How the web works - 2

Transmission Control Protocol / Internet Protocol

Domain Name System

How the web works - 2

Transmission Control Protocol / Internet Protocol

Domain Name System

Names and Numbers

Translating between names

www.google.com

To numbers

74.125.77.103

Problems with TCP/IP

Numbering system not intuitive

148. 87 .9 44

141. 241. 2. 11

Can't work out location, service or even remember easily

Problems with TCP/IP

Names much better

148. 87 .9 44	www. oracle. com
141. 241. 2. 11	pop3. kingston. ac. uk

Can work out location, service and remember easily

What happens when you browse?

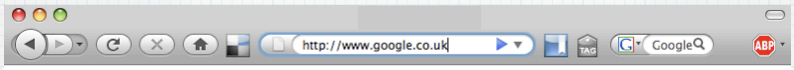

What happens when you browse?

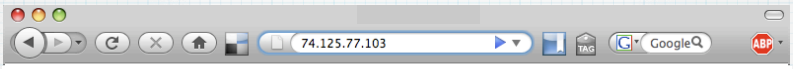

What happens when you browse?

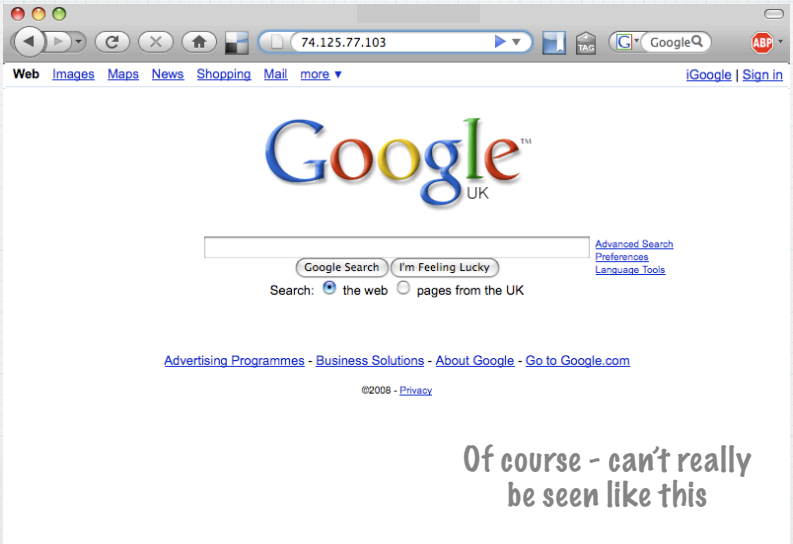

Of course - can't really be seen like this

Web Browser software

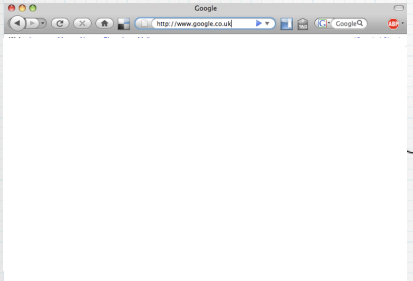

Server on the internet

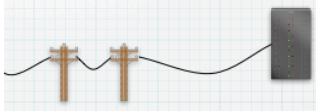

Web Browser software

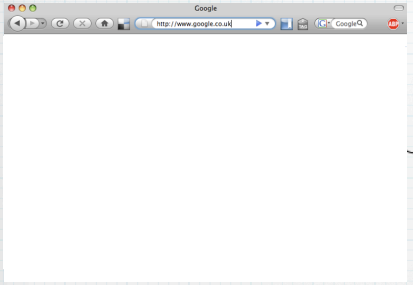

Server on the internet

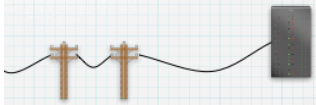

192.0.0.1 <----------> 148.87.9.44

Web Browser software

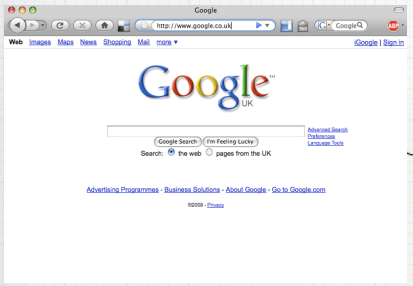

Server on the internet

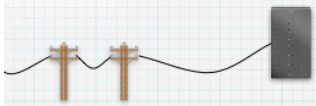

192.0.0.1 <----------> 148.87.9.44

Simplest way to implement the conversion

Use a text file with columns of names and IP numbers

```
localhost              127.0.0.1
www.oracle.com         148.87.9.44
```

A simple lookup process converts the name to number

Called the **hosts** file

This file still exists on machines today and is checked before the DNS/BIND system is used

Demo

- Changing a hosts file to associate a name with an IP number
- Using ping command to see if it works

Demo

- Changing a hosts file to associate a name with an IP number

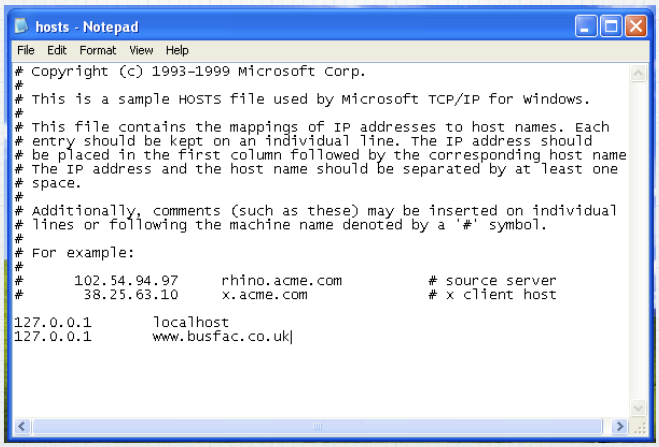

Demo

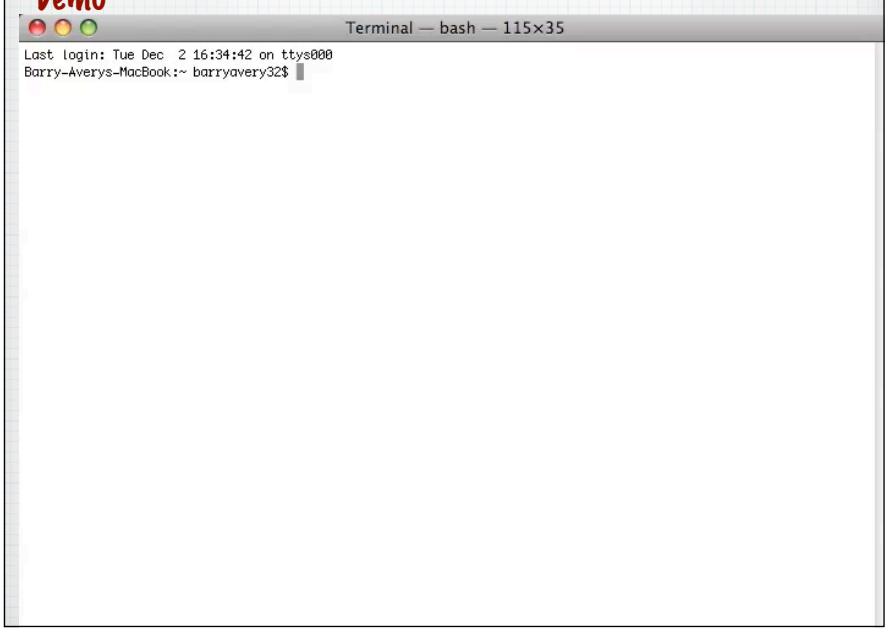

Works for very small number of machines

- Have one copy of the hosts file held centrally
- Regularly copied to other machines

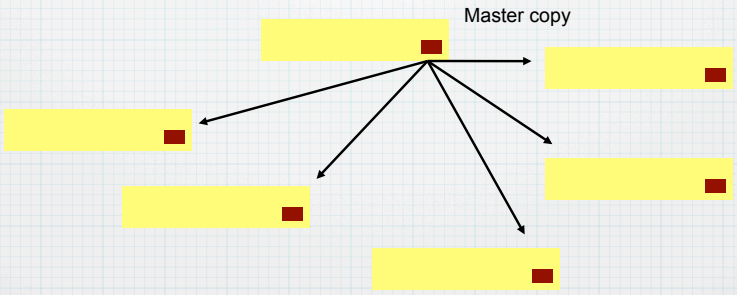

Obvious problems

- Doesn't scale well
- Synchronisation problems
- How do names outside the domain get managed
- Name collisions / management problems

So: DNS and BIND

Domain Name System	Standard
Berkeley Internet Name Domain	Software

So: DNS and BIND

A distributed database system initially invented by Paul Mockapetris

RFCs 882, 883, 1034, 1035 + others

BIND is the most popular software for DNS

DNS can be visualised as an inverted tree

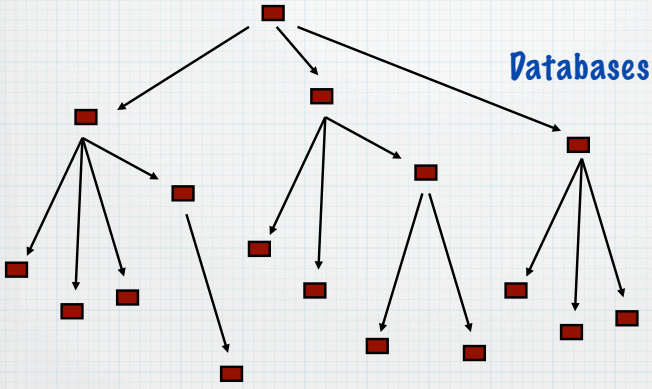

DNS can be visualised as an inverted tree

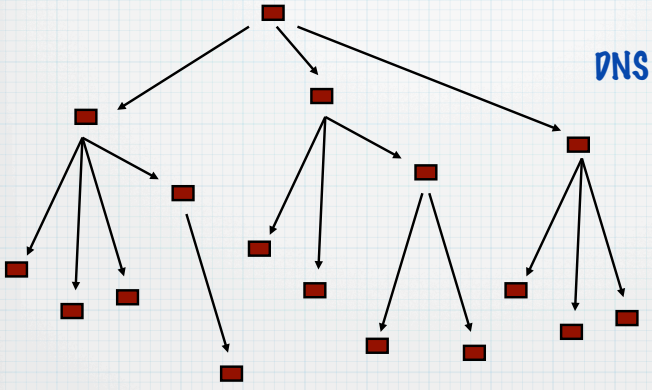

Other terms used

Top level DNS servers

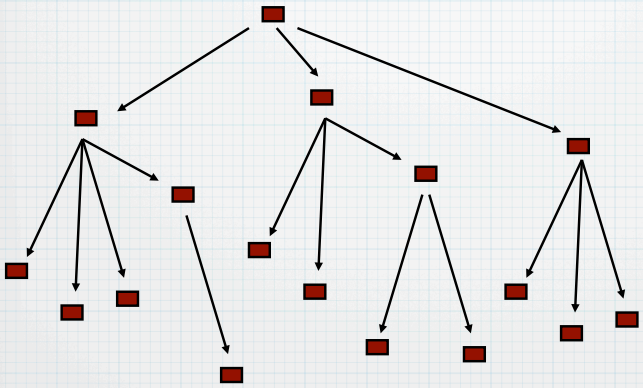

Other terms used

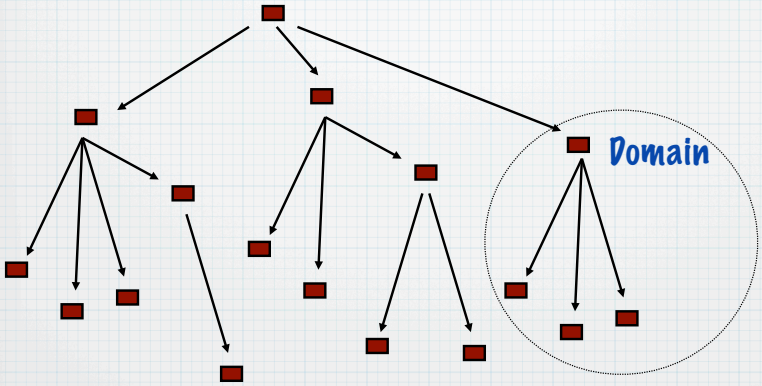

Other terms used

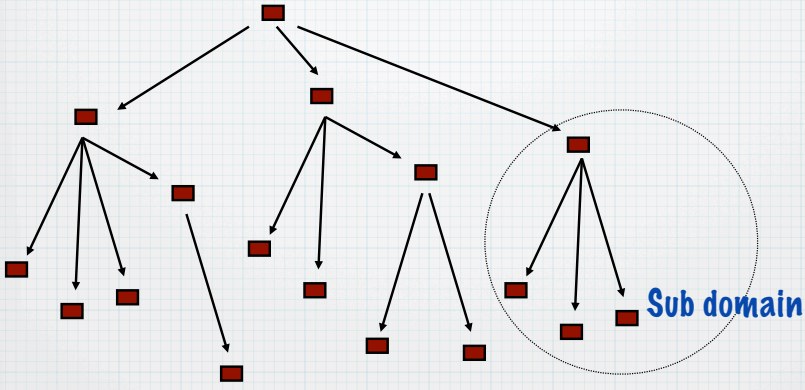

Other terms used

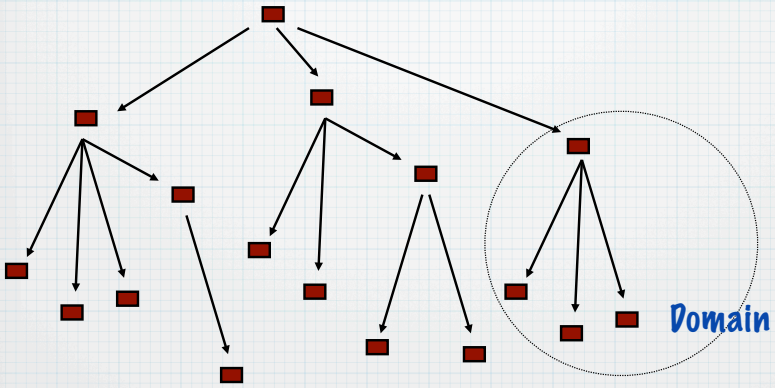

Domains and sub domains

academic domain .ac.uk

Run by JANET

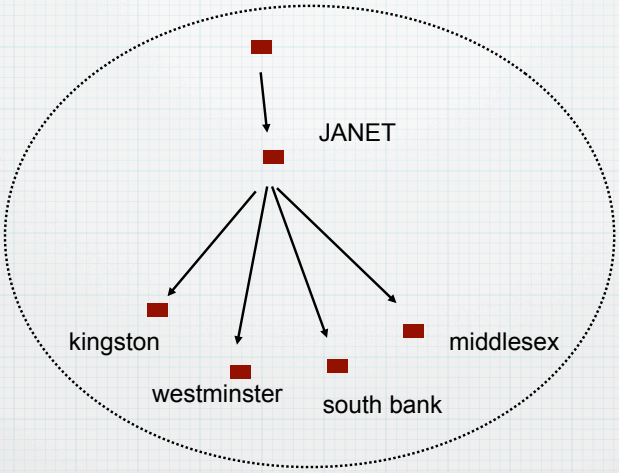

Domains and sub domains

kingston domain

kingston.ac.uk

Run by Kingston

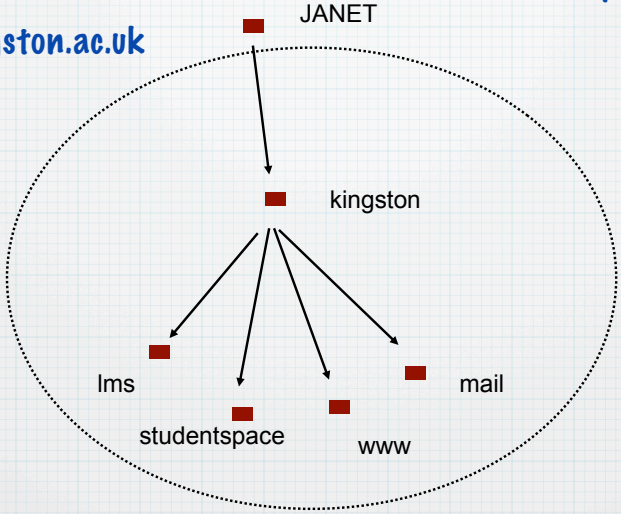

Note how names are formed

- names formed from right to left
- we control the names in our domains

- if we have a computer called bob
- if westminster have a computer called bob

No name conflict

bob.kingston.ac.uk

bob.westminster.ac.uk

So how does name resolution work? Local example

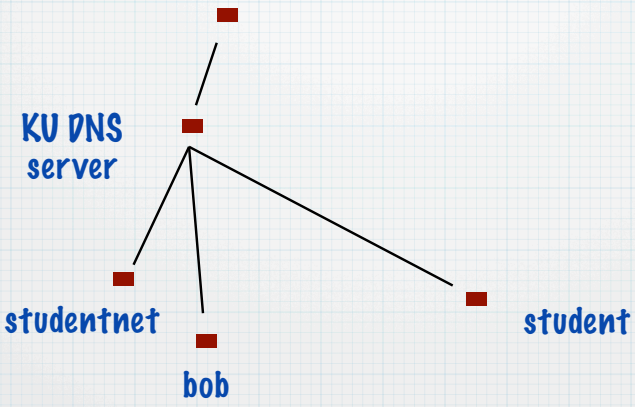

So how does name resolution work? Local example

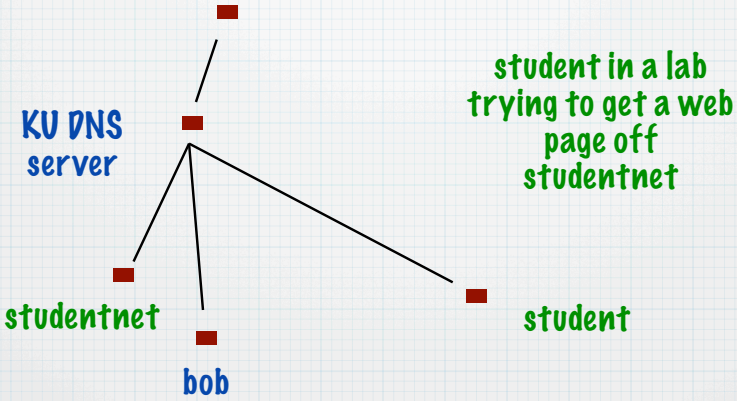

student in a lab trying to get a web page off studentnet

So how does name resolution work? Local example

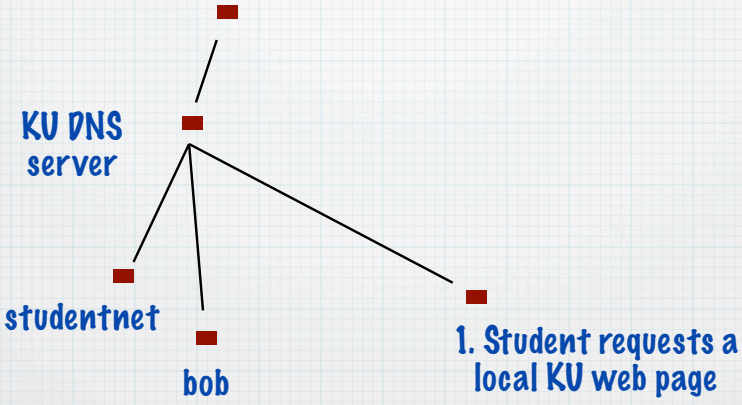

So how does name resolution work? Local example

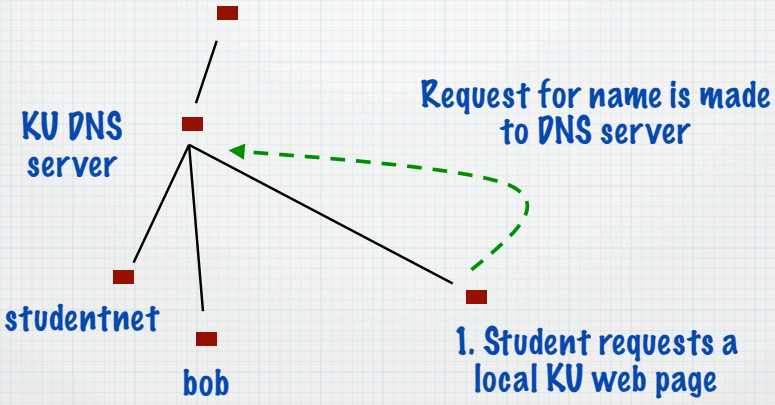

So how does name resolution work? Local example

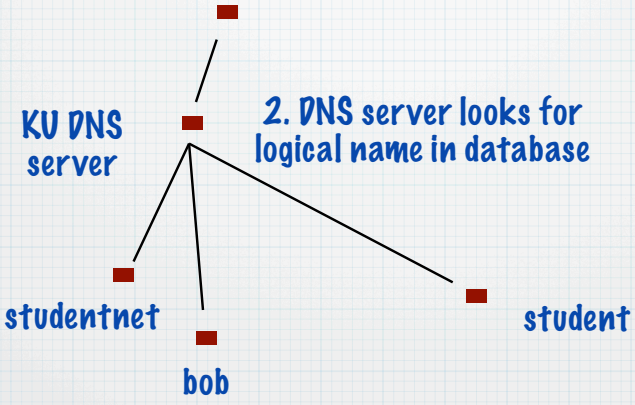

So how does name resolution work? Local example

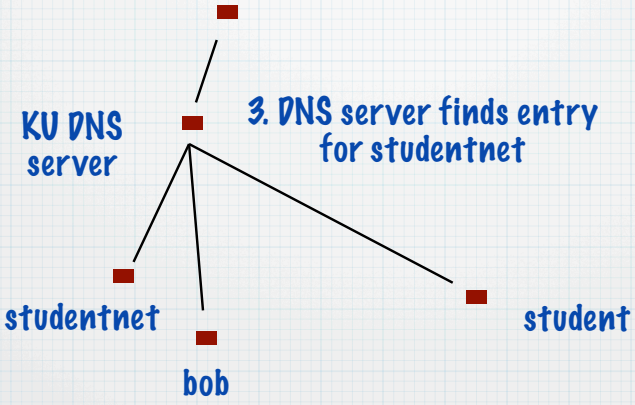

So how does name resolution work? Local example

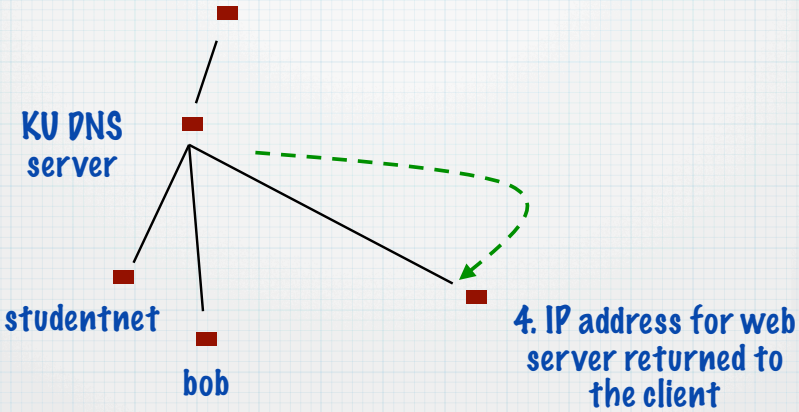

So how does name resolution work? Local example

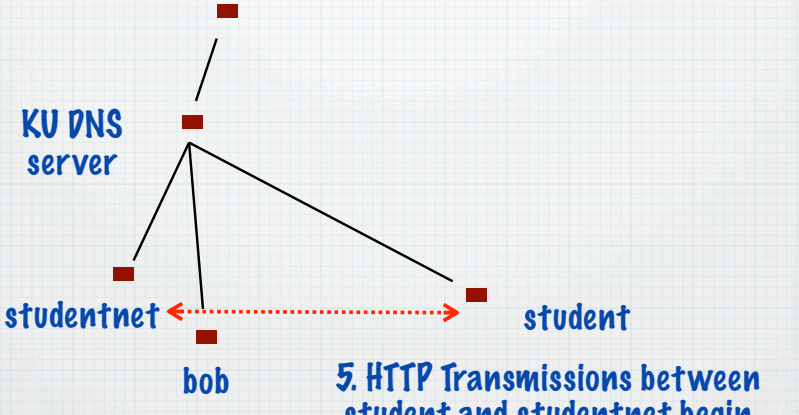

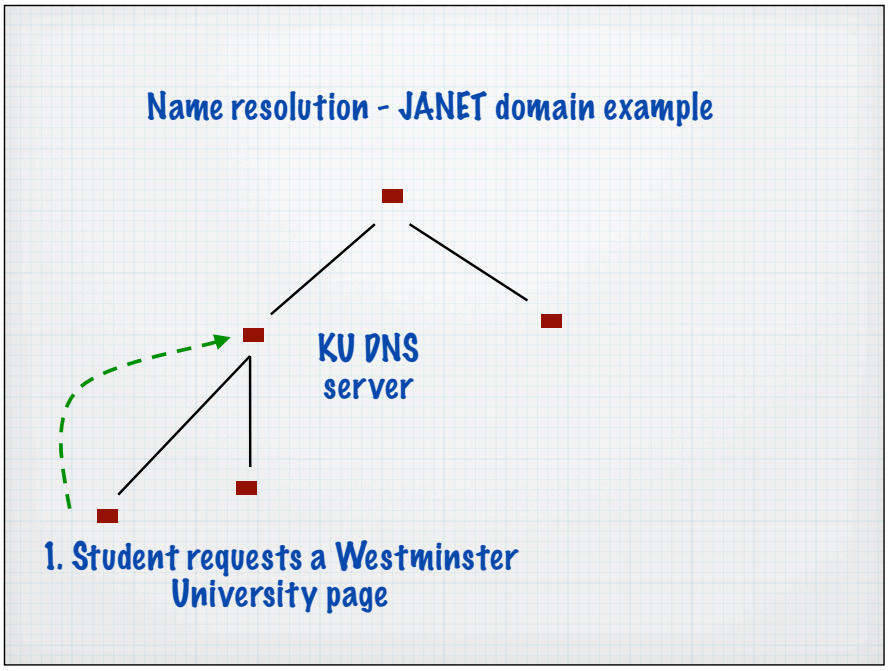
Name resolution - JANET domain example
KU DNS server
1. Student requests a Westminster University page

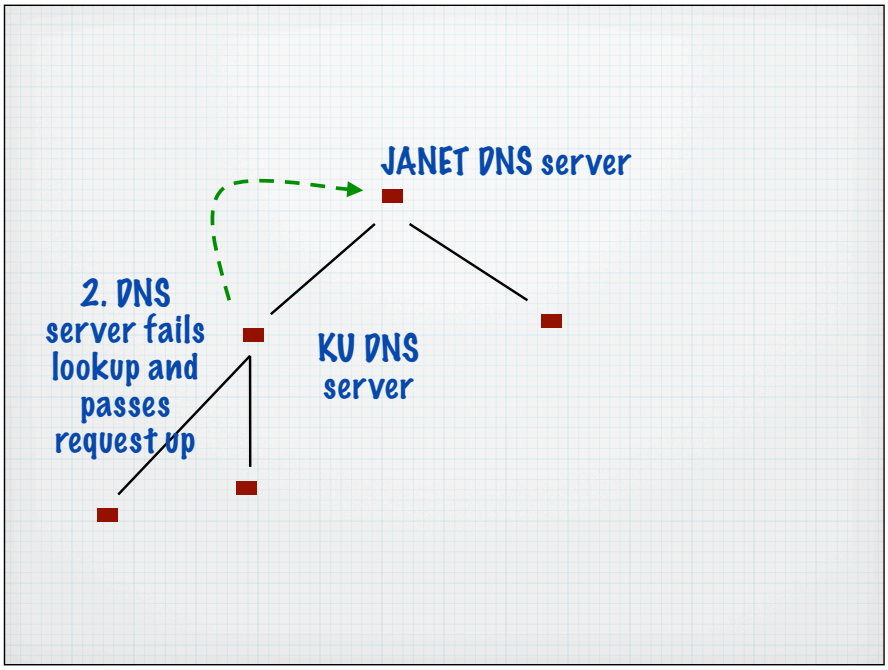
JANET DNS server
2. DNS server fails lookup and passes request up
KU DNS server

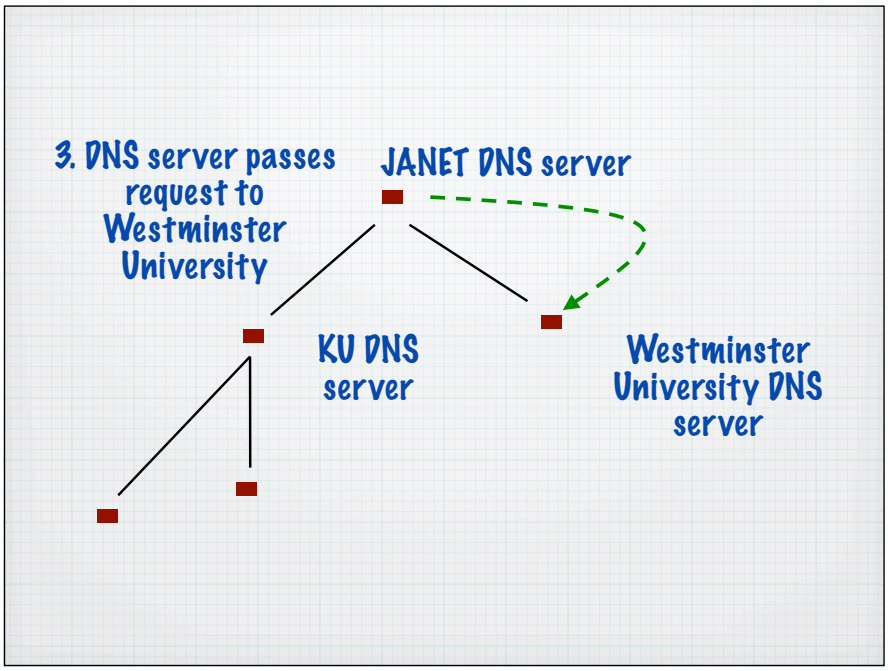
3. DNS server passes request to Westminster University
JANET DNS server
KU DNS server
Westminster University DNS server

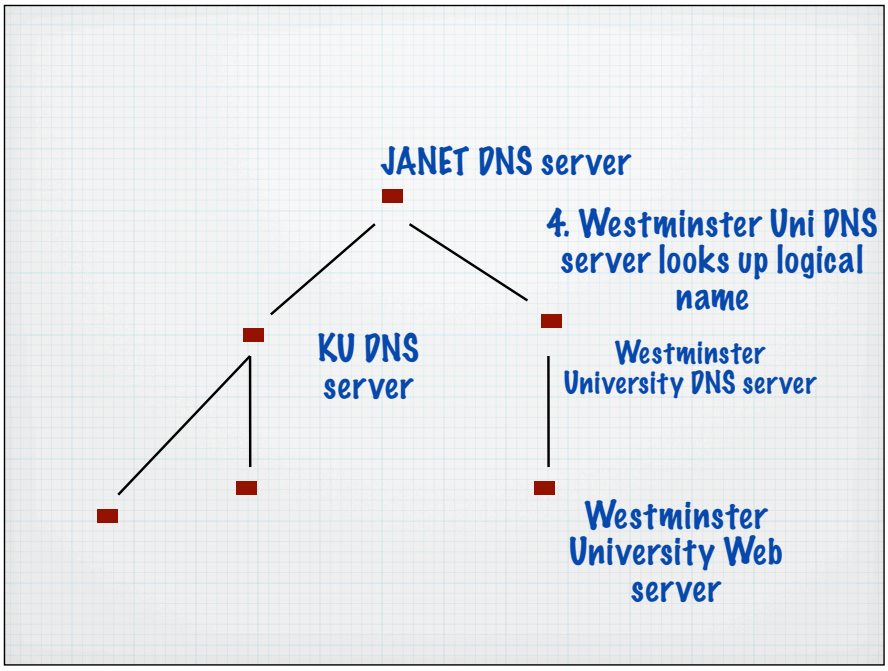
JANET DNS server
4. Westminster Uni DNS server looks up logical name
KU DNS server
Westminster University DNS server
Westminster University Web server

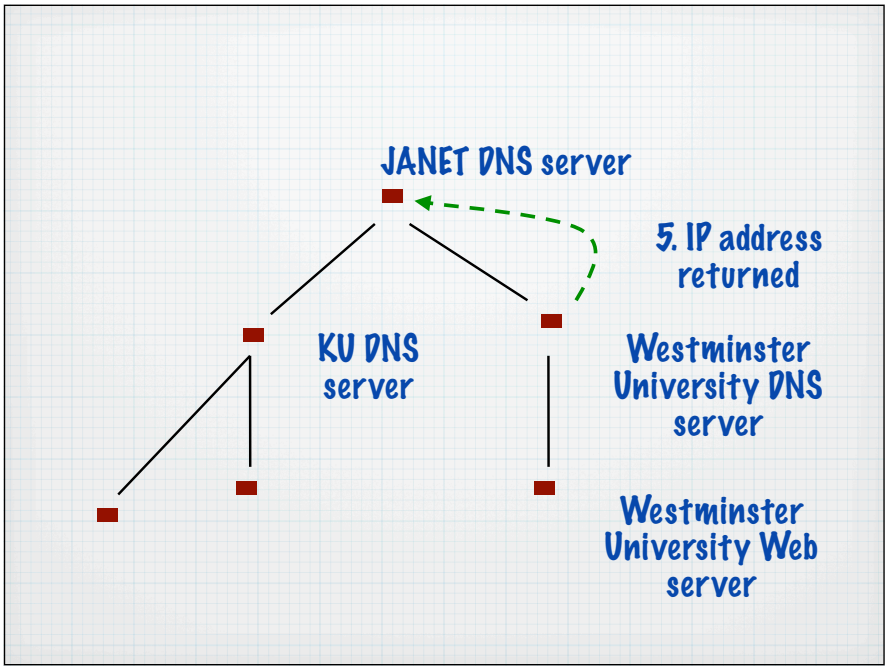
JANET DNS server
5. IP address returned
KU DNS server
Westminster University DNS server
Westminster University Web server

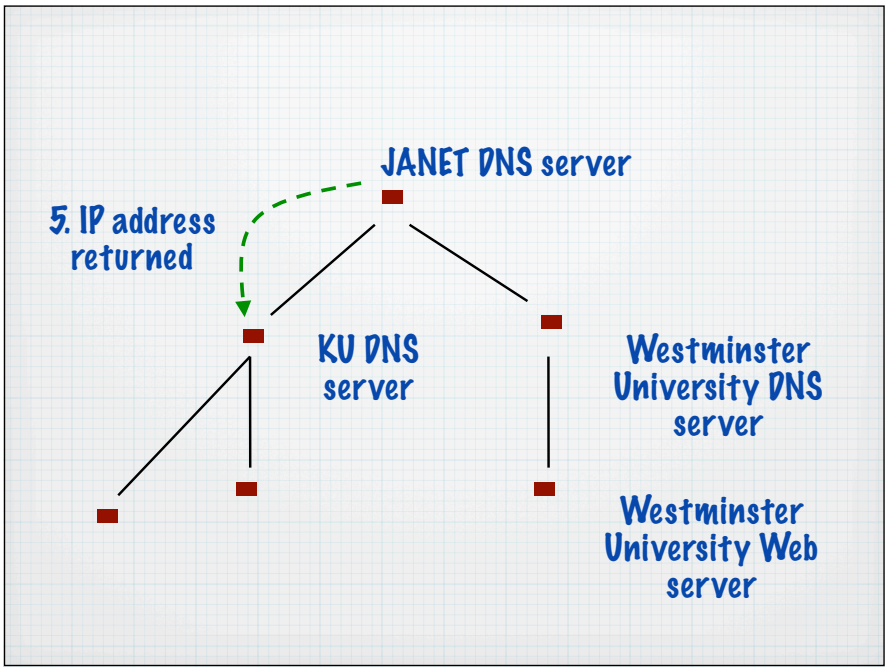
JANET DNS server
5. IP address returned
KU DNS server
Westminster University DNS server
Westminster University Web server

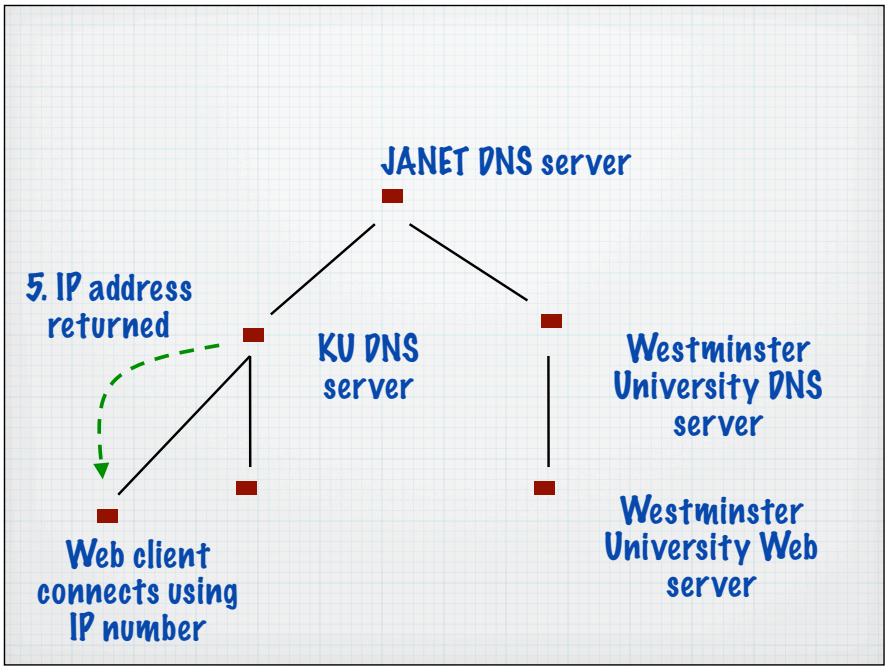
JANET DNS server
5. IP address returned
KU DNS server
Westminster University DNS server
Westminster University Web server
Web client connects using IP number

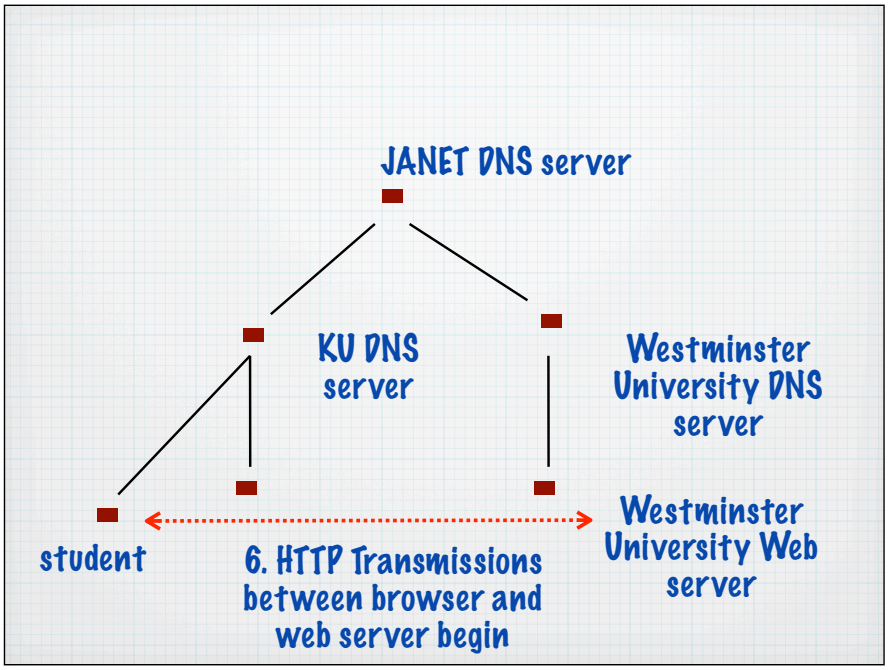
JANET DNS server
KU DNS server
Westminster University DNS server
Westminster University Web server
student
6. HTTP Transmissions between browser and web server begin

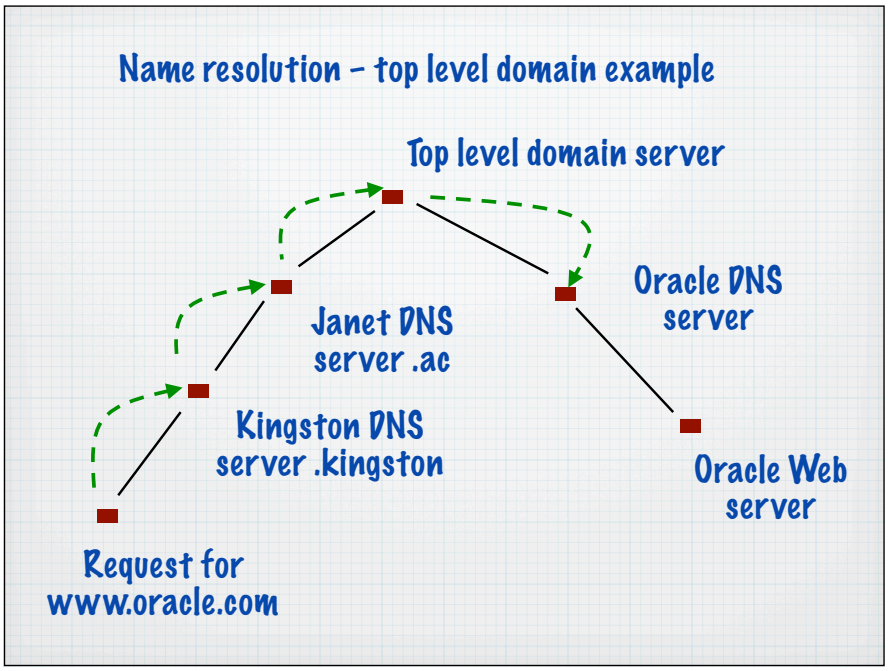
Name resolution – top level domain example
Top level domain server
Oracle DNS server
Janet DNS server .ac
Kingston DNS server .kingston
Oracle Web server
Request for www.oracle.com

Points to note

Queries only travel up the tree as far as needed

Each DNS server has link to next level up

Cache at each level can reduce the amount of traffic

In fact many top level servers (replicated to overcome failures)

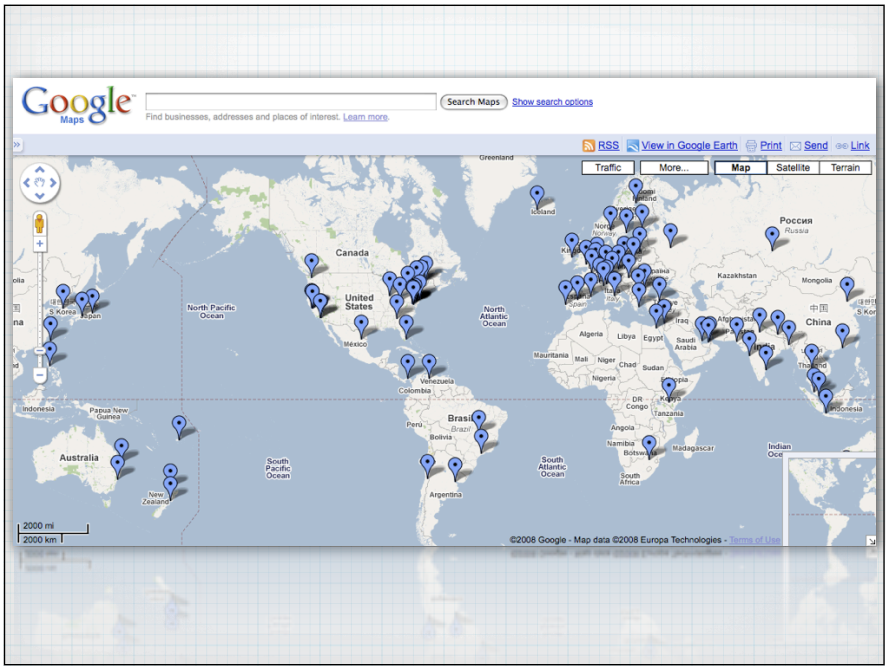

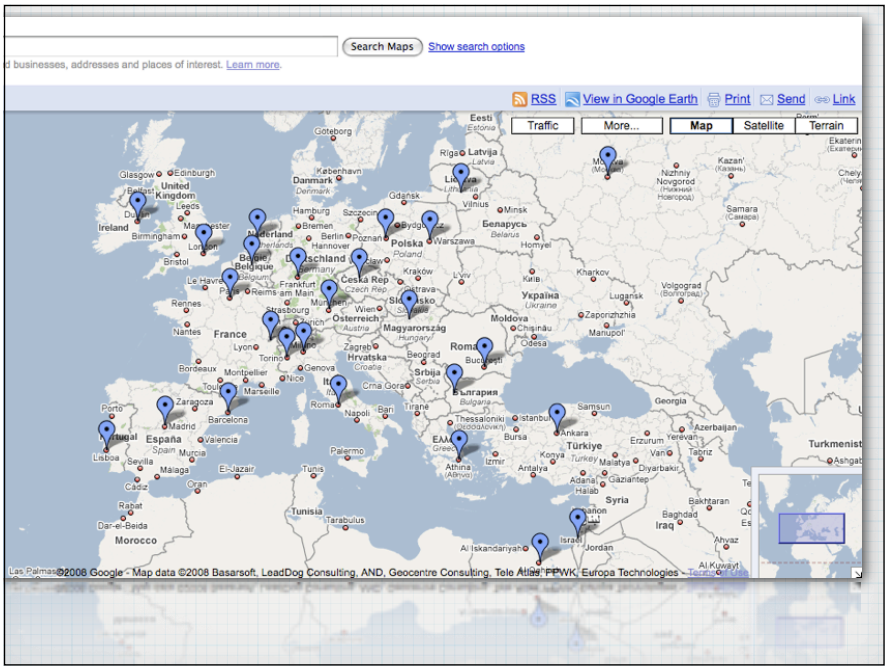

Things to try at home

traceroute name

Shows the route and time taken for the hops

Things to try at home

http://visualroute.visualware.com/

Visually shows the route and time taken for the hops