

CS 447 : Networks and Data Communications Programming Assignment #01

Total Points: 150

Assigned Date : Thursday, September 13, 2018

Due Date : Thursday, September 27, 2018 @ 12:29:59 p.m.

Correction (09/20/18) : See correction to Technical Requirements #10

Overview

Your first programming assignment is to **implement a basic client/server application** using the socket interface. There are several objectives of this assignment. These are:

- to get yourself familiarized working with the socket programming basics;
- to understand the ordering of the socket interface primitives;
- to get you exposed to linux system calls (if you already haven't);
- to gain a basic understanding of network protocols; and
- to set yourself up for the rest of the course.



Back Story

Captain Haddock is having all sorts of issues with the carrier pigeon system he uses when he is on sea voyages. Often, he is finding out that the messages he sends do not reach their intended recipients *in their original form* because the pigeons like to peck on the special corn-based paper Haddock uses. On the receiving side, Haddock has observed that, most times, the messages he receives are *late*. He also does not enjoy writing messages on behalf of his crewmates; right now, none of his crewmates can afford to own their own carrier pigeons.

Professor Calculus has a brilliant solution. “Let me introduce you to email, Haddock”, said Calculus. He promises Haddock that this new email solution will provide reliable delivery when sending emails, be capable of supporting more than one sender at the same time, and to have super speedy (although sometimes unreliable) delivery when receiving.

Technical Requirements

- For this assignment, you will need to write a client-server application to support Calculus’s email system.

2. Your sender → server interaction should follow the SMTP protocol, and must support the following SMTP commands: **HELO**, **MAIL FROM**, **RCPT TO**, **DATA**, and **QUIT**. Read Section 4.1 of the SMTP RFC found at <https://tools.ietf.org/html/rfc2821> for exact command specifications.
 - Closely related to the SMTP commands are the corresponding reply codes. Read Section 4.2.2 of the RFC #2821, select appropriate reply codes, and implement them. I anticipate you to find at least 5-6 reply codes necessary for your implementation. Justify your choice of reply codes in your report.
3. Email addresses should have the typical email format, i.e., should include the @ sign. For the purpose of simplicity in parsing, assume all send and receive email addresses are from the 447f18.edu domain. E.g. calculus@447f18.edu
4. Emails are written at the server, not at the sender. In other words, your SMTP interaction should not be a file transfer.
5. The sender → server interaction should run over TCP.
6. The sender → server interaction should support multi-threading; more than one client should be capable of sending emails at the same time.
7. Your server → receiver interaction should follow the HTTP/1.1 protocol and implement it's **GET** method. Read Section 4.3.1 of the HTTP/1.1 RFC found at <https://tools.ietf.org/html/rfc7231> for exact specifications.
 - The corresponding HTTP/1.1 reply codes (referred to as status codes in the RFC) are found in Section 6.1 of RFC #7231. Bare minimum, you must implement reply codes **200**, **400**, and **404**. Again, make sure to justify your use and choice of reply codes (and any additional reply codes you decided to use) in your report.
8. The server → receiver interaction should run over UDP. Multi-threading is not required for this interaction.
9. To simplify the implementation, we'll use a slightly modified **GET** request and response format for this project. Please find more information under the Logistics section below.
10. ~~Your client programs—both the sender and the receiver—should prompt to user to enter appropriate information, rather than making them type in the correct protocol commands; your programs should handle the correct protocol interaction internally. In other words, assume the user of your program is only interested in writing and retrieving emails, and have no knowledge of protocols and how they work. You, as the developer on the other hand, is well-versed in protocols, which would be reflected on your code. If this is not clear, make sure to talk to the instructor immediately.~~

Please ignore this requirement. Corrected on: 09/20/18 @ 2:10pm

Logistics



IP addresses/hostnames and port numbers should not be hard coded.

- Depending on how you decide to implement your server(s), your server executable will accept one (or two) command line argument to denote the corresponding listening port(s) as follows:
./server <tcp-listen-port> (<udp-listen-port>)
- Your client executable will accept two command line arguments as follows:
./client <server-hostname> <server-port>.

Your may assume that your client knows the server hostname/ip address and the correct port to connect to.

2. Your code must compile and run on a typical linux computer.
3. I will test at least 2 simultaneous client connections, once as two senders, and once as one sender + one receiver. Bare minimum, make sure this is covered in your testing plans.
4. All clients should exit gracefully. Server process is permitted to be forcefully killed.
5. Use the following file management strategy:
 - When your SMTP server fires up for the first time, make it programmatically (not manually) create a folder named **db** to store emails.
 - Programmatically create a new subfolder inside **db** for each new recipient that's mentioned in the **RCPT TO** command above.
 - Store emails as sequentially numbered files. E.g. first email to Haddock will be stored as `/db/haddock/001.email`
 - When a new receiver fires up for the first time, programmatically (not manually) create a folder under receiver's name to store retrieved emails.
6. Here's a sample .email file. Note the timestamp added by the server.

Answer

```
Date: Fri, 07 Sep 2018 13:04:20 -0500
From: <tintin@447f18.edu>
To: <haddock@447f18.edu>
Subject: The Last Unicorn

Dear Haddock,

Glad to hear that you found the last Unicorn. We are looking
forward to your safe return.

Yours truly,
Tintin and Snowy.
```



7. Here's a sample GET request.

Answer

```
GET /db/haddock/ HTTP/1.1
Host: <server-host-name>
Count: 1
```

Note: **Count** denotes the number of emails to download. Additionally, you may find reading the HTTP/1.1 Message Syntax and Routing RFC found at <https://tools.ietf.org/html/rfc7230> helpful for understanding.

8. Here's a corresponding successful GET response from the server. Store the response as a .txt file under the receiver's folder.

Answer

```
HTTP/1.1 200 OK
Server: <server-hostname>
```

Last-Modified: Fri, 07 Sep 2018 18:43:20 -0500
Count: 1
Content-Type: text/plain
Message: 1

Date: Fri, 07 Sep 2018 13:04:20 -0500
From: <tintin@447ss15.edu>
To: <haddock@447ss15.edu>
Subject: The Last Unicorn

Dear Haddock,

Glad to hear that you found the last Unicorn. We are looking forward to your safe return.

Yours truly,
Tintin and Snowy.

9. At the end of your implementation, you should be able to:
- Compile and run your code on a typical linux machine.
 - Run your server program(s) first.
 - Run one or more clients to connect to the server to send emails.
 - Run a receiver to retrieve email.
 - Exit the client(s) gracefully.

Instructions

- This is an individual assignment. **Do your own work.**
- **Start early!! Take backups of your code often!!**. Use of a version control software is highly recommended!
- Make sure to test your program properly before your final submission. It is **highly** recommended to test build and run your submission on the home server, **home.cs.siu.edu**.
- You may use any programming language of your choice out of C, C++, Java, Python or Go as these are all available on the home server. However, you **must** make sure that your code compiles and runs on a typical Linux machine. Absolutely **DO NOT** include executables with your submissions.
- A **Makefile** is mandatory. Whether or not your program needs to be compiled, have it echo instructions to run the program.
- Follow a good coding standard. Use the Google C++ coding standard found here <http://goo.gl/1rC1o>, if you don't already follow one.
- The report part of your solution must be produced using a word processor. \LaTeX is highly recommended but not a requirement.
- Your final report should be in **PDF** format. No exceptions.
- Any figures, graphs, plots, etc., should also be produced using appropriate computer applications. If using \LaTeX , the pgfplots package is very useful for making all sorts of graphs.
- The due date of this assignment is **Thursday, September 27, 2018 @ 12:29:59 p.m.** A dropbox will be opened for submission on Moodle.

Deliverables

A complete solution comprises of:

- A short report (max 5 pages) of the design and implementation of your system. Your report should include the followings:
 - Introduction
 - Design choices and protocol/reply codes used.
 - The output of a sample run (including screenshots where applicable).
 - Summary and Issues encountered (if applicable).
- A short README file with compilation and run instructions.
- A makefile to compile your code especially if it involves compiling multiple executables with flag options..
- A compressed tarball of the directory containing your source code, report, README, and makefile. **Absolutely do not** include executables, folders created by your programs, your version control repositories, or your test emails in this tarball. To create a compressed tarball of the directory `source`, use the following command: `tar -zcvf siue-id-pr1.tar.gz source/`.
e.g. `tar -zcvf tgame-pr1.tar.gz PR01/`.
- File formatting standards (**PDF, README, .email, .txt, .tar.gz**) will be strictly monitored and is subject to penalties.

Collaborating on ideas or answering questions is always encouraged. Most times, I find that you learn a lot from your peers. However, do not share/copy/duplicate code from others. If you use code found online, remember to site their source in your report. Issues related to academic integrity and plagiarism have **ZERO** tolerance.

Useful Resources

- Linux Man pages – found in all linux distributions
- Beej's Guide to Network Programming – A pretty thorough free online tutorial on basic network programming http://beej.us/guide/bgnet/output/print/bgnet_USLetter.pdf
- Simple Mail Transfer Protocol RFC #2821 <https://tools.ietf.org/html/rfc2821>
- Hypertext Transfer Protocol (HTTP/1.1): Semantics and Content RFC #7231 <https://tools.ietf.org/html/rfc7231>
- Hypertext Transfer Protocol (HTTP/1.1): Message Syntax and Routing RFC #7230 <https://tools.ietf.org/html/rfc2821>