Enterprise Digital Infrastructure – Assignment #2 – 22 & 28 April 2020

HTTP in action

Preliminary steps

Close **all** applications that could generate any network traffic.

Before each experiment

Remove the cookies and clear the cache and navigation history of the browser.

Open Wireshark and filter the traffic from/to your device.
Open a terminal on your device.
Open another terminal on your device and connect via **`ssh`** to your account on the server of my lab (i.e., **`pel.unipv.it`**, **`eva.unipv.it`**, **`poldo.unipv.it`**). Open the log files stored by the Apache Web server (**`tail -f /var/log/httpd/access_log`** or **`error_log`**).

1. Basic experiments
    a. Open a TCP connection on port 80 of the Web server where you have an account; use the **`telnet`** command to download **`/testfile.txt`** and **`/EDI`** using HTTP/1.0.
    b. Repeat the experiments using HTTP/1.1 (with/without header lines in the HTTP request message).
    c. Repeat experiments a. and b. using **`openssl`** on port 443.
    d. Download **`/EDI`** using Google Chrome and/or Mozilla Firefox browsers.

Analyze the HTTP requests and responses and discuss the differences. Analyze and discuss the impact of the browsers.

2. Open a TCP connection on port 80 or 443 of web servers hosting institutional/commercial websites of your choice and download the home page.

Did you succeed to download the home page? Analyze the headers sent along with the HTTP requests and responses. Does the web server provide details about itself (e.g., server used)? Did the server send any cookie? How many?

3. Measure the Page Load Times spent to download Web pages from the commercial/institutional websites including those used at point 2. (e.g., under Mozilla Firefox: Options -> Web Developer -> Network or under Google Chrome: More tools -> Developer tools).

Analyze and discuss the results. How many and what type of objects are embedded in the pages? Are the objects stored on the same Web server? How many TPC connections does the browser open? Did you get/send any cookie? How many? Did you get any third party cookie? When do they expire? What happens if you disable all cookies? Or third-party cookies? Can you forge the string associated with the User-Agent or Referer header field names? Did you notice any effect?

4. Test the caching policies of different websites by exploring the headers in the HTTP request and response messages, the content of the browser cache (in the navigation bar of Mozilla Firefox **`about:cache`** or Application tab in the Developer tool of Google Chrome) and the usage of Conditional requests .

Analyze and discuss the impact of the caching mechanisms on the network traffic and on the Page Load Times.

5. Modify the configuration settings of the browser with respect to parallelism (in the navigation bar of Mozilla Firefox **`about:config`**) and test the new configurations on different websites [remove cookies and the cache content before each download].

Analyze and discuss the impact of these changes on the network traffic and on the Page Load Times.

6. Open commercial/institutional websites using HTTP/1.1 and HTTP/2 using a browser and **`nghttp`**.

Analyze and discuss the differences and the impact on the Page Load Times.

7. Use the **`ab`** – Apache http server benchmarking tool and **`h2load`** to generate HTTP traffic towards different commercial/institutional websites using different types of configurations (**`man ab`** and **`h2load`** for details).

Analyze and discuss the results obtained under different configurations and for different websites.

The report should present and discuss the main findings of the experiments and highlight expected/unexpected behaviors.

The report – a **PDF file** not exceeding **two pages** – is due **May 4, 2020 23:59PM**.

The report should NOT include any screenshot.

Results of experiments referring to points 1 & 2 are not to be discussed in the report.