




Figure 37.1: TCP Segmentation

Segmentation brings benefits too?

Imagine you download a web page, which contains a graphics file called logo.gif. Its size is 146 KB. If TCP sends all the data in one TCP segment a single bit can get an error during transmission. This will result in TCP will have to resend the entire packet. Now let TCP sends logo.gif as 100 segments with 1460 bytes in each segment. If a single bit error occurred, only one segment would have to be re-sent. Thus, improving the network performance.

When TCP segments the application data, it uses the first byte of each segment for sequence number. TCP actually numbers the bytes, not the segments. This is shown next.




Figure 37.2: TCP segment the application Data

Topic 38: Basics of Internet Protocol (IP)

In this topic, we provide the basics of Internet Protocol (IP).

In a typical large network, many routers and LAN switches sit between a client PC and a server. **Routers** are networking devices that connect to multiple physical networks, such as multiple Ethernets. They make decisions about where to forward the data so that it reaches the correct destination.

Example

Assume Hannah on a LAN in Mason wants to get a web page from a web server <http://www.cisco.com> on a LAN in Cincinnati. This is shown next.




Figure 38.1: Topic 38: Basics of Internet Protocol (IP)

Routing is the complete process by which a computer sends the data, passing through all the routers and eventually arriving at the destination. Previous example shows only LANs, the routers could also be connected to a wide-area network (WAN).

Two Key functions

Forwarding is moving a packet from router's input to appropriate router output. Routing is the determination of the route taken by packets from source to destination. These ideas are shown next.




Figure 38.2: Routing with connected to wide area network (WAN)

IP provides logical communication between hosts. IP makes its best effort to deliver segments between communicating hosts, but it makes no guarantees. IP is said to be an unreliable service.




Figure 38.3: Transport Layer TCP, UDP

Topic 39: Working of IP and IP addresses

This topic explains the working of IP and usage of IP addresses.

Internet Protocol (IP) is the network layer protocol of TCP/IP model. IP defines addressing, forwarding and routing. The addressing details recommended by IP facilitate easy and efficient forwarding of IP packets. IP addresses are 32-bit binary numbers.

$$223.1.1.1 = \underbrace{11011111}_{223} \underbrace{00000001}_1 \underbrace{00000001}_1 \underbrace{00000001}_1$$

Each of the decimal numbers in an IP address is called an octet. Range falls in [0 255]. 2^{32} is equal to 4 billion possible IP addresses. Everyone should have a unique IP address to avoid confusion when trying to deliver data to that address. Each network interface on a computer needs an IP address. A network interface is simply a card that has a physical connector for some type of network. Ethernet NIC is one such example. The NIC takes care of TCP/IP network interface layer details, which are the equivalent of OSI Layers 1 and 2. Most end user computers also known as TCP/IP hosts have a single network interface. Devices that have more than one network interface have more than one IP address. Routers are such devices. IP defines a 20-byte long header, which includes a 4-byte source IP address and a 4-byte destination IP address. An IP packet includes the IP header, along with any data that follows the IP header. To send data from one computer to the other, a sender puts the destination computer's IP address into the destination IP address field and puts its own IP address in the source IP address. This is shown next.




Figure 39.1: IP address field

Next, we show the process of getting a home page. Application layer generates an HTTP GET request. TCP adds its header. The TCP segment gets encapsulated in an IP packet.




Figure 39.2: TCP segment gets encapsulated in an IP packet

HTTP is concerned about getting a web page to the user. TCP is concerned about segmentation, error recovery, and other things, on behalf of HTTP. IP provides an end-to-end delivery of packets service to TCP. Hannah types web server's name `http://www.fredsco.com` and corresponding IP address is 3.3.3.3. Names are more human friendly while addresses are more computer friendly.

Topic 40: How to run an IP Network?

This topic explains how an IP network runs.

The postal service provides a wonderful service. You put a letter into the mailbox, and it magically appears at the right address. By sending an IP packet into the network, the networking devices should collectively be able to forward the packet to the right destination. An example of an IP address is 1.1.1.1. There are 4 decimal numbers, between 0 and 255 inclusive, separated by periods. The structure and meaning of IP addresses tell us how IP addresses are used in an internetwork. To understand the concepts, let's take postal service as a reference. In addition to the number, street, town, and state, the postal service uses postal codes to make it easy

to sort the mail. Everyone in the same general geographical area such as a town has the same postal code. The only time the whole postal address is needed is when the letter gets to the final post office; the one that handles mail for that postal code.




Figure 40.1: IP network

All IP addresses on the same physical network such as one Ethernet LAN have a portion of their IP addresses in common. This common portion works like a postal code. Like postal codes, all the IP addresses on the same Ethernet are in the same general area, so routing can take advantage of that fact.




Figure 40.2: Same Ethernet are in the same general area

IP address grouping makes routing easy. Let's imagine the logic needed by R1: Packets whose destination begins with 1 should be forwarded to the left. Packets whose destination begins with 2 should be forwarded to the Ethernet on the right. Packets whose destination begins with 3 should be forwarded to R2 so that R2 can forward the packet. The routers do not have to know about every IP address in the

network. IP calls the group of IP addresses that share a common beginning to part of their addresses an IP network. The common portion is called network number and is used to represent a network. For an IP network number, the bits other than the common portion have all 0s. In previous example, the network numbers are 1.0.0.0, 2.0.0.0 and 3.0.0.0. A host cannot use a network number as an IP address in the network.

Topic 41: Classes of IP Networks

This topic describes the classes of IP networks.

Let's assume a simple internetwork with two routers and three IP networks.




Figure 41.1: Classes of IP Networks

How many IP addresses can you have in network 1.0.0.0? You can't use 1.0.0.0 and 1.255.255.255 – are reserved. Valid addresses: more than 16 million addresses. IP defines three sizes of networks as different classes. The three different network classes are called Class A, B, and C. All addresses in the same Class A, B, or C network have the same numeric value for the network portion of the addresses. The rest of the address is called the host portion of the address.

Any Network of This Class	Number of Network Bytes (Bits)	Number of Host Bytes (Bits)	Number of Addresses per Network
A	1 (8)	3 (24)	$2^{24} - 2$, or 16,777,214
B	2 (16)	2 (16)	$2^{16} - 2$, or 65,534
C	3 (24)	1 (8)	$2^8 - 2$, or 254

Let's assume the same previous internetwork. By using one IP network number for each physical network, allows easy routing.




Figure 41.2: Classes of IP Networks

Class A networks allow for a ton of IP addresses in a single IP network. Class B networks allow for a pretty large number, and Class C networks allow for a smaller number of host IP addresses. The people who made up IP addressing chose three sizes of networks because one size doesn't fit all companies and organizations. Larger companies can use Class A, medium-sized companies can use Class B, and small companies can use Class C.

Class	First Octet Range	Valid Network Numbers	Total Number of Networks of This Class
A	1 to 126	1.0.0.0 to 126.0.0.0	2^7-2 , or 126
B	128 to 191	128.1.0.0 to 191.254.0.0	$2^{14}-2$, or 16,382
C	192 to 223	192.0.1.0 to 223.255.254.0	$2^{21}-2$, or 2,097,150

Topic 42: IP Subnetting

In this topic, we explain the use of IP subnetting.

These days, single LANs with more than 1000 devices are extremely rare. Many IP addresses will go unused if you use a Class A or Class B network for the devices on a LAN. Subnetting provides a solution. Let's assume that a design requires three networks, and uses three different Class B networks. Each Class B network can host 65,534 addresses.




Figure 42.1: IP Subnetting

A subnet is just a subdivision of a larger Class A, B, or C network. Subnetting refers to the process whereby the engineer decides to create subnets. **Without Subnetting**, devices in the same Class A, B, or C network cannot be separated from each other by a router. Also, devices in different Class A, B, or C networks must be separated from each other by a router. **With Subnetting**, devices in the same subnet cannot be separated from each other by a router. Also, devices in different subnets must be separated from each other by a router. Subnetting allows a network engineer to configure network devices of a class B network for example, 150.1.0.0, such that they can think that the first 3 octets of the addresses identify the network. Following subnets can be created: all addresses that begin with 150.1.1, all addresses that begin with 150.1.2, all addresses that begin with 150.1.3 and so on.

An IP subnet number is used to represent a subnet. The subnet number has the same value in the first part of the number as all the host addresses, and 0s in the last part.




Figure 42.2: IP Subnetting

This internetwork of three Ethernet LANs only uses a part of Class B network 150.1.0.0. The same internetwork without subnetting fully uses three B classes (150.1.0.0, 150.2.0.0, and 150.3.0.0).

Figures and Material used for this topic have been adapted from Kurose and Ross's book with the title “Computer Networking: A Top-down Approach”, 6th Edition, 2013.

Topic 43: Network Address Translation (NAT)

This topic describes explain the concept of NAT.

Small office, home office (SOHO) subnets connect devices such as cameras, computers, smartphones, game consoles. Every IP-capable device needs an IP address. A SOHO needs a range of addresses that can be allocated by the ISP. With 32-bits, possible IP addresses are equal to 2^{32} or 4 billion. If the subnet grew bigger, a larger block of addresses would have to be allocated. But what if the ISP had already allocated the contiguous portions of the SOHO network's current address range? Network address translation (NAT) is a simpler approach to address allocation. It is defined in RFC 2663. It has found increasingly widespread use.

An Example

Assume a home network with 3 computers. A NAT-enabled router, residing in the home has an interface that is part of the home network. All four interfaces in the home network have the same subnet address of 10.0.0/24.




Figure 43.1: Network Address Translation (NAT)

There can be hundreds of home networks, many using the same address space, 10.0.0.0/24. Devices within a given home network can send packets using 10.0.0.0/24 addressing. How is addressing handled when packets are sent to or

received from the Internet? The NAT-enabled router does not look like a router to the outside world. Instead, it behaves to the outside world as a single device with a single IP address. All traffic leaving the home router for the Internet has a source IP address of 138.76.29.7, and all traffic entering the home router must have a destination address of 138.76.29.7. A NAT-enabled router hides the details of the home network from the outside world. If all datagrams arriving at the NAT router from the Internet have the same destination IP address, then how does the router know the internal host to which it should forward a given datagram?




Figure 43.2: NAT-enabled router

Motivations:

Local network uses just one IP address as far as outside world is concerned: 1) range of addresses not needed from ISP. 2) Addresses of devices in local network can change without notifying outside world. 3) ISP can change without changing addresses of devices in local network.

Figures and Material used for this topic have been adapted from Kurose and Ross's book with the title *"Computer Networking: A Top-down Approach"*, 6th Edition, 2013.

Topic 44: Dynamic Host Configuration Protocol (DHCP)

In this topic, we describe the Dynamic Host Configuration Protocol.

An organization obtains a block of IP addresses from an ISP. Then, it can assign individual IP addresses to the host and router interfaces in its organization. A system

administrator will typically manually configure the IP addresses into the router. Host addresses can be configured manually or Dynamic Host Configuration Protocol.

1- Manual configuration: manually done

- Windows: control-panel->network-> configuration-> tcp/ip-> properties
- UNIX: /etc/rc.config

2- DHCP: allows a host to dynamically obtain its IP address from network server when it joins network. Because a host gets connected automatically into a network, it is “plug-and-play” protocol.

Motivating Example1

Consider a student who carries a laptop from a dormitory room to a library to a classroom. In each location, he connects to a new subnet and hence will need a new IP address.

Motivating Example2

Consider a residential ISP that has 2,000 customers, but no more than 400 customers are ever online at the same time. The ISP does not need 2,048 IP addresses. Only get a block of 512 IP addresses. As the hosts join and leave, the DHCP server allocates an arbitrary address from its current pool of available addresses. Each time a host leaves, its address is returned to the pool.

DHCP is ideally suited to these situations, as there are many users coming and going, and addresses are needed for only a limited amount of time. DHCP is a client-server protocol. A client is typically a newly arriving host wanting to obtain network configuration information, including an IP address for itself.




Figure 44.1: Dynamic Host Configuration Protocol (DHCP)

Working of DHCP:

For a newly arriving host, DHCP protocol is a four-step process.

- 1- DHCP server discovery:** Host generates a UDP packet with destination port 67. Broadcast to all nodes attached to subnet.
- 2- DHCP server offer:** Server generates a message with proposed IP address for the client, the network mask, and an IP address lease time.
- 3- DHCP Request:** Client echoes back the configuration parameters.
- 4- DHCP ACK:** Server confirms the requested parameters.




Figure 44.2: Working of DHCP

Figures and Material used for this topic have been adapted from Kurose and Ross's book with the title *"Computer Networking: A Top-down Approach"*, 6th Edition, 2013.

Topic 45: Internet Control Message Protocol (ICMP)

In this topic, we explain the usage of ICMP.

There are three Major Components of Network Layer namely, IP protocol, routing component, and a facility to report errors in IP datagrams called ICMP.

ICMP is used by hosts and routers to communicate network-layer info to each other. Error messages: unreachable host, network, port, etc. Informational messages: echo request/reply. ICMP is often considered part of IP, but architecturally lies above IP. ICMP messages are carried in IP datagrams as IP payload. An **ICMP message** consists

of type, code and first 8 bytes of IP datagram that caused the ICMP message to be generated in the first place (sender can determine the datagram that caused the error).

Type	Code	description
0	0	echo reply (ping)
3	0	dest. network unreachable
3	1	dest. host unreachable
3	2	dest. protocol unreachable
3	3	dest. port unreachable
3	6	dest. network unknown
3	7	dest. host unknown
4	0	source quench (congestion control)
8	0	echo request (ping)
9	0	route advertisement
10	0	router discovery
11	0	TTL expired
12	0	bad IP header

Many tools obtain information about the Internet by sending probes and waiting for ICMP responses. For example, ping , traceroute etc.

Ping and ICMP

The ping program sends a datagram to a specified destination to test if it can be reached. It then reports the results of the probe by declaring whether the destination responds. The ping uses ICMP echo messages. When a user invokes it by giving the IP address or hostname, ping sends datagram that contains an ICMP type 8 code 0 message (echo request) to the specified host. The destination host, seeing an echo request, sends back a type 0 code 0 ICMP echo reply message.

Traceroute & ICMP

It allows us to trace a route from a host to any host in the world. It employs TTL expired and a destination port unreachable ICMP messages. The source sends a series of IP datagrams to the destination. Each carries a UDP segment with an unlikely UDP port number.




Figure45.1 Traceroute & ICMP

The **Stopping Criterion** creates an UDP segment eventually arrives at destination host. Destination returns ICMP “destination port unreachable” ICMP (type 3, code 3) message. When source gets this ICMP, it stops.

Figures and Material used for this topic have been adapted from Kurose and Ross’s book with the title “*Computer Networking: A Top-down Approach*”, 6th Edition, 2013.

Topic 46: End-to-End: Processing IP Packet

In this topic, we describe fragmentation i.e. processing of IP packet along the end-to-end route.

At the link-layer, an IP datagram is encapsulated within a link-layer frame. All link layer protocols cannot carry network-layer packets of same size e.g. Ethernet frames can carry up to 1500 bytes of data. The maximum amount of data a link-layer frame can carry is called Maximum Transmission Unit (MTU). The MTU of the link-layer protocol places a hard limit on the length of an IP datagram.

Problem

Each of the links along route between sender and destination use different link-layer protocols with different MTUs. When a router is about to forward an IP packet, it may find that its length is greater than MTU.

Solution

Fragment data in IP datagram into two or more smaller IP datagrams, and encapsulate each of these smaller IP datagrams in a separate link-layer frame.

IP Reassembly & Fragmentation

Each of these smaller datagrams is referred to as a fragment. Fragments need to be reassembled before they reach the transport layer at the destination. Both TCP and UDP are expecting to receive complete, un-fragmented segments from the network layer. Datagram reassembly is done in the end systems. Keep the network core simple.




Figure 46.1: IP Reassembly & Fragmentation

There are three fields in IP header. **Identification:** allows a destination host to determine which datagram a newly arrived fragment belongs to. A sender increments identification number for each datagram. **MF:** stands for more fragments and is set for all fragments except the last one. **Fragmentation offset:** tells where in the current datagram this fragment belongs.

Example:

Assume a datagram of 4000 (20 + 3980) bytes arrives at a router and must be forwarded to a link with MTU of 1500 bytes
1480 bytes in data field
offset = 1480/8

length	ID	MF	offset
=4000	=x	=0	=0

One large datagram becomes several smaller datagrams

length	ID	MF	offset
=1500	=x	=1	=0

length	ID	MF	offset
=1500	=x	=1	=185

length	ID	MF	offset
=1040	=x	=0	=370

$3980 - 1480 - 1480 = 1020$
bytes of data

Figures and Material used for this topic have been adapted from Kurose and Ross's book with the title "*Computer Networking: A Top-down Approach*", 6th Edition, 2013.

Topic 47: A Link State Routing Algorithm

In this topic, we explain the working of a link state routing algorithm.

When an IP packet arrives to a router, the router indexes a forwarding table and determines the link interface to which the packet is to be directed. Routing algorithms, operating in network routers, exchange and compute the information

that is used to configure these forwarding tables. The job of routing is to determine good paths from senders to receivers, through the network of routers. A “good” path is one that has the least cost, e.g. physical length of the link, link speed, or monetary cost associated with it.





Figure 47.1: A Link State Routing Algorithm

Link State routing algorithm uses global information such as network topology and all link costs are known. Each node broadcasts identities and costs of its attached links to all other nodes.

Pseudo-code:

```

1  Initialization:
2   $N' = \{u\}$ 
3  for all nodes v
4  if v adjacent to u
5  then  $D(v) = c(u,v)$ 
6  else  $D(v) = \infty$ 
7  Loop
8  find w not in  $N'$  such that  $D(w)$  is a minimum
9  add w to  $N'$ 
10 update  $D(v)$  for all v adjacent to w and not in  $N'$  :
11    $D(v) = \min( D(v), D(w) + c(w,v) )$ 
12 until all nodes in  $N'$  or  $N' = N$ 


```

Notation


$c(x,y)$: link cost from node x to y ; $= \infty$ if not direct neighbors. $D(v)$: current value of cost of path from source to destination v . $p(v)$: predecessor node along path from source to v . N' : subset of nodes

An Example:

Step	N'	$D(v), p(v)$	$D(w), p(w)$	$D(x), p(x)$	$D(y), p(y)$	$D(z), p(z)$
0	u	2, u	5, u	1, u	∞	∞
1	ux	2, u	4, x	2, x	∞	∞
2	uxy	2, u	3, y	4, y	4, y	∞
3	uxyv	2, u	3, y	4, y	4, y	∞
4	uxyvw	2, u	3, y	4, y	4, y	4, y
5	uxyvwz	2, u	3, y	4, y	4, y	4, y



Resulting shortest-path tree from u:



Resulting forwarding table in u:

destination	link
v	(u,v)
x	(u,x)
y	(u,x)
w	(u,x)
z	(u,x)

Figures and Material used for this topic have been adapted from Kurose and Ross's book with the title "*Computer Networking: A Top-down Approach*", 6th Edition, 2013.


Topic 48: Distance Vector Routing Algorithm

This topic explains the working of a distance vector routing algorithm.

Bellman-Ford Equation

If $dx(y) :=$ cost of least-cost path from x to y , then $dx(y) = \min \{c(x,v) + dv(y)\}$, where min is taken over all neighbors v of x .

Bellman-Ford Example



Clearly, $d_v(z) = 5$, $d_x(z) = 3$, $d_w(z) = 3$

B-F equation says:

$$\begin{aligned} d_u(z) &= \min \{ c(u,v) + d_v(z), c(u,x) + d_x(z), \\ &\quad c(u,w) + d_w(z) \} \\ &= \min \{ 2 + 5, 1 + 3, 5 + 3 \} = 4 \end{aligned}$$


Node that achieves minimum is next hop in shortest path. Thus, the solution to the Bellman-Ford eq. provides the entries in a node's forwarding table. Each node x maintains the following routing information:


- 1-For each neighbor v , the cost $c(x,v)$
- 2-Node x 's distance vector $D_x = [D_x(y): y \in N]$, where $D_x(y)$ = an estimate of cost of least-cost path from itself to y .
- 3- The distance vectors of each of its neighbors. i.e., $D_v = [D_v(y): y \in N]$ for each neighbor v of x

Basic idea:


From time-to-time, each node sends a copy of its distance vector to each of its neighbors. When a node x receives a new distance vector from any of its neighbors v , it saves v 's distance vector, and then uses the Bellman-Ford equation to update its own distance vector as follows. If x 's distance vector has changed as a result of this update step, x will then send its updated distance vector to each of its neighbors, which can in turn update their own distance vectors. The estimate $D_x(y)$ converges to the actual least cost $dx(y)$.

An Example:





The DV algorithm continues until no update messages are sent.



Topic 49: Working of Default Gateway Router

In this topic, we explain the working of a default gateway router.

Let's assume that Hannah wants to access the webpage www.example.com residing on a webserver. The internetwork consists of three different IP networks, each with one IP network number.




Figure 49.1: Working of Default Gateway Router

After IP software of Hannah's PC has built the IP packet that needs to be sent to the server, it needs to know where to send the packet first. To send the packet to R1, Hannah needs to know R1's IP address of the router interface that's connected to the same Ethernet as Hannah. R1 is termed as Hannah's default router/gateway. A PC's default router is simply the router to which that PC sends packets when the destination is in another network or subnet.




Figure 49.2: Software of Hannah's PC has built the IP

Routers typically have one IP address per physical interface. Routers have lots of interfaces of many different types, which are labeled with a name and a number, such as Ethernet1.

Topic 50: Address Resolution Protocol

In this topic, we explain how IP packets get encapsulated in Ethernet frames and how ARP works. Let's assume that Hannah's PC is connected to an internetwork that uses three different IP networks, each with one IP network number. She accesses www.example.com from a webserver.




Figure 50.1: Address Resolution Protocol

Two issues need attention. 1 - Hannah can't send an IP packet over an Ethernet, but she can send an Ethernet frame over an Ethernet. 2- The destination address field of IP header holds IP address of the destination web server. There is no place for "default gateway IP address" in the header. Hannah must have some other means to ensure that R1 receives the packet. **Solution of issue 1:** Encapsulation. Hannah encapsulates the IP packet in an Ethernet frame for transmission over the LAN. Ethernet places IP packet between a header and trailer. The Ethernet header contains source and destination Ethernet address fields. The trailer contains a frame check sequence field (FCS) to determine whether errors occurred during transmission.




Figure 50.2: Solution of issue 1




Figure 50.3: Solution of issue 1

Solution of issue 2: To deliver the packet to R1, Hannah includes R1's Ethernet address as the destination in the frame.




Figure 50.4: Solution of issue 2

Hannah learns the IP address of default router manually or via Dynamic Host Configuration Protocol (DHCP). To learn MAC address of R1, Hannah uses address resolution protocol (ARP).




Figure 50.5 Dynamic Host Configuration Protocol (DHCP)

Hannah generates an ARP broadcast "Hey, if this is your IP address, tell me your Ethernet MAC address." Everyone on the LAN gets ARP broadcast as switch forwards LAN broadcasts to all devices in network.




Figure 50.6: Ethernet MAC Address

Topic 51: Router's Routing Logic and Table

This topic describes the router's routing logic and contents of IP routing table.

Let's assume that Hannah's PC is connected to an internetwork that uses three different IP networks, each with one IP network number. She accesses www.example.com from a webserver.






Figure 51.1: Router's Routing Logic and Table

As R1 receives the packet, it needs to forward it to R2 next. This process can be broken down into 3 steps: 1. De-capsulation. 2. Decide where to forward the packet next. 3. Encapsulation.



1. De-capsulation: Check the incoming frame's FCS. If there are errors, discard it. Otherwise, remove header and trailer, leaving the original IP packet whose source and destination addresses are 8.1.1.1 and 130.4.3.3

2. Decide where to forward the packet next: A router makes a decision about where to forward the packet next by looking at the destination IP address of the IP packet. For routing to work well, the router needs to know how to reach the various IP networks and subnets in the internetwork. This info is contained in a routing table. RIP and OSPF are examples of routing protocols.



3. Encapsulation: At this point, R1 has an IP packet sitting in memory. To send the IP packet out on E2, R1 needs to encapsulate it in a frame. In the newly created Ethernet frame: source address is MAC address of my E2 interface. Destination address is the MAC address of R2. If not known, use ARP to find R2's MAC address – put it in a cache for future use.





Figure 51.2: Encapsulation

Logic of Final Router: R2 has the same logic as that of R1 except R2's routing table differs from R1's: R2 does not need to send the packet to another router, but it should instead send the packet directly to the web server.



From R2 to web server: Create a new Ethernet frame. Own Ethernet address on the outgoing interface as source MAC. Use ARP to learn web server's MAC and use it as

destination. Finally, the web server has the frame and extracts the packet. The routing of the packet is complete!

Topic 52: Routing with Subnets

This topic explains how a router's routing logic works with subnets.

Let's assume an internetwork of a single Class B IP network 150.1.0.0. This IP network is subdivided into three subnets namely, 150.1.1.0, 150.1.2.0, and 150.1.3.0. All hosts in a given IP subnet have the same value in the first 3 octets of their IP addresses. Next, we look at the routing tables.





Figure 52.1: Routing with Subnets

When Hannah sends a packet to the web server, destination address is 150.1.3.3. Hannah's logic works like it did without subnetting. Hannah's default gateway is R1. Hannah sends an ARP broadcast for IP address 150.1.1.250. R1 replies and Hannah sends a frame to R1 with the IP packet inside the Ethernet frame. R1 receives the frame, and if it is error free, it extracts the IP packet. R1 uses destination address and finds an entry for it in its IP routing table. R1 forwards packet out on its E2 interface to R2 next. Finally, when R2 receives the frame, it does the usual error check and extracts the IP packet. R2 finds an entry for destination IP address. R2 can forward the packet directly to the web server. Whether or not an internetwork uses subnetting, the basic routing logic remains the same. If Hannah sends a packet to someone who's on the same subnet, she can send it directly to the destination without involving default router.



Figures and Material used for this topic have been adapted from Kurose and Ross's book with the title *"Computer Networking: A Top-down Approach"*, 6th Edition, 2013.

Topic 53: Router Hardware Architecture

In this topic, we explain the router hardware components.

There are two key router functions: run routing algorithms (RIP, OSPF) and, forwarding datagrams from incoming links to appropriate outgoing links. A Generic router architecture is shown next.




Figure 53.1: Router Hardware Architecture

Input Ports

Input ports perform physical & data link layer functions of an incoming physical link at a router. They determine output port to which an arriving packet will be forwarded via the switching fabric.




Figure 53.2: Input Ports

Switching fabrics

Switching fabrics transfer packet from input buffer to appropriate output buffer. Rate at which packets can be transferred from inputs to outputs is called switching rate. The three types are shown next.





Figure 53.3: Switching fabrics

Output port



- buffering required when datagrams arrive from fabric faster than the transmission rate
- scheduling discipline chooses among queued datagrams for transmission

Figure 53.4: Output port

Where does queuing occur?

Packet queues can form at both input and output ports. The location and extent of queuing will depend on the traffic load, speed of switching fabric and line speed. As these queues grow large, router's buffer space will eventually be exhausted and packet loss will occur when no memory is available to store arriving packets. Suppose switch speed is N times faster than speed of line. Even then output buffering will still occur when multiple inputs send to same output. This scenario is shown next.




Figure 53.5: Packet queues

The Head-of-the-Line (HOL) blocking at Input Port is shown next.




Figure 53.6: The Head-of-the-Line (HOL) blocking at Input Port

Topic 54: Routing to Nearby Places

In this topic, we explain how a router learns routing table for networks/subnets that are directly connected to it.

At a given time, a router knows which of its physical interfaces are up and working. It knows the IP addresses used on each interface. Also, it knows about IP networks or subnets that are connected to its interfaces. The router can add a route to the subnet in its routing table.

Configuring a Router When you buy a brand new router, a network engineer connects to the router and tells the router which IP addresses to use. For each interface, use an IP Address and Subnet mask to figure out the range of valid IP addresses in each subnet.




Figure 54.1 Configuring a Router

After R1's interfaces are up, it knows about subnets: the subnet numbers, the outgoing interface it should use to forward packets to them. There is no need to send packets to another router for them. In our previous example, R1's Ethernet1 interface is connected directly to the same subnet as Hannah. As a result, R1 simply adds a route for each directly connected subnet to its routing table.




Figure 54.2: Configuring a Router

Routers always add routes for directly connected subnets and networks, as long as the interface is both configured and working. Also, they should include other routes in their routing tables. In previous example, R1 didn't have a route to subnet 150.1.3.0, so it could not forward a packet that was destined for IP address 150.1.3.3.

Static IP route: The network engineer can address this problem by configuring a static route on R1 for subnet 150.1.3.0, with outgoing interface Ethernet2, and next-hop router of 150.1.2.2.




Figure 54.3: Static IP route

Network topologies tend to change a lot. Static routes make it difficult to use all the possible routes to the same part of the network when you have multiple possible physical paths.

Topic 55: Dynamically Learning Routing Tables

This topic explains explain how routing table entries are learnt dynamically.

The most typical way a router learns all the rest of the routes in an internetwork, beyond just its directly connected routes, is by using a routing protocol. **Routing Protocols** define messages by which routers can exchange route information with other routers. If all the routers participate, all routers should have routes for all

subnets or networks. Let's assume an internetwork that consists of two routers and three subnets. Several steps that occur over time are:

- 1- Each router knows only its directly connected routes.
- 2- A router generates a routing protocol message (routing update) that contains information about IP networks and subnets and send it to another router.
- 3- The receiving router: a) puts that interface into the route in which it received the routing update. b) Also puts the IP address of the router that sent the routing update as the next-hop router.




Figure 55.1: Dynamically Learning Routing Tables

Next, we show R2 advertising its routes, with R1 learning a route to subnet 150.1.3.0.




Figure 55.2: Dynamically Learning Routing Tables

Topic 56: How to Pick the Best Route

In this topic, we describe how the best route gets selected.

Routing protocols not only help routers learn routes, but they also help routers learn the best routes to a destination when there is more than one way to get there. Let's assume an internetwork consisting of three routers. R1 can send packets to subnet 150.1.3.0 through R2 or through R3.




Figure 56.1: How the best route gets selected

Routing protocols define a number called a metric that's associated with each route in a routing update. That number represents how good or how bad that route is. When a router receives multiple routing updates, it might learn of multiple ways to reach a subnet. By looking at the metrics, the router can pick the best route.




Figure 56.2: The router can pick the best route

Imagine that someone turns off the power on the Ethernet switch between R1 and R2. That route would then fail and be unavailable. The routing protocol on R1 would remove the route from the routing table. At the same time, R1's routing protocol

would notice that another route exists the one through R3 and add that route to the routing table.

Figures and Material used for this topic have been adapted from Kurose and Ross's book with the title "*Computer Networking: A Top-down Approach*", 6th Edition, 2013.

Topic 57: Interior & Exterior Routing Protocols

In this topic, we explain the need for interior and exterior routing protocols.


Hierarchical Routing

Considering network to be flat, taking all routers to be identical and running one single routing protocol is not realistic. At least two reasons:

1-Scale: with hundreds of millions destinations can't store all destinations in routing tables. Similarly, routing table exchange would swamp links.

2-Administrative Autonomy: internet is a network of networks. Each network admin may want to control routing in its own network


Autonomous System (AS): Organize routers into regions. An AS is under the same administrative control. For example, operated by the same ISP or belonging to same company network. All routers within the same AS run the same routing protocol (intra-AS or interior routing protocol) and have information about each other. Routers in different AS can run different intra-AS routing protocols. To connect ASes to each other, one or more of the routers in an AS will have to be responsible for forwarding packets to destinations outside the AS – Gateway routers.



If a router in AS1 receives an IP packet destined outside of AS1, then it uses an inter-AS (exterior) routing algorithm 1- to learn which destinations are reachable through

AS2, which through AS3, and 2- propagate this information to all routers in AS1. This is shown next.

Forwarding table configured by both intra- and inter-AS routing algorithm



Well-known IP routing protocols are listed next.

IP Routing Protocols

Routing Protocol	Public or Proprietary?	Interior or Exterior?
Routing Information Protocol (RIP)	Public	Interior
Interior Gateway Routing Protocol (IGRP)	Proprietary	Interior
Open Shortest Path First (OSPF)	Public	Interior
Enhanced IGRP (EIGRP)	Proprietary	Interior
Border Gateway Protocol (BGP)	Public	Exterior

Topic 58: Introduction to Domain Name System

In this topic, we explain why a domain name system (DNS) is required.

People have many identifiers. For example, name, passport#, NTN#, NIC#. One identifier cannot be used everywhere. Can Hassan Hamid be used to recognize a tax payer? Similarly, Internet hosts have many identifiers. A "hostname", e.g., www.yahoo.com is human friendly but difficult for routers to process variable length alphanumeric characters. IP address (32 bit) is router friendly. If you need to call someone and you know the person's name but not the phone number, you can just look up the information in the phone book. It's simple and easy provided the info is present in the book. TCP/IP hosts can have the equivalent of a phone book in a file. The local host file contains a list of TCP/IP host computer names and their corresponding IP addresses. Assume Hannah types http://www.example.com in a

browser. Hannah's PC can send an IP packet to the web server, if its IP address is known. Look local host file (www.example.com, 150.1.3.3).




Figure 58.1: Domain name system (DNS)

A local host file cannot list the names and IP addresses of every server on the planet. If the pair changes, conveying this updated info to everyone is cumbersome. TCP/IP uses a domain name system (DNS) instead of local host file. DNS defines how to discover which names correspond to which IP address. DNS also defines the structure and format of TCP/IP host and names.

Figures and Material used for this topic have been adapted from Kurose and Ross's book with the title "*Computer Networking: A Top-down Approach*", 6th Edition, 2013.

Topic 59: Domain Name System (DNS) Servers

In this topic, we explain the need for different domain name system (DNS) servers.


DNS is an application-layer protocol that allows hosts to query the database and to resolve names (name, address). It is employed by other application-layer protocols such as HTTP, SMTP, and FTP.

Problems with centralized DNS

Single point of failure: If the DNS server crashes, so does the entire Internet. Traffic volume: a single server would have to handle all DNS queries. Distant centralized database: a single server cannot be "close to" all the querying clients –significant delays. Maintenance: a centralized database would have to be updated frequently for every new host. In order to deal with the issue of scale, the DNS uses a large

number of servers, organized in a hierarchical fashion and distributed around the world.

Distributed, Hierarchical Database




Interaction of Servers: an approximation

Let's assume a client wants IP address for `www.amazon.com`. 1- The client queries a root server to find com DNS server. 2- The client queries com DNS server (TLD) to get amazon.com DNS server. 3-client queries amazon.com DNS server (Authoritative) to get IP address for `www.amazon.com`.

Root Name Servers: in the Internet there are 13 root servers. **Top-level Domain (TLD) Servers** are responsible for com, org, net, edu, gov, and all top-level country domains, e.g.: uk, fr, ca, jp. **Authoritative Servers:** organizations with publicly accessible hosts (e.g. web servers) must provide publicly accessible DNS records that map the names of those hosts to IP addresses. **Local DNS Servers** do not strictly belong to hierarchy. Each ISP has one server which is also called "default name server". When host makes DNS query, query is sent to its local DNS server.

DNS name resolution example

- ✧ host at `cis.poly.edu` wants IP address for `gaia.cs.umass.edu`



DNS Caching: once (any) name server learns mapping, it caches mapping.

Topic 60: Working of Domain Name System

In this topic, we explain working of domain name system (DNS) inside and outside company.

DNS defines how to figure out names and IP addresses for the entire Internet, as well as inside a single site at a single company.

Inside the Company: Let's assume that Hannah's PC is inside the example.com corporation's enterprise network. She wants to view <http://www.example.com>.




Figure 60.1: Working of Domain Name System

For DNS to work well inside a company, someone must be responsible for supporting it. That work includes updating and changing the list of names and IP addresses.

Outside the Company: Let's assume that Hannah wants to access the <http://www.fredesco.com> website, which is located in Fredesco enterprise network.




Figure 60.2: Fredesco enterprise network

The DNS servers must work together as shown next.




Figure 60.3: The DNS servers

Figures and Material used for this topic have been adapted from Kurose and Ross's book with the title "*Computer Networking: A Top-down Approach*", 6th Edition, 2013.

Topic 61: DNS Resource Records

This topic describes the use of domain name system (DNS) Resource Records.

The domain name service (DNS) is a directory service, which translates hostnames into IP addresses. It is a distributed database implemented in a hierarchy of many DNS servers. The DNS distributed database is stored on servers in the form of resource records (RRs):

DNS Resource Record Format:

RR format: (name, value, type, TTL)

Time to live (TTL) – after this time a resource record becomes invalid

The meaning of **name** and **value** depend on **type**

The type field is defined as:

Type=A ▪ name is hostname ▪ value is IP address ▪ (relay1.bar.foo.com, 145.37.93.126, A)	Type=CNAME ▪ value is canonical name for the alias hostname Name. ▪ querying hosts learn the canonical name ▪ (foo.com, relay1.bar.foo.com, CNAME)
Type=NS ▪ name is domain ▪ value is hostname of authoritative name server that knows how to obtain IP addresses for hosts in this domain ▪ (foo.com, dns.foo.com, NS)	Type=MX ▪ value is the canonical name of a mail server associated with an alias name ▪ (foo.com, mail.bar.foo.com, MX)

Inserting resource records into DNS

Let's assume that Alice just created a new startup company "Network Utopia". She wants to register domain name networkuptopia.com.

1- She needs to contact a registrar (A registrar accredited by Internet Corporation for Assigned Names and Numbers (ICANN) verifies the uniqueness of the domain name, enters it into the database and collects a small fee for its services.)

2- She needs to provide registrar with names and IP addresses of her primary and secondary authoritative name servers.

3-Registrar inserts two RRs (NS and A) into the TLD com server:

- (networkuptopia.com, dns1.networkuptopia.com, NS)
- (dns1.networkuptopia.com, 212.212.212.1, A)

4 -She also has to make sure that Type A record for web server www.networkuptopia.com and Type MX record for her mail server mail.networkuptopia.com are entered into her authoritative DNS servers.

Now suppose Bob views www.networkuptopia.com. His host will first send a DNS query to his local DNS server, which will then contact a TLD com server. This TLD com server contains the NS and A RRs listed above and sends a reply back to Bob's local DNS server. This local DNS server then sends a query to 212.212.212.1, asking the type A record. This record provides the IP address, say 212.212.71.4, to Bob's host and now browser can initiate a TCP connection.

Topic 62: Introduction to Wide Area Network

In this topic, we describe what a wide area network (WAN) is.

A wide-area network (WAN) defines a type of a network, or part of a network, in which the devices are relatively far apart. A WAN is a network, or part of a network, for which the cabling must pass outside the property of one company. The distance might only be a few miles, or it might be thousands of miles! We next show a WAN.





Figure 62.1: Wide Area Network

You can run an Ethernet cross-over cable between two devices and the two devices can communicate with Ethernet. That works well, but can we use it for longer distances? If two routers in two cities need to forward packets to each other, they need some sort of physical medium over which to send the packets.



When R1 sends out an electrical signal over the cable, R2 needs to receive that same signal on the wires that it expects to receive data. Likewise, R1 needs to receive what R2 transmits. Because the physical interfaces on the routers use the same pins to transmit, the cable connects the twisted pair used for transmitting by R1 over to R2's receive pins, and vice versa. In short, all that the two routers need is a cable between them, with transmit and receive pairs of wires, so that they can send and receive anytime they want.

In the previous example, there are two problems stop you from using Ethernet:

- 1-You are not legally allowed to run a cable between Atlanta and Cincinnati.

- 2-Those who can run the cable namely, the telephone companies (telcos) do not lease or sell 450-mile Ethernet cross-over cables.

Topic 63: Different Aspects of WAN Link

In this topic, we describe different aspects of a WAN Link.

To connect two routers located let's say at Atlanta and Cincinnati (450 miles), a telephone company (telco) should offer:

- 1- A service between two routers that acts like a cable with four wires (two pairs) in it.
- 2- When a router sends on one pair, that pair is crossed to the other pair before it gets to the other end cable. That way, when one router sends on one pair, the other router receives on the other pair.
- 3- Your routers use a serial interface because the service requires to physically encode data a little differently than on an Ethernet.

Result is that two routers can send and receive data to and from each other. Essentially, a telco can lease you a 4-wire cable, or 4-wire circuit, between two points. Telcos have already run cables between almost every town and city. Telcos have offices, called central offices (COs), almost everywhere. In addition, telcos have the right-of-way, which is the legal right to dig up roads and put cabling in the ground.





Figure 63.1: Aspects of WAN Link

Switches in the CO are called telco switches, phone switches, or WAN switches. Telco needs to ensure path between Atlanta and Cincinnati:

- 1- A cable from office building in Cincinnati where other router sits and telco CO in Cincinnati.

2- Electrical signals between Atlanta CO and Cincinnati CO.



Topic 64: A Cross-over Cable versus Leased Circuit

In this topic, we describe differences between a cross-over cable and a leased circuit.


If you install an Ethernet cross-over cable between two routers by plugging the cable into Ethernet interface on the routers, the two routers can forward Ethernet frames to each other. Ethernet can support transmission speeds of 10 Mbps, 100 Mbps, 1000 Mbps (1 gigabit/second), and even 10 Gbps. If you install a WAN link between two routers, a telco creates something similar to a cross-over cable with some differences. With WAN links, many different transmission speeds are supported. Typically, WAN links use speeds that are multiples of 64 Kbps. For example, a telco offers a service called T1 line or T1 circuit whose speed is 1.5 Mbps, which is 24 times 64 Kbps. A telco can also offer links in multiples of 1.544 Mbps. You can also order even higher speeds that generally come in multiples of 51.84 Mbps.

Key points about speed of WAN links

Many different speeds are possible. You specify the speed when you order the leased line. You need to configure the routers to use the right speed.

How to control the WAN link speed

Serial links run at different speeds, and must be preconfigured. The channel service unit/data service unit (CSU/DSU) of a serial interface controls the speed. The CSU/DSU functions can be done with an external device or as a function of the serial interface card on the router. If you use a CSU/DSU that sits outside the router, you must configure the speed on that external device. If CSU/DSU is built into the router, you must configure the speed of the WAN link on the router.



A WAN Link Installation Plan

1. Contract with the phone company to provide a leased line at a certain speed.
2. Install a router at each site, near where the telco will run its cable.
3. Install an external CSU/DSU near the routers at each site, if you didn't buy routers with internal CSU/DSUs.
4. Configure the CSU/DSUs with the correct speed.
5. After the phone company runs the cables, install the cables into the CSU/DSU (external) or serial interface of the router (internal CSU/DSU) at each site.

Topic 65: Routers and WANS

This topic describes how WANs and Routers work together.

Routers use WAN links when they need to forward IP packets to a subnet at a remote site. End user devices at a company's office site use an Ethernet NIC and never directly connect to a WAN link. To send traffic to an IP host at another site, end user devices send the data to router. The router forwards the IP packet to another router at the other site, which then forwards the packet to the other IP host.




Figure 65.1: WANs and Routers work together

WAN data link framing

Let's assume that Hannah sends a packet to the `www.example.com` web server at IP address 150.1.3.3. The IP packet arrives to her default gateway namely R1. As R1 receives the Ethernet frame, it checks to see whether errors occurred. If no errors occurred, R1 extracts the IP destination address 150.1.3.3 and find a match in his routing table.




Figure 65.2: WAN data link framing

R1 needs to encapsulate the IP packet inside a data link layer frame. Two popular data link layer protocols for point-to-point WAN links are high-level data link control (HDLC) and Point-to-Point Protocol (PPP).




Figure 65.3: High-level data link control (HDLC) and Point-to-Point Protocol (PPP)

R1 needs to encapsulate the IP packet inside a data link layer frame. Two popular data link layer protocols for point-to-point WAN links are high-level data link control (HDLC) and Point-to-Point Protocol (PPP).

WAN data link addressing

Inside the HDLC and PPP header, there's only one address field, and it is 1 byte long. When R1 sends anything on this point-to-point WAN link, the only device that could possibly get the data is R2. So, although the address field exists, it doesn't really matter.

Differences bet. HDLC and PPP: HDLC was defined by the ITU in the 1970s. PPP was defined in RFC 1661 during the 1990s. PPP has more advanced features. Cisco uses a nonstandard version of HDLC. Cisco also conforms to the standard for PPP.

Topic 66: Frame Relay

In this topic, we explain the basics of Frame Relay.

Frame Relay is a wide area network (WAN) technology that uses one physical WAN link connected to each site, while allowing each site to send data to each other site. When you want to build a network to connect multiple remote sites, you can use a lot of serial links. Frame Relay requires less work, less new hardware and is more cost-effective solution compared to leased lines. A Frame Relay network acts like a big WAN switch, with routers connecting to it. To send data to another router, the sending router just needs to send a frame with the right address in it.

To make Frame Relay work, each router needs a physical cable between itself and a device called a Frame Relay switch. A Frame Relay switch is the equipment that understands Frame Relay and can forward traffic based on Frame Relay protocols. The telco uses Frame Relay switches in its local central offices (COs) that together switches the data to the correct sites. When a router physically connects to a Frame Relay switch, it is connecting to a Frame Relay service. The company that sells Frame Relay services is called a Frame Relay service provider. Often, a Frame Relay service

provider is also a telco. In most cases, that company works with other telcos to create the Frame Relay network.




Figure 66.1: Physical Parts of Frame Relay

Although the cabling and channel service unit/data service units (CSU/DSU) are the same as with a leased line, the telco does something different in the CO: It connects the cable to a Frame Relay switch. The provider's collective set of Frame Relay switches, along with the other equipment between them, form that provider's Frame Relay network. Frame Relay is a set of protocol specifications, all matching the functions of OSI Layer 2; data-link layer.

Topic 67: Frame Relay Switching

This topic explains the working of Frame Relay switching to multiple remote sites.

When a telco sells you a Frame Relay service, it forwards Frame Relay frames sent by one router to another router. Before a router can send a packet, it must add the correct data link header and trailer to the packet. Each Frame Relay header contains a single address field called a data-link connection identifier (DLCI). It is a 10-bit number, usually written as a decimal number between 0 and 1023. Each Frame Relay switch forwards the frame, based on the DLCI, through the network, until it gets to the router on the other side. Frame Relay switches must be configured to know where to forward frames with particular DLCIs in their headers. They cannot learn addresses automatically.




Figure 67.1: Cabling with Internal and External CSU/DSUs

To see the real advantage of Frame Relay over leased lines, let's look at an example with three Frame Relay sites. We further assume router R1 (Atlanta) wants to send data to both R2 (Cincinnati) and R3 (Boston). With Frame Relay, a router can use its single physical access link to forward traffic to multiple remote routers. This is shown next.




Figure 67.2: Three Leased Lines to Connect Three Routers

Topic 68: Virtual Circuits

In this topic, we explain the concept of virtual circuits.

Let's assume that there are three sites using Frame Relay network and frames are going from one site (R1) to the other two sites (R2 and R3).




Figure 68.1: Frame Relay Switching to Multiple Remote Sites

In Frame Relay lingo, the ability for R1 to send data to R2 over the Frame Relay network is called a virtual circuit (VC). A VC is like a leased circuit, in that exactly two devices can send and receive data using it. It is called virtual to contrast it with a physical leased circuit. As a VC is typically predefined to always be there, VCs are often called permanent virtual circuits (PVCs). It is shown next.




Figure 68.2: Frame Relay PVC Concepts

The service provider preconfigures all the required details of a PVC. Configuring the switches to forward frames with DLCI 102 to Cincinnati and frames with DLCI 103 to Boston. In the previous network, there is no PVC between R2 and R3. The choice of which sites need to have a PVC depends on where the network engineers think that traffic needs to flow in the network. When routers use Frame Relay, and there is a PVC between each pair of routers, the PVCs are in a full mesh. When not all routers have a PVC, it's called a partial mesh. A full mesh is shown next.