

## Solutions to Chapter 2 Problems

Q1. True or false

- a) A user requests a Web page that contains of some text and three images. For this page, the client will send one request message and receive four response messages.
- b) 2. Two distinct Web pages (for example, [www.mit.edu/research.html](http://www.mit.edu/research.html) and [www.mit.edu/students.html](http://www.mit.edu/students.html)) can be sent over the same persistent connection.
- c) 3. With non-persistent connections between browser and origin server, it is possible for a single TCP segment to carry two distinct HTTP request messages.
- d) 4. The Date: header in the HTTP response message indicates when the object in the response was last modified.
- e) 5. HTTP response messages never have an empty message body.

(This question is taken from Kurose & Ross's book, Chapter 2 Problem 1)

Solution:

- a) F
- b) T
- c) F
- d) F
- e) F

Q2.

- a) HTTP and FTP are both file transfer protocols and have many common characters; for example, they both run on top of TCP. What are the main differences between HTTP and FTP?
- b) How does SMTP mark the end of a message body?

(This question is taken from Kurose & Ross's book, Chapter 2, parts of Problem 14)

Solution:

- a) The main difference is that FTP uses two parallel TCP connections to transfer a file, a control connection and a data connection. The control connection is used for sending control information between the two hosts – information such as user identification, password, commands to change remote directory, etc. The data connection is used to actually send a file. Because FTP uses a separate control connection. FTP is said to send its control information out-of-band. HTTP, as you recall, sends request and response header lines into the same TCP connection that carries the transferred file itself. For this reason, HTTP is said to send its control information in-band.

- b) SMTP uses a line containing only a period to mark the end of a message body. HTTP uses “Content-Length header field” to indicate the length of a message body.

Q3. Referring to Example 2 (in Chapter 2 Examples and Homework Problems), suppose the HTML file references eight very small objects on the same server. Neglecting transmission times, how much time elapses with

- a) Non-persistent HTTP?  
b) Persistent HTTP?

(This question is taken from Kurose & Ross’s book, Chapter 2, parts of Problem 8)

Solution:

- a)  $RTT_1 + \dots + RTT_n + 2RTT_0 + 8 * 2RTT_0$   
b)  $RTT_1 + \dots + RTT_n + 2RTT_0 + 8 * RTT_0$

Q4. In this problem, we use the useful *dig* tool available on Unix and Linux hosts to explore the hierarchy of DNS servers. Recall a DNS server higher in the DNS hierarchy delegates a DNS query to a DNS server lower in the hierarchy, by sending back to the DNS client the name of that lower-level DNS server. First read the man (manual) page for *dig*, and then answer the following questions.

- a) Starting with a root DNS server (from one of the root server [a-m].root-servers.net), initiate a sequence of queries for the IP address for [www.cse.yorku.ca](http://www.cse.yorku.ca) by using *dig*. Show the list of the names of DNS servers in the delegation chain in answering your query.  
b) Repeat part a) for [www.google.com](http://www.google.com), [www.yahoo.com](http://www.yahoo.com).

(This question is taken from Kurose & Ross’s book, Chapter 2 Problem 19)

Solution:

The following delegation chain is used for [www.cse.yorku.ca](http://www.cse.yorku.ca)

a.root-servers.net

j.ca-servers.ca

dns11.ipns.yorku.ca (authoritative)

ns2.cse.yorku.ca

First command:

`dig +norecurse @a.root-servers.net any www.cse.yorku.ca`

;; AUTHORITY SECTION:

ca.	172800	IN	NS	j.ca-servers.ca.
ca.	172800	IN	NS	a.ca-servers.ca.
ca.	172800	IN	NS	l.ca-servers.ca.
ca.	172800	IN	NS	e.ca-servers.ca.
ca.	172800	IN	NS	c.ca-servers.ca.

ca.	172800	IN	NS	tld.isc-sns.net.
ca.	172800	IN	NS	sns-pb.isc.org.
ca.	172800	IN	NS	z.ca-servers.ca.
ca.	172800	IN	NS	k.ca-servers.ca.

;; ADDITIONAL SECTION:

a.ca-servers.ca.	172800	IN	A	199.253.251.2
a.ca-servers.ca.	172800	IN	AAAA	2001:500:80:5000::2
c.ca-servers.ca.	172800	IN	A	192.228.28.9
e.ca-servers.ca.	172800	IN	A	192.228.30.9
j.ca-servers.ca.	172800	IN	A	198.182.167.1
j.ca-servers.ca.	172800	IN	AAAA	2001:500:83::1
k.ca-servers.ca.	172800	IN	A	156.154.100.4
l.ca-servers.ca.	172800	IN	A	156.154.101.4
z.ca-servers.ca.	172800	IN	A	192.228.25.66
tld.isc-sns.net.	172800	IN	A	63.243.194.3
tld.isc-sns.net.	172800	IN	AAAA	2001:5a0:10::3
sns-pb.isc.org.	172800	IN	A	192.5.4.1
sns-pb.isc.org.	172800	IN	AAAA	2001:500:2e::1

Among all returned edu DNS servers, we send a query to the first one.  
dig +norecurse @j.cs-servers.ca any www.cse.yorku.ca

;; AUTHORITY SECTION:

yorku.ca.	86400	IN	NS	dns11.ipns.yorku.ca.
yorku.ca.	86400	IN	NS	dns21.ipns.yorku.ca.
yorku.ca.	86400	IN	NS	hanna.cac.washington.edu.

;; ADDITIONAL SECTION:

dns11.ipns.yorku.ca.	86400	IN	A	130.63.9.2
dns21.ipns.yorku.ca.	86400	IN	A	130.63.10.2

Among all three returned authoritative DNS servers, we send a query to the first one.  
dig +norecurse @dns11.ipns.yorku.ca any www.cse.yorku.ca

;; AUTHORITY SECTION:

cse.yorku.ca.	21600	IN	NS	ns2.cse.yorku.ca.
cse.yorku.ca.	21600	IN	NS	ns3.cse.yorku.ca.
cse.yorku.ca.	21600	IN	NS	ns1.cse.yorku.ca.

;; ADDITIONAL SECTION:

ns1.cse.yorku.ca.	21600	IN	A	130.63.97.28
ns2.cse.yorku.ca.	21600	IN	A	130.63.92.33
ns3.cse.yorku.ca.	21600	IN	A	128.100.1.1

The answer for [www.google.com](http://www.google.com) could be:

d.root-servers.net

d.gtld-servers.net

ns2.google.com(authoritative)

You may issue the command: dig [www.google.com](http://www.google.com) +trace

The answer for [www.yahoo.com](http://www.yahoo.com) could be:

h.root-servers.net

k.gtld-servers.net

ns1.yahoo.com

Q5. Suppose that your department has a local DNS server for all computers in the department. You are an ordinary user (i.e. not network/system administrator). Can you determine if an external Web site was likely accessed from a computer in your department a couple of seconds ago? Explain.

(This question is taken from Kurose & Ross's book, Chapter 2 Problem 21)

Solution

Yes, we can use dig to query that Web site in the local DNS server.

For example, "dig cnn.com" will return the query time for finding cnn.com. If cnn.com was just accessed a couple of seconds ago, an entry for cnn.com is cached in the local DNS cache, so the query time is 0 msec. Otherwise, the query time is large.