

COMPSCI 314 S1 C

DNS, Email, HTTP,
DHCP, Network Management

Domain Name System: name space

[Halsall section 8.2]

- The Domain Name System (DNS) maps host names (e.g. www.cs.auckland.ac.nz) to IP addresses (130.216.33.106)
- DNS is a distributed (tree-structured) database system
- Each node in the tree provides information for one domain, i.e. a set of hosts in a single network
- The root node is .
- Below that are top-level domains (TLDs), maintained by ICANN
 - ISO 3166 country codes ([ccTLDs](#))
 - ‘generic’ domains ([gTLDs](#)), e.g. com, net, org, biz, ...
- Next level down are ‘2nd-level’ domains
 - e.g. [microsoft.com](#), [ac.nz](#), [net.nz](#), ...
- Each ISP and Enterprise network runs its own *nameserver*
 - e.g. [auckland.ac.nz](#), [cs.auckland.ac.nz](#)


DNS: Resource Records

- Each nameserver keeps its data in a *zone* file
- The zone file records are DNS *Resource Records*
 - Each contains: *name type value*
- There are many record types: *[Halsall 8.2.2]*
 - A address IPv4 address for this host
 - AAAA IPv6 address IPv6 address for this host
 - NS nameserver name for this domain
 - MX mail exchange name “ “ “
 - ...
 - SOA authority data for this domain

DNS: looking up a domain name

- DNS looks up a *Fully Qualified Domain Name* (FQDN)
- A client sends a lookup query to a nameserver, e.g. using **dig** or **nslookup**. Nameserver will try to answer query from its own records. If it can't, it will start either a recursive or a non-recursive query
- Recursive query: DNS server will query higher-level nameservers on behalf of the client and return the result
- 'Higher-level' normally means walking down the tree from its root, but it may also mean asking the 'next-higher' nameserver
- DNS servers usually cache (temporarily store) records retrieved from other DNS servers – this reduces lookup traffic
- Cached records are 'non-authoritative,' genuine original records are 'authoritative'

DNS: Nameserver Hierarchy



DNS: Root Servers

- There are 13 root servers, A-M; e.g. f.root-servers.net, run by various different organisations
- There are also 13 gTLD servers, A-M, run by Verisign
- Many of the root servers are anycast
 - All instances use the same IP address and AS number
 - The global (BGP) routing system finds the ‘best’ instance for each user
 - For example, F root has 36 instances. Our closest one is at APE, floor 54 of the SkyTower
 - Likewise, I root has an instance at WIX in Wellington
- Anycast servers share the request load, and make the DNS more resilient against attacks

SMTP – Simple Mail Transfer Protocol

- SMTP is used to send e-mail
- Very widely used
- Uses port 25
- An SMTP connection is known as a ‘session’
- All commands in clear text!
- Unauthenticated!

SMTP at work

- Client connects to *server* port 25 using TCP
- HELO
- *250 mail.compsci314.com Hello a-client.com [130.216.197.66], pleased to meet you*
- MAIL From: jane@bloggs.com
- *250 jane@bloggs.com ... Sender OK*
- RCPT To: joe@bloggs.com
- *250 joe@bloggs.com ... Recipient OK*
- DATA
- *354 Enter mail, end with '.' on a line by itself*
- Subject: SMTP is soooo cool!
Howzit goin, Joe?
- .
- *250 VBB09405 Message accepted for delivery*
- QUIT
- *221 mail.compsci314.com closing connection*

SMTP problems

- SMTP does not authenticate its user!
- SMTP does not ask for username and password before accepting a mail for delivery
- Anyone can connect to an SMTP server and send spam e-mails from fake senders!
- This is a major problem on the Internet now – most of the e-mail traffic is now spam
- SMTP traffic can be eavesdropped on – not encrypted!

Anti-spam strategies used by network operators

- Source filtering: Accept connections only from within own network or with secondary authentication (e.g. NetLogin) – no ‘open relay’
- Log connections with IP address, time, etc.
- Optional: Don’t accept mails with unknown user/domain names for delivery

Aside: Anti-spam strategies you can use

- Don't put your e-mail address on the web or into a newsgroup
- Don't run an SMTP server (sendmail) on your computer – if you do, make sure it won't accept outside connections
- Never reply to spam e-mail
- Don't 'unsubscribe' yourself from spam e-mails. This only confirms that your address is live!

Email: User Interface

- Network Administrator runs an email server, i.e. a Mail Transfer Agent (MTA)
- User runs an email client, i.e. a Mail User Agent (MUA)
- User composes email and sends it to her MTA, MTA sends it on to other MTAs using SMTP
- Incoming mail arrives at user's MTA. She retrieves it from there to her MUA using a protocol such as IMAP
- Web-based email systems now provide a common, easily-accessible MUA

POP – Post Office Protocol

- Client connects to *POP3 server* on port 110
- *+OK POP3 server pop3.foobar.com ready*
- USER joebloggs
PASS iluvjane
- *+OK 1 message(s)*
- LIST
- *+OK 1 message(s)*
1 4578
.
- RETR 1
- *+OK 4578 octets*
From: jinx@foo.com
To: joe@bloggs.com
...
- QUIT
- *+OK Goodbye!*

POP3 shortcomings

- Plain TCP connection is not secure enough in many cases – can use secure POP to deal with this
- No compression of mail data (=long downloads)
- Not very flexible if we wish to receive mail on multiple machines – IMAP addresses this problem

HTTP – HyperText Transfer Protocol

- Used for the communication between web clients (often browsers) and web servers
- Client connects to web server using TCP, generally on port 80. (Other ports are also used, especially for secondary servers on a machine)
- Client sends an HTTP request and receives an HTTP response
 - HTTP 1.0 closes the TCP connection
 - HTTP 1.1 keeps connections open, allowing further objects to be transferred through them

BOOTP: RFC 951

- ARP finds a host's IP address, but ARP is a **link-layer** (layer 2) protocol, it can't be routed
- BOOTP (the Bootstrap Protocol) is an application protocol using **UDP transport**:
 - Requesting host sends a BOOTP request packet containing the client's MAC address to port 67 at the IP broadcast address
 - BOOTP server looks up the client MAC address in a database, then sends response back to client port 68
 - Server may add client to its ARP cache, otherwise it broadcasts the reply

BOOTP relay

- Administrator can maintain database of MAC addresses and their IP addresses
- Simple networks may use a separate BOOTP server for each subnet (e.g. City and Tamaki)
- Larger networks can run a single server, and configure routers to relay BOOTP datagrams
 - Router receives a BOOTP request, forwards it to a specific BOOTP server
 - Router receives BOOTP response, sends it to requesting client

DHCP: RFC 2131

- An extension of BOOTP to carry other data besides client IP address, e.g. netmask, default gateway, nameserver
- Includes ability to allocate addresses *dynamically*
 - Administrator maintains list of permanently-allocated IP addresses for fixed hosts, and range of addresses for other hosts
 - Client is **leased** (allocated) an address from the range for a specified lease time. Leases may be extended or revoked
- DHCP (and BOOTP) check for clients using same IP address


Network Management

- ‘Management’ (IETF-style) means
 - Configuring devices
 - Monitoring device operation
 - Changing device configurations
 - Keeping track of configurations as they change
- Simple (late 1980s) view
 - Each device keeps configuration in a ‘**Management Information Database**’ (MIB)
 - **SNMP** allows a manager to interact with a device via its MIB
- Current (2000s) view
 - Need to manage large sets of devices as a coherent whole

Management Information Base

- Management of a network involves reading and setting many network-related values, as we evaluate the network performance and adjust its operation
- The management values are held in an extensive database using a structure agreed between Internet and OSI network management, known as the *Management Information Base* or *MIB*
- We can monitor a device's behaviour by watching values in the MIB, e.g. bytes in/out of an interface
- The MIB is a hierarchical structure, with a name identified by the sequence of node indices as the tree is traversed from the root to the node

MIB




All internet Management variables start with
1.3.6.1.2. ... or iso.org.dod.internet.mgmt

First levels of Management Information Base tree

- There are three basic trees, from an unnamed root node.
- The one of interest is ‘internet,’ which is managed by ‘dod’ (the Department of Defense), which is an ‘org’ (Organisation) known to the International Standards Organisation ‘iso’

MIB (2)



Namespace for variables within Internet MIB-I (RFC 1156)

- MIB-I is ‘for Network Management of TCP/IP-based internets’
- It’s *interface* variables are very useful, e.g.
ifInOctets and **ifOutOctets**

Simple Network Management Protocol (SNMP)

The SNMP protocol provides the user interface to the MIB and is in some respects an extension of the MIB.

Based on a request/response, or fetch/store paradigm, in which the manager may request values for variables or supply values; actions may occur as a side-effect of writing values.

Command	Meaning
get-request	Fetch a value from a specific variable
get-next-request	Fetch the value following that fetched by earlier get-request
set-request	Store a value in a specific variable
get-response	Reply to a get operation
trap	Message triggered by an event

SNMP (2)

- SNMP operations are atomic, meaning that all of the actions of a message must occur, or none will occur; any error causes the whole request to be abandoned
- The problem of scanning through tables is solved by the *get-next-request* which has an identifier for a known item in a table, but triggers a *get-response* corresponding to the next table entry. A *get-next-request* to the table itself returns the first entry
- SNMP is usually implemented as a datagram protocol, using UDP. SNMP response packets serve as Acks to set- or get- requests
- Each command is given a 32-bit sequence number which is returned in replies
- Replies can be associated with requests, or repeated commands ignored

Details from here on are not examinable, but principles should be known.

Fortunately, SNMP is supported by a suite of public-domain routines which do most of the message formatting and interpretation for the user, so that even people working at the detailed level seldom need to know all the gory details.

SNMP(3)

- The GetRequest-PDU for example contains a 32-bit *requestID* which is essentially a sequence number of the request to match requests and responses, error indicators and a list of objects for which values are wanted.
- The messages tend to be complex. The next slide contains the lines —

```
A0      1C      02      04      05      AE      56      02
getreq  len=28  INTEGER  len=4  |----- request id -----|
```

A0	identifies a Get-Request command
1C	length of the command
02	an integer follows (known to be the request ID)
04	the length of the integer
05 AE 56 02	the 4 bytes of the integer value

SNMP(4)

```

    30      29      02      01      00
SEQUENCE len=41 INTEGER      len=1 vers=0

    04      06      70 75 62 6C 69 63
string   len=6      p  u  b  l  i  c

    A0      1C      02      04      05 AE 56 02
getreq   len=28 INTEGER len=4 |-- request id --|

    02      01      00 02      01      00
INTEGER  len=1      00 INTEGER len=1 err.index

    30      0E      30      0C      06      08
SEQUENCE len=14 SEQUENCE len=12 objectid len=8


    2B      06 01 02 01 01 01 00
1.3      6  1  2  1  1  1  0

    05      00
null     len=0
```

Example of SNMP message – get-request for sysDescr (1.3.6.1.2.1.1.1)

Monitoring a link with MRTG

MRTG reads SNMP variables and plots them on web pages, for example:



Max **In**: 8280.3 kb/s (8.3%) Average **In**: 4017.3 kb/s (4.0%) Current **In**: 3500.9 kb/s (3.5%)
Max **Out**: 50.8 Mb/s (50.8%) Average **Out**: 25.1 Mb/s (25.1%) Current **Out**: 22.7 Mb/s (22.7%)

NeTraMet: a Network Traffic Meter

A system of:

- *Meters*, to collect measurements on the attached networks
- *Meter readers*, to collect data from meters (meters & readers may be many to many)
- *Managers*, to coordinate and process data from readers
- Communication (configuration and reading) is by SNMP

