Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

Distributed Computing

FS 2017 Prof. R. Wattenhofer

Computer Engineering II

Solution to Exercise Sheet Chapter 4

1 Quiz Questions

a) A user provides his credentials, i.e., login information. The server sets a cookie with a session ID associated with the server side information about the user, i.e., the session data. Basic Auth, i.e., including the username and password in the URL, is not commonly used since it shows a popup that cannot be integrated into the look and feel of the webpage. Furthermore, copy-pasting a URL somewhere may result in the credentials being leaked.

b) It may be ambiguous whether the client is using encrypted or unencrypted requests. The age of the protocol is not the cause since it is trivial to introduce an upgrade command which bootstraps the encrypted communication, e.g., `STARTTLS` for SMTP upgrades a plaintext connection after the initial handshake to use Transport Layer Security (TLS).

c) No. Nowadays most users deliver outgoing mails to their provider's outgoing mailserver which then takes care of forwarding the mail to the destination. This adds authentication since the outgoing mailserver can now certify that the mail was sent by a legitimate user and it may attempt delivery multiple times should the destination mailserver not be reachable.

d) Yes, the server may have a cached response to the resolution request and respond directly without having to perform the entire resolution. Should the server not have a cached response the overhead of using an intermediate server is minimal since the ISP is on the communication path to the authoritative nameserver anyway.

e) There are at least 3 connections being opened:

- The first connection is opened to retrieve `http://google.ch/`. This page redirects (status code 301) to `http://www.google.ch/` which may not be served by the same server.
- The second connection is opened to retrieve `http://www.google.ch/` (port 80) redirects to `https://www.google.ch/` (port 443) which is not on the same port.
- The third connection actually retrieves the content of the webpage, but additional connections may be opened to retrieve referenced resources in parallel.

2 Send me a comment

a) Depending on the website you visited there will be an initial connection to retrieve the HTML page followed by a cascade of resources referenced in the HTML page. Since each of the resources may themself reference some other resources it is not uncommon to see multiple waves of parallel requests.

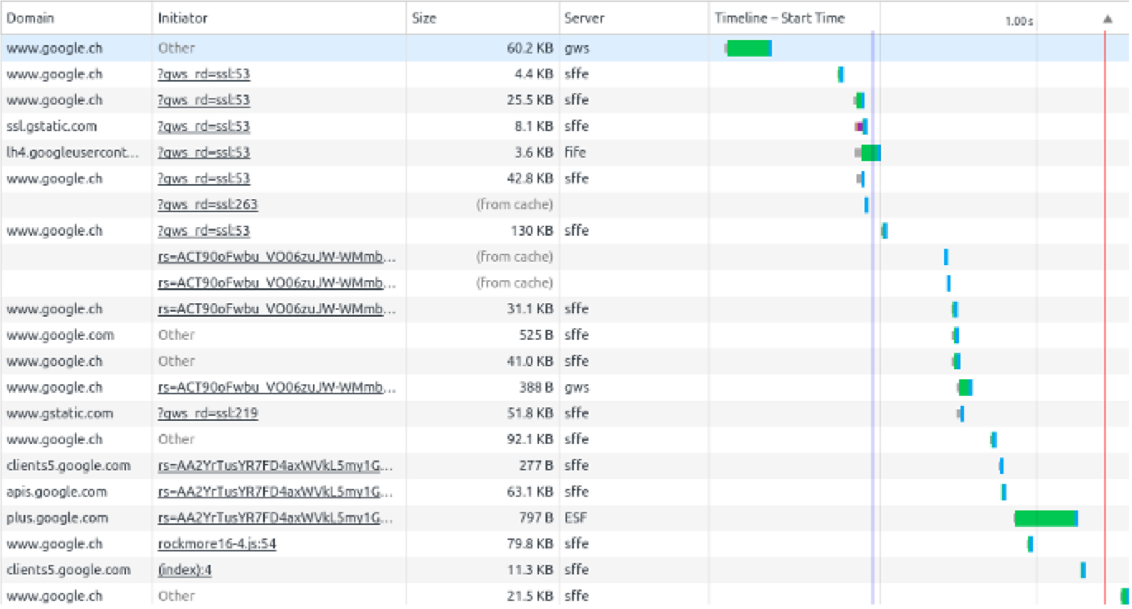

Domain	Initiator	Size	Server	Timeline – Start Time
www.google.ch	Other	60.2 KB	gws	
www.google.ch	?gws_rd=ssl:53	4.4 KB	sffe	
www.google.ch	?gws_rd=ssl:53	25.5 KB	sffe	
ssl.gstatic.com	?gws_rd=ssl:53	8.1 KB	sffe	
lh4.googleusercont...	?gws_rd=ssl:53	3.6 KB	fife	
www.google.ch	?gws_rd=ssl:53	42.8 KB	sffe	
	?gws_rd=ssl:263	(from cache)		
www.google.ch	?gws_rd=ssl:53	130 KB	sffe	
	rs=ACT90oFwbu_VO06zuJW-WMmb...	(from cache)		
	rs=ACT90oFwbu_VO06zuJW-WMmb...	(from cache)		
www.google.ch	rs=ACT90oFwbu_VO06zuJW-WMmb...	31.1 KB	sffe	
www.google.com	Other	525 B	sffe	
www.google.ch	Other	41.0 KB	sffe	
www.google.ch	rs=ACT90oFwbu_VO06zuJW-WMmb...	388 B	gws	
www.gstatic.com	?gws_rd=ssl:219	51.8 KB	sffe	
www.google.ch	Other	92.1 KB	sffe	
clients5.google.com	rs=AA2YrTusYR7FD4axWVkL5my1G...	277 B	sffe	
apis.google.com	rs=AA2YrTusYR7FD4axWVkL5my1G...	63.1 KB	sffe	
plus.google.com	rs=AA2YrTusYR7FD4axWVkL5my1G...	797 B	ESF	
www.google.ch	rockmore16-4.js:54	79.8 KB	sffe	
clients5.google.com	(index):4	11.3 KB	sffe	
www.google.ch	Other	21.5 KB	sffe	

Figure 1: The waterfall of requests when opening a simple webpage.

b) The following is an example HTTP request and response:

```
# telnet virt13.ethz.ch 80
Trying 82.130.102.226...
Connected to virt13.ethz.ch.
GET / HTTP/1.1
Host: virt13.ethz.ch

HTTP/1.1 200 OK
Transfer-Encoding: chunked
Date: Wed, 09 Mar 2016 15:15:54 GMT
Content-Type: text/html
Server: TwistedWeb/15.5.0

249

<html>
  <head>
    <title>Disco Comments</title>
    <link href="http://getbootstrap.com/dist/css/bootstrap.min.css" rel="stylesheet" />
  </head>
  <body lang="en">
    <header class="navbar navbar-static-top">
      <div class=container>
        <div class=navbar-header>
          <div class=navbar-brand>Disco Comments</div>
        </div>
      <div>
    </header>
    <div class=container>
      <div class="col-md-12">
```

```
30          <dl>
31            <dt id="mid-0">Disco</dt>
32            <dd>Welcome to the Disco Comment page</dd>
33          </dl>
34        </div>
35      </div>
36    </body>
37  </html>
38  0
39
40  Connection closed.
```

The request (lines 4-6) specifies that HTTP/1.1 is to be used, hence the connection will remain open for a few seconds before the server closes it due to a timeout. The response header (lines 7-12) returns some information about the server and metadata about the response. Notice that the Transfer-Encoding is set to `chunked` which means that the server returns the response in multiple chunks, each prefixed with the number of bytes (as hexadecimal numbers) on a separate line. In this case we have two chunks, one of 249_{hex} bytes starting at line 13 and another chunk of 0_{hex} bytes, indicating the end of the response, starting at line 38.

c) As mentioned we need to construct a `PUT` request to resource `/`. The following sends the message `hello world` to the server.

```
 1  # telnet virt13.ethz.ch 80
 2  Trying 82.130.102.226...
 3  Connected to virt13.ethz.ch.
 4  PUT / HTTP/1.1
 5  Host: virt13.ethz.ch
 6  Content-Length: 11
 7
 8  hello world
 9  HTTP/1.1 201 Created
10  Transfer-Encoding: chunked
11  Date: Wed, 09 Mar 2016 15:30:37 GMT
12  Content-Type: text/html
13  Server: TwistedWeb/15.5.0
14
15  18
16  Comment added with ID 3
17
18  0
```

The request now includes a `Content-Length` option which tells the server how many bytes to read. The request spans lines 4-8, including the payload on line 8.

3 Send me a mail

a) Simply sending a mail to `somebody@virt13.ethz.ch` will not work because when searching for the responsible mailserver for the domain `virt13.ethz.ch` using DNS no `MX` record is returned hence the sender is unable to find the responsible server.

b) The following is a valid SMTP session:

```
 1  telnet virt13.ethz.ch 25
 2  Trying 82.130.102.226...
 3  Connected to virt13.ethz.ch.
 4  220 cd652f72069f NO UCE NO UBE NO RELAY PROBES ESMTP
 5  EHLO pc-5305.ethz.ch
 6  250 cd652f72069f Hello 82.130.102.226, nice to meet you
 7  MAIL FROM: someone@student.ethz.ch
 8  250 Sender address accepted
 9  RCPT TO: recipient@virt13.ethz.ch
10  250 Recipient address accepted
11  DATA
12  354 Continue
13  This is a test message
14  that spans multiple lines
15  .
16  250 Delivery in progress
17  QUIT
18  221 See you later
19  Connection closed by foreign host.
```

This is an interactive session with the SMTP server, lines starting with a number are sent by the server. The client needs to identify itself using the `EHLO` command on line 5. On lines 7 and 9 we specify sender and recipient, notice that the sender needs to be specified before the recipient. The `DATA` command is then used to start the content transfer, which is terminated on line 15 with a single dot on a line.

4 DNS

4.1 Getting Started

a)
- If you are outside the ETH network: `dig disco.ethz.ch` or `dig +domain=ethz.ch disco`
- If you are within the ETH network: `dig disco.ethz.ch`

b) CNAME stands for Canonical Name. The entry `disco.ethz.ch. 101 IN CNAME disco01-srv.ethz.ch.` should be read as: `disco.ethz.ch` is an alias for the CNAME (true name) `disco01-srv.ethz.ch.` Using CNAMEs, we can point several aliases to the same CNAME, which in turn points to one IP address. Therefore, if an IP address changes, one only has to change one entry (namely the A or AAAA record).

c) A stands for **A**ddress. An A record maps a DNS name to an IPv4 address. AAAA records are used for IPv6 addresses.

d) There are four mail servers in the `ethz.ch.` domain: `phil[1-4].ethz.ch`. You can use the command `dig mx ethz.ch` to get this result. The number in front of the DNS names indicates a preference for the mail servers. If there are several mail servers, it chooses the one with the lowest number. In our case, all mail servers have the same preference number and will be treated equally, i.e., the mail server is chosen uniformly at random.

e) `disco.` is a Fully-Qualified Domain Name (FQDN) and therefore not in the ethz.ch. domain. `disco` on the other hand is a relative domain name, and can therefore be found under `disco.ethz.ch.`.

4.2 DNS Queries

a) Typing `dig` will return all root servers.

b) `dig +norec @a.root-servers.net disco.ethz.ch`. The root servers do not know `disco.ethz.ch`. However, they know the name servers of the `ch.` top-level domain. So, the root servers cannot give us an answer directly, but they return a list of name servers that might know the answer.

c) `dig +norec @a.nic.ch. disco.ethz.ch` and `dig +norec @ns1.ethz.ch disco.ethz.ch`. Then, you can find the IP address of `disco.ethz.ch` in the answer section.

d) Using the command `dig disco.ethz.ch`, we ask our local resolver for the IP address of `disco.ethz.ch`, which probably does not know it yet. The local resolver will then automatically perform the steps we just did by hand, and return the result to the client. And of course, the IP addresses that were returned are the same.

4.3 DNS Caching

a) Just issue any `dig` request, and look at the bottom of the output. The `SERVER` field tells you the IP address of your standard resolver. This could for example be the router in your home network.

b) Since we are issuing a non-recursive query, the default name server can only return the result to us, if it has previously been cached. Therefore, if the name server's reply does not contain an *ANSWER SECTION*, it means that the query was not cached. If we then perform a standard recursive query, we get a response with an *ANSWER SECTION*. Now we can issue another non-recursive query, which will then also provide us with a response with an *ANSWER SECTION* containing the IP address we were looking for, since the query has now been cached.

c) Cached queries are of course much faster. In our case, the recursive query without cache-hit took 53ms, and the iterate query with cache-hit only 5ms.