

Network Services

Hands-On UNIX System Administration DeCal

Lab 8 — 14 March 2011

Due 28 March at 6:10 PM

1 Final project

Final project proposals are **due in hard copy at the start of next class**, 28 March (after spring break). Please arrange to stick around after the lecture portion of class is complete that day — I'll be meeting with each group to discuss your proposals during lab. Your proposal should include the following:

- **Elevator pitch.** Convince me that your project is awesome. What will it do, and why should I care? What problem will it solve, or how will it make my life as a UNIX systems administrator easier?
- **Summary.** How do you plan to implement your project? Include requests for additional resources — e.g., additional VMs, different operating systems, custom DNS records, ports forwarded, etc.; I'll attempt to honor any reasonable requests.
- **Timeline.** List major milestones necessary to complete your project, how long each should take, and by when you plan to have finished each milestone. You'll be presenting your project to the class on 25 April, so plan accordingly.

Bonus points¹ if your proposal looks professional (hint: if you're a CS or maths major, you really should learn how to use \LaTeX).

2 DNS queries

If you're doing this exercise on an Instructional server, note that `host` and `dig` are installed but are not in your `PATH`. If you're using your virtual server, you should install the `host` and `dnsutils` packages.

1. Use `host` to look up the MX record for Berkeley.EDU. Show the syntax you used.
2. What's an NS record, and what are Berkeley.EDU's NS records? Where are the servers Berkeley's NS records point to actually located, and why might IST have set things up that way?

¹Well, not really — this is a pass/not-pass class, after all. But you *do* want to be awesome, don't you?

3. Use `dig` to perform the same queries as in questions (1) and (2), and show the syntax you used. Then run `dig` with no arguments — what do you see in the “ANSWER SECTION”? What are these servers?

3 HTTP

(You’ll want to install `netcat` on your virtual server before continuing.) In this part, you’ll be writing a simple shell script that downloads the contents of an arbitrary URL, saving the result to the file “output.” Assume that the URL your script is passed is in the format

```
http://www.example.com/path/to/document
```

(don’t worry about escapes or special characters). Here’s what your script should do:

1. Use `cut` to extract the domain name (`www.example.com`) and the path (`/path/to/...`). Bear in mind that a valid path can contain an arbitrary number of slashes.
2. Use `nc` to connect to the server on port 80. (You can pipe input into `netcat`’s `stdin`.) An HTTP request looks like this:

```
GET /path/to/document HTTP/1.1
Host: www.example.com
(newline)
```

Use `echo` and the “enable interpretation of backslash escapes” option — your HTTP request, on one line, will look something like

```
GET /path/to/document HTTP/1.1\nHost: www.example.com\n\n
```

(note the *two* concluding newlines).

3. Pipe this input into a read loop. We’re not interested in the HTTP headers, which are terminated by a blank line, so you should discard lines until you find an empty one. (To check if you’ve found an empty line, test whether it’s equal to the output of `echo -e '\r'` — try comparing it to `$(\r)`.) Once you’ve found a blank line, `echo` every remaining line.
4. Save the output of this read loop to a file called “output.”

Before you begin, try out `nc` and `echo -e` on their own and make sure you understand how they work.

Extra for Experts™!

Once you’re done, if you’re feeling adventurous, you might try...

- The filename “output” is not very descriptive. Make your script save its output somewhere more useful.

- Look at the Wikipedia article on HTTP status codes. The first HTTP header your script receives includes a status code — common ones include “200 OK,” “301 Moved Permanently,” “404 Not Found,” and so on. Build in support for a few of these status codes — maybe follow redirects (look at the Location: header), or print a warning when you encounter a 404.
- If you’re a shell-scripting ninja, you finished the entire lab in ten minutes flat, and you’re bored, read the appropriate RFCs and write a bash-based IRC client.