

Network Programming Tools

Richard Ha

Agenda

Administrivia

Proxy Lab Recap

HTTP Overview

Networking Tools

Debugging Tips & Techniques

Administrivia

Proxy lab due date extension!

New due date is Aug. 6, 2015 @ 11:59pm

Will cover changes to handout in a few minutes.

New due date overlaps with final exam dates so be sure to finish the lab early.

Administrivia

Proxy lab due date extension!

New due date is Aug. 6, 2015 @ 11:59pm

Will cover changes to handout in a few minutes.

New due date overlaps with final exam dates so be sure to finish the lab early.

Lab Grading

Shell lab grades are done, please read annotated code.

Malloc grades soon to follow. (Heapchecker, style)

Administrivia

Final Exam

First round final will take place Thursday+Friday.

Recitation for final exam review will be held on Tuesday.

If you want a question answered / topic covered in the final exam review, send an e-mail titled “FINAL RECITATION QUESTION” to the staff list.

We will pick a list of questions to cover for the recitation, as much as we can fit.

Agenda

Administrivia

Proxy Lab Recap

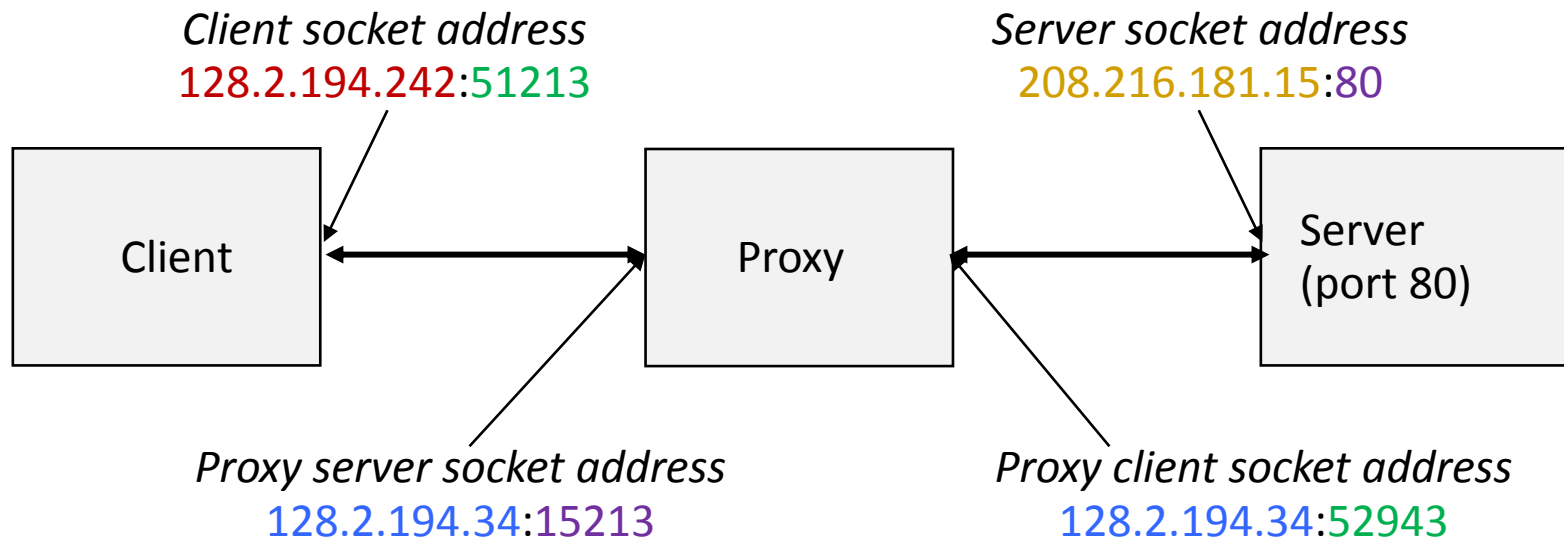
HTTP Overview

Networking Tools

Debugging Tips & Techniques

What is a Proxy?

A proxy is a go-between a client and a server



It is both a server AND a client!

Proxy Requirements

HTTP/1.0

Make and receive HTTP/1.0 requests & responses.

Convert all HTTP/1.1 requests to HTTP/1.0

Proxy Requirements

HTTP/1.0

Make and receive HTTP/1.0 requests & responses.

Convert all HTTP/1.1 requests to HTTP/1.0

Concurrency

Needs to be able to handle multiple clients at once.

Proxy Requirements

HTTP/1.0

Make and receive HTTP/1.0 requests & responses.

Convert all HTTP/1.1 requests to HTTP/1.0

Concurrency

Needs to be able to handle multiple clients at once.

Caching

Needs to temporarily cache web objects in a shared cache amongst threads.

Proxy Requirements

Robustness

Malformed requests and responses are a fact of real life, and not all of them are intentional or malicious.

Your proxy needs to be able to anticipate and deal with malformed requests and responses without crashing, segfaulting, erroring, exiting, or otherwise prematurely terminating.

Servers are expected to run (mostly) forever without human intervention. Your proxy is no exception.

Agenda

Administrivia

Proxy Lab Recap

HTTP Overview

Networking Tools

Debugging Tips & Techniques

Structure of a HTTP Request

HTTP requests are divided into two parts:

The request line

The headers

\r\n

Structure of a HTTP Request

Request Line:

<Request Method> <Request-URI> <HTTP-Version>\r\n

There are nine different HTTP methods. We only care about GET.

Examples:

GET / HTTP/1.0\r\n

GET <http://cs.cmu.edu/~213> HTTP/1.0\r\n

Structure of a HTTP Request

Headers:

The headers section consists of a list of header name and value pairs.

<HeaderName>: <HeaderValue>\r\n

Example:

Host: cs.cmu.edu:80\r\n

User-Agent: Mozilla/5.0 (X11; Linux x86_64; rv:10.0.3)

Gecko/20120305 Firefox/10.0.3\r\n

Structure of a HTTP Request

Headers:

The change in the lab handout was related to the Accept and Accept-Encoding headers.

Handout originally stated you should insert/modify the headers.

Now, just treat those two headers like most other headers and don't modify/change them. Just pass them as-is if they're present.

Structure of a HTTP Response

HTTP Responses are split into three sections:

Status Line

Headers

\r\n

Entity-Body

Structure of a HTTP Response

Status Line

<HTTP-Version> <Status Code> <Reason>\r\n

Example:

HTTP/1.0 200 OK\r\n

HTTP/1.0 302 Moved Permanently\r\n

HTTP/1.0 404 Not Found\r\n

Structure of a HTTP Response

Headers

Same format as request headers.

If there is an entity body that comes with the response (e.g. HTML page, image, etc.), then there will be a Content-Length header.

Structure of a HTTP Response

Content Body

Just a stream of bytes in the designated encoding.

Could be text, could be binary data, could be compressed, could be not compressed, etc.

Agenda

Administrivia

Proxy Lab Recap

HTTP Overview

Networking Tools

Debugging Tips & Techniques

Networking Tools

Just knowing the lab specifications and the HTTP/1.0 spec is not enough to do a lab.

As with any development effort, you need debugging and diagnostic tools to help you catch bugs and errors and identify where the problem is occurring.

Thankfully, Shark machines have plenty of tools to use.

Also included some in the lab handout.

Networking Tools

- Tiny server
- cURL / wget
- Telnet
- Netcat (nc)
- Valgrind
- Your browser(s)
- And many more!

Tiny server

There is a tiny HTTP server included in your proxy lab handout.

```
rkh@lemonshark:/tmp/proxylab-handout$ ls
csapp.c  driver.sh      Makefile      port-for-user.pl  README
csapp.h  free-port.sh  nop-server.py proxy.c          tiny
```


Tiny server

There is a tiny HTTP server included in your proxy lab handout.

It is a 200-line web server written in C.

Invoked with: `./tiny <port>`

Serves everything inside `./`

Comes with an example adder CGI script

Consult the README for details.

cURL / wget

Command-line programs that fetch a given web URL.

e.g. `$ curl http://lemonshark.ics.cs.cmu.edu:15213/godzilla.jpg`

will fetch `godzilla.jpg` from the tiny server.

`curl` has a `--proxy` option that lets you specify a HTTP proxy that it can make requests through

telnet

Client application for the TELNET protocol

Telnet client can conveniently be used to communicate and interact with text-protocol servers like HTTP and IRC servers.

```
rkh@lemonshark:/tmp/proxylab-handout$ telnet lemonshark.ics.cs.cmu.edu 15213
Trying 128.2.220.17...
Connected to lemonshark.ics.cs.cmu.edu (128.2.220.17).
Escape character is '^]'.
_
```

netcat (nc)

Like the cat program except with networking capabilities.

Can pipe strings or files to a target server and port.

```
rkh@lemonshark:/tmp/proxylab-handout$ echo -e "GET /home.html HTTP/1.0\r\n\r\n" | nc lemonshark.ics.cs.cmu.edu 15213
HTTP/1.0 200 OK
Server: Tiny Web Server
Connection: close
Content-length: 120
Content-type: text/html

<html>
<head><title>test</title></head>
<body>

Dave O'Hallaron
</body>
</html>
```

netcat (nc)

Can also invoke as a listen server on a given port.

e.g. `$ nc -l 15213`

Easily usable in shell scripts to run tests.

Very versatile tool.

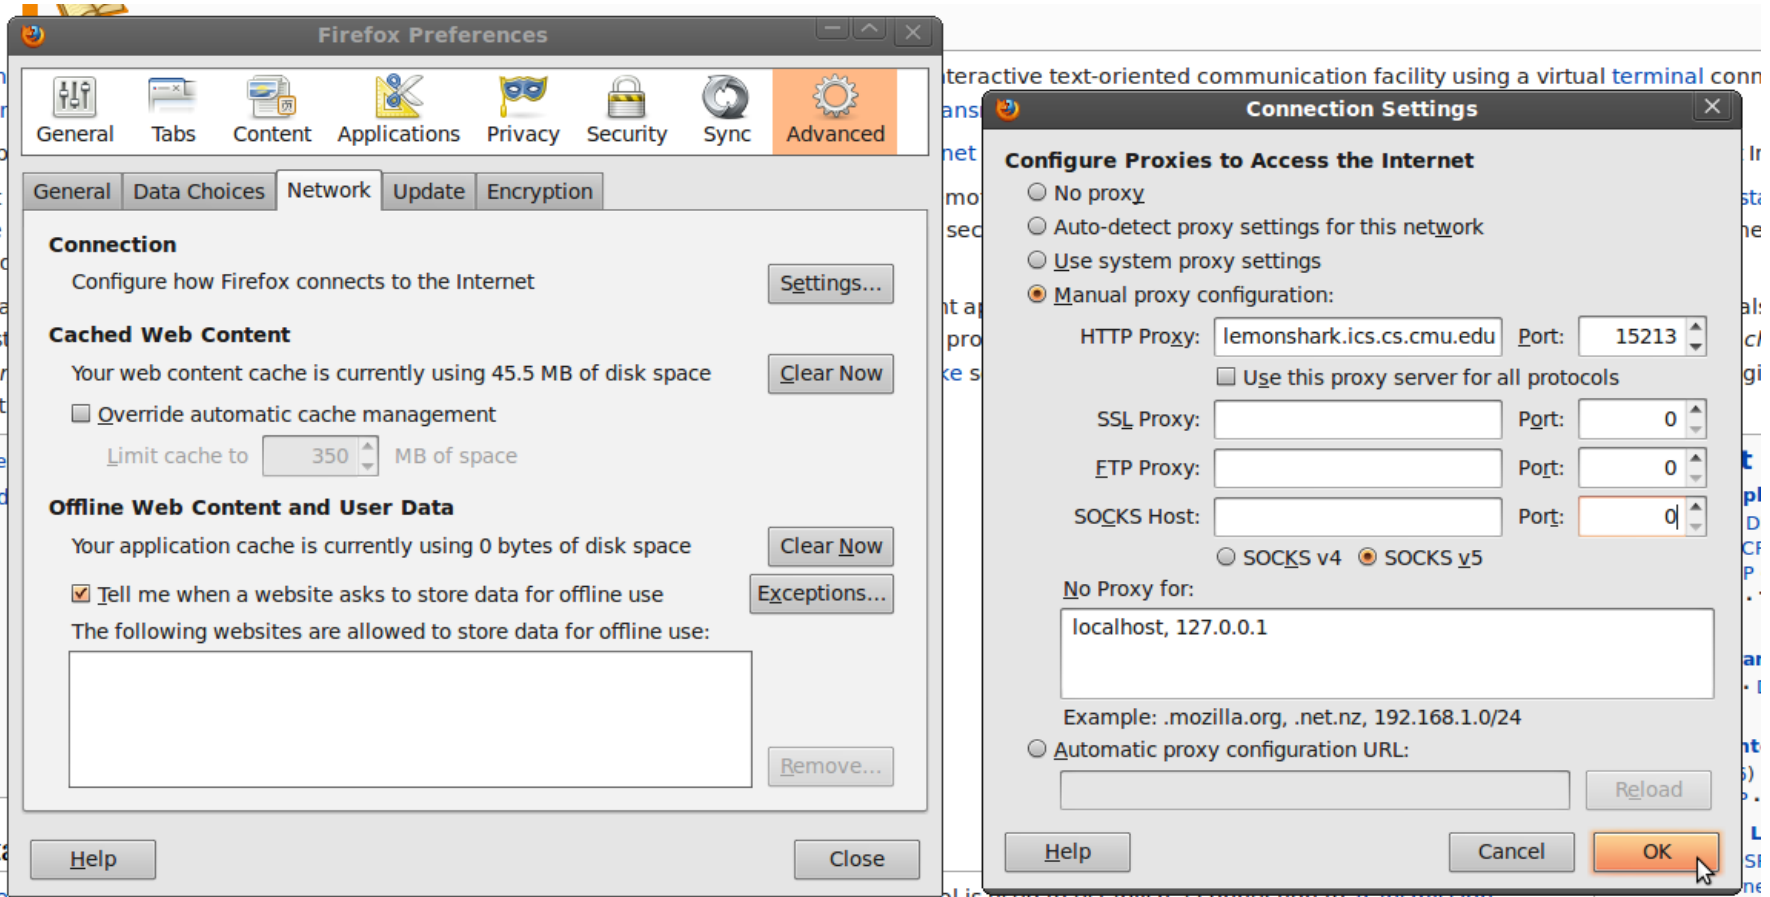
Valgrind

Very powerful memory debugging tool.

Useful in spotting memory leaks, illegal memory accesses, etc.

Your browser

Most modern browsers have a HTTP proxy option you can specify. Very useful in seeing if your proxy actually works with real-life usage.



Your browser

Most modern browsers also come with a slew of networking analysis and debugging tools. You can log and inspect the requests and responses your browser makes and receives to and from your proxy and/or other web servers.

Try Ctrl+Shift+I for developer tools in Chrome or Ctrl+Shift+K for the web console in Firefox.

The screenshot displays the browser's developer tools interface. The 'Network' tab is active, showing a list of requests. The 'Inspect Network Request' panel is open, showing details for a request to `http://reddit.com/`. The request method is GET, and the status code is HTTP/1.1 301 Moved Permanently. The request headers include:

- User-Agent: Mozilla/5.0 (Windows NT 6.1; WOW64; rv:39.0) Gecko/20100101 Firefox/39.0
- Host: reddit.com
- Connection: keep-alive
- Accept-Language: en-US,en;q=0.5
- Accept-Encoding: gzip, deflate
- Accept: text/html,application/xhtml+xml,application/xml;q=0.9,*/*;q=0.8

The response headers include:

- Transfer-Encoding: chunked
- Server: cloudflare-nginx
- Location: http://www.reddit.com/
- Date: Sun, 02 Aug 2015 18:00:38 GMT
- Connection: keep-alive
- cf-ray: 20fbaf7d3040880-IAD

A received cookie is also shown: `__cfduid: d425eae017ed0ab77b93e3bd9737834b1438538438`. The background shows the browser's console with a list of network requests, including the main page request and various static assets.

Agenda

Administrivia

Proxy Lab Recap

HTTP Overview

Networking Tools

Debugging Tips & Techniques

Debugging Tips & Techniques

- Always go over your proxy's output with a fine-toothed comb.
- Use nc or tiny to display exactly what your proxy is sending.
- Try to avoid blindly using string functions like strcpy() and strlen() when handling data. Use strncpy() and strlen() instead.
- Try to use large port numbers (e.g. 15213-65535) when running your proxy and tiny server(s) locally on the Shark machines. Smaller numbers might be reserved and already be in use by something else.

Debugging Tips & Techniques

- Make sure your underlying cache data structure implementation is correct before worrying about synchronisation. Maybe write a separate test program that includes/uses your cache library.
- Consider making your own testing framework. Save valid requests/responses your proxy breaks on to test on the proxy again later.
- Remember to include a method to monitor program flow while debugging race conditions and deadlocks.

Questions?

GOOD LUCK