

COMP312-09A Communications and Systems Software

Lecture 4 – DNS, HTTP, Apps
Pages 588 to 611 – Tanenbaum 4th Ed.

Matthew Luckie
mluckie@cs.waikato.ac.nz



© THE UNIVERSITY OF WAIKATO • TE WHARE WAIKATO O WAIKATO

DNS

- We know a fair amount of DNS from COMP202-08B
 - Root servers
 - DNS uses UDP by default
 - Top level domains (TLDs), second level domains (SLDs)
 - Iterative protocol
 - TTLs
- This lecture will add a couple of details to that
- Pages 579-588 in Tanenbaum 4th Ed.

© THE UNIVERSITY OF WAIKATO • TE WHARE WAIKATO O WAIKATO

© March 2009

2

DNS record types

- **A:** Address.
 - Maps name to IP address
- **PTR:** Pointer.
 - Maps IP address to name
- **MX:** Mail exchanger.
 - Host willing to accept mail
- **NS:** Name server
 - Who will answer DNS requests for specified domain
- **TXT:** Text record
 - Arbitrary text
- **CNAME:** Canonical name
 - Maps alias to a single true name

© THE UNIVERSITY OF WAIKATO • TE WHARE WAIKATO O WAIKATO

© March 2009

3

- Client has name: **www.amazon.com**

1. Ask root server: what is the IP address for **www.amazon.com**?

- reply: I don't know. Try
.com NS a.gtld-servers.net
a.gtld-servers.net A 192.5.6.30
or 15 other servers they also specify

2. Ask .com name server: IP address for **www.amazon.com**?

- reply: I don't know
amazon.com NS udns1.ultradns.net
amazon.com NS udns2.ultradns.net
udns1.ultradns.net A 204.69.234.1
udns2.ultradns.net A 204.74.101.1

3. Ask amazon.com NS: IP address for **www.amazon.com**?

- Reply: ask **www.amazon.com NS ns-923.amazon.com**
ns-923.amazon.com A 72.21.204.20

4. Ask ns-923.amazon.com NS: IP address for **www.amazon.com**:

- Reply: **www.amazon.com A 72.21.210.250**

© THE UNIVERSITY OF WAIKATO • TE WHARE WAIKATO O WAIKATO

© March 2009

4

You can play along at home

- `dig www.amazon.com @a.root-servers.net`
- `dig www.amazon.com @a.gtld-servers.net`
- `dig www.amazon.com @udns1.ultradns.net`
- `dig www.amazon.com @ns-923.amazon.com`

Correct at March 9th 2009

• Use:

- `dig <name to look up> <optional server to query>`
- If server to query is not specified, answer will come from name server your machine has configured
- Answer is likely to be cached from previous lookup if querying local name server

© THE UNIVERSITY OF WAIKATO • TE WHARE WAIKATO O WAIKATO

© March 2009

5

```
[mluckie@sorcererer mjl]$ dig www.google.com
```

Answer comes from local name server

```
;; QUESTION SECTION:
;www.google.com.                IN      A

;; ANSWER SECTION:
www.google.com.                568515  IN      CNAME    www.google.com.
www.l.google.com.              280    IN      A          2085.171.104
www.l.google.com.              280    IN      A          2085.171.147
www.l.google.com.              280    IN      A          2085.171.99
www.l.google.com.              280    IN      A          2085.171.103

;; AUTHORITY SECTION:
google.com.                    83534   IN      NS          ns.google.com.
google.com.                    83534   IN      NS          ns.google.com.
google.com.                    83534   IN      NS          ns.google.com.
google.com.                    83534   IN      NS          ns.google.com.

;; ADDITIONAL SECTION:
ns1.google.com.                163922  IN      A          21639.32.10
ns2.google.com.                163922  IN      A          21639.34.10
ns3.google.com.                163922  IN      A          21639.36.10
ns4.google.com.                163922  IN      A          21639.38.10
```

© THE UNIVERSITY OF WAIKATO • TE WHARE WAIKATO O WAIKATO

© March 2009

6

HTTP

- We know a fair amount of HTTP from COMP202-08B
 - Hypertext Transfer Protocol
 - Protocol spoken by web browsers
 - Assignment 2 was to implement a basic web server
 - Uses TCP
 - Reliable, byte-stream, connection oriented transport protocol
 - Well known port 80
 - request/response protocol
- Pages 611 to 662 Tanenbaum 4th Ed.
 - Though 652 to 662 are the most important
 - Won't be considering HTML etc, there are whole other courses on that

HTTP

- `http://www.wand.net.nz/~mluckie/`
- URL consists of
 - Protocol (`http`)
 - name of machine where resource is found (www.wand.net.nz)
 - File (`/~mluckie/`)

HTTP

- `http://www.wand.net.nz/~mluckie/`
- Protocol steps follow from URL
 1. Parse the URL for its components
 2. Ask DNS for IP address of `www.wand.net.nz`
 3. DNS replies with `130.217.250.15`
 4. Make a TCP connect to port 80 on `130.217.250.15`
 5. Request `/~mluckie/`
 6. Server sends file associated with this resource (`index.html`)
 7. Close TCP connection
 8. Display page
 9. Fetch all images and media in page, using steps 1-8

HTTP GET

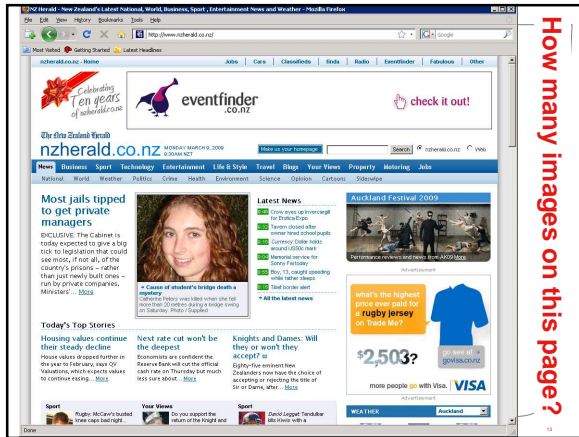
```
GET /~mluckie/ HTTP/1.1
Host: www.wand.net.nz
User-Agent: Mozilla/5.0 (X11; U; FreeBSD i386; rv:9.0.5)
Accept: text/html,*/*;q=0.8
Accept-Language: en-us,en;q=0.5
Accept-Encoding: gzip,deflate
Accept-Charset: ISO-8859-1,utf-8;q=0.7,*;q=0.7
Keep-Alive: 300
Connection: keep-alive
Pragma: no-cache
Cache-Control: no-cache
```

HTTP/1.1 200 OK

```
HTTP/1.1 200 OK
Date: Sun, 08 Mar 2009 20:38:04 GMT
Server: Apache/2.0.54 (Debian GNU/Linux)
Last-Modified: Mon, 15 Dec 2008 20:52:54 GMT
ETag: "6c3810b-965-6dbb580"
Accept-Ranges: bytes
Content-Length: 2405
Keep-Alive: timeout=15, max=99
Connection: Keep-Alive
Content-Type: text/html; charset=ISO-8859-1
```

Main Problem with HTTP 1.0

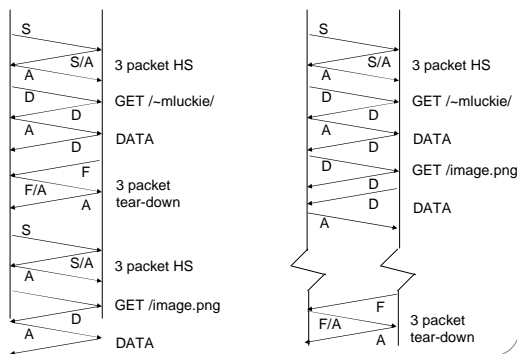
- Protocol: each request uses a separate TCP connection
 - This worked fine when a web page was marked-up text
- Problem:
 - These days web pages contain lots of images in them
 - Each image requires a separate TCP connection
 - 3-way handshake
 - Data
 - 3 or 4-packet disconnect
 - Overhead is magnified for people living far from most of the web's content. i.e. NZ to US round-trip-times.
 - TCP connections would rarely get out of slow-start
 - i.e. would go slow, rarely using the available capacity in the network



HTTP/1.1

- Persistent connections
 - Establish TCP connection
 - Send a request, get a response
 - Send further requests and get further responses
 - Good for fetching images to be displayed inline
 - Keep TCP connection open for a short period (about a minute)
 - Good for fetching other web pages from the same server when the user clicks a link
- Pipelining
 - Send multiple requests simultaneously over single TCP connection


HTTP 1.0 vs 1.1



Further Performance Enhancements

- Browser Caching
 - HTTP headers provide information on when a file was last modified, information on if the data should be cached, etc
 - HTTP protocol support for fetching If-Modified-Since
- Proxy Server Caching
 - Server whose job it is to cache files
 - Multiple clients can use, some benefit to be had in addition to browser caching if clients tend to visit the same sites
 - Useful for authentication and billing
 - Client has shorter TCP retransmission domain

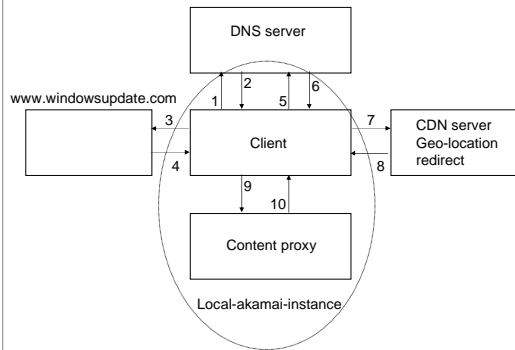
TCP retransmission domain



Content Delivery Networks

- Company pays ISPs to host servers which replicate large websites
 - ISPs will happily take this service because it makes them money, gives their users better user-experiences
 - Content providers (windows update, etc) pay CDN to host their website
 - CDN needs to have some way to map a user to their closest replica, i.e. IP to geo-location service
- Akamai largest CDN example

Example



© THE UNIVERSITY OF WARWICK • TE WAIWAE WAIWANGA O WAIWAKATO

19 March 2009

19

Summary

- DNS has multiple record types
- HTTP good example of protocol which can be optimised in ways designers had not first thought about
 - Persistent connections
 - Pipelining
 - Proxy servers
 - CDNs
- For a protocol to be worth optimising, it must first get critical mass
 - Bittorrent another example of protocol where significant research into optimisation is taking place.

© THE UNIVERSITY OF WARWICK • TE WAIWAE WAIWANGA O WAIWAKATO

19 March 2009

20