



Communication Networks

Prof. Laurent Vanbever

Solution: Exercise 11 – HTTP & Video Streaming

11.1 Setting up a video streaming service (Exam Style Question)

In this task, you will set up a brand new video streaming service.

- a) You first create a website (<http://www.yourvideo.ch>) where visitors can buy subscriptions. To improve the performance of your website, you want to make it small enough that it fits in one single IPv4 packet. Assuming the maximum size of an IP packet is 1500 bytes, how many bytes are remaining for the application layer at most? Explain your answer.

Solution: HTTP usually runs over TCP, but could also be run over UDP. Therefore, we provide the answer both for TCP and UDP. Within the available 1500 bytes, we also need to pack the header for IP and TCP/UDP. The IP header has a minimum length of 20 bytes. The TCP header also has a minimum length of 20 bytes, whereas the UDP header only needs 8 bytes.

- $1500 - 20 - 20 = 1460$ bytes (IP, TCP)
- $1500 - 20 - 8 = 1472$ bytes (IP, UDP)

- b) Can you use all the available bytes that you computed above for the source code of the website? Explain your answer.

Solution: No, the HTTP headers also need some space.

- c) After a user bought a subscription, you want to show her the list of all available videos. As you know, the HTTP protocol is stateless. Which technique can you use to maintain state about whether a user is logged in or not? Mention the name of the technique, where the state is stored and how your web server knows whether a user is logged in upon receiving a HTTP request from her.

Solution: To maintain user-state, we use a technique called “cookies”. These cookies are stored at the client. With every HTTP request, the client sends the cookie within the header to the server.

- d) Unfortunately, you now realize that the website and the 10 images that it contains do not fit in a single IP packet and you therefore need to think of other ways to improve the page load speed. You come up with the following solution: Instead of storing everything on the same server, you add 10 additional servers to your data center and store each image on one of them (i.e., image 1 is on 123.0.0.101/image1.jpg, image 2 is on 123.0.0.102/image2.jpg and so on). Each server has a high bandwidth Internet connection.

Compared to the initial solution (everything on the same server), does this new approach have an impact on the number of established connections if the client uses HTTP 1.0 (also referred to as naive HTTP in the lecture) or if the client uses HTTP 1.1? Explain your answers.

Solution: When HTTP 1.0 is used, the additional servers have no impact on the number of connections as the protocol anyway requires you to establish a new connection for each element (e.g., image).

HTTP 1.1 uses a persistent connection that is setup once and can be used to fetch all elements. Hence, distributing the elements across multiple servers will lead to an increase in the number of connections as they cannot be transmitted over that one common connection.

Now that you've finished your website, you can finally start adding videos that your customers can stream.

- e) You heard that your competition provides the videos in multiple resolutions and bit-rates. Is this a good idea for your case too, even if you do not use a CDN? Explain your answer.

Solution: Yes, it is a good idea as your customers can request the video chunks at the resolution and bit-rate that is best for their situation: network conditions, screen size, etc. This has nothing to do with the use of a CDN. A CDN just helps you bring your content closer to your customer independent of the resolutions and bit-rates offered.

- f) Assuming you decided to also provide the videos in different resolutions and bit-rates. Explain how the player knows which versions are available as well as how and when it selects the best version?

Solution: The player knows which versions are available thanks to the "Manifest". It selects the version based on the buffer occupancy and will reconsider its decision for every chunk anew.

- g) One day customers tell you they often observe "rebuffering". Explain what this is and describe one situation in which it can occur.

Solution: Rebuffering happens when the buffer of the player is empty and there are no more chunks locally available to be shown. This happens often when the connection is too slow and/or the player request too large chunks.