

Introduction to TCP/IP

INFO 341

1

Objectives

- What was the original name of the Internet?
- Who built the Internet?
- What are the layers of the Internet Reference Model?
- What are the main Internet address classes?
- What is subnetting?
- What is supernetting?
- BOOTP & DHCP are two examples of what?
- What does traceroute do?
- What is the role of DNS in the network?
- What are some routing protocols?
- What is routed, RIP, OSPF?

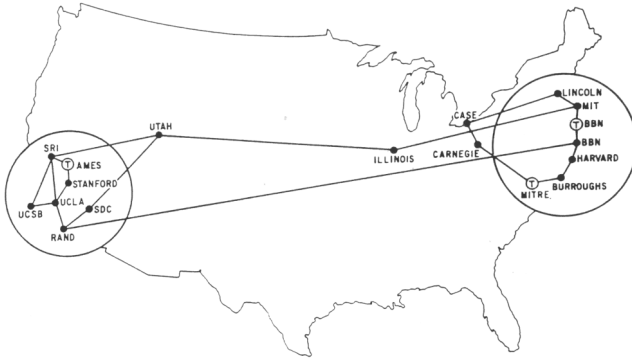
2

Internet & TCP/IP

- Transmission Control Protocol/Internet Protocol
 - Development began in the early 1970's
 - Largely funded by the defense department through the Advanced Research and Project Agency (ARPA) as ARPANet
 - Motivation was to interconnect different computers located on different physical networks located over a large geographical area (like the entire US)
 - Researchers working at major universities on defense research projects need to collaborate and be able to share access to mainframes, share files, etc.

3

The Original ARPANet



MAP 4 September 1971

4

Additional Design Considerations

- Wanted the network to be able to scale well, to potentially connect a very large number of nodes and different network types
- Wanted the network to be able to continue functioning even if a particular segment failed or there was unreliable communications
 - Consider again this is being funded by defense department – a segment could fail not just because equipment failed, but they wanted to consider the possibly of a military action taking out a city or region

5

TCP/IP Protocol Suite

- What is a protocol?
 - An end-to-end agreement about how to communicate
 - TCP/IP is not the only protocol used on LANs. Others include IPX, Netbios, DecNet, AppleTalk, others
 - Multiple protocols can exist at the same time on a LAN although some Network Administrators only allow TCP/IP
- TCP/IP is called a “suite” of protocols because it includes many different protocols at different layers
 - The OSI 7 layer model was actually developed before TCP/IP
 - TCP/IP was designed with a 5 Layer Internet Reference Model

6

IP (Internet Protocol) Addressing

- Each “host” or device must have a unique “IP Address” (NAT firewalls let us relax this)
 - In the current version of IP (IPv4), these addresses are 4 octets long or 32 bits
 - To make them easy for humans to remember they are represented by 4 octets of decimal numbers, separated by a “.”
 - For example: 152.2.81.1 or 128.95.220.25
 - “dotted quad” address

10

Back to Binary Again...

Consider the following possible IP address:

10000000	11010000	1100100	1100111
128	208	100	103

Which is written as: 128.208.100.103

Given this format:

The smallest value of any octet could be 00000000 or 0 decimal

The largest value of any octet could be 11111111 or 255

11

Special addresses...

- 0 and 255 have special meanings in IP addresses (using standard netmasks)
 - 0 and 255 are reserved for special purposes
 - Example: 128.208.100.255 is a broadcast address for the 128.208.100.x network with a netmask of 255.255.255.0
- Devices have addresses from 1-254

12

Address Classes

- Each 32 bit address is actually divided into 2 different fields
 - The NetID portion of the address identifies the network that a host is connected to
 - The HostID portion of the address gives each node on a given network a unique identifier
- When the addressing scheme was devised it was assumed that there would be a few networks with a very large number of hosts, a moderate number of networks with an intermediate number of hosts, and a large number of networks with a small number of hosts
 - Different "address classes" were designated for these scenarios

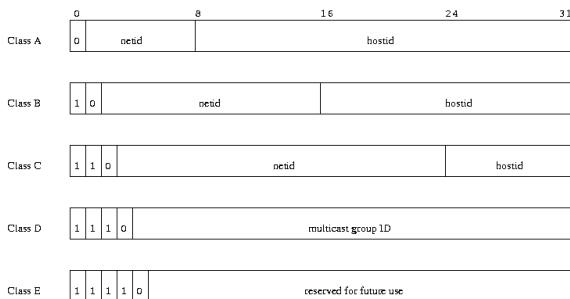
13

Address Classes - A, B, C

- Class A addresses support:
 - 16 million hosts on each of 127 networks
- Class B addresses support:
 - 65,000 hosts on each of 16,000 networks
- Class C addresses support:
 - 254 hosts on each of 2 million networks

14

IP Address Classes



Class A: 127.0.0.1 and below
 Class C: 192.0.0.0 to 223.255.255.255
 Class E: 240.0.0.0 to 247.255.255.255

Class B: 128.0.0.0 to 191.255.255.255
 Class D: 224.0.0.0 to 239.255.255.255
 Remaining 248.0.0.0 to ????

15

Reserved IP Addresses

- 127.0.0.1 – “loopback interface”, the local machine
- 10.x.x.x – class A private IP networks
- 169.254.x.x - self-assigned IP addresses when a DHCP server cannot be reached (a.k.a. APIPA)
- 172.16.0.0 – 172.31.255.255 – class B private networks
- 192.168.x.x – class C private networks

These special addresses can be used for testing – under normal circumstances routers will not pass packets with these addresses

16

How are IP numbers assigned?

- If you want to setup a private TCP/IP network not connected to the global Internet, you can use what you want – best to use one of private network numbers when you do it however
- If you want to connect to the Internet, your organization must be assigned an address space that is unique to use
- Private ISPs (Internet Service Providers) coordinate with IANA (Internet Assigned Numbers Authority) to provide organizations with IP addresses
- If you are in a large organization or university, it is likely that your organization already has an address space assigned and your address will fall within that range
 - You may have to work with your organization’s central networking group to obtain a network address for your department or division

17

UW Example

- The University of Washington has been assigned the following Class B network addresses
 - 128.208.x.x
 - 128.95.x.x
 - 140.141.x.x
 - 140.142.x.x
- The UW can manage those addresses itself and allocate “subnets” in that pool to various departments
 - 128.208.100.x has been allocated to the iSchool
 - We can allocate the 254 possible addresses to hosts in the iSchool

18

Assigning Numbers to Machines

- Once you have a range of numbers you can use, an administrator decides how to allocate addresses on the local network
- This can be done “statically” where one machine is assigned a certain number to use all the time, or
- “Dynamically” where each machine requests an IP address from a pool of shared addresses when they boot up. Each time they start they may get a different IP address
- In Windows and Macintosh the actual assignment of the address is done in the Network Control Panel (unfortunately you must be “Administrator” to see this in Windows 2000/XP)
- In Unix these values are typically entered in configuration files (and of course you have to be “root” to change those values)

19

Static Assignments

- Always used with “servers” or computers that other people need to connect to on a recurring basis
 - Why?
Examples?
- Until relatively recently, most computers connected to the Internet had static assignments and hence the “crisis” several years ago with people worrying about running out of IP numbers for each device
- Routers, switches, printers etc. should all have static IP addresses - they are infrastructure

20

Dynamic Assignments

Two Mechanisms (ignoring RARP):

- BOOTP - Bootstrap Protocol
- DHCP – Dynamic Host Configuration Protocol

21

[BOOTP]

- The “Bootstrap” Protocol
 - On boot-up the machine broadcasts a request for an IP address
 - If a BOOTP server is listening on that network, it will look-up in its database the MAC address of the machine making a request and assign it an IP address
 - Note: this requires a database to be maintained of MAC addresses and IP addresses
 - So while we say BOOTP is a dynamic address assignment protocol, is it really?
 - BOOTP was also used to boot a diskless client entirely from the network from a network image on a bootp server

22

[DHCP]

- DHCP - Dynamic Host Configuration Protocol
 - Similar to BOOTP and actually an extension to the BOOTP protocol
 - On boot-up the machine broadcasts a request for an IP address (FF-FF-FF-FF-FF-FF and 255.255.255.255)
 - If a DHCP server is listening for that network, it will assign an IP address - a process called leasing
 - DHCP can be configured to just give an address at random from a pool of available addresses, or it can be setup to supply the same address all the time to a particular MAC address like BOOTP

23

[Additional DHCP info]

- DHCP uses the concept of a lease that can be renewed or that can timeout after a specified time period
- DHCP is the dominant protocol used today for dynamic assignment of addresses – it “won”. BOOTP is rarely used today
- DHCP does not require a static list of MAC addresses
- Most commercial ISP’s use DHCP for address assignment. You typically have to pay extra for a static address. One day IPv6 will change this?

24

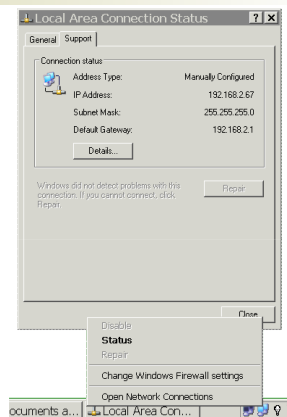
Advantages of DHCP

- DHCP allows many more devices to connect to the Internet without each device requiring it's own IP address.
 - For instance, a commercial service provider like Earthlink, can have a pool of IP addresses for all their users. They only need one address for each modem in the dial-in pool rather than one address for every subscriber to their service
 - This advantage starts to go away however when you consider "always on" users like DSL or cable modem users that don't power down their machines
- DHCP makes it easier to setup your machine if you move from place to place. In a static assignment, your address is very much tied to the physical topology of your network.
 - Again at UW, the iSchool network was assigned the address 128.208.100.x
 - CS at UW has the address 128.95.1.x. So what happens if I have a laptop and want to move from the iSchool network to the CS network?
- Another good example is the WLAN of HP's corporate visitor center, where guests with wired or wireless laptops are given one-hour IP address leases

25

Checking your IP Setup

- In Win95/98 and Win2k/XP (if you have permission) open the Network Control Panel
 - Or in Win2K/XP at a command prompt type: `ipconfig /all`
 - Or right click the LAN icon in the system tray ----->



26

What we saw...

- The machines in this room get their IP address via DHCP
- In addition to the IP address for that machine, ipconfig displayed a number of other IP addresses including:
 - Default gateway
 - Subnet mask
 - Primary Name server
 - Secondary Name server

We will discuss each a bit later today...

27

Delivering packets on a local Ethernet

- We have an IP address for each device on our LAN
- We also had said that each device has a MAC address that uniquely identifies the NIC in that device
- If we want to deliver a packet from one device on our LAN to another and we specify an IP address, how does that packet that was addressed using an IP address get delivered to the proper Ethernet MAC address?
 - Example: ping 128.208.100.56
 - How does the Ethernet card know where to deliver that packet?
 - Recall that Ethernet frames use MAC addresses, not IP addresses
 - Also recall that the IP address on any particular NIC can change (an administrator can change a static address, or it can change by itself if the address is dynamic)

28

ARP – Address Resolution Protocol

- ARP provides a mechanism for one device to discover the MAC address of another device given an IP address
- We saw ARP last week...
- Suppose we want to ping a device on another network such as a PC on the UW Computer Science network.
 - Assume that PC has an IP address 128.95.1.56
 - Do we just sent out an ARP broadcast looking for the device?

29

ARP isn't the answer

- ARP doesn't work because the PC we want is not on our local network
- How do we know what addresses are on our local network and which are not?
 - Recall that we said IP address are 32 bits long and that they are divided into (2) pieces
 - The NetID that identifies the network
 - The HostID that identifies the host

30

[Two pieces really?]

- When we wrote an address we wrote it like: 128.208.100.103
 - Using that notation how can you tell what part of that address is the network and part identifies the host? – No!
- Today on the Internet we use something called “subnet addressing” as a way to easily identify which part of the address is for the network, and which part is for the host
- To use subnet addressing, each network device is configured with not only with an IP address, but also with a “subnet mask”
- The subnet mask is a 32 bit value in which 1’s represent the network portion of the address and 0’s represent the host portion
- The subnet mask is written in dotted decimal notation just like the IP address is written (for example: 255.255.255.0)

31

[Subnet mask example]

- Consider the iSchool network. We have been assigned the 128.208.100.x network. Hosts on our network have addresses in the range:

128.208.100.1 – 128.208.100.254

The first 3 octets represent the “network” the last octet represents a host. We said we use 1’s to represent the network and 0’s to represent the host, so we have a subnet mask that looks like:

11111111 11111111 11111111 00000000

or in decimal, 255 . 255 . 255 . 0

32

[Using subnet masks]

- Nodes perform an “and” operation using their address and the subnet mask to determine what network they are on
- Example: My address 128.208.100.103, subnet mask 255.255.255.0

Address: 10000000 11010000 01100100 01100111

Mask: 11111111 11111111 11111111 00000000

And result: 10000000 11010000 01100100 00000000
128 208 100 0

We have computed that we are on the 128.208.100.x network

33

[Destination addresses]

- Similarly, nodes perform an “and” operation on destination host addresses and the subnet mask to determine if that destination is on their network or another network.

Consider the destination address 128.208.95.56

```
Address:      10000000 11010000 01011111 00111000
Subnet Mask: 11111111 11111111 11111111 00000000
And result:   10000000 11010000 01011111 00000000
              128      208      95       0
```

This destination is on the 128.208.95.x network, I am on the 128.208.100.x network, these are different networks

34

[Intuitively obvious?]

- At first glance you might think the “and” operation is a lot of work, in that intuitively you can easily tell what network you are on
- True when the subnet masks fall on the octet boundaries, but this is not always the case
- Consider a small company that has been given a Class C subnet address of 150.199.10.x
 - They have 4 divisions in their company that are in 4 different physical locations. Routers connect the 4 different networks together
 - If they used the standard 255.255.255.0 subnet mask they could only have one network with up to 254 hosts on that one network
 - They need to represent 4 different networks using a combination of an IP address and a subnet mask – what do they do?

35

[Subnetting...]

- In this case, we need to represent 4 different networks all in the 150.199.10.x space
 - How many binary digits does it take to represent 4 possibilities?

2 digits – 00, 01, 10, 11

So our subnet mask needs to be 2 bits longer like this:

```
11111111 11111111 11111111 11000000 or
      255      255      255      192
```

36

[Supernet example addresses]

If you had two class C addresses

150.199.6.x and 150.199.7.x
by using a 255.255.254.0 subnet mask you create a single network with valid addresses for hosts from 150.99.6.1 – 150.99.7.254

As expected, 150.99.6.0 can not be used as it is reserved and 150.99.7.255 is “broadcast”

Thus 150.199.6.0/24 and 150.199.7.0/24 equals 150.199.6.0/23, this is CIDR notation

CIDR = Classless Inter Domain Routing

40

[Generally speaking]

- While it is possible to use “unusual” subnet masks, network managers generally avoid them because they add complexity to the addressing
- These “unusual” subnet masks implement what is called ‘classless’ addressing (recall CIDR)
- Most subnet masks fall on the boundaries, so you will most frequently see 255.255.255.0 or 255.255.0.0 as subnet mask values
- Netmasks can be messy, hence the use of CIDR notation to simplify it.

41

[Subnet mask summary]

- The subnet mask in conjunction with an IP address allows a network device to know if a packet is destined for the local network, or if it has to go to a different network via a router
- Back to pinging a CS machine again...
 - We use the destination address and our subnet mask to know that this address is not on our network.
 - Hence we do not “arp” for it, but instead we must send the packet to our “default router”

42

Default Router (Gateway)

- Whenever you see the word “gateway” when TCP/IP is being discussed – think Router
- The “Default Router” or “Default Gateway” is where packets not on your network are sent
- Recall that a router is used to interconnect different networks. The “Default Gateway” is the router that is connecting your network to other networks
- Routers insure that traffic (including broadcast) that is between devices on your network stays on your network
- Routers only pass data along when it needs to go another network
- The default router IP address is configured again in the Network Control Panel in Windows and Mac, or in a configuration file on Unix.
- If a machine receives it's IP address via DHCP, the DHCP server tells the machine what the IP address of the default gateway is when the machine is assigned it's IP address

43

Routers themselves have default routers

- Example packet arrives at the router and the router says 128.95.1.56... is the 128.95.1.x network connected to one of my local interfaces?
 - If so we will send it to that network
 - If not the router tries to determine if it already knows a route to get to that destination
 - Routers run software to determine optimal routes and they exchange those routes with other routers using routing protocols like RIP, OSPF, or BGP
 - If the router doesn't know the route – the packet is sent to the router's default gateway (another router up in the hierarchy) for it to figure out
 - The process repeats
- To get from one network to another, a packet may traverse multiple routers

44

Traceroute

- Traceroute is a simple utility that can be used to show the routers that a packet traverses to get to a particular destination
- The route may change from time to time!
Why?
- In Unix, Linux, and Mac OS X use the command:
 - traceroute ipaddress
- In Windows, at a Command prompt:
 - tracert ipaddress

45

[In the lab, you can try ...]

- Start, Run, cmd to get to Command prompt
 - tracert www.ischool.washington.edu
 - tracert www.cs.washington.edu
 - What do we see from this info?
CS and the iSchool must both be connected to the same router!
 - tracert www.washington.edu
 - Did you all see the same routes?
 - What does this tell us?
 - tracert ils.unc.edu (a machine at the Univ of NC)
 - Notice the word "abilene"? That's the academic Internet 2 network
 - tracert www.cnn.com – what happens?

46

[IP Names and Numbers]

- In several of our demos we have used addresses like www.washington.edu rather than numbers like 140.142.3.7
- From an end user's perspective IP numbers are often hidden and names are used instead
 - Why are names preferable to numbers?
- The use of names rather than numbers can be done two different ways
 - Configuration files on each machine like /etc/hosts in Unix that have an IP number name pair. The device consults this table whenever a name is used to determine the address.
 - Obviously this method does not scale well! But still useful occasionally ie. In Lab1 you were able to ping using a name rather than a number because the names and address were pre-entered into a local table
 - Use of a distributed system called DNS (the Domain Name Service)

47

[More on DNS...]

- Domains are hierarchical
- Top level domains
 - .com, .edu, .org – **what else?**
 - Some new top level domains have recently been announced...like .info, .museum
- From the top level we may see a hierarchy like:
 - washington.edu or ischool.washington.edu or pc3.ischool.washington.edu

48

Important Note

- Domain names imply a logical hierarchy of addresses, but that hierarchy does not necessarily apply to the physical numerical addresses

For example, I could have names like:

pc1.demostore.com and pc2.demostore.com

pc1 could have an IP address of 128.208.100.150 and pc2 could be 150.63.20.2

Just because the IP name suggests that these hosts are on the same network, does not mean that they really are. In fact it is important that the names NOT map directly into the IP numbers. Why?

Consider a business with multiple networks or a company with multiple physical locations

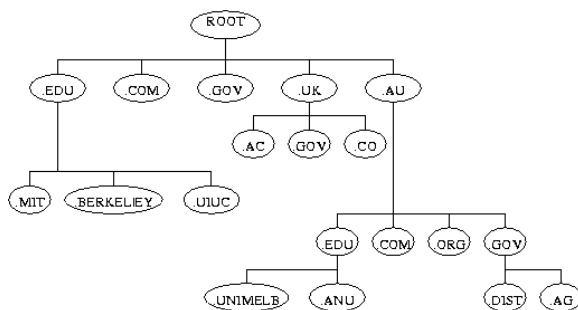
49

The logical hierarchy

- There is a reason for the hierarchy
 - DNS is a distributed system, it allows management to be delegated as needed
 - Individual organizations can run their own DNS servers and change names as needed – there is no central authority that they must inform to make changes
 - DNS consists of root servers at the top of the hierarchy that do not know all the names, but who to ask for each name
 - Each top level DNS server (ie .edu, .com) knows which DNS server in that domain to ask for a name/IP number

50

Portion of the DNS Hierarchy



51

[Getting a name]

- Just like with numbers, names have to be unique and they must be registered
 - Note that names don't have to have 4 parts like an IP address does, e.g. cnn.com is a valid DNS name
- Currently Network Solutions is one of the companies responsible for assigning .com, .org. and .net addresses.
 - www.networksolutions.com
 - Note the ability to register a name for a fee, also note the "whois" database lookup feature and check washington.edu and uw.edu. Interesting!

52

[Registration Issues (start here)]

- Can you register any name?
 - Well...people have tried and the issue has often ended up in the courts if there is a conflict
- Issues with NSI itself
 - a possible monopoly organization
 - now there are multiple registrars, not just one
- Recently efforts have been made to make the new top level domains accurately reflect those organizations that are part of it (for example .museum)
 - You actually do have to be a museum to get a .museum name

53

[Trying out DNS]

- Normally, your machine queries the DNS server automatically, you don't even know it is happening
- In Unix, Linux, Mac OS X, and Win2K/XP there is a tool called "nslookup" that you can use to manually query a DNS server for an IP address. Note nslookup is deprecated on the Unices and Linuxes, in favor of ...
- Also explore the "dig" command

54

[nslookup Sample]

- At a command prompt:
 - First do `ipconfig /all` again and note what it says is our DNS server's address. Note there is a primary and secondary – why do we have two?
 - Now type `nslookup`, and then try:
 - `argo.ischool.washington.edu` then `junk.ischool.washington.edu`
 - Try `128.208.100.103`
 - Try `www.cs.washington.edu`
 - Try connecting to it using the **IP number** in your Web browser
 - `www.cnn.com`
 - Very interesting, what is this showing? Can it work the other way, can one machine have more than one name? (check `fileserv.ischool.washington.edu` and `mail.ischool.washington.edu`)
 - Why might this be useful?

55

[DNS Summary]

- The domain name service is designed to take an Internet name that you enter (like www.cnn.com) and return to you the IP the number
- DNS servers are distributed throughout the Internet. Many organizations maintain their own DNS servers and these servers exchange their records with others on a periodic basis
 - With static IP assignments, a network administrator in your organization enters the information into a table or database that lists the individual machines and their name(s).
- Without a DNS server, your machine will not be able to connect to other machines using an Internet name. BUT...you could connect if you knew the IP number

56

[Layer 3]

- Routers & Routing are Layer 3 items
- Routing Protocols figure out routes
 - Solve the routing problem by allowing routers to communicate with each other
 - They facilitate the discovery of additional routing connections

57

[Examples]

- Some IP Routing protocols
 - Border Gateway Protocol (BGP)
 - Open Shortest Paths First (OSPF)
 - Routing Information Protocol (RIP)
 - Interior Gateway Routing Protocol (IGRP)


58

[Interior & Exterior Routing]

- Interior routing is for routes among routers that you control - within your autonomous system (AS)
- Exterior routing is for the larger global networks

59

[Interior/Exterior Diagram]



60

[IGP/EGP]

- The most common EGP
 - BGP (Border Gateway Protocol)
 - Specifically between large systems
 - BGP4 is used today
- Common IGPs
 - OSPF
 - RIP

61

[OSPF]

- OSPF is an open standard
 - IETF (Internet Engineering Task Force)
 - RFCs (Request For Comment) 1131, 1247, 1583, 2328

62

[OSPF Mechanics]

- Hierarchical Assignment
 - Each router is assigned a level in a hierarchy
 - The number of levels in an organization define the administrative area (responsibility) for each router
 - There must be at least 2 levels in the hierarchy
 - Logically this corresponds to the Internet (outside the organization) and LAN (inside the organization)
 - Only routers at the same level in the hierarchy exchange data

63

[OSPF Mechanics]

- Route Determination
 - Routers at the same level in the hierarchy exchange Link-State Advertisements (LSAs)
 - The LSA indicates a link is 'live' – after some time without an LSA the link is assumed to be 'dead'
 - Routers use the LSAs to build a graph representation of the known 'networks'
 - The Shortest-Paths algorithm is used to determine the best route to a known location

64

[Why use OSPF?]

- Hierarchical strategy
 - Good for scalability, can handle many routers in an organization
 - Limits scope of LSAs
 - High administrative overhead and complex to implement, only high end devices typically support OSPF
- Because of the overhead, smaller organizations often pick RIP

65

[RIP]

- The most simple of the Interior Gateway Protocols (IGP)
- One of the first TCP/IP routing protocols
- Originated in BSD Unix
 - routed or gated
- Has become a standard over time
 - two different versions, RIP1 and RIP2

66

[RIP Mechanics]

- Each RIP router periodically broadcasts its whole route table on each active network interface
- When a RIP message is received, the router compares its internal table with the route table it received
- If the router that sent the RIP message knows a shorter path to a network, accounting for the cost of sending to that router, then the new route replaces the old route
- Any route longer than 16 hops on the network is considered unreachable (4-bit field limits this routing metric a.k.a. hop count)
- The shortest of the available routes to the same destination has the lowest hop count metric and is thus the favored route.
- If no RIP broadcasts are heard on a RIP-enabled interface, then the router disables that interface

67

[Why use RIP?]

- Quick implementation (just turn it on)
- Simple to run, but...
 - No security, broadcasts go to everyone
 - When you have a large number of routers running RIP, the RIP messages get to be large and it can use a lot of bandwidth
 - RIP doesn't scale well due to limit of 16 hops
 - RIP doesn't support variable length subnet masks!
 - RIP can't calculate routes based on throughput, reliability, or delay – just number of hops

68

[Summary]

- Brief history of the Internet, origins of TCP/IP
- TCP/IP Layering model
- Supernetting & Subnetting
- IP Names, DNS
- Routing basics, EGP, IGP

69
