

Quiz 2 Solution

Your Name:

9/16/2014

1. (4 points) A company is trying to assess the benefit of deploying a web proxy on their internal network. The company has a 10 Mbps link connecting it to the Internet. The link load is currently 95% just to accommodate web traffic. From measurements, the company has determined that on average requests to individual websites are roughly of the same size, but that 80% of users' requests target only 100 websites. The remaining 20% of users' requests are for webpages that are rarely common across users. The 100 most popular websites together store a total of about 500 Gbytes worth of webpage data with all pages equally likely to be requested. However, 75% of those webpages are dynamic pages that have content which continuously changes, *i.e.*, will typically be different even for two consecutive requests.

Assuming that the company decides to setup a web proxy that only caches the “static” pages of the 100 web sites it has identified as popular targets, how big should the proxy cache be and what fraction of requests will it be able to successfully serve? What will then the traffic load be on the company's 10 Mbps link?

The static pages make-up 25% of the 500 Gigabytes worth of webpages on the 100 most popular websites, i.e., for a total of 125 Gigabytes. Similarly, those static webpages will be targeted by approximately 25% of the requests aimed at those popular websites for a total of, therefore, 20% of all page requests.

Given that those latter requests will now be served directly from the cache, the load on the 10 Mbps access link will decrease by 20% from 95% to $0.8 \cdot 95\% = 76\%$.

2. (2 points) Assume that instead of using a “.” on a single line, the SMTP protocol opted for the combination “:;” on a single line as its end of mail message. Is this sufficient to safely indicate the end of any mail message or should additional steps be specified. If no additional steps are needed, why, and if you think some are required, what would they be?

*Unfortunately and even if unlikely, the combination “:;” could still occur on a single line as part of a user message. To avoid this, we can use the same “character stuffing” approach that SMTP uses, namely, whenever the line of a user message starts with the combination “:;”, the mail client should replace it with two such combinations in sequence, *i.e.*, “:;;:”. Conversely, on receipt of a message, the server should remove the second copy of any such combination when found at the start of a line (if it is the only such combination at the start of the line, then it is the end of message command).*

3. (4 points). Suppose a server with IP address 156.12.34.56 executes the following lines of java.

```
ServerSocket sock = new ServerSocket();  
sock.bind(InetSocketAddress("156.12.34.56", 4576));  
Socket connsock = sock.accept();  
InetAddress x = connsock.getInetAddress();
```

Now, suppose a host with IP address 72.35.64.85 executes the following lines.

```
Socket sockA = new Socket();  
sockA.bind(InetSocketAddress("72.35.64.85", 4576));  
sockA.connect(InetSocketAddress("156.12.34.56", 14333));
```

At this point, how many sockets are there at the server?

1 socket, namely, the ServerSocket sock, since the connect request from host 72.35.64.85 is to a port that the server is not listening on.

A short while later, another host with IP address 23.44.55.66 executes the following lines.

```
Socket sockB = new Socket();  
sockB.bind(InetSocketAddress("23.44.55.66", 55621));  
sockB.connect(InetSocketAddress("156.12.34.56", 4576));
```

At this point, how many sockets are there at the server?

2 sockets, the server and client sockets for the connection to host 23.44.55.66.

How many port numbers are being used at the server?

1 port number (both sockets are associated with port 4576).

What is the value of the variable x at the server?

$x = 23.44.55.66$, the address of the client.