15-441: Computer Networks
Homework 2

Assigned: Sep 28, 2011
Due: Oct 13, 2011 1:30 PM in class
Lead TA: Charles Guo <rang@andrew.cmu.edu>

1 Tools

1. In this section, you will learn a couple of practical tools: **route**, and **ifconfig**, and **netstat**. Below is a very brief description of what they do. For more detailed information, check the man pages (**Note on unix.andrew route and ifconfig are located under /sbin/.** Either add this to your PATH or use the full path to run the commands).

   route The `route` command can be used to view and manipulate the IP routing table.

   netstat is a tool that can be used to display network connections, routing tables, interface statistics and many other things.

   ifconfig The `ifconfig` command is used to configure a network interface and to display the status of the currently active interfaces.

   (a) Run the `route` command. What is the use of the entry with netmask 0.0.0.0?

   (b) Suppose a malicious attacker runs "route del" to delete the routing entry corresponding to 128.2.13.128/26 on one of the unix.andrew.cmu.edu machines. Now you try to run "ping www.cnn.com". Do you expect the ping to get through? Harry Bovik, the local networking guru argues that "Well, there is still a default route, so the ping should go through!" Harry's pessimistic alter ego suggests otherwise, but does not have a reason. Should Harry believe his alter ego or dismiss it as a case of unjustified pessimism? Give a short 1-2 line answer why.

   (c) What is the command to view the routing table of your machine using netstat? What is the command to only show IP addresses and not host names in the routing table?

   (d) How can you use netstat to find out what the network interfaces of your machine are? What is the MTU of your Ethernet interface?

   (e) Run ifconfig. What information does the field "Mask" give you?

   (f) What happens if you run ifconfig and configure an interface to be in promiscuous mode?

2 Destination-Based Forwarding

2. [Sec. 4.3.1] A forwarding table for a router in a network using CIDR is given below.

Address prefix	Next hop
200.0.0.0/5	A
192.0.0.0/2	B
128.0.0.0/2	C
0.0.0.0/0	D

(a) If the router receives a packet with destination 199.42.13.37, what will the next hop be?

(b) If the router recieves a packet with destination 255.255.255.255, what will the next hop be?

(c) Suppose I add the entry 128.25.0.0/16 to the routing table. Will lookups still yield a single forwarding entry? Please explain.

(d) What actually happens to a packet that only matches 0.0.0.0/0?

3 ARP

3. Because there are both network-layer addresses (IP-addresses) and link-layer addresses (i.e. LAN addresses), there is a need to translate between them. For the Internet, this is the job of the address resolution protocol (ARP). The command `arp` allows you to view the *ARP table* which contains the mapping of IP addresses to LAN addresses on the local network.

   Run `arp` on your machine. Is there an entry in the ARP table for every host on the same LAN? Find out how entries enter and leave the ARP table and give a brief description. How can you use this information to determine the Ethernet address of a host on the same LAN that is currently not present in the ARP table of your machine?

4 BGP Tomfoolery

[Sec 4.1.2] Suppose a hacker obtains control of all the BGP-speaking routers in several different Autonomous Systems (ASes). Our hacker has each AS "hijack" several IP blocks. That is, each AS under his or her control announces via BGP that it owns IP blocks for which it does not. For example, our hacker has AS (CMU) announce a one-hop path to the IP block 18.0.0.0/8 (MIT).

4. Assuming that the AS graph still converges to a stable state, can this attack cause routing loops to form? Explain why or why not.

5. Suppose the ASes under attack are identified. Can other ASes change their routing policies to ensure that their traffic still reaches the hijacked IP blocks? Explain.

6. In response to this attack, suppose all ASes agree to check a central registry for IP block ownership before a path is considered valid. That is, whenever an AS receives a route to a prefix P, it checks that the last AS in the route actually owns P. For example, upon receiving a path to 18.0.0.0/8 (MIT), an AS will check that the last AS in the route is 3 (MIT). Can a hacker still hijack IP address blocks belonging to ASes he or she does not control? (i.e., can he or she cause traffic destined to those IP blocks to be routed to the ASes he controls?) Explain.

7. Suppose a solution was devised where IANA hosted a server on the Internet that was able to validate all AS paths. Assume that this server is always trustworthy and paths are valid if and only if the server says so.

   True or false: With this solution an AS can always check the validity of a BGP path advertisement it receives.

5 Routers and Bridges

[Sec 3.1] A small company has the network topology shown below. In this topology: Three Ethernet segments (L1, L2, L3) are interconnected with a transparent bridge B and router R. Nodes E, F, I, and J are endsystems. Endsystems do not forward packets.

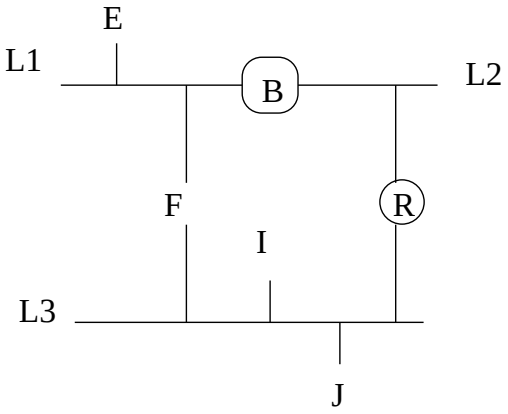

8. In this problem, use the following notation: the MAC address of node E is e (lower case), its IP address is E (UPPER CASE). If a node has multiple IP or MAC addresses, use the index of the segment to differentiate between them. For example, host F has two different IP addresses, F[1] and F[3], since it is on two LANs. The first line in the table provides an example entry for a packet from E to F on link L1.

   Suppose host E opens a connection to the Web server J. In the following table, give the MAC and IP source and destination addresses of the SYN packet of the TCP connection as it is observed traveling on links L1, L2, and L3 (i.e. fill in the next 3 lines).

Location	MAC Source	MAC Dest	IP Source	IP Dest
E to F on link L1	e	f[1]	E	F[1]
E to J on L1				
E to J on L2				
E to J on L3				

9. The administrator has allocated class C addresses to each of the routed IP subnets. All the networking parameters of the different nodes in the system are statically configured. At some point, the company wants to move server J from link L3 to link L1.
   (a) What changes are needed to the configuration of J?
   (b) What changes are needed to the configuration of R? (Don't worry about DNS for any of these questions).
   (c) What other nodes need to be updated?

10. At a later time endsystem E is moved to link L2. What nodes must have their configuration changed and how?

Because of high network utilization in L3, the administrator reduces network load in L3 by segmenting it.

11. What kind of device should be used if the administrator wishes to minimize his work? Give a reason for your answer.

12. What kind of device should be used if a significant portion of the network traffic is broadcast (primarily ARP traffic) on the LAN? Give a reason for your answer.

6 Using 'dig' to Understand DNS

13. In this question you will use the unix utility 'dig' to explore the contents of DNS messages. Please use dig on unix.andrew.cmu.edu.

    The format of a dig request is simple. Just type: dig www.princeton.edu to perform a look-up for that DNS name. As you now know, DNS requests can do more than just ask for the IP address corresponding to a single DNS name. Type dig princeton.edu ANY to see DNS records of all types that are associated with the domain 'princeton.edu'.

    (a) What IP address did the computer you are logged into contact to make the DNS request? Where do you think this server is located?

    (b) List all of the different types of records received as a result of your query. For each record, explain its purpose, using one of the entries provided in the reply as a concrete example.

    (c) Note that some of the names in the reply are not in the domain 'princeton.edu'. Use the DNS names and/or 'traceroute' to find the general location of one of these servers. Where is it? Given the type of record, why would Princeton do this?

    (d) Use dig to find the names of two non-local servers you could contact in the process of identifying the nameserver for the domain 'cnn.com' (assume no DNS information is cached anywhere).

    (e) Use dig to find the TTL for the DNS mappings of 'www.cnn.com' and 'www.cs.stanford.edu'. What are they? If your boss asks you to provide two positive and two negative effects of having a short DNS TTL for the company's e-commerce site, what would you say?

7 IPv6

14. [Sec. 4.1.3]

    (a) IPv6 allows for 2^{128} addresses. That is about 2^{95} addresses for every person alive. Why was IPv6 designed to have so many addresses? Please identify two specific reasons.

    (b) IPv4 contained a checksum while IPv6 does not. Please explain why the designers of IPv6 decided the checksum was not needed. Provide at least 2 concrete reasons why they felt they could or should drop this feature.

8 RIP

15. [Sec. 3.3.2] For this problem, consider the following network topology:

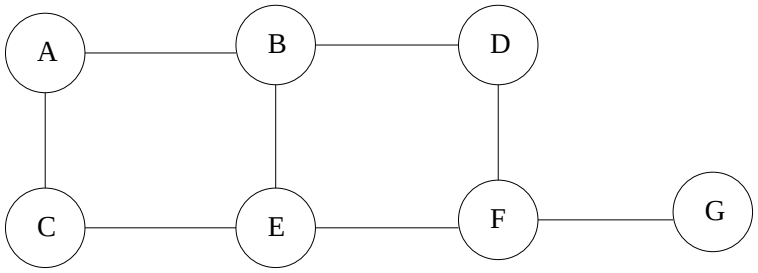

Suppose that the routers are running RIP, and are all turned on simultaneously. Recall that in RIP, the weight assigned to each link is one. You can also assume that if RIP identifies multiple paths of equal cost, it picks one randomly.

The initial routing table at node A is:

Destination	Cost	Next Hop
B	1	B
C	1	C
D	∞	—
E	∞	—
F	∞	—
G	∞	—

Fill in the following tables to show the initial routing table at node F:

Destination	Cost	Next Hop
A		
B		
C		
D		
E		
G		

16. Now show the contents of the routing table after each iteration of the algorithm:

(a) After iteration 1:

Destination	Cost	Next Hop
A		
B		
C		
D		
E		
G		

(b) After iteration 2:

Destination	Cost	Next Hop
A		
B		
C		
D		
E		
G		

17. In some failure situations, the administrator notices that it takes an exceptionally long time for the routing protocol to stabilize in this network.

(a) What problem with RIP is the cause?

(b) The administrator is told that BGP does not suffer from this problem. What prevents BGP from having this problem?

9 ping vs traceroute

18. Instead of traceroute, you can also use ping to find the route that packets take to a given network host. It does so by using the IP Record Route option. Try the following two commands:

ping -R www.cmu.edu

traceroute www.cmu.edu

What is the difference in their output? Now try the following:

ping -R www.google.com

Do you get the desired output? Explain your observation. Lastly, try running the following command a few times.

traceroute www.whitehouse.gov

What do you notice about the results you get from different runs?

10 DNS Redirection

[Sec 9.3.1] Harry Bovik is working on a web site that has multiple replicated servers located throughout the Internet. He plans on using DNS to help direct clients to their nearest server replica. He comes up with a hierarchical scheme. Harry has divided his server replicas into three groups (east, west and central) based on their physical location. A typical query occurs as follows:

- When a client makes a query for `www.distributed.hb.com`, the root and `.com` name servers are contacted first. It returns the name server (NS) record for `ns1.hb.com`. The TTL of this record is set to 1 day.
- The `ns1.hb.com` name server is then queried for the address. It examines the source of the name query and returns a NS record for one of {`east-ns, central-ns, west-ns`}`.distributed.com`. The choice of which name server is based on where `ns1` thinks the query came from.
- Finally, one of {`east-ns, central-ns, west-ns`}`.distributed.com`. is contacted and it returns an address (A) record for the most lightly loaded server in its region.

Answer the following 3 questions based on this design.

19. Harry's name server software has only two choices for TTL settings for A and NS records - 1 day and 1 minute. Harry chooses the following TTLs for each record below:

    1. NS record for {`east-ns, central-ns, west-ns`}`.distributed.com` - 1 day TTL.
    2. A record for {`east-ns, central-ns, west-ns`}`.distributed.com` - 1 day TTL.
    3. A record returned for the actual Web server - 1 minute TTL.

    Briefly explain why Harry's choices are reasonable, or why you would have made different choices.

20. In general, name resolution systems map names based on the name and context. In this particular case, what are ***TWO*** items of context that the name resolution uses?

21. Harry's Web site is especially popular among CMU students. The CMU network administrator estimates that there is one access from CMU every 3 minutes. Each access results in the application resolving the name `www.distributed.hb.com`. Assume the following:

    - No other DNS queries are made in CMU
    - All CMU clients use the same local name server.
    - This local name server is mapped to the `east-ns` region.
    - Web browsers do not do any caching on their own.

How many accesses per hour will be made to the following name servers to resolve these CMU queries? Explain your calculation.

1. The Root Servers
2. `ns1.hb.com`
3. `east-ns.distributed.com`