=1 ttl=63 time=2.223 ms
64 bytes from 192.168.0.1: icmp_seq=2 ttl=63 time=1.979 ms
64 bytes from 192.168.0.1: icmp_seq=3 ttl=63 time=2.298 ms
```

A successful ping proves your network stack (software) is working and your network interface card (hardware) is working and your cable is working and your local area network is working. It means you do not have to test localhost.

If ping fails to get a response from your gateway, it typically means that your gateway has failed, or the cable between you and your gateway has failed. It can also mean that your gateway is not accepting ping requests, but this would be unusual because ping is such a valuable troubleshooting tool.

31.2.4 Step 3a: ping Your Neighbors

By neighbors, we mean the other devices on your local area network. We do not mean devices beyond your LAN.

Sometimes you cannot identify your gateway from the **ipconfig** information, and your guesses are not successful. It may be time to try the broadcast address.

pinging the broadcast address is often a good way to find what other IP addresses are in use, but it should be used with care. See **nmap** in section 33.3 (page 242) for another useful tool for mapping out the network.

http://en.wikipedia.org/wiki/Smurf_attack has more on pings that involve broadcast addresses, the dangers that they can pose, and the reasons they may not work.

Exam Question 393 (p.342)**:** What is a smurf attack?
Acceptable Answer: A broadcast ping with a fake source address.

Exam Question 394 (p.342)**:** What is a broadcast ping with a fake source address called?
Acceptable Answer: A smurf attack

Since about 1999, Smurf attacks have not been a real threat because broadcast pings seldom get past the first router. All replies will come from within the local area network. But out of an abundance of caution, some computers will not reply to broadcast pings for fear they are Smurf attacks.

Below is a ping to a broadcast address. Notice two things. First, we are getting back responses that are labeled (DUP!). That is because every device in the local area network is responding, including ourself.

Second, we are getting back the IPv4 addresses of all devices on the network. In this case, they are *.126, *.1, and *.130.

Exam Question 395 (p.342)**:** What information can a broadcast ping provide?
Acceptable Answer: (a) IP addresses and (b) latencies of all devices within your LAN.

Judging by the times listed, the first one is coming from ourself, at 0.108 ms. The second one is from the router, at 1.205 ms. The third one is from some other device on the LAN. In this case, it is an iPod Touch with its Wi-Fi activated. It took much longer to respond.

```
macosx> ping 192.168.1.255
PING 192.168.1.255 (192.168.1.255): 56 data bytes
64 bytes from 192.168.1.126: icmp_seq=0 ttl=64 time=0.108 ms
64 bytes from 192.168.1.1: icmp_seq=0 ttl=64 time=1.205 ms (DUP!)
64 bytes from 192.168.1.130: icmp_seq=0 ttl=64 time=9.390 ms (DUP!)
64 bytes from 192.168.1.126: icmp_seq=1 ttl=64 time=0.142 ms
64 bytes from 192.168.1.1: icmp_seq=1 ttl=64 time=1.162 ms (DUP!)
64 bytes from 192.168.1.130: icmp_seq=1 ttl=64 time=109.981 ms (DUP!)
```

31.2.5 Step 3b: ping Your Neighbors

Sometimes you cannot run **ipconfig**. It may be time to try the global broadcast address: 255.255.255.255.

We first introduced this in section 18.1.1 (page 108).

Review section 31.2.4 (page 229). The same warnings apply here.

Exam Question 396 (p.342)**:** What IP address is used for global broadcast?
Required Answer: 255.255.255.255

With the global broadcast address, ping provides much the same information as pinging the local broadcast address, but without requiring you to know your local address. Very handy, when it works.

Routers will generally stop global broadcasts, or any broadcasts, and restrict them to the local area network. That is because of the mischief they could cause in a world-wide Internet.

```
macosx> ping 255.255.255.255
PING 255.255.255.255 (255.255.255.255): 56 data bytes
64 bytes from 192.168.1.126: icmp_seq=0 ttl=64 time=0.116 ms
64 bytes from 192.168.1.1: icmp_seq=0 ttl=64 time=3.577 ms (DUP!)
64 bytes from 192.168.1.130: icmp_seq=0 ttl=64 time=17.144 ms (DUP!)
64 bytes from 192.168.1.126: icmp_seq=1 ttl=64 time=0.160 ms
64 bytes from 192.168.1.1: icmp_seq=1 ttl=64 time=1.084 ms (DUP!)
64 bytes from 192.168.1.130: icmp_seq=1 ttl=64 time=42.293 ms (DUP!)
```

31.2.6 Step 4: ping Something Beyond

Conventional wisdom says to do this fourth, but there is no harm in doing it first, except maybe wasted time.

If you can successfully ping something substantially beyond your local area network, you have proved that your network is up and all the pieces in between are working, at least some of the time.

Exam Question 397 (p.342)**:** What things does a successful ping to something beyond your LAN prove?

Acceptable Answer: 1. Your network stack (software) is working, 2. your network interface hardware is working, 3. your network connection (cable) is working, and 4. your gateway router is working.

That is a lot to know. That is very good information.

But normally if you are pinging things, it is because you have failed to reach a web site that you can normally reach. You are trying to figure where the fault lies so you can fix it.

So, we start by pinging things that are close, and work our way out to things that are farther away.

The router is close. We did that in step 3.

Next, we might try the cable or dsl modem. If you know the IP address, give it a try. Here the local cable modem is at 192.168.100.1. (That is not an uncommon value.)

```
macosx> ping 192.168.100.1
PING 192.168.100.1 (192.168.100.1): 56 data bytes
64 bytes from 192.168.100.1: icmp_seq=0 ttl=62 time=3.312 ms
64 bytes from 192.168.100.1: icmp_seq=1 ttl=62 time=3.208 ms
64 bytes from 192.168.100.1: icmp_seq=2 ttl=62 time=3.421 ms
```

Beyond that, you might know the name of a web site provided by your ISP. It is probably close to you, in network terms. Give it a try.

Beyond that, you might have some well-known sites that are always up and running. On a typical home network, you might choose google.com.

Following is a successful ping to google.com. The host machine is running Mac OS X. Ignoring the first ping, which is often non-typical, the later pings are still much longer than the 3 ms we saw above. Instead they are around

75 ms, or 75,000 millionths of a second. It takes far more time than the local router because the Internet with multiple hops is involved.

```
macosx> ping google.com
PING google.com (74.125.224.178): 56 data bytes
64 bytes from 74.125.224.178: icmp_seq=0 ttl=52 time=143.750 ms
64 bytes from 74.125.224.178: icmp_seq=1 ttl=52 time=79.072 ms
64 bytes from 74.125.224.178: icmp_seq=2 ttl=52 time=87.303 ms
64 bytes from 74.125.224.178: icmp_seq=3 ttl=52 time=66.365 ms
```

Chapter 32

Intermediate Power Tools

Contents

32.1 traceroute

traceroute is like **ping** on steroids. Pick a distant destination (or a close one), and **traceroute** will tell you the identities and ping times (latencies) for each hop along the way.

It can be used to discover where the bottlenecks or routing errors might be occurring.

http://en.wikipedia.org/wiki/Traceroute provides an overview of **traceroute** (Unix) and **tracert** (Windows).

Exam Question 398 (p.343)**:** What two things does traceroute report?
Acceptable Answer: (a) The list of routers between source and destination. (b) The latency for each router.

traceroute does this by manipulating the TTL values on the ping packets.

The **TTL** (time to live) parameter controls how far a packet will go before it expires. Each time a packet is forwarded through a router, the TTL value is decreased by one (decremented). If the value ever reaches zero, the

packet "expires," is dropped, and an error message is returned to the original sender. This message identifies the router at which the packet expired.

Exam Question 399 (p.343)**:** What does TTL stand for?
Required Answer: time to live

Exam Question 400 (p.343)**:** What is the purpose of TTL?
Acceptable Answer: TTL stops infinite loops.

For normal packets, TTL should never reach zero. For packets that are being improperly routed in a loop, TTL prevents infinite loops. traceroute makes clever use of this built-in defense mechanism.

Exam Question 401 (p.343)**:** How does traceroute use TTL?
Acceptable Answer: It sends packets with varying TTL counts. As each packet dies, its location of death is reported.

Error messages from dropped packets are sent using **ICMP**, the Internet Control Message Protocol.

In the following example from my home, we can see the first two hops are through non-routable addresses within my residence.

Hop 3 is the ISP.

Hops 4 through 7 are through Road Runner, the company that provisioned my line to my ISP, Earthlink. Hop 4 is on Oahu. The other three are socal, Southern California.

Hop 5 has two entries. This indicates that several packets, each with a TTL of 5, died at different places. The route from 4 to 6 has at least two alternative paths.

Hop 8 is tbone. The bone part suggest that this is a backbone segment at the core of the Internet. The lax part indicates Los Angeles.

From hop 9 through hop 12, the routers are not named. They only have IP addresses. Hop 9 also has two entries, indicating two routes between 8 and 10.

```
macosx> google.com
traceroute: Warning: google.com has multiple addresses; using 74.125.224.178
traceroute to google.com (74.125.224.178), 64 hops max, 52 byte packets
 1  192.168.1.1 (192.168.1.1)  1.799 ms  0.907 ms  0.863 ms
 2  192.168.0.1 (192.168.0.1)  2.321 ms  1.843 ms  1.810 ms
 3  user-10cm4g1.cable.mindspring.com (64.203.18.1)  12.768 ms  23.945 ms  12.387 ms
```

```
 4  ge0-0-0-1-oahuhimili-rtr2.hawaii.rr.com (24.25.225.145)  24.747 ms  26.682 ms  12
 5  te0-5-0-3-tustca1-rtr3.socal.rr.com (24.25.230.130)  77.599 ms
    te0-4-0-3-tustca1-rtr3.socal.rr.com (24.25.230.134)  75.164 ms  72.756 ms
 6  be25-tustca1-rtr1.socal.rr.com (66.75.161.50)  71.294 ms  69.243 ms  71.997 ms
 7  107.14.19.30 (107.14.19.30)  73.264 ms  111.126 ms  111.937 ms
 8  ae-1-0.pr0.lax10.tbone.rr.com (66.109.6.131)  98.640 ms  96.452 ms  129.582 ms
 9  72.14.197.157 (72.14.197.157)  230.857 ms  226.836 ms
    72.14.198.73 (72.14.198.73)  90.543 ms
10  64.233.174.238 (64.233.174.238)  281.218 ms  267.655 ms  284.535 ms
11  72.14.236.11 (72.14.236.11)  136.589 ms  145.517 ms  141.934 ms
12  74.125.224.178 (74.125.224.178)  78.874 ms  79.888 ms  89.720 ms
```

In this next example, we traceroute from an Ubuntu box on campus to the same IP address used in the first example.

Hops 1 through 6 are on campus (byuh.edu), with 2 through 4 being non-routable internal IP addresses.

Hops 7 and 8 are by way of twtelecom.net, one of the companies that acts as an ISP to campus.

Hop 10 is to the 64.233.174.* network, which also appears in the traceroute from home. After that, the route from school is very similar to the one from home.

```
ubuntu> traceroute 74.125.224.178
traceroute to 74.125.224.178 (74.125.224.178), 30 hops max, 60 byte packets
 1  gateway.cis.byuh.edu (216.228.254.1)  0.545 ms  1.123 ms  1.386 ms
 2  10.11.224.21 (10.11.224.21)  0.448 ms  0.466 ms  0.708 ms
 3  10.0.224.14 (10.0.224.14)  0.700 ms  1.026 ms  1.295 ms
 4  10.224.224.18 (10.224.224.18)  0.659 ms  0.984 ms  1.253 ms
 5  216.228.251.225 (216.228.251.225)  0.925 ms  1.229 ms  1.478 ms
 6  216.228.251.218 (216.228.251.218)  0.570 ms  0.618 ms  0.607 ms
 7  64-128-3-221.static.twtelecom.net (64.128.3.221)  2.577 ms  2.635 ms  2.623 ms
 8  pao1-pr1-ge-1-0-0-0.us.twtelecom.net (66.192.243.98)  62.600 ms  62.600 ms  62.58
 9  209.85.240.114 (209.85.240.114)  87.781 ms  91.144 ms  91.462 ms
10  64.233.174.206 (64.233.174.206)  71.119 ms  71.110 ms  71.099 ms
11  64.233.174.189 (64.233.174.189)  71.059 ms  71.054 ms  71.069 ms
12  72.14.236.11 (72.14.236.11)  71.308 ms  71.307 ms  71.480 ms
13  74.125.224.178 (74.125.224.178)  70.989 ms  71.190 ms  71.183 ms
```

The variation in routes is due to an important characteristic of the Internet:

redundancy. It is very common, especially near the core or backbone of the Internet, to find several alternate paths forward. This is done partly for load balancing, and partly to avoid having a single point of failure that could stop all communication.

It is similar to the way streets and highways occur in heavily populated areas. If one road is closed or congested, other roads can often lead to the same destination.

In the fringes of the network, very close to the source or destination computer, it is common to have only a single route forward. In a case like that, of the route stops working, the outer computers are simply cut off from the Internet until the route is restored.

32.2 ftp

ftp stands for File Transfer Protocol, and provides for (insecure) file transfer with another computer.

http://en.wikipedia.org/wiki/Ftp has more.

Exam Question 402 (p.343)**:** What does FTP stand for?
Required Answer: file transfer protocol

Exam Question 403 (p.343)**:** Is FTP considered to be secure? Why?
Acceptable Answer: No. Traffic (data) is not encrypted.

Because traffic is in the clear, if it is intercepted it can be understood by anyone listening in.

ftp is a very old protocol, one of the first implemented in the Internet. It is widely available.

ssh, discussed in section 33.1 (page 239), presents a commonly used and secure alternative to **ftp**.

Exam Question 404 (p.343)**:** Which is more secure, ssh or ftp?
Required Answer: ssh

Exam Question 405 (p.343)**:** Which is more widely available, ssh or ftp?
Required Answer: ftp

32.3 telnet

The **telnet** command makes an insecure shell (command line) connection with another computer. It allows the user (or another program) to carry out a character-based interaction on any port.

http://en.wikipedia.org/wiki/Telnet has more.

Exam Question 406 (p.343)**:** Is telnet considered to be secure? Why?
Acceptable Answer: No. Traffic (data) is not encrypted.

Because traffic is in the clear, if it is intercepted it can be understood by anyone listening in.

Specifically, when you type in a user name and a password, that information is transmitted in the **clear** from your current computer to the computer you are trying to reach. Anyone listening in can **sniff** your password from the network traffic.

telnet is a very old protocol, one of the first implemented in the Internet. It is widely available.

ssh, discussed in section 33.1 (page 239), presents a commonly used and secure alternative to **telnet**.

Exam Question 407 (p.343)**:** Is ssh considered to be secure? Why?
Acceptable Answer: Yes. Traffic (data) is encrypted.

Exam Question 408 (p.343)**:** Which is more secure, ssh or telnet?
Required Answer: ssh

Exam Question 409 (p.343)**:** Which is more widely available, ssh or telnet?
Required Answer: telnet

Exam Question 410 (p.343)**:** What port does telnet normally use?
Required Answer: 23

Port 23 leads to telnetd, the telnet daemon, which generally provides shell access to the server.

Exam Question 411 (p.343)**:** What is telnetd?
Acceptable Answer: telnet daemon

Exam Question 412 (p.343)**:** What is a daemon?
Acceptable Answer: A program that interacts mostly with other programs and little if any with end users.

Server and daemon usually mean the same thing.

Daemon is an old-fashioned spelling of demon.

Exam Question 413 (p.343)**:** In a program name, what does the suffix d usually mean?
Acceptable Answer: daemon

telnet can also be used to connect to any other port and talk to any other daemon on a server. Secure connections encrypt their traffic, and are virtually impossible to do by hand using telnet, but technically they can be done also. It just requires the mental speed of a computer to know what to type next.

Because telnet can talk to any port, it can be used to debug communication failures in services such as email (through smtp, the simple mail transfer protocol).

Chapter 33

Advanced Power Tools

Contents

These tools are a bit more advanced than the basic tools of the previous chapter. They are more advanced in the sense that fewer people know them, or they rely on advanced features such as encryption.

ssh is introduced in 33.1 (page 239).

dig is introduced in 33.2 (page 240).

nmap is introduced in 33.3 (page 242).

Wireshark is introduced in 33.4 (page 243).

33.1 ssh

The first s in ssh stands for secure. All data between two computers is encrypted if it goes through an ssh connection.

ssh provides a **secure shell** (command line) connection with another computer. It also has the ability to securely copy files. In fact, ssh provides several other capabilities, making it an entire suite of tools.

http://en.wikipedia.org/wiki/Secure_Shell has more.

Exam Question 414 (p.343)**:** What port does ssh normally use?
Required Answer: 22

Exam Question 415 (p.343)**:** What does ssh stand for?
Required Answer: secure shell

Exam Question 416 (p.343)**:** How does ssh establish a secure connection?
Acceptable Answer: It uses a public-key system like RSA to create a shared secret.

Once the connection is established, ssh uses faster methods of encryption to carry the rest of the traffic.

Section 24.1 (page 154) discusses public and private key systems.

ssh also provides tools such as ssh-keygen that allow you to create your own public and private key pairs.

33.2 dig

The **dig** command is a DNS (domain name system) lookup utility. It traverses the DNS system and reports the IP address or addresses of the requested domain name. It also reports the authority by which it came to that conclusion.

http://en.wikipedia.org/wiki/Domain_Information_Groper has more information about the dig command.

The **nslookup** command is similar to the **dig** command.

Exam Question 417 (p.344)**:** Which is more current, nslookup or dig?
Required Answer: dig

nslookup has been deprecated in favor of dig.

Exam Question 418 (p.344)**:** What does deprecated mean?
Acceptable Answer: being phased out

Deprecated means the old way of doing things, and that the old way is in the process of being replaced by something better. When something has been deprecated, it should not be used or relied upon. Something better has replaced it. But just removing the old thing would frustrate people and cause other programs to break. Therefore, deprecated things are often kept

around for many years. This gives people time to update their use of the old item.

Deprecated does not mean that it has **been** phased out. It is not gone yet. It is still usable. But it is **being** phased out, and you should assume that some day it will simply be gone.

Things are deprecated because the newer thing is substantially better, and there will be no loss when the old thing is gone.

Deprecated things still often show up on exams, partly because the exam has not been updated, and partly because some historical knowledge can be helpful.

Here is an example using dig against n101.doncolton.com.

```
ubuntu> dig n101.doncolton.com

; <<>> DiG 9.7.0-P1 <<>> n101.doncolton.com
;; global options: +cmd
;; Got answer:
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 4200
;; flags: qr rd ra; QUERY: 1, ANSWER: 1, AUTHORITY: 2, ADDITIONAL: 0

;; QUESTION SECTION:
;n101.doncolton.com.            IN      A

;; ANSWER SECTION:
n101.doncolton.com.     10120   IN      A       69.89.31.217

;; AUTHORITY SECTION:
doncolton.com.          78192   IN      NS      ns1.bluehost.com.
doncolton.com.          78192   IN      NS      ns2.bluehost.com.

;; Query time: 1 msec
;; SERVER: 216.228.255.4#53(216.228.255.4)
;; WHEN: Fri Mar 25 21:58:52 2011
;; MSG SIZE  rcvd: 97
```

The answer section includes n101.doncolton.com. 10120 IN A 69.89.31.217 which tells us that the IP address of n101.doncolton.com is currently 69.89.31.217.

33.3 nmap

Broadcast ping is discussed in section 31.2.4 (page 229). There we see a way that frequently works to discover what other devices are sharing the same local area network with you.

nmap goes this one better. In addition to finding the IP addresses of those other devices, it can probe to find the ports that are open. An open port means a service that is being provided.

http://en.wikipedia.org/wiki/Nmap has more.

http://nmap.org/download.html has free downloads for Windows, Mac, and Linux.

Here is an example scan of my home network.

The first host is my router. It supports telnet, dns, and http.

The second host is the laptop computer from which I issued the command. It supports several more services.

The third host is an Apple iPod Touch. The only service it supports is one for synchronization.

```
macosx> nmap 192.168.1.0/24

Starting Nmap 5.00 ( http://nmap.org ) at 2011-03-26 18:45 HST
Interesting ports on 192.168.1.1:
Not shown: 997 closed ports
PORT   STATE SERVICE
23/tcp open  telnet
53/tcp open  domain
80/tcp open  http

Interesting ports on 192.168.1.126:
Not shown: 993 closed ports
PORT     STATE SERVICE
21/tcp   open  ftp
80/tcp   open  http
88/tcp   open  kerberos-sec
139/tcp  open  netbios-ssn
445/tcp  open  microsoft-ds
```

```
548/tcp  open  afp
3306/tcp open  mysql

Interesting ports on 192.168.1.130:
Not shown: 999 closed ports
PORT      STATE SERVICE
62078/tcp open  iphone-sync

Nmap done: 256 IP addresses (3 hosts up) scanned in 10.96 seconds
```

But there's more. Using the `-A` command-line option, **nmap** can probe the ports with various special requests and can analyze their responses. The result of the analysis is a more detailed look into the nature of that machine.

Here is an example scan of my home router.

```
macosx> nmap 192.168.1.1 -A

Starting Nmap 5.00 ( http://nmap.org ) at 2011-03-26 18:55 HST
Interesting ports on 192.168.1.1:
Not shown: 997 closed ports
PORT   STATE SERVICE VERSION
23/tcp open  telnet  DD-WRT telnetd 24 std (c) 2007 NewMedia-NET GmbH
53/tcp open  domain  dnsmasq 2.40
80/tcp open  http    Linksys wrt54g DD-WRT firmware http config
|_ html-title: Treehouse5 - Info
Service Info: OS: Linux; Device: WAP

Service detection performed. Please report any incorrect results at http://nmap.org/s
Nmap done: 1 IP address (1 host up) scanned in 10.52 seconds
```

33.4 Wireshark

Wireshark is a software tool that allows you to examine network traffic (for example, packets) that is visible to the computer you are using. It can be a great aid in debugging network problems. It is also a great way to gain an understanding of networking.

http://en.wikipedia.org/wiki/Wireshark has more.

http://www.wireshark.org/ is the official web site for Wireshark. Free downloads are available there.

Exam Question 419 (p.344)**:** What does Wireshark do?
Acceptable Answer: It captures and reports network traffic (data) as the traffic passes through an interface.

Exam Question 420 (p.344)**:** How does Wireshark get data?
Acceptable Answer: It captures data that is passing through a nearby interface.

Wireshark displays the internal structure of the data it has captured. This helps you educate yourself.

Wireshark shows you exactly what is happening on the wire (interface), including user names and passwords that were sent in **clear text**."

Wireshark is readily available for Microsoft Windows and Mac OS X. It can also be compiled from source code for other platforms.

Wireshark is GPL-based free and open source software.

Wireshark was formerly called **Ethereal**.

Networks tend to be busy, mostly with broadcast traffic. Wireshark will capture everything it hears. You will want to filter it down to just the interesting parts.

The exact steps to filter are beyond the scope of this book. But just know that it can be done, in case you need it.

Unit VIII

Switching

Chapter 34

Overview of Switching

Contents

34.1 Topologies

Topology means the way devices are connected among themselves to create a network. Who is your neighbor, or neighbors?

Star: The most popular topology for wired local area networks is called the star topology. At the center of the star is a switch, and at the end of each ray is a computer. It is popular because of simplicity and fast speeds.

Bus: Years ago the most popular topology for wired networks was the bus topology because of its low cost. A single wire goes from computer to computer, like a road around an island. Omnibus means "for all," or shared, or combined. It is the main topology for wireless networks because devices all share the same radio frequency bandwidth.

Mesh: This is the most popular topology for connecting several LANs. Each node of the network can connect to a few or many other nodes. The human brain uses a mesh topology where nerve cells can connect to many other nerve cells. High performance parallel processing computers use mesh to connect the individual CPU cores.

Exam Question 421 (p.344)**:** List in any order the three main network topologies in use today.
Acceptable Answer: star, bus, mesh

34.2 Collision Domains

Keeping costs low is always important. Voltage on copper is one of the least expensive types of physical media to utilize. And using a bus topology also saves money.

To communicate in this setting, a computer would transmit its message as voltages on the shared wire. Other computers would sense the voltages and recognize them as a message. They would make a copy of the message and then decide whether it was for them or not.

But there could be collisions. What if two computers talk at almost the same time? New rule: a computer would first listen to the wire. If the wire was quiet, it would start sending its message. It would also listen just in case some other computer started talking at about the same time. It would sense the voltages on the wire. If another machine started talking, the voltages would overlap each other creating higher peaks than normal. Those higher peaks would be recognized as a collision, and the computer would stop talking. The other computer would also detect the collision and it would stop talking too.

Exam Question 422 (p.344)**:** What is a collision?
Acceptable Answer: two devices talking at the same time on the same media

The solution was to back off and wait some random amount of time and try again. This generally ended the conflict. One computer would randomly talk first. The other would wait and then talk next. If there was another collision, they would wait again, this time longer.

The shared wire is called a collision domain. It is the set of places where two devices are not allowed to talk at the same time. It is also the set of places

where every computer can hear every message from every other computer.

Exam Question 423 (p.344)**:** What is a Collision Domain?
Acceptable Answer: shared physical media

The physical medium includes the full extent to which signals can propogate, including passing through layer 1 hubs and repeaters.

Normally we worry about collision domains that have lots of hosts sharing the same medium, but technically it is still a collision domain even if there are only one or two hosts.

Exam Question 424 (p.344)**:** How do you break up a Collision Domain?
Acceptable Answer: insert switches to make separate segments

In wireless (Wi-Fi) settings, collision domains can be large. This is a problem.

In wired settings, collision domains are less of a problem. That is because switches are commonly used instead of hubs, and full-duplex, star-topology wiring is commonly used instead of bus-topology wiring.

34.3 Half Duplex

We use the words "half duplex" to say that a channel can send or receive but not both at the same time. For example, a residential driveway is typically half duplex.

Exam Question 425 (p.344)**:** What does half duplex mean?
Acceptable Answer: can send or receive but not at the same time

We use the words "full duplex" to say that a channel can go both ways at the same time. For example, a two-lane road is full duplex.

Exam Question 426 (p.344)**:** What does full duplex mean?
Acceptable Answer: can send and receive at the same time

With bus topology, communication could only be half duplex because there was really only one channel.

With star topology, communication can be faster because fewer devices are on each channel, and with proper wiring full duplex becomes possible.

34.4 Dividing Collision Domains

As computers got faster their communication needs also got bigger and the wires got more busy. They often reached a point where there was so much traffic that you would get a traffic jam. Think of a room with two people talking. Now add more people until you can barely carry on a conversation. That was the collision domain problem.

The next step forward was to separate large collision domains into smaller pieces. You could split it into two pieces by using a bridge. The bridge would listen to each message and copy it to the other segment and send it out again. Bridges broke down collision domains and they also extended the range of the network. Without a bridge, the maximum wire length was something like 100 meters. But with a bridge, you could get 100 meters per segment.

If you have a smart enough bridge, it remembers the MAC address of each device on each side of the bridge. If the destination was on the other side of the bridge, it would copy it across. If not, it would not copy it. This cut down on network traffic on both sides.

An important feature of communication is the so-called 80-20 rule, which applies to many things in life and nature. This rule says that 80 percent of your communication will be with 20 percent of your neighbors. Obviously the numbers 80 and 20 are not precise, but they indicate the direction that things happen.

If you could segregate the devices wisely, you could make sure that the 20 percent of popular neighbors were on the same network segment as yourself.

Instead of using a bridge, you could use a switch and have lots of segments.

34.5 Counting Collision Domains

A popular test question that is easy to grade is to show a network diagram and ask how many collision domains there are.

Switches and bridges are the dividing points.

There is one slightly tricky issue: the full-duplex wired connection.

Each side of a UTP wire creates an independent communications channel. (1,2 goes to 3,6 and vice versa.) There is only one transmitter and one

receiver in each channel, so no collisions are possible. Each side of the channel is a separate collision domain.

Because of these separate channels, UTP wire supports full duplex communication. Devices can talk and listen (upload and download) at the same time.

By tradition, and on every test I have seen, a full-duplex communication link counts as one collision domain even though it is really two.

34.6 From Bus To Star

As things got faster, network contention became a real problem.

Switches were recognized as the solution. And lots of rewiring happened.

Computers gave up their coax cable connections and replaced them with 8P8C (RJ45) jacks. And wiring went from the bus-topology model to the star-topology model, also called home-run wiring. Each computer had its own private highway between itself and the switch. No contention. And each highway could be a two-way road, full duplex. Each computer could send and receive at the same time.

Switches have an internal fabric called the back plane that does all the cross connections between the various ports. A good switch can copy messages from A to B at the same time as it is copying messages from C to D.

As a result, all the collision domains disappeared. Okay, they still existed, but they were irrelevant. The traffic jams did not happen on the wires. They happened at the switches.

34.7 Broadcast Domains

Direct communication is used when we know the exact address of our destination. Sometimes we do not know the address. Examples include DHCP and ARP.

Exam Question 427 (p.344)**:** List in any order two protocols that use broadcast in a LAN.
Acceptable Answer: dhcp, arp

DHCP: DHCP stands for **Dynamic Host Configuration Protocol**. When a machine first connects to the network, it normally makes a DHCP request to find out its own IP address, subnet mask, gateway address, and other useful details. The DHCP request is a broadcast.

Exam Question 428 (p.344)**:** What does DHCP stand for?
Required Answer: dynamic host configuration protocol

ARP: ARP stands for **Address Resolution Protocol**. ARP starts with a logical address (a layer 3 IP address) and asks for a physical address (a layer 2 hardware address, like a MAC address).

Exam Question 429 (p.344)**:** What does ARP stand for?
Required Answer: address resolution protocol

Exam Question 430 (p.344)**:** What does ARP do?
Acceptable Answer: convert ip address to mac address

There is another protocol that is sometimes mentioned: **RARP**. It is the **Reverse Address Resolution Protocol**. It broadcasts a physical address and asks for the logical address. It is obsolete, largely replaced by **DHCP**.

Exam Question 431 (p.344)**:** What is a Broadcast Domain?
Acceptable Answer: all devices that can hear the same broadcast

A broadcast domain is the extent to which broadcasts are sent. It includes the media and the machines connected to the media. It is almost always a single local area network.

Layer 2 broadcasts are sent to MAC address FF:FF:FF:FF:FF:FF and are expected to be received by all devices on the same LAN.

Layer 3 broadcasts are sent to IP address 255.255.255.255 and are currently expected to be received by all devices on the same LAN or small group of closely related LANs.

In the early days of the Internet, Layer 3 broadcasts were received by all devices everywhere on the Internet. Nowadays they are filtered by routers and not passed along to neighboring LANs outside of the organization.

Normally broadcasts pass through switches and are filtered by routers.

Exam Question 432 (p.344)**:** How do you break up a Broadcast Domain?
Acceptable Answer: insert routers to make separate lans

Chapter 35

Plan B: Redundancy

Contents

We like redundancy. Redundancy is having a "plan B" ready to go in case our current plan stops working. Normally we plan for equipment failure including wiring failure and switch/router failure. Wires get cut by construction equipment or gnawed through by rodents. Switches and routers lose power or wear out. We want to be back up and running in a few seconds before anybody notices, not in a few days when there are lots of unhappy people. That requires redundancy.

No Single Point of Failure: We want to avoid having any point in our network that, when it dies, brings down a large part of our network. Such a point is called a choke point or a "single point of failure." To avoid that we create alternate routes through the network. When our preferred route disappears, we can use one of the alternates.

Loops: But alternate routes create loops. That is the price we need to pay for redundancy. (Loops can also be created accidentally.) And loops can be a problem.

Imagine that computer A sends a broadcast. It goes to its switch, call it S1. S1 is smart enough to send it out on all the other ports, but not back to A. Down one of those lines is another switch, S2, and down another line is S3. And for redundancy S2 and S3 also have a direct connection to each other.

So the message goes from A to S1, and from S1 to both S2 and S3. And from S2 to S3 (again). And from S3 to S2 (again). Soon the message is running in circles.

Broadcast Storm: We call this situation a broadcast storm. It will happen whenever switches are connected in a loop. Switches cannot remember every message they have seen before. They simply pass every broadcast forward to all the other switches.

The solution is to break the loop by disconnecting one of the wires so the network becomes a "tree" without loops. Guess who had to do that? The network administrator. The broadcast storm ends. And you lose your redundancy.

35.1 Spanning Tree Allows Redundancy

The tree thing is very important. What if switches could be made to automatically disable certain links, turning any net into a tree? Then redundancy could exist in switching networks. In that case, if a link went down the network could heal itself. Redundant links make the network stronger and more reliable.

The solution is an algorithm called the **Spanning Tree Protocol** (**STP**). This protocol is a set of steps followed by switches to let them do just what we want: automatically disable redundant links until they are needed.

Exam Question 433 (p.345)**:** What does STP stand for?
Acceptable Answer: spanning tree protocol

35.2 Convergence

Convergence: Original STP can take 30 to 50 seconds to converge. During this time all routine traffic must wait. Rapid STP takes about 6 seconds to converge.

Exam Question 434 (p.345)**:** What is convergence?

Acceptable Answer: when devices agree

Here we are talking specifically about switches, but the exact same concept applies to routers.

Each switch has some concept of what the network looks like. If all the switches have that same concept, things run smoothly. If not, loops can happen and messages may not get delivered.

When network changes occur, knowledge about the change does not immediatly reach all the switches that need to know. Those immediately connected find out first. From them it gradually rolls out from switch to switch until everybody knows. Sometimes the change is small and sometimes the change is huge.

As the knowledge rolls out, switches make decisions and come to understandings. Eventually they all come to the same understanding. When that happens, we say the network has converged.

Exam Question 435 (p.345)**:** About how many seconds does it take STP to converge?
Acceptable Answer: 30

Wording: Switches and bridges are the same thing. Generally a bridge is a two-port switch, and a switch is a multi-port bridge. STP uses the two terms to mean the same thing.

35.3 Step One: Switches Elect A Leader

The first thing STP does is elect a root bridge. It becomes the center of everything, the root of the tree. Links that would create loops get blocked.

Exam Question 436 (p.345)**:** Which switch is the root bridge?
Acceptable Answer: lowest bridge id

The switch with the lowest bridge ID will become the root. The bridge ID has two parts: bridge priority and MAC address. Priority is the important thing. MAC address is only used as a tie-breaker when two switches have the same priority. The bridge priority defaults to 32768 but can be changed by the network administrator. Because the network administrator can set the bridge priority, they can control the root bridge election.

Normally the network administrator wants the switch that is closest to the router to be the root bridge.

35.4 Step Two: Switches Identify Best Paths

After the root bridge is agreed upon, each other switch must decide which of its ports provides the best path to that root bridge. Best is defined as having the least cost, which means the fastest speed. That port becomes the root port for the switch.

Exam Question 437 (p.345)**:** Which port is the root port?
Acceptable Answer: least cost path to root bridge

The network administrator can manually disable a port. Disabled ports are never part of the spanning tree. The other ports can become members of the spanning tree.

We do not really want to manually disable any ports. We want the STP protocol to disable them as needed. Then, as needed, STP can re-enable ports to heal a broken network. Manually disabled ports cannot be re-enabled by STP. They must be manually re-enabled. This means a network administrator must get directly involved.

Exam Question 438 (p.345)**:** List in any order classic STP's four port state options.
Acceptable Answer: forwarding, learning, listening, blocking

Under the classic STP protocol, each non-disabled port of each switch is in one of the following four states:

Forwarding: This port is a member of the tree. All ports on access routers are forwarding ports.

Learning: This port has decided to become a member of the tree but does not have enough information to be useful, so it just listens and learns about the network.

Listening: This port is deciding whether or not to become a member of the tree. It has not yet decided.

Blocking: This port decided that joining the tree would create a loop.

35.5 Step Three: Traffic Resumes

As mentioned above, during convergence all routine traffic must wait. When convergence is achieved, routine traffic resumes.

35.6 RSTP Is Faster

Because classic STP has slow convergence, it is deprecated. It has been replaced by the **Rapid Spanning Tree Protocol** (**RSTP**). RSTP will fall back to classic STP if necessary.

Exam Question 439 (p.345)**:** What does RSTP stand for?
Acceptable Answer: rapid spanning tree protocol

Exam Question 440 (p.345)**:** About how many seconds does it take RSTP to converge?
Acceptable Answer: 6

RSTP protocol achieves faster convergence by having more roles and states. Roles include root, designated, alternate, and backup. States include discarding, learning, and forwarding.

https://en.wikipedia.org/wiki/Spanning_Tree_Protocol provides an in-depth article.

Unit IX

Routing

Chapter 36

Review of Routing

Contents

Routing is one of the most crucial aspects of modern networking. Researchers are constantly working to improve speeds and reduce convergence times. Our goal in this unit is to understand the fundamentals of routing. This foundation will help you, the reader, understand and work with any routing protocols or problems you may encounter.

We have previously mentioned routing a number of times. In this chapter we review that information.

36.1 Routing Packets Between LANs

In section 3.3 (page 21) we introduced routing. We said:

A computer that belongs to more than one local area network can act as a router, and can pass messages between those LANs.

The moving of a message from one LAN to an adjacent LAN is called a hop. It is normal for messages to make ten or more hops before they reach their

destination.

Exam Question 441 (p.345)**:** What is a hop?
Acceptable Answer: moving a packet from one lan to the next lan

The router must send the packet on its next hop toward its ultimate destination. To do this, each router must belong to two or more LANs.

Routers talk to their neighboring routers. They learn what networks each one can reach. This enables them to pick the best direction for the next hop.

36.2 Anatomy of a Hop

In chapter 9 (page 56) we further developed the concept of hops, where messages are passed from one local area network to another.

In section 9.4 (page 59) we said that a hop is the activity performed by a router when it receives a packet on one local area network, and passes it across to a different local area network, hopefully one step closer to its final destination.

In section 9.5 (page 59) we said that routers must figure out a reasonable path forward for each packet or the packet will not be delivered.

36.3 Time To Live

Sometimes routing mistakes are made. There is a time-to-live mechanism that prevents infinite loops when those mistakes happen. When a packet starts its journey, the source computer "pays postage" by setting the TTL counter in the packet header to some number of allowed hops. Each time a packet makes a hop, the router decrements the TTL counter. When the counter reaches zero, the router drops the packet instead of forwarding it. The router also sends back a message to the original source telling that the packet got dropped.

Exam Question 442 (p.345)**:** What does TTL stand for?
Required Answer: time to live

Exam Question 443 (p.345)**:** What is the purpose of TTL?
Acceptable Answer: stop infinite loops

36.4 Trace Route

Because TTL reports the location of death, a clever use of this information gives us a list of routers through which a packet will travel to its destination. This information can be helpful in debugging problem situations. It can be used to discover where the bottlenecks or routing errors might be occurring.

The Trace Route command (**traceroute** or **tracert**) sends several packets with a TTL value of 1. As those die, the location of death is sent back to the trace route program. It sends several more with a TTL value of 2 and records their deaths. It continues increasing the TTL value until the final destination is reached.

http://en.wikipedia.org/wiki/Traceroute provides an overview.

Exam Question 444 (p.345)**:** List the two main things that traceroute reports.

Acceptable Answer: identity and latency for each router in the path

Traceroute also lists the number of hops to get to each router, and the percentage of packet loss (expected responses not received).

Chapter 37

Netmasks and Addressing

Contents

Internet Protocol (IP) addresses are logical addresses rather than physical addresses. MAC addresses are the physical addresses used in Ethernet.

IP addresses consist of a number of bits. The first part of the address identifies the network or subnet. The rest of the address identifies the host. The location of the dividing line is called the netmask and has changed over the years.

In section 27.4 (page 183) we looked at pre-1981 addressing. We talked briefly about core routers and hand-crafted routing tables.

37.1 Constant Netmasks

With pre-1981 addressing, the 32-bit IP address was divided into two parts: the network and the host. The network part was always 8 bits. The host part was always 24 bits.

37.2 Implicit Netmasks: Classful Addressing

Classful Addressing came next. Instead of always having 8 bits for the network, the number depended on the address class. The netmask itself was not stated, but it was implied by the first few bits of the IP address.

Exam Question 445 (p.345)**:** Class A addresses start with what bit(s)?
Required Answer: 0

Exam Question 446 (p.345)**:** Class A addresses have what netmask?
Acceptable Answer: /8

Class A corresponds to IP addresses from 0.x.x.x through 127.x.x.x.

With Class A addresses, the network portion is 8 bits and the host portion is 24 bits. This is exactly the same as the legacy pre-1981 system, making the Classful Addressing approach backward compatible with the pre-1981 system.

Exam Question 447 (p.346)**:** Class B addresses start with what bit(s)?
Required Answer: 10

Exam Question 448 (p.346)**:** Class B addresses have what netmask?
Acceptable Answer: /16

Class B corresponds to IP addresses from 128.x.x.x through 191.x.x.x.

With Class B addresses, the network portion is 16 bits and the host portion is 16 bits.

Exam Question 449 (p.346)**:** Class C addresses start with what bit(s)?
Required Answer: 110

Exam Question 450 (p.346)**:** Class C addresses have what netmask?
Acceptable Answer: /24

Class C corresponds to IP addresses from 192.x.x.x through 223.x.x.x.

With Class C addresses, the network portion is 24 bits and the host portion is 8 bits.

There are two other classes. Class D addresses start with the bits 1110. It is for multicast use. Class E addresses start with the bits 1111. It is for experimental use.

37.3 Explicit Netmasks: Classless, CIDR, VLSM

In section 28.1 (page 191) we looked at CIDR, Classless Inter-Domain Routing. CIDR started about 1993 as a way to consolidate groups of networks into a single entry in the routing table. CIDR made the network masks explicit and thereby more flexible.

This approach is not backward compatible with classful addressing because it requires an additional field, the explicit netmask (or subnet mask). Going classless required updates to routing protocols. For example, RIP, which we will discuss in chapter 41 (page 281), handles classful addresses but not classless addresses. RIP version 2 handles classless addresses.

Exam Question 451 (p.346)**:** What does CIDR stand for?
Acceptable Answer: classless inter-domain routing

In chapter 29 (page 202) we looked at VLSM, Variable Length Subnet Masking, and how it was made possible by the widespread use of explicit subnet masks.

Exam Question 452 (p.346)**:** What does VLSM stand for?
Acceptable Answer: variable length subnet mask

Chapter 38

Types of Routers

Contents

38.1 Core, Distribution, and Access

The Internet is like the road system of a large country.

Core Routers: There are major highways that carry traffic for long distances. These highways are like the backbone of the Internet, and where the

highways intersect we have major interchanges. These interchanges are like core routers.

Distribution Routers: There are major roads that connect the neighborhoods of each city. These are like the distribution layer of routers in a business or large organization.

Access Routers: There are city streets that connect the houses within the neighborhoods. These are like the access layer of routers that handle traffic for individual offices, workgroups, or homes.

These three kinds of routers exist because the workload requirements differ. Fundamentally all three kinds of routers are the same. But as built, they have different strengths and weaknesses.

Access routers mostly do things like NAT and only simple routing. They are optimized for that. We look at them in the remainder of this chapter.

Distribution routers do quite a bit of actual routing and also do things like QoS (Quality of Service). We look at distribution routers in chapter 39 (page 272).

Core routers do high-speed heavy-duty actual routing. They are very expensive. They use routing protocols that are beyond the scope of this book.

Exam Question 453 (p.346)**:** List in any order the three layers of the Cisco router model.
Required Answer: core, distribution, access

Switches: These are like the intersections in a neighborhood.

In this chapter we consider the role of access routers.

38.2 Autonomous Systems

The Internet is divided into autonomous systems. An **autonomous system** (**AS**) is roughly the collection of networks belonging to a large organization. Network administrators of that organization manage all the routers within the autonomous system.

Exam Question 454 (p.346)**:** What does AS stand for?
Acceptable Answer: autonomous system

The routing protocols used within an autonomous system are called **interior gateway protocol**s (**IGP**s). In this unit we are mainly concerned with

IGPs. IGPs are further divided into distance-vector and link-state protocols. One well-known distance-vector example is **RIP**. Some well-known link-state examples are **EIGRP**, **OSPF**, and **IS-IS**.

IGPs are used mostly by distribution layer routers.

The routing protocols used between autonomous systems are called **exterior gateway protocol**s. The most important example is **BGP**, the Border Gateway Protocol.

38.3 Living On The Edge

Access routers form the outer edge of the Internet. These routers tend to be very simple and inexpensive. They reach downward to a single local area network with a few or many host computers and one or more switches. They reach upward to a more powerful distribution router that handles many access routers.

Because they are on the edge their routing decisions are very simple. They spend most of their time doing NAT rather than actually routing. They are often called gateway routers. Gates only have two sides: the inside and the outside.

Home and small office routers fall into this category. They only talk to one other router which is their "uplink." Their downlink is to a switch that is the root bridge (and often the only bridge) of a local area network of end users. Home routers also typically provide Wi-Fi, DHCP, and NAT. A home router usually has one uplink and four downlinks.

When a computer uses DHCP to find out about its network, one piece of information it gets is the gateway address. That is the address of the router that serves as the gateway to networks beyond the local one.

38.4 User Sends A Packet

When the source computer has information to send the information goes down the OSI stack from application to presentation to session to transport to the network layer. The information is divided into a series of packets. At this point the netmask (or subnet mask) becomes important.

Inside or Outside? The source computer looks at the IP address of its destination. It uses its own subnet mask to identify its own subnet. It checks to see whether the destination is in the same subnet.

Exam Question 455 (p.346)**:** How can we tell if two machines are in the same LAN?
Acceptable Answer: network portion of both addresses is the same

38.4.1 Sending Inside

If the destination is within the same LAN as the source, the packet is sent directly, without the use of a router. The source computer uses **ARP** if necessary to discover the MAC address of the destination. It uses that to build a frame for the packet and it sends the bits out on the physical medium.

This would be the case for things like a local network printer, or a home media server, or video surveillance cameras. As we get more into the Internet of Things (IoT), there will be more things in the local area network.

Exam Question 456 (p.346)**:** When does a computer send a frame directly to its destination?
Acceptable Answer: destination is in same lan

38.4.2 Sending Outside

We don't pay for home internet so we can print to our local printer. We want content from websites or Netflix or Google. And that is outside of our local area network.

If the destination is not within the same LAN as the source, the packet is sent to the router using the gateway address. The source computer uses **ARP** if necessary to discover the MAC address of the gateway. It uses that to build a frame for the packet and it sends the bits out on the physical medium.

Exam Question 457 (p.346)**:** When does a computer send a frame directly to the gateway?
Acceptable Answer: destination is not in same lan

Exam Question 458 (p.346)**:** What does a gateway do?
Acceptable Answer: provide access to other networks

Exam Question 459 (p.346)**:** If the router's IP address is 1.2.3.4, what is

the most likely value for the gateway address in that LAN?
Acceptable Answer: 1.2.3.4

The gateway address sent by DHCP is actually just the router's IP address in that LAN.

38.4.3 The Role of NAT

After the traffic arrives at the router, typically **network address translation** (**NAT**) takes place because the inside computers do not have routable addresses, but the router itself does. So the router makes a note in its NAT table telling which device and which port is sending the message. It then substitutes its own IP address and makes up a port number, slaps them on the packet, and sends them through the uplink.

Exam Question 460 (p.346)**:** What does NAT stand for?
Required Answer: network address translation

Exam Question 461 (p.346)**:** What does NAT do?
Acceptable Answer: replace one ip address with another

Normally (but not always) NAT is replacing a non-routable address with a routable address.

38.5 User Receives A Packet

NAT acts as a simple firewall to protect computers within the LAN.

When a client starts a conversation with someone outside of the LAN, a record is made in the NAT table and responses can be accepted.

38.5.1 Continuing A Conversation

After we send a packet to Google or whatever, we expect a response. The response will come to the router through the uplink.

An outsider can continue a conversation that was started by the local computer if the NAT table has a matching entry.

When a message comes from outside, the destination port number is checked. If the port number is in the NAT table, the address and port number get updated and the message gets sent forward.

38.5.2 Starting A Conversation

Normally an outsider cannot start a conversation with a local computer.

Port Forwarding: The exception is when a local computer wants to be a server. Maybe it does this to host a multi-player game. Maybe it does this to host a Skype conversation. To do this, it must get around the NAT table restriction. The router must be configured explicitly to do port forwarding. Port forwarding says "I am a server. If a message comes in for my port send it along to me and I will handle it."

38.5.3 Defeating Hackers

If the port number is not forwarded or in the NAT table, the message is dropped. In this way NAT acts as a simple firewall. Because there are 65536 possible port numbers, but only a small number of those are in the NAT table, the hacker's odds of guessing correctly are pretty small.

Exam Question 462 (p.346)**:** How does NAT defend against attacks on local computers?
Acceptable Answer: outsiders cannot start conversations with locals

38.6 Daisy Chaining Access Routers

Children sometimes pick flowers and weave them into necklaces. Daisies are popular for this activity, and the weaves are sometimes called daisy chains.

That term has been adopted into networking and other parts of the computing world. Daisy chaining means linking one thing to another, to another, to another, possibly many links long.

In our present discussion, we can say that it is possible to have a collection of access routers, just like branches or rootlets of a tree. The key feature is that there is only one path up through each router in that network. Loops are not possible. And access routers work great.

The decision of whether to daisy chain several access routers together, or to have a distribution router talk to several access routers, is really more of an economic decision. Distribution routers are more expensive and powerful than access routers. Sometimes an access router is good enough.

With a cascade of several access routers, it could be that a packet from

outside reaches the first access router, and that router updates the address and port number, sending it along its way. The router presumes (but does not really care) that the packet is going directly to its final destination. But what if that destination is not the final destination after all? What if it is yet another access router with yet another NAT table? It could happen. In fact, it is fairly common.

The truth is, we don't know, and probably cannot easily tell, how many times NAT has been applied to a packet. We just know that like the breadcrumbs in the Hansel and Gretel fairy tale, the routers provide a path that leads back to the ultimate client.

38.6.1 Two Router Systems

Let's say you want to share your Internet connection. Perhaps you have roommates or guests and you, being a nice guy, are happy to help them out.

But at the same time you know that devices inside your LAN are not subject to the firewall provided by NAT. Somebody could hack your computer. You want to avoid that.

There is a two-router solution. You can place one access/gateway router at the front door of your system. It talks to the Internet, and it provides a connection to guests and roommates. It provides the shared network.

It also provides a connection to your private subnet. That is where the second access/gateway router comes in. Your computer sits inside the private subnet. And your private devices, like your laser printer and your media server sit inside your private subnet.

Some routers support guest networks automatically. You may not need two routers to achieve the same results.

38.6.2 What Could Go Wrong?

Two things could go wrong.

First, you might put your private network closer to the front door, and try to make the shared space in the second network. This has a nice feel to it, but it does not work. Everything in the innermost network can reach everything in the middle network.

Second, if you get two routers that are basically the same, they will have the

same default addressing. But for routing to happen the subnets must have different network numbers. That means that somebody must configure one of the routers to have a different IP address range. And anytime you are doing configuration, it is scary to the average person.

Exam Question 463 (p.346)**:** What happens when both LANs for the router have the same network number?

Acceptable Answer: traffic will not get through

Chapter 39

Distribution Routers

Contents

A distribution router makes choices about where to send each packet. It would typically have many ports linking it to both access routers and to other distribution routers.

39.1 Routing Tables

To handle the choices, the router has a table. It could be constructed by hand or it could be automatically generated.

Basically, the table makes a list of all possible IP addresses, and for each IP address, it tells whether the traffic should be sent through port 1, port 2, port 3, or whatever.

Obviously with 4 billion IPv4 addresses, you would not build a table that lists each of them explicitly. Instead, we do them in groups based on prefix. If the prefix matches the desired destination, then it goes out that port.

For example, if all IP addresses that start with 2.4.6 are to go through port 3, that entry would be in the routing table. It would look something like this:

Address	Mask	Port
2.4.6.0	24	3
0.0.0.0	0	1

The 24 means to check the first 24 bits of the destination, and if it matches the first 24 bits of 2.4.6.0, then send it through port 3.

The 24 corresponds to /24 in CIDR (classless inter-domain routing) format. It corresponds to a netmask with 24 1s which is 255.255.255.0.

The 0.0.0.0/0 route in the table is the default route. Anything that has not been matched yet will be compared to the first "zero" bits of 0.0.0.0 (which is guaranteed to match) and will be sent through port 1.

39.2 Route Aggregation

Maybe all the traffic for 2.4.7 should also go through port 3. The table might look like this instead:

Address	Mask	Port
2.4.6.0	24	3
2.4.7.0	24	3
0.0.0.0	0	1

The table can be pretty big. Its size is really only limited by the available memory in the router and the speed of the processor in the router. But both of those things cost money, so they tend to be small in ordinary routers.

Dotted Quad	Binary
2.4.6.0	00000010.00000100.0000011**0**.00000000
2.4.7.0	00000010.00000100.0000011**1**.00000000

Notice that the first 23 bits of both networks are identical. Only the last bit differs. Because the IP addresses are similar and the port is the same, so we can aggregate the routes like this:

Address	Mask	Port
2.4.6.0	23	3
0.0.0.0	0	1

The 23 means to check the first 23 bits of the destination, and if it matches the first 23 bits of 2.4.6.0, then send it through port 3.

Route aggregation is the hugely important, and pretty obvious, solution to

avoiding having 4 billion entries in our routing table.

Exam Question 464 (p.347)**:** List the other names for aggregation.
Acceptable Answer: summarization, supernetting

39.3 Longest First

When a packet comes in, the router looks at the destination address and makes a decision about which line in the routing table should be used.

The rule is to use the line with the longest prefix. Thus, a /24 line would be used before a /23 line, and the /0 line would be used last of all.

Chapter 40

Routing Table Example

Contents

Access routers do not need complex routing tables. They do not converse among themselves. They have no redundancy. If an access router goes out, a workgroup drops off the network. They may be down for a few hours. The solution is to throw in a new router and everything is back up and running. Being a single point of failure for a workgroup is fine.

Above the workgroup level, when many workgroups are involved, it becomes valuable to have redundancy to avoid having any single point of failure. With all that redundancy the routing tables for distribution layer routers can become pretty complex so we want them to be managed automatically, not by hand. We do this by using Routing Protocols to manage the routing tables.

The first widely used routing protocol was RIP, the Routing Information Protocol. It came out in 1988.

Exam Question 465 (p.347)**:** List the two big advantages of RIP.
Acceptable Answer: widely supported, easy to configure

The basis of this protocol is two things. First, each router knows what

networks it is directly connected to. Second, each router can share what it knows with its neighbors.

This is a simple-minded but very effective approach. And because it has been around so long, many routers speak this language.

40.1 Network Diagram

We will use a network diagram to illustrate what we are talking about. Here are some symbols.

The cloud symbol represents the Internet.

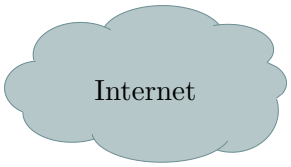

This round thing with arrows is the industry-standard symbol that represents a router. We will generally label our routers with the letter R followed by a number.

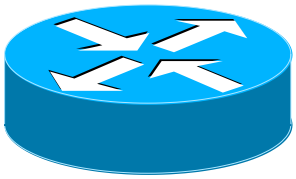

This square thing with arrows is the industry-standard symbol that represents a workgroup switch. We will generally label our switches with the letter S followed by a number.

The cloud with computers inside it represents a workgroup of computers and other devices that are all in the same local area network.

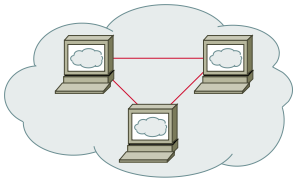

40.2 Typical Access Network

A typical access network has an access router with a built-in switch. Here we show it connected on the left to the cloud through the ISP (internet service provider) and on the right to the local area network of office or household computers.

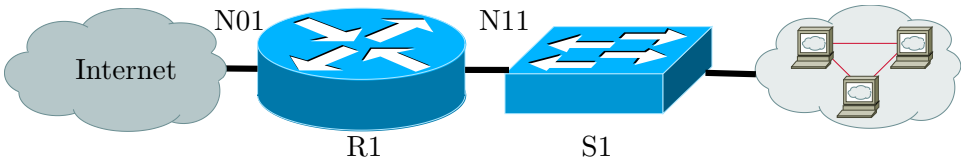

We have two local area networks here: N01 and N11. The N01 network is owned by the ISP. The uplink port of the R1 router belongs to the N01 network. The N11 network is owned by the office or household. The downlink port of the R1 router belongs to the N11 network. The S1 switch belongs to the N11 network. All devices in the workgroup belong to the N11 network.

Access Routing Table

Every network needs an IP address range and a network (or subnet) mask. Let's give the N11 network this address: 1.1.1.0/24. For the N01 network we will just use 0.0.0.0/0, the route that matches everything. 0.0.0.0/0 is also called the default route, or the route of last resort. The R1 router has the following routing table.

Address	Port
0.0.0.0/0	0
1.1.1.0/24	1

When any traffic arrives at the router, it will consult its table.

The router always uses the most narrowly specified route that it can. And the bigger the CIDR number, the more narrowly specified the route is. In this case the first choice is /24.

Traffic with a destination in 1.1.1.0/24 will be sent through port 1.

All remaining traffic has a destination in 0.0.0.0/0 and will be sent through port 0.

40.3 Two Access Networks

Here we have two access routers.

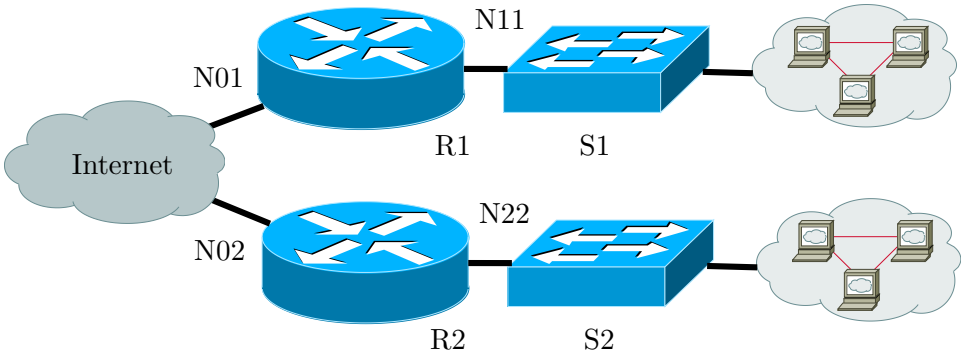

This is pretty similar to the previous network. We have four local area networks here: N01, N11, N02, and N22. The N01 and N02 networks are owned by the ISP. The N11 and N22 networks are each owned by an office or household.

The uplink port of the R1 router belongs to the N01 network. The downlink port of the R1 router belongs to the N11 network. The S1 switch and its associated workgroup belong to the N11 network.

The uplink port of the R2 router belongs to the N02 network. The downlink port of the R2 router belongs to the N22 network. The S2 switch and its associated workgroup belong to the N22 network.

When traffic wants to go from the N11 network to the N22 network, it must go through the R1 router, then through the Internet, then through the R2 router, and then into the N22 network.

40.4 Two Distribution Networks

Now we will go beyond access routers. We will let the routers talk to each other directly without going through the Internet. For this to happen, the routers must have more than just two ports and the owners of R1 and R2 must create a physical connection between the routers. This router-to-router connection is in its own local area network, a two-device network.

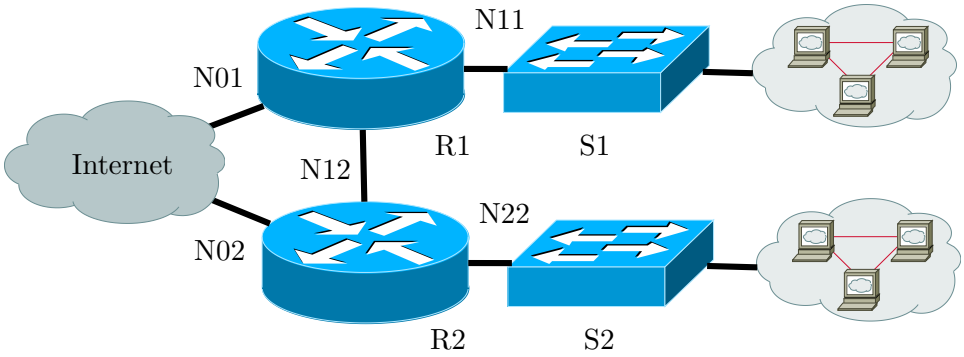

Again, this is similar to the previous example. But here we have five local area networks, including a new one labeled N12.

When traffic wants to go from the N11 network to the N22 network, it must go through the R1 router, then it can go directly through the R2 router, and then into the N22 network. This cuts down on traffic to the Internet. ISPs often bill based on the amount of traffic that goes through their connection, so having a direct connection can be a cost savings.

To make this happen, we must update the routing tables.

Let's assume that the N11 network is 1.1.1.0/24.

Let's assume that the N22 network is 2.2.2.0/24.

For the R1 router, the new routing table will look something like this:

Address	Port
0.0.0.0/0	0
1.1.1.0/24	1
2.2.2.0/24	2

When any traffic arrives at the router, it will consult its table. For traffic that has a destination matching 1.1.1.0/24 (which is the N11 network) it will send the traffic through port 1. For traffic that has a destination matching 2.2.2.0/24 (which is the N22 network) it will send the traffic through port 2. Otherwise, for traffic that has a destination matching 0.0.0.0/0, it will send the traffic through port 0.

40.5 Direct Knowledge

Routers have to know directly about each LAN to which they are connected.

Access routers may learn about their uplink through DHCP provided by their uplink. They learn about their downlink by being manually configured by a network administrator.

Distribution routers typically learn all their direct knowledge from configuration information provided by the network administrator.

What about the 2.2.2.0 entry in the table above? It is not direct knowledge.

Address	Port
2.2.2.0/24	2

R1 has no direct knowledge of the 2.2.2.0 network. This information was given to it by the network administrator as a hand-coded "static route."

When networks are small, it is very reasonable to use hand-crafted static routes to tell the routers how to do their jobs. But networks get big and network administrators wisely want to avoid tedious and error-prone aspects of their jobs. That is why routing protocols were invented. We want to insert the 2.2.2.0 entry into the routing table automatically, not by hand. That is called "dynamic routing."

Chapter 41

RIP: Routing Information Protocol

Contents

Building upon our example from the previous chapter, we now come to RIP, the Routing Information Protocol, the first major routing protocol. Because it is first, it is important to understand.

Exam Question 466 (p.347)**:** What does RIP stand for?
Acceptable Answer: routing information protocol

With RIP, the routers talk to each other. Every 30 seconds or so, each router sends a message to all of its direct neighbors giving them a copy of its routing table. These messages are called advertisements.

Those neighbor routers will look at the routing table to discover local area networks that can be reached by going through the first router.

In this way, R2 learns that through R1 it can reach 1.1.1.0/24, and R1 learns that through R2 it can reach 2.2.2.0/24.

These are our networks:

Network	Address
N12	9.1.2.0/30
N11	1.1.1.0/24
N22	2.2.2.0/24

Before the first advertisement, the routing tables look like this:

R1 Routing Table	
Address	Port
0.0.0.0/0	0
1.1.1.0/24	1
9.1.2.0/30	2

R2 Routing Table	
Address	Port
0.0.0.0/0	0
2.2.2.0/24	1
9.1.2.0/30	2

After the first advertisement, with the new information merged in, the routing tables look like this:

R1 Routing Table	
Address	Port
0.0.0.0/0	0
1.1.1.0/24	1
9.1.2.0/30	2
0.0.0.0/0	2
2.2.2.0/24	2
9.1.2.0/30	2

R2 Routing Table	
Address	Port
0.0.0.0/0	0
2.2.2.0/24	1
9.1.2.0/30	2
0.0.0.0/0	2
1.1.1.0/24	2
9.1.2.0/30	2

In these tables we are using **bold type** to show the information each router has just learned from the advertisement. Notice that the 9.1.2.0/30 entries are duplicates of something already in the table so they are not actually re-added. Instead, we have this.

R1 Routing Table	
Address	Port
0.0.0.0/0	0
1.1.1.0/24	1
9.1.2.0/30	2
0.0.0.0/0	2
2.2.2.0/24	2

R2 Routing Table	
Address	Port
0.0.0.0/0	0
2.2.2.0/24	1
9.1.2.0/30	2
0.0.0.0/0	2
1.1.1.0/24	2

When we receive an advertisement, we merge the networks into our routing table so long as the advertisement provides a different route (port number)

than something we already knew about.

We do not keep the second copy of 9.1.2.0/30 port 2 because we already have one with the same destination and port.

But we do keep the second copy of 0.0.0.0/0 because one uses port 0 and the other uses port 2, so they are not the same.

Convergence: Over time each router learns everything it can about the whole network, and the advertisements no longer contain any new information. At that point, we say the network has converged.

41.1 Distance-Vector Routing

Distance-Vector Routing makes routing decisions based on distance. In the case of RIP, distance is measured in hops.

Redundancy: Each router now has two routes to 0.0.0.0/0. Having several routes is called redundancy. It is a good thing because if one of the routes fails the other route can be used. Maybe a tree limb breaks and falls on a utility line and our communication line snaps. Maybe a backhoe digging somewhere cuts the line. Or a squirrel does. Or maybe someone simply unplugs a wire.

But when we have both routes, which route is best? If R1 receives a packet destined for 0.0.0.0/0 it cannot tell. We want each packet to take the fastest route available.

Distance Vector: RIP solves this problem by keeping track of the number of hops involved. This number is called the distance vector.

Here are the new starting tables. They include the number of hops required to reach the destination.

R1 Routing Table		
Address	Port	Hops
0.0.0.0/0	0	1
1.1.1.0/24	1	1
9.1.2.0/30	2	1

R2 Routing Table		
Address	Port	Hops
0.0.0.0/0	0	1
2.2.2.0/24	1	1
9.1.2.0/30	2	1

After the first advertisement, the routing tables look like this.

R1 Routing Table		
Address	Port	Hops
0.0.0.0/0	0	1
1.1.1.0/24	1	1
9.1.2.0/30	2	1
0.0.0.0/0	2	2
2.2.2.0/24	2	2

R2 Routing Table		
Address	Port	Hops
0.0.0.0/0	0	1
2.2.2.0/24	1	1
9.1.2.0/30	2	1
0.0.0.0/0	2	2
1.1.1.0/24	2	2

After the second advertisement, the routing tables look like this.

R1 Routing Table		
Address	Port	Hops
0.0.0.0/0	0	1
1.1.1.0/24	1	1
9.1.2.0/30	2	1
0.0.0.0/0	2	2
2.2.2.0/24	2	2
1.1.1.0/24	2	3

R2 Routing Table		
Address	Port	Hops
0.0.0.0/0	0	1
2.2.2.0/24	1	1
9.1.2.0/30	2	1
0.0.0.0/0	2	2
1.1.1.0/24	2	2
2.2.2.0/24	2	3

At this point the routes have converged and future advertisements will not result in any changes unless the network itself changes.

If R1 gets a packet destined for 3.3.3.3, it checks its routing table. It has a couple of /24 addresses but none of them match. So it falls through to the default routes. It has two of those:

R1 Routing Table		
Address	Port	Hops
0.0.0.0/0	0	1
0.0.0.0/0	2	2

It can send the packet through port 0 with an expected cost of one hop. Or it can send the packet through port 2 with an expected cost of two hops. Since it wants the fastest route, it sends the packet through port 0.

If the route through port 0 fails, R1 will remove it from its table. Then the only route remaining will go through port 2, so that is the route it will select.

Problem solved. We have redundancy to give us options when part of the network has failed, and we have a way to pick between identical destinations based on the distance cost.

41.2 Naïve Routing Problem

Let's see what happens when R1 loses its connection to 1.1.1.0.

These are the routing tables after convergence, but before R1 loses its connection to 1.1.1.0/24.

R1 Routing Table		
Address	Port	Hops
0.0.0.0/0	0	1
1.1.1.0/24	1	1
9.1.2.0/30	2	1
0.0.0.0/0	2	2
2.2.2.0/24	2	2
1.1.1.0/24	2	3

R2 Routing Table		
Address	Port	Hops
0.0.0.0/0	0	1
2.2.2.0/24	1	1
9.1.2.0/30	2	1
0.0.0.0/0	2	2
1.1.1.0/24	2	2
2.2.2.0/24	2	3

These are the routing tables we have after the port 1 link went down. All port 1 routes are removed. (There was only one such route.)

R1 Routing Table		
Address	Port	Hops
0.0.0.0/0	0	1
9.1.2.0/30	2	1
0.0.0.0/0	2	2
2.2.2.0/24	2	2
1.1.1.0/24	2	3

R2 Routing Table		
Address	Port	Hops
0.0.0.0/0	0	1
2.2.2.0/24	1	1
9.1.2.0/30	2	1
0.0.0.0/0	2	2
1.1.1.0/24	2	2
2.2.2.0/24	2	3

R1 thinks it has a route to 1.1.1.0/24 through port 2. What R1 does not know is that the route will never work. Let's see what actually would happen.

R1 receives a packet that should go to 1.1.1.0 and sends it through its port 2 to R2.

R2 receives a packet that should go to 1.1.1.0 and sends it through its port 2 back to R1.

(repeat)

The traffic would just bounce back and forth endlessly. This is a loop. Problems like this are why TTL was invented.

You might think that in today's modern world, such a thing could never happen. But the truth is that it still does happen. The examples are a bit

more complex, but it does still happen.

And that is where Time To Live (**TTL**) comes into play. Each time the packet is handed off, the counter is decremented (decreased by one). When the counter reaches zero, the packet is mercifully dropped.

RIP is not just simple distance-vector routing. It is a bit more sophisticated than that. In this chapter, we look at some of these important features in the RIP protocol.

41.3 Split Horizon

The first trick is fairly simple, actually. It is called Split Horizon. The concept here is that anything the R2 router learned from the R1 router should never be advertised back to the R1 router. It would be redundant because R1 already knows it, and R1 is one step closer to the truth.

So, we only advertise things we learned from other sources.

Here are the new starting tables. They include the source port number through which we learned the information. P = port, H = hops, Src = source. We use * to indicate direct knowledge.

R1 Routing Table			
Address	P	H	Src
1.1.1.0/24	1	0	*
9.1.2.0/30	2	0	*
0.0.0.0/0	0	0	*

R2 Routing Table			
Address	P	H	Src
2.2.2.0/24	1	0	*
9.1.2.0/30	2	0	*
0.0.0.0/0	0	0	*

After the first advertisement, we have these updated routing tables.

R1 Routing Table			
Address	P	H	Src
1.1.1.0/24	1	0	*
9.1.2.0/30	2	0	*
0.0.0.0/0	0	0	*
2.2.2.0/24	2	1	2
0.0.0.0/0	2	1	2

R2 Routing Table			
Address	P	H	Src
2.2.2.0/24	1	0	*
9.1.2.0/30	2	0	*
0.0.0.0/0	0	0	*
1.1.1.0/24	2	1	2
0.0.0.0/0	2	1	2

So far, so good. Now the big test. What happens after the second advertisement?

R1 Routing Table			
Address	P	H	Src
1.1.1.0/24	1	0	*
9.1.2.0/30	2	0	*
0.0.0.0/0	0	0	*
2.2.2.0/24	2	1	2
0.0.0.0/0	2	1	2

R2 Routing Table			
Address	P	H	Src
2.2.2.0/24	1	0	*
9.1.2.0/30	2	0	*
0.0.0.0/0	0	0	*
1.1.1.0/24	2	1	2
0.0.0.0/0	2	1	2

There is no change at all in the routing tables.

Problem solved. Thank you Split Horizon.

So, can we get rid of the hop count? After all, we never used it with the example we just did.

Not yet.

The problem gets pushed off but does not go away. If we have **THREE** routers, each connected to the other two, we can still get a routing loop.

We still want to keep Distance-Vector information because we might have more than two routers. If we have R1, R2, and R3 all connected to each other, R1 can tell R2 about LAN X, and R2 can tell R3 about LAN X, and R3 can tell R1 about LAN X.

Split Horizon cuts out the first-order duplication of information, but as the network gets more complex, it is not powerful enough. Still, it is very useful.

http://en.wikipedia.org/wiki/Split_horizon has more.

41.4 Thrashing

RIP needs to avoid thrashing. What is thrashing?

Here is a real-world example. For a while I maintained the master phone list for a group of about 200 people. Over time, phone numbers changed. Sometimes I would get a report from someone that person A's phone number was wrong. Let's say I delete it from the list. Then person B helpfully says, oh, I have a phone number for person A. Here, add it to your list. But it is the same wrong number that I had before. That is thrashing.

Thrashing is also called **flapping**.

Exam Question 467 (p.347)**:** What is thrashing?
Acceptable Answer: repeatedly doing and undoing something

In my phone number example I solved the problem by keeping the wrong number but marking it as wrong. RIP does much the same thing.

When RIP discovers that a route has failed, it poisons the routing table entry. It does this by setting the distance metric to infinity (actually 16, the maximum value allowed). That means: yeah, this route existed in the past but now it is broken. If you advertise this route, I will ignore you.

41.5 Route Poisoning

Exam Question 468 (p.347)**:** In RIP, what is route poisoning?
Acceptable Answer: setting an infinite hop count to prevent use of a route

When a router determines that a route has failed, it does route poisoning by sending an infinite distance (16 for RIP) to all its neighbors.

Exam Question 469 (p.347)**:** In RIP, what is poison reverse?
Acceptable Answer: sending route poison back to the router where you learned it

Normally **split horizon** prevents a router from repeating anything it learned back across the interface where it learned it. This is to prevent routing loops. Poison reverse is an exception to this rule.

If we have (A) to (B) to (C), then A sends routing information to B, and B sends it along to C but not back to A. However, with poison reverse, if A poisons a route, A sends that information to B, and B sends it along to C and also back to A.

Is poison reverse necessary? No. Split horizon is enough. Poison reverse is optional but it is always safe and can speed up convergence.

41.6 RIP Timers

But what if the route that was poisoned really comes back? With phone numbers that was unlikely. If a number stops working, it usually stays that way. But with networking a link may go down for a while due to power outage or broken wire and later it gets fixed.

We need a way to avoid thrashing and yet let routes be reinstated.

To do this RIP uses four timers: Update, Invalid, Holddown, and Flush.

Update Timer: Usually 30 seconds, when it expires, the router sends an update to all neighboring routers, giving the routing information that it knows.

Invalid Timer (Expiration Timer): Usually 180 seconds. For each route in the table, this timer is reset to 180 when an update comes in that reaffirms the route. If no update reaffirms the route by the time this timer counts down to zero, the router concludes that the route is invalid.

Flush Timer: Usually 240 seconds. After being marked invalid, the router advertises to its neighbors that the route is invalid. This is called **route poisoning**. It does this for 60 seconds (two update cycles). But if no update reinstates the route, it is dropped from the routing table when the flush timer reaches zero.

Holddown Timer (Cisco): Usually 180 seconds. When a route update advertises a higher distance than it had before, a holddown goes into effect. This higher distance could be valid, but it often means the network is not stable. To restore stability, no updates are accepted for this route until the holddown expires. This lets **route poisoning** do its job.

41.7 Holddown

Exam Question 470 (p.347)**:** In RIP, what is holddown?
Acceptable Answer: the time during which a route cannot be reinstated so route poisoning can work

Holddown is the time during which route poisoning is free to propogate. After that time, it is assumed that all routers would have received the poison. Only a router with direct knowledge about a link would feel authorized to report that link as being up.

The holddown time starts when a route is poisoned: reported as being unreachable. During the holddown, the router refuses to believe any reports which say the route is actually up. This is because some routers may have not yet received the unreachable message, so they may still be falsely advertising the route as valid.

For RIP, the default holddown time is 180 seconds.

http://en.wikipedia.org/wiki/Holddown has more.

41.8 RIP Summary

RIP is the first widespread routing protocol. It is a distance vector routing protocol. It is widely supported and easy to configure. However, it is not widely used.

Exam Question 471 (p.347)**:** List the two big disadvantages of RIP.
Acceptable Answer: slow convergence, does not scale well

RIP is very chatty. It sends lots of update messages consisting of lots of information. As routers are added to the network, the amount of chatting goes up and the messages get larger. For this reason, it does not scale up well.

RIP uses a collection of timers to keep track of the reliability of the information in its routing table.

RIP communicates using UDP port 520.

http://en.wikipedia.org/wiki/Routing_Information_Protocol has more.

41.9 RIPv2

RIP, aka RIPv1, does not support explicit subnet masks. It just works with the implicit network masks that ruled the day before CIDR became popular. RIP2, aka RIPv2, fixes that problem. RIP2 supports CIDR and explicit subnet masking and makes other improvements.

Chapter 42

Link-State Routing

Contents

Link-State routing is better than Distance-Vector routing.

In this chapter we look at the overall link-state routing approach, and briefly mention some well-known link-state routing protocols. We will speak in generalities because each routing protocol can differ in small ways from the overall link-state approach.

Distance-Vector Spreads Rumors: With distance-vector routing, the whole routing table is passed along, so hearsay is passed along as though it were truth. For the most part it works, but it breaks down because tables fill up with information that was heard from a friend of a friend of a friend, and the origins are lost in the mists of time. It is hard to correct misinformation: convergence is slow.

On the other hand, distance-vector is simple. With Distance-Vector each router keeps a single table that holds everything it knows. Simple. Just one.

42.1 Link-State Advertisements

Link-State Keeps It Pure: Link-State routing gets back to basics. Each router says what it directly knows, and does not say anything that it learned from other sources. There is no friend of a friend of a friend rumor mongering. Just reliable knowledge about direct connections.

With Link-State each router keeps a separate chunk of data called the LSA (Link-State Advertisement) for each and every router. If there are 12 routers, each router eventually has 12 LSAs, one for itself and one for each other router. Each LSA is small. These LSAs taken together are called the Link-State Database (LSDB)

Exam Question 472 (p.347)**:** What does LSA stand for?
Acceptable Answer: link-state advertisement

Exam Question 473 (p.347)**:** What does LSDB stand for?
Acceptable Answer: link-state database

Each LSA contains the router ID, its list of direct connections (its topology), and a sequence number. For each connection it can have a cost factor to show how desirable it is. The sequence number goes up by one each time the topology changes (when a direct link is added or deleted).

Sending Advertisements: Each router sends its own LSA to each of its direct neighbors. This happens from time to time and whenever the topology changes.

Receiving Advertisements: When an LSA is received, the sequence number is checked. If the sequence has not changed, nothing else happens.

New Advertisements: If the sequence number is different, the LSA is passed along to all the direct neighbors. This causes an instant flooding of the whole network with the new advertisement. Every router will have the latest version of every LSA.

This flooding makes convergence much faster than with the distance-vector model. Because convergence is faster, this makes link-state better than distance-vector.

Exam Question 474 (p.347)**:** Which routing method is better: link-state or distance-vector?
Acceptable Answer: link-state

One experiment in 2001 showed that distance-vector (RIP) took about ten

times as long to converge as link-state (OSPF), 138s vs 15s. Of course the actual convergence time would depend on the actual before/after topography of the network.

Split horizon is important for distance-vector routing but is not important for link-state routing because the sequence number keeps the number of messages from getting out of hand.

42.2 Routing Tables

One way that old-fashioned distance-vector is better than link-state is that it puts less of a computing load on the individual router. With distance-vector routing, once you get the routing table integrated, your work is done.

Link-state can require a more powerful router. With link-state routing, once you get your LSA advertisements, your work is just starting. You still do not have a routing table. You must build it. There are well-known computer algorithms to help you build the table, but it is still a separate step.

Each router will recompute its routing table when it receives a new LSA. While the new LSA is flooding the network, there is a period of time that routers will have differing views of the network. During this time routing loops can form. Generally this time is short.

42.3 EIGRP

Exam Question 475 (p.347)**:** What does EIGRP stand for?
Acceptable Answer: enhanced interior gateway routing protocol

EIGRP is an improved version of IGRP, the Interior Gateway Routing Protocol. It supports classless networking. IGRP did not support classless networking.

Exam Question 476 (p.347)**:** What is the biggest advantage of EIGRP?
Acceptable Answer: fast convergence

Exam Question 477 (p.348)**:** What is the biggest disadvantage of EIGRP?
Acceptable Answer: limited to cisco equipment

Proprietary means owned by. EIGRP is owned by Cisco and they have limited its use to Cisco equipment only. They want to force people to stay

with Cisco equipment.

Also, because the protocol is owned by one company, that company can continue to make improvements to the protocol without getting anybody else's permission.

42.4 OSPF

Exam Question 478 (p.348)**:** What does OSPF stand for?
Acceptable Answer: open shortest path first

OSPF is almost, but not quite, as efficient as EIGRP. It is open source. It works on both Cisco and non-Cisco routers.

Exam Question 479 (p.348)**:** What is the biggest advantage of OSPF?
Acceptable Answer: works on non-cisco and cisco

42.5 IS-IS

Exam Question 480 (p.348)**:** What does IS-IS stand for?
Acceptable Answer: intermediate system to intermediate system

IS-IS is widely used.

42.6 Further Study

https://en.wikipedia.org/wiki/Link-state_routing_protocol has more on link-state routing.

https://en.wikipedia.org/wiki/Enhanced_Interior_Gateway_Routing_Protocol has more on EIGRP.

https://en.wikipedia.org/wiki/Open_Shortest_Path_First has more on OSPF.

https://en.wikipedia.org/wiki/IS-IS has more on IS-IS.

Unit X

IPv6

Chapter 43

IPv6 Addressing

Contents

43.1 Why Should Anyone Care?

Networking professionals need to understand the IPv6 protocol and be able to think critically about the differences between it and the traditional IPv4 protocol.

Understanding begins with knowing why it was invented in the first place, and also includes why it is still not being used by everyone.

IPv6 (version 6) was approved in 1998 as a replacement for the 15-year-old IPv4 (version 4) protocol (established 1983). But as of 2016, nearly 20 years later, IPv4 is still dominant. IPv6 is set to overtake as IPv4 runs into problems, but it could take a long time.

In favor of IPv4:

(a) It works almost everywhere.

(b) Networking professionals already understand it.

In favor of IPv6:

(a) It is more efficient (faster) than IPv4.

(b) The rate of adoption has been increasing.

(c) It has lots more IP addresses.

Popularity: Currently (December 2016), network traffic typically happens using IPv4, but IPv6 is increasingly being supported as an option.

Google allows connections to itself using both IPv4 and IPv6, and Google publishes statistics showing the percentage of usage that happens through IPv6. You can see the current numbers here:

https://www.google.com/intl/en/ipv6/statistics.html

0.5% January 2012
1.0% January 2013 (100% increase)
2.5% January 2014 (150% increase)
5.0% January 2015 (100% increase)
9.0% January 2016 (80% increase)
13.5% December 2016 (50% increase)

43.2 If IPv6 Is So Great, Why Is IPv4 Still Used?

There are groups of people that say IPv6 is the wave of the future and anyone who clings to IPv4 is misguided and stuck in the past. Maybe. If you just read articles from the same "echo chamber," you may end up with a mistaken impression of reality. Be alert.

There is a saying that you cannot teach an old dog new tricks. People tend to stick with whatever has worked in the past, and they are reluctant to learn new things unless there is a substantial benefit.

So far, there is no substantial benefit to using IPv6 instead of IPv4. Sure, IPv6 is an improvement, but IPv4 still works and will continue to work far into the future. What's the big deal, then?

For you, the student of networking, I have a saying: "In the land of the blind, the one-eyed is king." And with so many people resisting the move to IPv6, if you know even a little about it, you become less blind, and when IPv6 comes up, you will be king.

So, let's get you ahead of the adoption curve.

43.3 Am I Already Using IPv6?

This is actually a pretty interesting question. You may be using IPv6 without even knowing it. Before we go any farther, here are a few things you can try so you will know where you stand right now.

First, visit http://test-ipv6.com/ to see whether your device can communicate using IPv6 at this moment. For me, the first time I tried it the answer was no, but after making a few small adjustments (mentioned below) the answer turned into yes.

Next, to see what is going on in your browser, install an IPv6 add-on. For example, install IPvFox for Firefox or IPvFoo for Chrome. (I installed IPvFoo for Chrome.) This will show you whether the webpage you are viewing arrived on the IPv6 channel or on the IPv4 channel. It will also show you where the sub-parts of that webpage came from.

Most people are just using IPv4 at this point (December 2016).

43.4 Alternate Paths

Communication through the Internet passes through, among other things, layer 3 of the OSI model, which is the network layer.

Exam Question 481 (p.348)**:** At which OSI layer does IPv6 operate?
Required Answer: 3 or network

IPv4 and IPv6 provide two alternate ways of passing data through the Internet. Both operate at OSI protocol layer 3, the Network layer. Because of the OSI multi-layer approach, the change-over to IPv6 does not require

changes in any of the other protocol layers. We just swap one mechanism (IPv4) for another (IPv6).

In order to use IPv4, both endpoints and all the routers in between must support the IPv4 protocol. This is the norm.

Similarly, in order to use IPv6, both endpoints and all the routers in between must support the IPv6 protocol. It is a complete replacement for IPv4.

IPv6 and IPv4 are not compatible with each other. They do not need to be. They exist side by side. They are different languages. Specifically, an IPv6-only device cannot talk to an IPv4-only device, nor vice versa. Instead, IPv6 can only talk to IPv6 and IPv4 can only talk to IPv4.

Exam Question 482 (p.348)**:** Are IPv6 and IPv4 compatible with each other?
Required Answer: no

But we do have bi-lingual devices, devices that can speak in both languages. We identify devices as being IPv4-only, IPv6-only, and dual-stack.

Exam Question 483 (p.348)**:** What does native dual stack mean?
Acceptable Answer: both ipv4 and ipv6 are supported by the sender or receiver

43.5 What Devices Are Involved?

There are four devices (or groups of devices) on your path to communication: Self. Gateway. Internet. Server.

Self: The first device is the one you are directly using. Your laptop. Your cell phone. Buried in its operating system there are network drivers, also known as the network stack. Your device can be IPv4-only, IPv6-only (theoretically), or dual stack. Modern devices are dual stack. Older devices may be IPv4-only.

Gateway: The second device is your gateway router. Typically this is the source of your IP address, which you learn through DHCP, the dynamic host configuration protocol. Often these have an IPv6 setting that is available, and often this setting is turned off. You may need to turn it on for IPv6 to work for you.

Normally devices are assigned an IP address through the DHCP configuration process. As each device begins to use the network, it requests informa-

tion about that network and its place in it. Normally this is provided by a DHCP server.

Potential Bottleneck: I found that my router was limiting me to IPv4 until I made a small change in its configuration. Then, suddenly, I was fully able to use IPv6.

Internet: The next group of devices is the routers of the Internet, starting with your ISP and ending with the ISP of your target (the server you are trying to communicate with). Since June 6, 2012 the Internet Backbone fully supports IPv6 so you should have no problems here.

There is actually an interesting trick at work here, too. Some parts of the Internet already run IPv6-only. To allow IPv4 traffic to pass, they do a thing called tunneling, where they temporarily encapsulate the IPv4 traffic inside an IPv6 shell, send it across, and strip off that shell when it reaches the other end of the backbone.

Server: The last device is the server you are trying to reach.

Many servers have been upgraded to be dual-stack. Such servers support both IPv4 and IPv6 for their traffic. Here are some well-known examples (as of December 2016): google, youtube, facebook, yahoo, wikipedia, linkedin, and netflix.

But there are still many websites that do not support IPv6 yet (as of December 2016). Here are some well-known examples: amazon, reddit, twitter, ebay.

Web hosting has still not caught IPv6 fever. When I checked (December 2016) several reviews of web hosting providers, nowhere was IPv6 even mentioned. One explanation is that IPv4 addresses are becoming scarce, so web hosting providers can charge extra for them, whereas IPv6 addresses are plentiful.

Potential Bottleneck: I run several websites and my web hosting provider does not support IPv6. Someday probably, but not today.

43.6 Getting an IPv6 Address

Under IPv4 there are two main ways addresses get assigned: static and dynamic.

Static (also called manual) means that someone enters the address into the

device and the device remembers it. When it boots up and connects to the network, it already knows its address. It just starts using it.

Dynamic (also called stateful) is usually done through DHCP. The DHCP server presents an address for the host to use.

Under IPv6 we can also have static and dynamic addresses. And there is one more that is popular: SLAAC (also called stateless).

Exam Question 484 (p.348)**:** What does SLAAC stand for?
Acceptable Answer: stateless address auto configuration

With SLAAC, the device uses the first 64 bits provided by the host network and creates its own last 64 bits. It just needs to verify (using the DAD protocol) that nobody is already using that address.

Exam Question 485 (p.348)**:** What does DAD stand for?
Acceptable Answer: duplicate address detection

43.7 Privacy and Security

Marketers and law enforcement have been both excited and frustrated with the advent of IPv6. Those who were excited believed that with all these IPv6 addresses it would be easier to track potential customers or criminals. That is because NAT made it difficult to track people under IPv4. But with IPv6 there is no need for NAT.

It is true that IPv6 lets each device have its own IP address instead of sharing one behind some form of NAT. But it is also true that IPv6 lets each device have a new IP address every day (or as often as it wants).

So, if you are into marketing or law enforcement and if you want to track who did what, and you are planning to use IPv6 addresses to tell you, you will probably be disappointed. You should be looking at other things like cookies and super cookies.

43.8 IoT: The Internet of Things

With many devices starting to use the Internet for communication, we have a growing phenomenon that is popularly called the Internet of Things.

Exam Question 486 (p.348)**:** What does IoT stand for?

Acceptable Answer: internet of things

Examples include intelligent light bulbs, thermostats, refrigerators, door locks, and coffee makers. It is not hard to imagine many devices that could join the IoT.

With IPv4 there is a problem of not having enough addresses for IoT things. There is less than one IPv4 address per person living on the earth today. With IPv6 that problem goes away. There is more than one IPv6 address for each atom on the face of the earth.

43.9 Technical Details

The primary reason for creating IPv6 was the obvious future that showed IPv4 would run out of addresses eventually. But once it was clear that IPv6 would be needed, it was important to fix as many problems as possible.

Addresses: IPv4 has an address space of about 4.3 billion (4.3×10^9) usable addresses. This is because IPv4 addresses are 32 bits long. Each time you add a bit, you double the number of possible addresses. 2^{32} is about 4 billion.

To get more addresses, we have to have more bits. How many should we have? After much deliberation, the committee decided to have 128 bits, four times as many as IPv4. That makes 2^{128} possible addresses, which is about 64 billion billion billion billion.

Exam Question 487 (p.348)**:** Which has a larger address space, IPv4 or IPv6?
Acceptable Answer: IPv6

Exam Question 488 (p.348)**:** In IPv4 how many bits does an address have?
Required Answer: 32

Exam Question 489 (p.348)**:** In IPv6 how many bits does an address have?
Required Answer: 128

Nobody really likes to type out individual bits. Instead we write them in groups. In IPv4 we do not write out 32 individual bits. Instead we write out four groups of 8 bits. In IPv6 we do not write out 128 individual bits.

Instead we write out eight groups of 16 bits.

Exam Question 490 (p.348)**:** In IPv4 how many groups of bits does an address have?
Required Answer: 4

Exam Question 491 (p.348)**:** In IPv4 how many bits are in each group?
Required Answer: 8

Exam Question 492 (p.348)**:** In IPv6 how many groups of bits does an address have?
Required Answer: 8

Exam Question 493 (p.348)**:** In IPv6 how many bits are in each group?
Required Answer: 16

The number base used to write each group is different now. Converting between binary and base 10 is a pain, but base 10 is familiar. Converting between binary and hex is easy, but hex is unfamiliar. For IPv6 they decided to go with easy.

Exam Question 494 (p.348)**:** In IPv4 each group of bits is written in what number base?
Required Answer: 10

Exam Question 495 (p.349)**:** In IPv6 each group of bits is written in what number base?
Required Answer: 16

The group separator is different now. For IPv4 a dot is used. For IPv6 a colon is used.

Exam Question 496 (p.349)**:** In IPv4 address groups are separated by what character?
Required Answer: .

Exam Question 497 (p.349)**:** In IPv6 address groups are separated by what character?
Required Answer: :

43.10 Address Abbreviation

IPv6 addresses can be abbreviated in two ways.

(a) leading zeros in each group can be omitted.

(b) if one or more entire groups are just zero, they can be replaced with `::` (the double colon).

For example, here is an IPv6 address:

0000:0000:0000:0000:0000:0000:0000:0001

We can remove leading zeros. That gives us this:

0:0:0:0:0:0:0:1

We can replace groups of zeros with ::, as follows:

::1

In fact, the zero (aka the "unspecified") address can be simply written as:

::

Here is another example:

2001:db8:0:0:0:0:0:1234

It could be written like this:

2001:db8::1234

Incidentally, the ::1 address in IPv6 corresponds to 127.0.0.1, the **localhost** address, in IPv4. The localhost address is also called the **loopback** address.

Exam Question 498 (p.349)**:** In IPv4 what is the localhost address?
Required Answer: 127.0.0.1

Exam Question 499 (p.349)**:** In IPv6 what is the localhost address?
Required Answer: ::1

Exam Question 500 (p.349)**:** In IPv6 what is the unspecified address?
Required Answer: ::

The unspecified address is the zero address.

Exam Question 501 (p.349)**:** In IPv6 what is the link local network address?
Required Answer: fe80::/64

fe80/64 is 1111 1110 1000 and the rest of the 64 bits are zero.

Link local is similar to self-assigned 169.254.x.x addressing in IPv4.

Exam Question 502 (p.349)**:** In IPv6 what is the all-host multi-cast address?
Required Answer: ff02::1

The all-host multi-cast address is the IPv6 equivalent of the local network or subnet broadcast address.

43.11 Subnets

In IPv4, a Class A address has 8 bits for network, and CIDR or a netmask is used to express how many of the remaining 24 bits are used for subnet. Up to 22 can be used.

In IPv4, a Class B address has 16 bits for network, and CIDR or a netmask is used to express how many of the remaining 16 bits are used for subnet. Up to 14 can be used.

In IPv4, a Class C address has 24 bits for network, and CIDR or a netmask is used to express how many of the remaining 8 bits are used for subnet. Up to 6 can be used.

In IPv6, there are no classes. There are 48 bits used for network and 16 (or more) bits are used for subnet. The remaining 64 bits (or less) are used to identify the host. If more than 16 bits are needed for subnet, CIDR is used to express how big the subnet is.

Exam Question 503 (p.349)**:** In IPv6 how many bits are reserved for the network?
Required Answer: 48

Exam Question 504 (p.349)**:** In IPv6 how many bits are reserved for the subnet?
Required Answer: 16

Exam Question 505 (p.349)**:** In IPv6 how many bits are reserved for the host?
Required Answer: 64

Unit XI

Cisco IOS

Chapter 44

Cisco IOS

Contents

Cisco is the dominant brand of networking equipment in the world (2016), which makes it the standard for networking commands.

IOS is the Cisco device operating system.

44.1 Internetworking

Exam Question 506 (p.349)**:** What is a Collision Domain?
Acceptable Answer: It is the physical layer medium within which one or more hosts might possibly contend for transmission time.

The physical medium includes the full extent to which signals can propogate, including passing through hubs and repeaters.

Exam Question 507 (p.349)**:** How do you break up a Collision Domain?
Acceptable Answer: Divide media segments by using switches or bridges (not hubs or repeaters).

Collision domains still exist and are common in wireless settings.

In wired settings, collision domains are less of a problem. That is because switches are commonly used instead of hubs, and because full-duplex, star-topology wiring is commonly used instead of bus-topology Ethernet wiring.

Each side of a full duplex UTP wire actually represents an independent communications channel. (1,2 goes to 3,6 and vice versa.) There is only one transmitter and one receiver in each channel, so no collisions are possible. Each side of the channel is a separate collision domain, so there are two collision domains.

In spite of this, the wire between a switch and a computer is normally counted as just being a single collision domain.

Exam Question 508 (p.349)**:** What is a Broadcast Domain?
Acceptable Answer: It is the extent to which broadcasts are sent.

Layer 2 broadcasts are sent to MAC address FF:FF:FF:FF:FF:FF and are expected to be received by all devices on the same LAN.

Layer 3 broadcasts are sent to IP address 255.255.255.255 and are currently expected to be received by all devices on the same LAN. In the early days of the Internet, they were received by all devices everywhere on the Internet. Nowadays they are mostly filtered by routers and not passed along to neighboring LANs outside of the organization.

Typical examples of broadcast include DHCP query and ARP/RARP.

So, generally broadcasts pass through switches and are filtered by routers.

Exam Question 509 (p.349)**:** How do you break up a Broadcast Domain?
Acceptable Answer: Divide into separate LANs connected by routers.

Exam Question 510 (p.349)**:** When would you use a Straight cable?
Acceptable Answer: PC to Switch, Switch to Router

For modern equipment, AUTO-MDIX works well. That means the equipment does not need you to pick the right kind of cable. You can use a straight cable or a crossover cable. It senses the cable you picked and configures itself to use that cable properly. That works for newer devices. For older devices, AUTO-MDIX does not work.

Exam Question 511 (p.349)**:** When would you use a crossover cable?
Acceptable Answer: When connecting like devices, such as PC to PC, Switch to Switch, or Router to Router.

You would also use a crossover to connect a PC directly to a Router.

Older devices may need a **crossover** cable.

Many newer devices can detect each other better and can use either straight cables or crossover cables to communicate. Most newer devices, starting about 1998, use **Auto-MDIX** so they do not require a crossover cable. Remember, it takes maybe ten years for everyone to start making and using new equipment, so you may still run into older equipment that requires crossover cables, and you may be asked about it on exams.

Exam Question 512 (p.349)**:** When would you use a Rolled cable?
Acceptable Answer: When managing a switch or router, you use a rolled cable between your PC serial port and the device's console port.

Exam Question 513 (p.350)**:** How does Flow Control work?
Acceptable Answer: The receiver notifies the sender when to start and stop sending.

There are several ways flow control can be done.

You should also be able to do number base conversions between binary, octal, hex, and decimal. This is usually easiest by converting to binary first, and then converting the result to the target number base.

44.2 IOS and SDM

Exam Question 514 (p.350)**:** Give the Cisco IOS command to enter privileged mode.
Required Answer: enable

Exam Question 515 (p.350)**:** Give the Cisco IOS command to leave privileged mode.
Required Answer: disable

Exam Question 516 (p.350)**:** Give the Cisco IOS command to leave user mode.
Required Answer: logout

Exam Question 517 (p.350)**:** Give the Cisco IOS command to enter configuration mode.
Required Answer: conf t

`conf t` is short for `configure terminal`.

Exam Question 518 (p.350)**:** What does the Cisco IOS 'co?' command

do?
Acceptable Answer: list all commands that start with 'co'

Exam Question 519 (p.350)**:** What does the Cisco IOS 'show history' command do?
Acceptable Answer: It displays all recently entered commands.

Exam Question 520 (p.350)**:** What does the Cisco IOS 'show version' command do?
Acceptable Answer: It displays basic information about your hardware and operating system.

44.3 Managing Cisco

You need to understand the four main places where things are stored: RAM, ROM, NVRAM, and Flash. You need to know basically what is stored in each place.

Nothing is "stored" in RAM, but when the router is operational, the running configuration is in RAM.

ROM cannot be changed, so it holds only the most basic things, things that would never change, such as the power-on self-test, the ROM monitor, and the mini-IOS.

Flash is removable. It holds the IOS.

NVRAM is like Flash, but is not removable. It holds the startup configuration, which includes the enable secret (password).

The configuration register is very small. I suspect it is stored in NVRAM along with the startup configuration, but I am not sure.

Exam Question 521 (p.350)**:** What does RAM stand for?
Required Answer: random access memory

Exam Question 522 (p.350)**:** What does ROM stand for?
Required Answer: read-only memory

Exam Question 523 (p.350)**:** What does NVRAM stand for?
Acceptable Answer: non-volatile random-access memory

Exam Question 524 (p.350)**:** What does POST stand for?

Acceptable Answer: power-on self-test

Exam Question 525 (p.350)**:** Where is the POST stored? (ram, rom, flash, nvram)
Required Answer: rom

Exam Question 526 (p.350)**:** What does TFTP stand for?
Required Answer: trivial file transfer protocol

Exam Question 527 (p.350)**:** Where is the ROM monitor stored? (ram, rom, flash, nvram)
Required Answer: rom

Exam Question 528 (p.350)**:** Where is the mini-IOS stored? (ram, rom, flash, nvram)
Required Answer: rom

The mini-IOS is like the ROM monitor, but the mini-IOS is more powerful and complex.

Exam Question 529 (p.350)**:** Where is the IOS stored? (ram, rom, flash, nvram)
Required Answer: flash

Exam Question 530 (p.350)**:** Where is startup-config stored? (ram, rom, flash, nvram)
Required Answer: nvram

Exam Question 531 (p.350)**:** Where is running-config stored? (ram, rom, flash, nvram)
Required Answer: ram

Exam Question 532 (p.350)**:** What does the configuration register control?
Acceptable Answer: boot sequence

Exam Question 533 (p.350)**:** What does the 0x prefix mean?
Acceptable Answer: hex digits follow

Exam Question 534 (p.350)**:** What does the 0 prefix mean?
Acceptable Answer: octal digits follow

Exam Question 535 (p.350)**:** What does 'copy run start' do?
Acceptable Answer: It copies the running configuration to where the startup configuration is stored. This saves any changes you made to the

configuration.

Exam Question 536 (p.350)**:** Where is the 'enable secret' password stored? (ram, rom, flash, nvram)
Required Answer: nvram

Exam Question 537 (p.351)**:** Why is BREAK important?
Acceptable Answer: It lets you interrupt the boot process, enter the ROM monitor, and change the configuration register even without the enable secret password.

Exam Question 538 (p.351)**:** Explain 0x2102.
Acceptable Answer: It is the normal configuration register setting.

Exam Question 539 (p.351)**:** Explain 0x2142.
Acceptable Answer: It is the configuration register setting to avoid loading the startup configuration, which allows password recovery.

44.4 IP Routing

In this chapter, we learn about static routing and RIP-based dynamic routing. We also learn the IOS commands to configure them.

Exam Question 540 (p.351)**:** What does 'sh ip route' do?
Acceptable Answer: show the IP routing table

By seeing the routing tables, we can verify that things are set up properly.

Exam Question 541 (p.351)**:** What does the frame destination address point to?
Acceptable Answer: next receiver in same lan

The frame destination address is a layer 2 physical (MAC) address.

For a packet that has a destination beyond the current LAN, the frame destination address will be the address of the next router in its path.

Exam Question 542 (p.351)**:** What does the packet destination address point to?
Acceptable Answer: final destination ip address

The packet destination addrtess is a layer 3 logical (IP) address.

Exam Question 543 (p.351)**:** What does an 'arp' request return?

Acceptable Answer: MAC address (of the next hop)

Exam Question 544 (p.351)**:** What does 'conf t' mean?
Acceptable Answer: configure via terminal

Exam Question 545 (p.351)**:** What does 'int fa0/1' mean?
Acceptable Answer: interface fastethernet 0/1

Exam Question 546 (p.351)**:** In the 'ip address x y' command, what are x and y?
Acceptable Answer: x is the IP address of the interface, and y is its netmask.

Exam Question 547 (p.351)**:** What does 'no shut' do?
Acceptable Answer: It reverses a shutdown command, either implicit or explicit.

Exam Question 548 (p.351)**:** What does 'config-if' mean?
Acceptable Answer: You are in interface configuration mode.

The major mode is 'config' (configuration), and the minor mode is 'if' (interface).

Exam Question 549 (p.351)**:** What does 'ip route 1.1.1.0 255.255.255.0 5.6.7.8' mean?
Acceptable Answer: 1.1.1.0 is the destination network. 255.255.255.0 is its netmask. 5.6.7.8 is the next hop IP address.

Exam Question 550 (p.351)**:** What does 'ip route 0.0.0.0 0.0.0.0 5.6.7.8' mean?
Acceptable Answer: 5.6.7.8 is the default gateway (or gateway of last resort).

Configure all interfaces on a router using these IOS commands: `interface`, `ip address`, `description`, `no shutdown`. Repeat for all interfaces.

`sh ip route` shows the routing table.

`ping` verifies connectivity.

Unit XII

Other Things

Chapter 45

Helping Your Friend Set Up A Router

Contents

Here is a scenario. Let's say your friend is trying to set up his new wireless router. He knows you took a class in networking, so he is asking for help. You want to help him. What do you do?

Here are some steps you can follow to be of assistance.

45.1 Some Assumptions

Let's assume the router is new and actually works.

Let's assume your friend has a wired connection to the Internet, and that connection will be feeding into the new router.

Let's assume that you can get ahold of a copy of the manual for the router.

Let's assume that you have a Cat5 cable you can use to connect directly to the router.

45.2 Do Some Research

Get a copy of the manual for that router. If you have the paper copy that came with the router, that is fine. Otherwise, maybe there is a CD-ROM that came with the router, and it may have a PDF of the manual on it. That is fine too. If all else fails, use a search engine like Google to find the manual and download a copy of it to your laptop.

We can assume that your friend has not changed the username or password for his router, but if he has, find out what they are.

You will want to find the default username and password for the router.

If all else fails, use a search engine (like Google) and search for "default username" and mention his router. You should be able to find the username and password.

45.3 Connect To The Router

The best way to do this is by using a Cat5 cable. Turn off your laptop's Wi-Fi and go directly through the cable.

If you go through Wi-Fi instead, there is a risk that you will not be connecting to your friend's router. Maybe you will connect with someone else's router. Better to be safe.

45.4 Find Your IP Address

The router probably has DHCP running, and will issue an IP address to your laptop. Use ipconfig or something like it to discover your own IP address, as assigned by the router.

It will probably be something like 192.168.1.100.

You also want to find the IP address of your gateway. It is probably the same thing as your own IP address, but the last quad will be 1. (It could be something else, but 1 is the most common.)

It will probably be something like 192.168.1.1.

45.5 Browse To Your Gateway

Using your favorite browser, type in the IP address of the router.

`http://192.168.1.1` for example.

With any luck, you should be a login screen asking you for a username and password.

Type in the username and password that you discovered with your research above. If it works, good. If not, reset the router and try again.

The instructions for resetting the router should be in the manual. Often it involves pressing a button for ten seconds, or something like that.

45.6 After You Log In

Admin Password: Work with your friend to select a new password. Pick something good. Write it down. In fact, it's okay to write it on a sticky note and tape it to the side or bottom of the router. After all, anyone that has physical access to the router can reset it anyway, so writing the password down is not a big deal.

Wi-Fi SSID: Work with your friend to select a Service Set Identifier to be the name of his Wi-Fi access point. Personalize it. Don't leave it as "linksys" or something else generic.

Wi-Fi Password: Work with your friend to select a password that others can use when they connect to his Wi-Fi access point. **THIS SHOULD NOT BE THE SAME AS THE ADMIN PASSWORD.** You should probably make it something easy to remember so he can tell his other friends when they visit.

Go through the other settings and adjust anything that seems important to you.

Make sure you save your changes. Typically this will cause the router to reboot itself, or ask you for the new password.

If you decided to change the network numbers, you may need to reconnect using a new IP address.

45.7 Try A Wireless Connection

Unplug the Cat5 cable that you have been using.

Turn on your laptop's Wi-Fi.

Look for the SSID that you just created.

Attach to it.

When it asks for a password, put in the new Wi-Fi password.

It should let you in, and you should be able to start browsing the Internet. Congratulations! You are done. Collect your fee (or dinner, or whatever) and escape before he thinks of anything else to ask for.

If not, you may need to backtrack, and redo some or all of the steps above, maybe even doing a reset again.

Unit XIII

Appendix

Appendix A

Test Bank

Test Bank

1: (p.4) What are target skills?

2: (p.4) What are basic skills?

Unit I: Networking Basics

Unit I, Chapter 1: Exploring The Web

3: (p.10) What's the difference between the Web and the Internet?

4: (p.11) What does URL stand for?

5: (p.11) What does URI stand for?

6: (p.11) List in any orderthe three most common parts of a URL.

7: (p.11) In networking, what is a protocol?

8: (p.12) What does HTTP stand for?

9: (p.12) List in any order five popular protocols.

10: (p.12) Does capitalization matter with domain names?

11: (p.13) Does capitalization matter with URLs?

12: (p.13) What is the structure of a domain name?

13: (p.13) Does capitalization matter in the path portion of a URL?

14: (p.13) What is the normal structure of the path?

15: (p.14) What does big-endian mean? Give an example.

16: (p.15) What does little-endian mean? Give an example.

17: (p.16) When is little-endian better?

18: (p.16) What does mixed-endian mean? Give an example.

Unit I, Chapter 2: Parts of the URL

19: (p.17) What is syntax?

20: (p.17) In a URL, where does @ (at) go?

21: (p.17) In a URL, where does : (colon) go?

22: (p.18) In a URL, where does ? (question mark) go?

23: (p.18) In a URL, where does = (equals) go?

24: (p.18) In a URL, where does & (ampersand) go?

25: (p.18) In a URL, where does ; (semi-colon) go?

26: (p.18) In a URL, where does + (plus) go?

27: (p.18) In a URL, where does % (percent) go?

28: (p.19) In a URL, where does # (hash) go?

Unit I, Chapter 3: How The Internet Works

29: (p.21) What does LAN stand for?

30: (p.21) What protocol do most LANs use for communication?

31: (p.22) In networking, what does IP stand for?

32: (p.22) What is a software port?

33: (p.23) What is Survivability?

34: (p.23) Why does the Internet avoid centralization?

Unit I, Chapter 4: Domain Names and DNS

35: (p.25) What does DNS stand for?

36: (p.26) What service does DNS provide?

37: (p.26) Is there any special meaning to the order of the parts in a domain name? If so, what?

38: (p.27) What does TLD stand for?

39: (p.28) What is Cyber Squatting?

40: (p.28) Can someone own a domain name?

41: (p.30) Is byuh.doncolton.com controlled by BYUH? Why or why not?

Unit I, Chapter 5: DHCP: Host Configuration

42: (p.32) What does DHCP stand for?

43: (p.32) What is a host?

44: (p.33) What is configuration?

45: (p.33) What does dynamic mean?

46: (p.33) What does static mean?

47: (p.34) What does DHCP provide? Include a specific example.

48: (p.34) How does a typical laptop computer discover its own IP address?

49: (p.34) How does a typical server computer discover its own IP address?

Unit II: OSI Model

Unit II, Chapter 6: The OSI Model

50: (p.38) What is layer 7 of the OSI model?

51: (p.38) What layer number is the Application layer of the OSI model?

52: (p.38) What is layer 6 of the OSI model?

53: (p.38) What layer number is the Presentation layer of the OSI model?

54: (p.39) At which OSI layer is encryption / decryption?

55: (p.39) At which OSI layer is data compression?

56: (p.39) What is layer 5 of the OSI model?

57: (p.39) What layer number is the Session layer of the OSI model?

58: (p.39) What is layer 4 of the OSI model?

59: (p.39) What layer number is the Transport layer of the OSI model?

60: (p.40) What does MTU stand for?

61: (p.40) What is the typical value for MTU (in bytes)?

62: (p.40) At which OSI layer are software ports?

63: (p.40) What does TCP stand for?

64: (p.40) Which protocol provides for guaranteed delivery of information?

65: (p.40) At which OSI layer is TCP?

66: (p.41) Which protocol provides for fast (but not guaranteed) delivery of information?

67: (p.41) What does UDP stand for?

68: (p.41) At which OSI layer is UDP?

69: (p.41) At which OSI layer do we find segments?

70: (p.41) What is the Protocol Data Unit at OSI layer 4?

71: (p.41) What is layer 3 of the OSI model?

72: (p.41) What layer number is the Network layer of the OSI model?

73: (p.42) At which OSI layer is the Internet?

74: (p.42) At which OSI layer are Wide Area Networks?

75: (p.42) What does WAN stand for?

76: (p.42) At which OSI layer is IP Addressing?

77: (p.42) In networking, what does IP stand for?

78: (p.42) At which OSI layer is Logical Addressing?

79: (p.42) At which OSI layer does a router operate?

80: (p.42) At which OSI layer does a gateway operate?

81: (p.42) At which OSI layer is Network Address Translation?

82: (p.42) At which OSI layer is Port Address Translation?

83: (p.42) At which OSI layer do we find packets?

84: (p.42) What is the Protocol Data Unit at OSI layer 3?

85: (p.43) What is layer 2 of the OSI model?

86: (p.43) What layer number is the Data Link layer of the OSI model?

87: (p.43) What does MAC stand for?

88: (p.43) How does a typical laptop computer discover its own MAC address?

89: (p.43) What does NIC stand for?

90: (p.43) At which OSI layer is the Local Area Network?

91: (p.43) At which OSI layer is MAC Addressing?

92: (p.43) At which OSI layer is Physical Addressing?

93: (p.43) At which OSI layer is Ethernet?

94: (p.43) At which OSI layer does a switch operate?

95: (p.43) At which OSI layer does a bridge operate?

96: (p.43) What is a multi-port bridge called?

97: (p.43) What is a two-port switch called?

98: (p.44) At which OSI layer do we find frames?

99: (p.44) What is the Protocol Data Unit at OSI layer 2?

100: (p.44) What is layer 1 of the OSI model?

101: (p.44) What layer number is the Physical layer of the OSI model?

102: (p.44) At which OSI layer is wireless signal?

103: (p.44) At which OSI layer is coaxial cable?

104: (p.44) At which OSI layer is cat5 cable?

105: (p.44) At which OSI layer is fiber-optic cable?

106: (p.44) What is a multi-port repeater called?

107: (p.44) What is a two-port hub called?

108: (p.44) At which OSI layer does a hub operate?

109: (p.44) At which OSI layer does a signal repeater operate?

110: (p.45) At which OSI layer do we find bits?

111: (p.45) What is the Protocol Data Unit at OSI layer 1?

Unit II, Chapter 7: IP Addressing Preview

112: (p.49) What is a software port?

Unit II, Chapter 8: Converting Between Bases

113: (p.51) Convert binary 11110100010000 to octal.

114: (p.52) Convert binary 11111111011001011001 to hex.

115: (p.52) Convert octal 16471 to binary.

116: (p.53) Convert hex 64209 to binary.

117: (p.53) Convert decimal 162 to binary.

118: (p.54) Convert binary 10001000 to decimal.

Unit II, Chapter 9: Anatomy of a Hop

119: (p.59) What is a hop?

120: (p.60) What does TTL stand for?

121: (p.60) What is the purpose of TTL?

Unit II, Chapter 10: Address Sharing (NAT)

122: (p.62) What does having a routable address mean?

123: (p.62) What does having a non-routable address mean?

124: (p.63) List in any order the five non-routable IP address blocks

125: (p.63) What does NAT stand for?

126: (p.63) What does PAT stand for?

127: (p.64) What does MITM stand for?

128: (p.64) Explain Man In The Middle.

129: (p.66) In NAT, how does the router remember the original sender?

130: (p.67) What is garbage collection?

131: (p.67) How long do NAT address pool entries last?

132: (p.68) What is a keep-alive?

Unit II, Chapter 11: Peer to Peer with NAT

133: (p.72) What does DMZ stand for?

Unit III: Home Networking

Unit III, Chapter 12: Home Network Components

134: (p.76) What does ISP stand for?

135: (p.76) What is a typical broadband download speed in megabits per

second (2013, Worldwide)?

136: (p.76) What is a typical broadband upload speed in megabits per second (2013, Worldwide)?

137: (p.76) What is bandwidth?

138: (p.76) What is throughput?

139: (p.76) List in either order the two measures of network speed.

140: (p.76) List in any order the three measures of network speed.

141: (p.77) What is latency?

142: (p.77) What does demarc stand for?

143: (p.77) Why is the demarc important?

144: (p.77) What does modem stand for?

145: (p.78) What does WAN stand for?

146: (p.78) What does a modem do?

147: (p.79) What does 8P8C stand for?

148: (p.79) In RJ45, what does RJ stand for?

149: (p.80) What is the technical term for a connection that can use either straight-through or crossover cables?

150: (p.80) What does UTP stand for?

151: (p.80) What does UTP do?

152: (p.80) What does Cat 5 stand for?

153: (p.80) Which is better quality, cat5 or cat6?

154: (p.80) What is troubleshooting?

155: (p.81) What two times should you connect your computer directly to a modem?

Unit III, Chapter 13: Home Router

156: (p.83) List in any order the five services a typical home router provides.

157: (p.83) What does a firewall do?

158: (p.83) What does a gateway do?

159: (p.83) What does DHCP do?

160: (p.83) What does NAT do?

161: (p.83) What does Wi-Fi do?

162: (p.84) What does WAP stand for?

163: (p.84) What does WLAN stand for?

164: (p.84) How fast is 802.11b Wi-Fi in Mb/s (theoretical max)?

165: (p.84) How fast is 802.11g Wi-Fi in Mb/s (theoretical max)?

166: (p.84) How fast is 802.11n Wi-Fi in Mb/s (theoretical max per channel)?

167: (p.85) How many connections can a Wi-Fi access point handle?

168: (p.85) How fast (in Mb/s) is a wired connection?

169: (p.85) What is duplex (in general)?

170: (p.86) What is half duplex?

171: (p.86) What is full duplex?

172: (p.86) What benefits does a switch provide?

173: (p.87) Name five components of a typical home Internet system.

Unit III, Chapter 14: Selecting the Pieces

174: (p.89) List in any order the six categories of ISP.

175: (p.91) What does UPS stand for?

176: (p.91) What is a hotspot?

177: (p.93) What does default mean?

178: (p.93) List in either order the two router configuration default values you should not keep.

179: (p.93) List in any order the two passwords a home router normally has.

Unit III, Chapter 15: Making Your Own Cat5 Patch Cable

180: (p.99) What is the maximum length (in meters) for Cat5 cabling?

181: (p.100) With T568 wiring, are the striped wires odd or even?

182: (p.100) With T568 wiring, are the solid-color wires odd or even?

183: (p.100) With T568A wiring, what color goes in slots 1 and 2?

184: (p.101) With T568B wiring, what color goes in slots 1 and 2?

185: (p.101) With T568A wiring, what color goes in slots 3 and 6?

186: (p.101) With T568B wiring, what color goes in slots 3 and 6?

187: (p.101) With T568A wiring, what color goes in slots 4 and 5?

188: (p.101) With T568B wiring, what color goes in slots 4 and 5?

189: (p.101) With T568A wiring, what color goes in slots 7 and 8?

190: (p.101) With T568B wiring, what color goes in slots 7 and 8?

Unit III, Chapter 16: Network Speed

Unit III, Chapter 17: Servers

Unit III, Chapter 18: Troubleshooting the Network

191: (p.108) What is a Global Broadcast Ping?

192: (p.108) What is the command to do a Global Broadcast Ping?

193: (p.109) With Global Broadcast Ping, who is the first responder?

194: (p.109) With Global Broadcast Ping, who is the second responder?

Unit IV: Wireless Networking

Unit IV, Chapter 19: Wi-Fi Configuration

195: (p.118) What is 802.11?

196: (p.118) What does GHz stand for?

197: (p.118) What are the two main Wi-Fi frequency ranges?

198: (p.119) What 802.11b Wi-Fi channels exist (in the USA)?

199: (p.119) What is Channel 196?

200: (p.119) Which Wi-Fi band gets better distance, 2.4 or 5.0?

201: (p.120) Which Wi-Fi band has less competition, 2.4 or 5.0?

202: (p.120) Which Wi-Fi band has more usable channels, 2.4 or 5.0?

203: (p.120) What is a site survey?

204: (p.121) In networking, what does sniff mean?

205: (p.121) What 802.11b channels are commonly usable (in the USA)?

206: (p.121) Why are many Wi-Fi channels not used?

207: (p.121) What does SSID stand for?

208: (p.121) What is the purpose of the SSID?

209: (p.121) How many characters long can an SSID be?

210: (p.122) Why are some SSIDs hidden?

211: (p.122) Does hiding your SSID improve security?

212: (p.123) List in any order the three Wi-Fi security methods that are commonly used.

213: (p.123) What does WEP stand for?

214: (p.124) What does WPA stand for?

215: (p.124) Which is better: WEP or WPA?

Unit IV, Chapter 20: Wi-Fi Antennas and Signal Strength

216: (p.125) What is the legal maximum Wi-Fi signal (in milliwatts) in the USA?

217: (p.127) What does dBm stand for?

218: (p.127) What is the difference between dBm and dB?

219: (p.128) What does SNR stand for?

220: (p.128) In what units is SNR measured?

221: (p.128) What is the minimum SNR (in dB) needed for a usable connection?

222: (p.129) What is the typical range (in meters) for Wi-Fi signals?

223: (p.129) For typical Wi-Fi, how much signal (in dB) can be used up before the SNR is too low for useful communication?

224: (p.129) For typical Wi-Fi, how much signal (in dB) is lost per ten meters of open air?

225: (p.130) For typical Wi-Fi, how much signal (in dB) is lost per interior wall (plaster-board, wooden studs)?

226: (p.130) For typical Wi-Fi, how much signal (in dB) is lost per exterior wall (wood, brick, cement block, metal studs)?

227: (p.130) For typical Wi-Fi, how much signal (in dB) is lost per floor (thick plywood, support beams)?

228: (p.131) Name at least three typical indoor obstacles that affect Wi-Fi signal strength.

229: (p.131) Name at least three typical outdoor obstacles that affect Wi-Fi signal.

230: (p.131) What conflict happens with 802.11b networks?

Unit V: Security

Unit V, Chapter 21: Passwords

231: (p.133) Why are weak passwords a significant problem in networks?

232: (p.136) List in any order the four measures of password quality.

233: (p.138) What is the problem with short passwords?

234: (p.138) What is the problem with long passwords?

235: (p.138) What is a dictionary attack?

236: (p.138) What is the problem with dictionary passwords?

237: (p.139) What do I recommend for a password?

238: (p.140) What is a high-value password?

239: (p.140) What is a low-value password?

240: (p.141) Does it matter if a low-value password is easy to guess?

241: (p.141) If several high-value passwords are the same is that okay?

242: (p.141) If several low-value passwords are the same is that okay?

243: (p.141) Are password managers a good thing?

244: (p.142) List up to three problems with changing passwords frequently.

245: (p.142) What is the problem with changing passwords rarely?

Unit V, Chapter 22: Security Protocols

Unit V, Chapter 23: Authentication

246: (p.146) List in any order the four types of things used to prove identity (four single words).

247: (p.147) What is multi-factor authentication?

248: (p.147) Is it multi-factor authentication if you have both a password and a security question?

249: (p.147) What does 2FA stand for?

250: (p.147) What is single sign-on?

251: (p.148) What is a hacker?

252: (p.148) Is hacking bad?

253: (p.148) What does black hat mean?

254: (p.148) What does white hat mean?

255: (p.148) What is pen testing?

256: (p.150) Is http considered to be secure? Why?

257: (p.150) Is https considered to be secure? Why?

258: (p.150) What does SSL stand for?

259: (p.151) What does TLS stand for?

260: (p.151) How does SSL protect confidentiality of a TCP connection?

261: (p.152) What are symmetric keys?

262: (p.152) What does rot13 stand for?

263: (p.152) How does rot13 work?

Unit V, Chapter 24: Public Key Systems

264: (p.155) Who knows Alice's public key?

265: (p.155) Who knows Alice's private key?

266: (p.155) Whose key, and which key do you use to send a private message to Bob?

267: (p.155) What is the purpose of encrypting a message?

268: (p.156) What is the purpose of signing a message?

269: (p.156) Whose key, and which key do you use to sign a message?

270: (p.156) How does signing prove authorship?

271: (p.157) How can Bob send a private, authenticated message to Alice?

272: (p.158) What do public-key systems make possible?

273: (p.158) Why is RSA special?

274: (p.159) What is a prime number?

275: (p.160) Why are prime numbers used in encryption?

276: (p.160) What does the RSA private key consist of?

277: (p.160) What does the RSA public key consist of?

278: (p.160) If RSA is so great, why are other things used?

Unit V, Chapter 25: Firewalls

279: (p.162) What is an Outside Threat?

280: (p.163) What is a botnet?

281: (p.163) For what three things are botnets commonly used?

282: (p.163) What does DDOS stand for?

283: (p.164) What is a zombie?

284: (p.164) What does PWN stand for?

285: (p.164) What is an Inside Threat?

286: (p.164) What two things does server mean?

287: (p.166) How can firewalls defend against network attacks on clients?

288: (p.166) How can firewalls defend against network attacks on servers?

289: (p.167) How does DDOS defeat firewall protection for servers?

290: (p.167) What does DMZ stand for?

291: (p.168) What service does DMZ provide?

292: (p.168) What service does port forwarding provide?

293: (p.168) How can sharing your Wi-Fi be dangerous?

Unit VI: IPv4 Addressing

Unit VI, Chapter 26: Number Bases

294: (p.174) What does a dotted quad number look like?

295: (p.174) When an IPv4 address is written in x.x.x.x format, the possible value for x range from 0 to what?

296: (p.177) Describe base 2.

297: (p.177) Describe base 8.

298: (p.177) Describe base 10.

299: (p.177) Describe base 16.

300: (p.177) Describe base 60.

301: (p.178) Describe base 64.

302: (p.178) Describe base 256.

Unit VI, Chapter 27: IPv4 Addresses: Advanced

303: (p.179) What is an octet?

304: (p.179) How many bits in a byte?

305: (p.179) How many bits in a nybble?

306: (p.180) The number 0755 is assumed to be in what number base?

307: (p.180) The number 755 is assumed to be in what number base?

308: (p.180) The number 0x755 is assumed to be in what number base?

309: (p.181) What are the two meanings of kilo?

310: (p.181) What are the two meanings of meg?

311: (p.181) What are the two meanings of gig?

312: (p.186) For 19.19.19.19, what class is it?

313: (p.186) For 199.199.199.199, what class is it?

314: (p.187) What does a net mask look like?

315: (p.187) In a net mask, what do the 1s mean?

316: (p.187) In a net mask, what do the 0s mean?

317: (p.188) What is the IPv4 special address range for the local network?

318: (p.188) What is the Class A Private Address Range?

319: (p.188) In the Class A Private Address Range, how many (classful) networks are there?

320: (p.188) In the Class A Private Address Range, what is the first IP address?

321: (p.188) In the Class A Private Address Range, what is the first usable host address?

322: (p.188) In the Class A Private Address Range, what is the last IP address?

323: (p.188) In the Class A Private Address Range, what is the last usable host address?

324: (p.188) What is the local host Private Address Range?

325: (p.188) What is the Link Local (APIPA) Private Address Range?

326: (p.189) How are Link Local (APIPA) addresses assigned?

327: (p.189) What is the Class B Private Address Range?

328: (p.189) In the Class B Private Address Range, how many (classful) networks are there?

329: (p.189) In the Class B Private Address Range, what is the first IP

address?

330: (p.189) In the Class B Private Address Range, what is the first usable host address?

331: (p.189) In the Class B Private Address Range, what is the last IP address?

332: (p.189) In the Class B Private Address Range, what is the last usable host address?

333: (p.189) What is the Class C Private Address Range?

334: (p.189) In the Class C Private Address Range, how many (classful) networks are there?

335: (p.189) In the Class C Private Address Range, what is the first IP address?

336: (p.190) In the Class C Private Address Range, what is the first usable host address?

337: (p.190) In the Class C Private Address Range, what is the last IP address?

338: (p.190) In the Class C Private Address Range, what is the last usable host address?

339: (p.190) What is the IPv4 global broadcast address?

Unit VI, Chapter 28: IPv4 Addresses: Classless

340: (p.191) What does CIDR stand for?

341: (p.192) What does CIDR notation look like?

342: (p.192) What is the other name for slash notation?

343: (p.192) What is the other name for CIDR notation?

344: (p.192) For 199.199.199.199, what is the default Net Mask in CIDR and

dotted quad notation?

345: (p.194) For /10, what is the subnet block size?

346: (p.195) For /23, what is the subnet block size?

347: (p.195) For /28, what is the subnet block size?

348: (p.196) For 255.255.248.0, what is the subnet block size?

Unit VI, Chapter 29: VLSM

349: (p.202) What does VLSM stand for?

Unit VI, Chapter 30: Ports

350: (p.209) What is a software port?

351: (p.209) Software port numbers range from 0 up to what number?

352: (p.210) What is port 80 normally used for?

353: (p.210) What port does http normally use?

354: (p.210) What does http stand for?

355: (p.210) What is port 443 normally used for?

356: (p.210) What port does https normally use?

357: (p.210) What does https stand for?

358: (p.210) What is port 21 normally used for?

359: (p.210) What port does ftp normally use?

360: (p.210) What does ftp stand for?

361: (p.211) What is port 22 normally used for?

362: (p.211) What port does ssh normally use?

363: (p.211) What does ssh stand for?

364: (p.211) What is port 23 normally used for?

365: (p.211) What port does telnet normally use?

366: (p.211) What is port 25 normally used for?

367: (p.211) What port does smtp normally use?

368: (p.211) What does smtp stand for?

369: (p.211) What is port 53 normally used for?

370: (p.211) What port does dns normally use?

371: (p.212) What does dns stand for?

372: (p.212) What is port 123 normally used for?

373: (p.212) What port does ntp normally use?

374: (p.212) What does ntp stand for?

375: (p.213) What does a client do?

376: (p.213) What does a server do?

377: (p.214) What two things does server mean?

Unit VII: Power Tools

Unit VII, Chapter 31: Basic Power Tools

378: (p.222) What does intermittent mean?

379: (p.222) What is another word for intermittent?

380: (p.222) What does a successful ping prove?

381: (p.222) What does a failed ping prove?

382: (p.222) What one statistic does ping report?

383: (p.223) Using ping to troubleshoot, conventional wisdom says you should ping what first?

384: (p.223) Using ping to troubleshoot, conventional wisdom says you should ping what second?

385: (p.223) Using ping to troubleshoot, conventional wisdom says you should ping what third?

386: (p.223) Using ping to troubleshoot, conventional wisdom says you should ping what fourth?

387: (p.225) What is localhost?

388: (p.225) What is the IPv4 address of localhost?

389: (p.225) What things does a successful ping to localhost prove?

390: (p.226) How can you (the user) find your IP address?

391: (p.227) What things does a successful ping to your own LAN address prove?

392: (p.228) What things does a successful ping to your neighbor prove?

393: (p.229) What is a smurf attack?

394: (p.229) What is a broadcast ping with a fake source address called?

395: (p.229) What information can a broadcast ping provide?

396: (p.230) What IP address is used for global broadcast?

397: (p.231) What things does a successful ping to something beyond your LAN prove?

Unit VII, Chapter 32: Intermediate Power Tools

398: (p.233) What two things does traceroute report?

399: (p.234) What does TTL stand for?

400: (p.234) What is the purpose of TTL?

401: (p.234) How does traceroute use TTL?

402: (p.236) What does FTP stand for?

403: (p.236) Is FTP considered to be secure? Why?

404: (p.236) Which is more secure, ssh or ftp?

405: (p.236) Which is more widely available, ssh or ftp?

406: (p.237) Is telnet considered to be secure? Why?

407: (p.237) Is ssh considered to be secure? Why?

408: (p.237) Which is more secure, ssh or telnet?

409: (p.237) Which is more widely available, ssh or telnet?

410: (p.237) What port does telnet normally use?

411: (p.237) What is telnetd?

412: (p.237) What is a daemon?

413: (p.238) In a program name, what does the suffix d usually mean?

Unit VII, Chapter 33: Advanced Power Tools

414: (p.240) What port does ssh normally use?

415: (p.240) What does ssh stand for?

416: (p.240) How does ssh establish a secure connection?

417: (p.240) Which is more current, nslookup or dig?

418: (p.240) What does deprecated mean?

419: (p.244) What does Wireshark do?

420: (p.244) How does Wireshark get data?

Unit VIII: Switching

Unit VIII, Chapter 34: Overview of Switching

421: (p.247) List in any order the three main network topologies in use today.

422: (p.247) What is a collision?

423: (p.248) What is a Collision Domain?

424: (p.248) How do you break up a Collision Domain?

425: (p.248) What does half duplex mean?

426: (p.248) What does full duplex mean?

427: (p.250) List in any order two protocols that use broadcast in a LAN.

428: (p.251) What does DHCP stand for?

429: (p.251) What does ARP stand for?

430: (p.251) What does ARP do?

431: (p.251) What is a Broadcast Domain?

432: (p.251) How do you break up a Broadcast Domain?

Unit VIII, Chapter 35: Plan B: Redundancy

433: (p.253) What does STP stand for?

434: (p.253) What is convergence?

435: (p.254) About how many seconds does it take STP to converge?

436: (p.254) Which switch is the root bridge?

437: (p.255) Which port is the root port?

438: (p.255) List in any order classic STP's four port state options.

439: (p.256) What does RSTP stand for?

440: (p.256) About how many seconds does it take RSTP to converge?

Unit IX: Routing

Unit IX, Chapter 36: Review of Routing

441: (p.259) What is a hop?

442: (p.259) What does TTL stand for?

443: (p.259) What is the purpose of TTL?

444: (p.260) List the two main things that traceroute reports.

Unit IX, Chapter 37: Netmasks and Addressing

445: (p.262) Class A addresses start with what bit(s)?

446: (p.262) Class A addresses have what netmask?

447: (p.262) Class B addresses start with what bit(s)?

448: (p.262) Class B addresses have what netmask?

449: (p.262) Class C addresses start with what bit(s)?

450: (p.262) Class C addresses have what netmask?

451: (p.263) What does CIDR stand for?

452: (p.263) What does VLSM stand for?

Unit IX, Chapter 38: Types of Routers

453: (p.265) List in any order the three layers of the Cisco router model.

454: (p.265) What does AS stand for?

455: (p.267) How can we tell if two machines are in the same LAN?

456: (p.267) When does a computer send a frame directly to its destination?

457: (p.267) When does a computer send a frame directly to the gateway?

458: (p.267) What does a gateway do?

459: (p.267) If the router's IP address is 1.2.3.4, what is the most likely value for the gateway address in that LAN?

460: (p.268) What does NAT stand for?

461: (p.268) What does NAT do?

462: (p.269) How does NAT defend against attacks on local computers?

463: (p.271) What happens when both LANs for the router have the same network number?

Unit IX, Chapter 39: Distribution Routers

464: (p.274) List the other names for aggregation.

Unit IX, Chapter 40: Routing Table Example

465: (p.275) List the two big advantages of RIP.

Unit IX, Chapter 41: RIP: Routing Information Protocol

466: (p.281) What does RIP stand for?

467: (p.287) What is thrashing?

468: (p.288) In RIP, what is route poisoning?

469: (p.288) In RIP, what is poison reverse?

470: (p.289) In RIP, what is holddown?

471: (p.290) List the two big disadvantages of RIP.

Unit IX, Chapter 42: Link-State Routing

472: (p.292) What does LSA stand for?

473: (p.292) What does LSDB stand for?

474: (p.292) Which routing method is better: link-state or distance-vector?

475: (p.293) What does EIGRP stand for?

476: (p.293) What is the biggest advantage of EIGRP?

477: (p.293) What is the biggest disadvantage of EIGRP?

478: (p.294) What does OSPF stand for?

479: (p.294) What is the biggest advantage of OSPF?

480: (p.294) What does IS-IS stand for?

Unit X: IPv6

Unit X, Chapter 43: IPv6 Addressing

481: (p.298) At which OSI layer does IPv6 operate?

482: (p.299) Are IPv6 and IPv4 compatible with each other?

483: (p.299) What does native dual stack mean?

484: (p.301) What does SLAAC stand for?

485: (p.301) What does DAD stand for?

486: (p.301) What does IoT stand for?

487: (p.302) Which has a larger address space, IPv4 or IPv6?

488: (p.302) In IPv4 how many bits does an address have?

489: (p.302) In IPv6 how many bits does an address have?

490: (p.303) In IPv4 how many groups of bits does an address have?

491: (p.303) In IPv4 how many bits are in each group?

492: (p.303) In IPv6 how many groups of bits does an address have?

493: (p.303) In IPv6 how many bits are in each group?

494: (p.303) In IPv4 each group of bits is written in what number base?

495: (p.303) In IPv6 each group of bits is written in what number base?

496: (p.303) In IPv4 address groups are separated by what character?

497: (p.303) In IPv6 address groups are separated by what character?

498: (p.304) In IPv4 what is the localhost address?

499: (p.304) In IPv6 what is the localhost address?

500: (p.304) In IPv6 what is the unspecified address?

501: (p.304) In IPv6 what is the link local network address?

502: (p.304) In IPv6 what is the all-host multi-cast address?

503: (p.305) In IPv6 how many bits are reserved for the network?

504: (p.305) In IPv6 how many bits are reserved for the subnet?

505: (p.305) In IPv6 how many bits are reserved for the host?

Unit XI: Cisco IOS

Unit XI, Chapter 44: Cisco IOS

506: (p.307) What is a Collision Domain?

507: (p.307) How do you break up a Collision Domain?

508: (p.308) What is a Broadcast Domain?

509: (p.308) How do you break up a Broadcast Domain?

510: (p.308) When would you use a Straight cable?

511: (p.308) When would you use a crossover cable?

512: (p.309) When would you use a Rolled cable?

513: (p.309) How does Flow Control work?

514: (p.309) Give the Cisco IOS command to enter privileged mode.

515: (p.309) Give the Cisco IOS command to leave privileged mode.

516: (p.309) Give the Cisco IOS command to leave user mode.

517: (p.309) Give the Cisco IOS command to enter configuration mode.

518: (p.309) What does the Cisco IOS 'co?' command do?

519: (p.310) What does the Cisco IOS 'show history' command do?

520: (p.310) What does the Cisco IOS 'show version' command do?

521: (p.310) What does RAM stand for?

522: (p.310) What does ROM stand for?

523: (p.310) What does NVRAM stand for?

524: (p.310) What does POST stand for?

525: (p.311) Where is the POST stored? (ram, rom, flash, nvram)

526: (p.311) What does TFTP stand for?

527: (p.311) Where is the ROM monitor stored? (ram, rom, flash, nvram)

528: (p.311) Where is the mini-IOS stored? (ram, rom, flash, nvram)

529: (p.311) Where is the IOS stored? (ram, rom, flash, nvram)

530: (p.311) Where is startup-config stored? (ram, rom, flash, nvram)

531: (p.311) Where is running-config stored? (ram, rom, flash, nvram)

532: (p.311) What does the configuration register control?

533: (p.311) What does the 0x prefix mean?

534: (p.311) What does the 0 prefix mean?

535: (p.311) What does 'copy run start' do?

536: (p.312) Where is the 'enable secret' password stored? (ram, rom, flash,

nvram)

537: (p.312) Why is BREAK important?

538: (p.312) Explain 0x2102.

539: (p.312) Explain 0x2142.

540: (p.312) What does 'sh ip route' do?

541: (p.312) What does the frame destination address point to?

542: (p.312) What does the packet destination address point to?

543: (p.312) What does an 'arp' request return?

544: (p.313) What does 'conf t' mean?

545: (p.313) What does 'int fa0/1' mean?

546: (p.313) In the 'ip address x y' command, what are x and y?

547: (p.313) What does 'no shut' do?

548: (p.313) What does 'config-if' mean?

549: (p.313) What does 'ip route 1.1.1.0 255.255.255.0 5.6.7.8' mean?

550: (p.313) What does 'ip route 0.0.0.0 0.0.0.0 5.6.7.8' mean?

Unit XII: Other Things

Unit XII, Chapter 45: Helping Your Friend Set Up A Router

Unit XIII: Appendix

Unit XIII, Chapter A: Test Bank

Index