15-441: Computer Networks Spring 2010
Homework #3

Due: Mar 23^{rd} 2010, in class
Lead TA: Daegun Won (`daegunw@andrew.cmu.edu`)

1 DNS Tools

In the following questions, you will learn to use two useful tools for querying for Domain Name System (DNS) information: `nslookup` and `dig`.
The `nslookup` program queries Internet domain name servers (DNS). Entering the command nslookup will give you the name of the name server your system knows and its IP address. The list of name servers used by a Unix machine can usually be found by looking at the file */etc/resolv.conf*. Read the man page for nslookup and answer the following questions.

1. What is the IP address of andrew webmail server (`webmail.andrew.cmu.edu`)?
2. When you send mail to `somebody@steelers.com`, which machine does the mail go to? What are the

machines used to process mail sent to somebody@cs.cmu.edu?

dig is another program that allows you to query DNS servers. For the purpose of this question, you should use the following format to invoke dig

```
dig +norecurse @<name.of.dns.server> <record-type><domain-name>
```

where

- <name.of.dns.server> is the hostname of the DNS server you wish to query such as A.ROOT-SERVERS.NET
- <record-type> is the type of DNS record you wish to retrieve, such as ANY, MX, etc.
- <domain-name> is the name of the host or domain you seek information on.

The DNS is a distributed architecture that uses hierarchical delegation. At the top of the system are the root name servers, which know which DNS server is responsible for each second-level domain (such as cmu.edu). If you send a root server a query for a particular machine, you will receive a reply listing the servers that have been delegated authority for that machine's second-level domain. It is common for a large domain such as cmu.edu to further delegate to departmental or workgroup DNS servers, which you discover by querying the second level servers.

In order to discover the chain of delegation in use at Akamai, run a series of NS queries for a1794.x.akamai.net. You may wish to start with any of the 13 root servers ([a-m].root-servers.net), and you should continue your sequence of queries until you stop getting new delegations (in some domains this is indicated by a DNS server returning you a delegation pointing to itself, and in other domains this is indicated by a DNS server returning you a SOA record instead).

As an example here is the delegation chain for aol.com:

Server queried	NS delegates to
B.ROOT-SERVERS.NET	A.GTLD-SERVERS.NET
A.GTLD-SERVERS.NET	DNS-01.NS.AOL.COM, DNS-0[267].NS.AOL.COM
DNS-01.NS.AOL.COM	DNS-01.NS.AOL.COM

3. Generate the delegation chain for a1794.x.akamai.net. Present your results in the table form shown above. Each NS query will typically return two or more answers: choose among them at random. If you

query a server and get a timeout, choose an alternate server.

DNS is also used for reverse lookup, i.e. to translate IP addresses into hostnames. Again, the database is distributed in a hierarchical fashion, with a wrinkle. The most-specific part of a domain name is on the left (i.e. `ux1` in `ux1.sp.cs.cmu.edu`), but the reverse is true of IP addresses (i.e. in 128.2.198.101, 128 is top-level, 128.2 is Carnegie Mellon in general and 128.2.198 belongs to the Computer Science Department. Thus, address-to-name mappings are discovered by reversing the bytes of the IP address and making queries in a special domain. To turn 128.2.198.101 into a hostname, various servers are sent queries seeking PTR records for `101.198.2.128.in-addr.arpa`. The first query would be:

```
dig @b.root-servers.net PTR 101.198.2.128.in-addr.arpa
```

You will know when you're done when your query gives you back a PTR record in addition to (or instead of) NS record.

Note that you have to reverse the bytes in the address, i.e. start with the lowest level byte, e.g.,

```
dig @b.root-servers.net PTR 36.0.26.18.in-addr.arpa
```

Server Queried	NS delegates to
`B.ROOT-SERVERS.NET`	`STRAWB.MIT.EDU`
`STRAWB.MIT.EDU`	returns the PTR record

The hostname is `mintaka.lcs.mit.edu.`

4. Find out the TTL for the DNS record of `mintaka.lcs.mit.edu` at `strawb.mit.edu`. What's the default TTL (i.e. the TTL when the record is first cached) for this record? Show your process and any results of command you used.

5. Fill in a table like the one above showing a query chain for the IP address 128.2.201.61.

2 DNS Redirection

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.

6. Harry's name server software has only two choices for TTL settings for A and NS records - 12 hours and 1 minute. Harry chooses the following TTLs for each record below:

   1. NS record for {`east-ns, central-ns, west-ns`}`.distributed.com` - 12 hours TTL.
   2. A record for {`east-ns, central-ns, west-ns`}`.distributed.com` - 12 hours 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.

7. 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?

8. 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 in average 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`

3 UDP vs TCP

9. Determine what transport layer protocol you would use for each of the following applications and then give as many reasons why as you can identify:
   (a) File Transfer
   (b) Live video streaming
   (c) DNS server
   (d) Web browser
   (e) VoIP