ECE 435 – Network Engineering Lecture 7

Vince Weaver

http://web.eece.maine.edu/~vweaver

vincent.weaver@maine.edu

19 September 2017

Announcements

- HW#3 was Posted

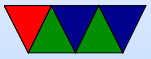

HW#2 Review

- C code will be discussed next time
- browser
  - Error 404 – not found
  - Error 418 – RFC 2324 coffee protocol (I'm a teapot)
  - Error 451 – Unavailable For Legal Reasons / Ray Bradbury
- http header
  - Apache 2.2.2
  - "w3 total cache" – wordpress accelerator?

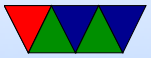

(wordpress is a PHP/mysql content management system)
  - X-pingback – enable pingbacks when linked by wordpress
- Can you connect with telnet? No, way more complex, even discounting encryption

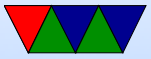

e-mail

- Been around since more or less start of networks
- ARPANET, Ray Tomlinson credited with first modern e-mail around 1971, decided to use '@' char
- UNIX mail, just a mail spool on your computer. Could use command line "mail" to send it. /var/spool/mail/username
- biff to interrupt you when mail came in (used to be exciting)
- mbox vs maildir. mbox format, tell each new e-mail via

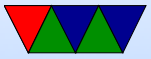

From:. So has to be escaped, you'll see this sometimes. Locking

- Want to send machine-to-machine e-mails. Various ways to do this. UUCP, etc.
- UUCP bang paths

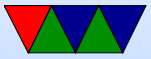

SMTP vs x.400

- As with OSI layer, the big formal ISO definition was made but the hacked-together SMTP won out.
- x.400 much better in many ways
  built-in security
  could tell you once e-mail was delivered
  can send binary files without hacks
- x.400 had horrible e-mail addresses
  C=country, A=administrator (like ISP?can be blank), P= Private Domain, etc

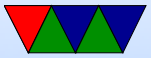

C=US;A=;P=UMaine;O=ECE;S=Weaver;G=Vince;

- x.400 actually used a lot in some situations. Microsoft exchange did for a while

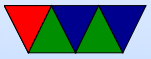

Internet e-mail

- Send/store, can wait on server (as opposed to an instant-message type system where both users have to be active)
- Compose message, send to outgoing server
- deliver to mailbox, collected
- user@host.network
- can often leave off subhost, looks up mailserver for domain via DNS

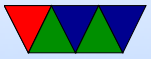

e-mail process

- Sender machine: MUA (mail user agent) sends by SMTP (simple mail transport protocol) to
- MSA (mail submission agent) which determines the destination to send to if not local.
- The MSA uses DNS to look up mail server for destination, then sends it to the receiving MTA (mail transfer agent)
- The final receiving MDA (mail delivery agent) puts into file/mailbox for user

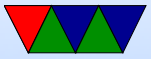

- Receive MUA on local machine via POP3/IMAP
- MUA – editor (optional) mutt/pine/thunderbird/outlook
  Often these days replaced by browser app
  can you use telnet as MUA?
- MTA – sendmail/qmail/postfix
  speaks SMTP. sendmail was standard, has more or less incomprehensible config setup
- MDA – fetchmail? deliver mail tp mailbox. Possibly just a single file, can also be series of directories
- MCA – retrieve e-mail via IMAP or POP

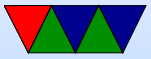

e-mail layout

- envelope
- header/body
- RFC 822/2822/5322
- originally plain 7-bit ASCII, anything more needs MIME and other extensions
- Headers
  - From: / Date: are required
  - Reply-to:
  - Received: (each transfer agent adds in)

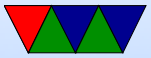

- Return-path:
- To:
- CC: (carbon copy)
- BCC: (blind carbon copy)
- Message-In:
- In-Reply-To:
- Subject:
- X-

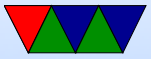

MIME

- How do you send Unicode/8-bit ASCII (accents) or Chinese/Japanese
- How do you attach audio/images?
- Multipurpose Internet Mail Extensions (RFC-1341)
- Backwards compatible
- Message headers:
  - MIME-Version:
  - Content-Description:
  - Content-Id:

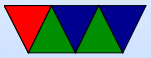

- Content-Transfer-Encoding:
- Content-Type:

- Encodings:
  - regular e-mail 7-bit ASCII lines less than 1000 chars
  - Same, but 8-bit
  - base64 – groups of 24 bits broken into 4 6-bit parts, each a legal ASCII. A=0 B=1 then lower case digits, + /
  - quoted-printable – 7-bit ASCII but higher characters encoded with = sign (hex digits) equal sign =3D
  - multipart

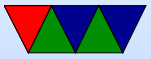

e-mail nettiquette

- Signature, 4 lines 76(?) chars (why?)
- No top-posting!
- Quoting
- Linux kernel rules. Text only. No attachments. No MIME. no line-wrapping. Include patch as text.
- Eternal September / September that never ended!

  sdate

  `Tue Sep 8785 11:59:14 EDT 1993`

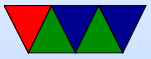

SMTP – simple mail transfer protocol

- RFC 821 in 1982
- connect port 25. Text. All commands 4 chars (no one remembers why)

  `S: 220 maine.com SMTP service ready`
- `HELP`
- `HELO a.com`

  `S: 250 maine.com says hello to a.com`

  There is an extended SMTP. You can detect by sending `EHLO` instead

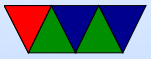

- `MAIL FROM: <xyz@maine.edu>`

  `S: 250 sender ok`
- `RCPT TO: <abx@maine.edu`

  `S: 250 recipient ok`
- `DATA`

  Put data. . on line by itself is end

  `S: 250 message accepted`
- `QUIT`

  `S: 221 maine.edu closing connection`
- Respond with 3-digit code
  - 2xx = successful

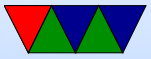

- 3xx = flow control problem
- 4xx failed
- 5xx error in command

- In theory supposed to keep retrying to send for up to 4days

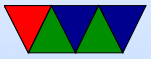

POP/IMAP

- POP (post office protocol) – download mail to local machine which handles
  port 110
- IMAP (internet message access protocol) – manipulate mail on server
  gmail is IMAP. tags are really imap "folders". Can actually download local (I do).

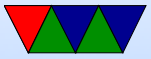

e-mail body

- What happens if try to start line with “From”? Try it.
- Useful to check headers for things like SPAM, phishing attacks
- Signatures (4 lines/80col?)

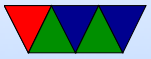

SPAM/other

- In early days, "open relays" if an e-mail came in the server would take mail from anyone and try to deliver it to anyone. Not a good idea (spammers)
- Origin of term SPAM?
  First commercial spam March 5, 1994 Law Firm, Green Card Spam
- Spam/Virus filtering (joke of getting viruses via e-mail)
- procmail sorting
- mail spoofing (What's to stop you from putting someone

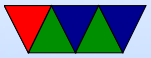

else's address at FROM? how can you catch this?)

- SPAM countermeasures
  - On the sysadmin side, make sure systems are secure. Many ISPs block outgoing port 25
  - SPF records in DNS, say which machines in your network are allowed to send e-mail. Downside, if user has bought a domain and uses it but the ISP doesn't support SPF.
  - Not posting your e-mail, intentionally mixing up your e-mail so address harvesters have trouble getting it. Downside? Things like + in e-mail address?

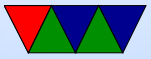

- Challenge/response. Need to ACK before e-mail goes through. No one likes this.
- DNS black lists, lists of known spamming sources
  Some people block whole countries or all cable-modem connections
- Strict SMTP implementations. Spammers don't always implement their mail senders well.
- Greylisting – delay delivery of the mail by a few minutes (with a 400 response). Most legitimate servers will retry, a lot of spam software doesn't bother.
- Filtering, blocking keywords/all-caps

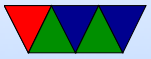

False positives?

e-mails with chunks of books in them, crazy characters

- Bayesian filtering – auto learning. Sometimes can see this in headers

- Vacation Messages
- Mailing lists

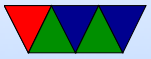

e-mail security

- SSL encrypted connection to SMTP server (usually plain text) SSMTP
- SMTP end to end still unencrypted
- Can use PGP (pretty good privacy) to encrypt e-mails, practically no one does this

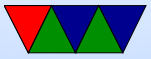

Domain Name System (DNS)

- Why do we need it? Send e-mails to vince@192.168.8.1? What if server moves?
- Hierarchical distributed database
- RFC 1034, 1035
- Maps hostnames to IP addresses
- In early days NIC.arpa has a "HOSTS.TXT" file you downloaded occasionally with all known machines. Didn't really scale.
- /etc/hosts is a relic of this, usually checked first

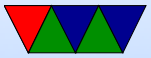

- On Linux this is configured via `/etc/nsswitch.conf`

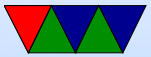

Domain Names

- Which ones can you name? .com/.org/.gov/.edu/.net/.mil
- Country codes (.us/.uk/.ie etc)
- Huge expansion in the last few years (.horse)
- Owner of a domain can subdivide, i.e. eece.maine.edu
- How do you buy them? Used to be fairly expensive and only for two years at a time from a single registrar. Not so much anymore.
- whois will show you info on who owns

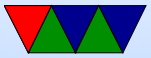

Name Rules

- Can have 127 levels, each 1-63 chars.
- Usually total name cannot exceed 253 chars.
- LDH (letters,digits,hyphens, cannot start with hyphen, not all numbers)
- Case-insensitive
- International names: "punycode". Trouble, why? Foreign letters that look like ASCII ones.

- punycode – snowman example `http://xn--n3h.net/`

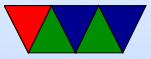

- First commercial name 15 March 1985 symbolics.com example.com set aside (why be careful with your example names?)
- Shortest? g.cn. Various one-letter domains (like x.org) but they were later reserved.

- Typosquatting, domain squatting, copyrighted names, etc.

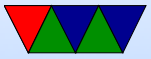

DNS Server

- Listens on port 53, usually UDP
- Zone records
- 5-tuple, NAME TTL CLASS TYPE VALUE
  - TTL (how long to cache)
  - Class (usually IN for internet)
  - Type and RDATA (resource data)
  - Common types
    - SOA – start of authority (parameters) primary source, e-mail of admin, etc

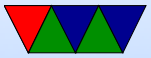

- A – IPv4 address of host (32bit int) `linux.deater.net 86400 IN A 1.2.3.4` can have multiple and be cycled through round-robin
- AAAA – IPv6
- MX – Mail exchange (can have multiple, specify priority)
- NS – name sever (name server for this domain)
- CNAME – Canonical name, allows aliases
- PTR – alias for IP, for reverse lookup `4.3.2.1.in-addr.arpa`
- HINFO – cpu and OS type (text) (uncommon)

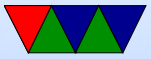

    - TXT – raw ASCII text
- Can you store other things in records? Text adventure? File transfer? Tunneling (iodine?)

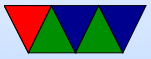

DNS client

- Name resolver, translate from ASCII (still?) name to IP addr

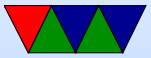

DNS Lookup Example

- Basically: application calls a library (resolver) with the hostname.
  gethostbyname() in socket examples
  This sends UDP to local DNS server, which figures out the address (possibly recursively) and returns the address to caller.
- Details
  - First check `/etc/nsswitch.conf` which might say to check `/etc/hosts` and maybe NIS/LDAP first

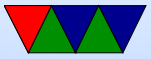

- Query via UDP local nameserver (`/etc/resolv.conf`)
- If the local is the official nameserver, get *authoritative response* (from responsible zone)
  the alternative is a cached response
- If local DNS server doesn't know about it, it has to ask up the chain.
- If not known, query "root" server. So if looking up weaver-lab.eece.maine.edu will ask root, which will direct to .edu DNS server
- Used to be (still?) 13 root servers, mostly in US
- Recursive query It will not likely know but will know

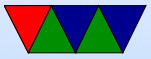

about maine.edu, so ask that one, which will ask eece.maine.edu, etc, and then passed back

- Result is then cached along the way (TTL) caching up to 68 years (or none at all). Why low values? Why can that be bad?
- Caching also means usually the root server does not have to respond to each request
- As the response is passed back through it will be cached along way.

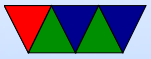

Reverse DNS request

- Given IP address, how can you find the name?
- Linux can use the "host" command.
- For IPv4, there is special in-addr.arpa domain
- To look up 1.2.3.4, lookup 4.3.2.1.in-addr.arpa
- It will iterate down. This gets trickier now with non-contiguous IP allocations.
- Similar thing for IPv6 using ip6.arpa

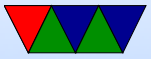

Other DNS Notes

- Can also do non-recursive (iterative)
- breaking out of circular queries
- Packet format
- BIND/named
- dig / nslookup tools
- Security issues?

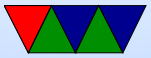

USENET

- Usenet, September that never ended

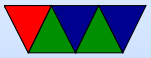

Other common protocols we won't cover

- Legacy (inetd): echo, chargen, discard, time, finger (.profile, .plan), qotd, systat, write, talk
  why no longer supported? security? lack of interest?
- Messaging:
  - IRC – internet relay chat
  - AIM/ICQ/MSN etc
  - unix talk/write
  - MUDs, talkers
- IPP – printer protocol (CUPS, lpd, jetdirect)

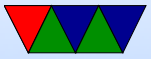

- backup software
- syslog
- telephony
  - skype
  - facetime
  - VOIP
  - ASTERISK
- ntp – network time protocol
- LDAP/Authentication
- Network Attached Storage/Fileservers
  - NFS

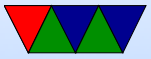

  - Samba/CIFS
  - andrewfs (afs)
- Databases: mysql
- Distributed/Torrent sites
- Distributed computing (SETI@Home)