Mapping Human Readable Names To Internet Addresses

- Simplest technique uses a table of name/address pairs
  - For example, the old Unix file "`/etc/hosts`"
- Problems with host tables
  - Storage cost (each host has table proportional to number of hosts)
  - Distribution cost
  - Maintaining consistency
  - Limited ability of local authority to control names and addresses
  - Does not scale well

Domain Name System (DNS)

- Provides automated mapping from human-readable machine names to Internet addresses
- Supports hierarchical name space
- Allows local autonomy in assigning names
- Uses distributed name lookup
- Requires on-line name servers that use the Internet to exchange requests and responses
- Uses tree-shaped topology of name servers
- Achieves efficiency through locality and caching
- Uses local name server as primary contact
- Lives at UDP/TCP port 53

Domain Name Assignments (Top Level Domains)

Domain	Meaning
COM	Commercial organizations
EDU	Educational institutions
GOV	Government institutions
MIL	Military groups
NET	Major network support centers
ORG	Other organizations
INT	International organizations
ARPA	Temporary ARPANET domain
country code	Countries (US included)

Illustration Of The Domain Name Hierarchy

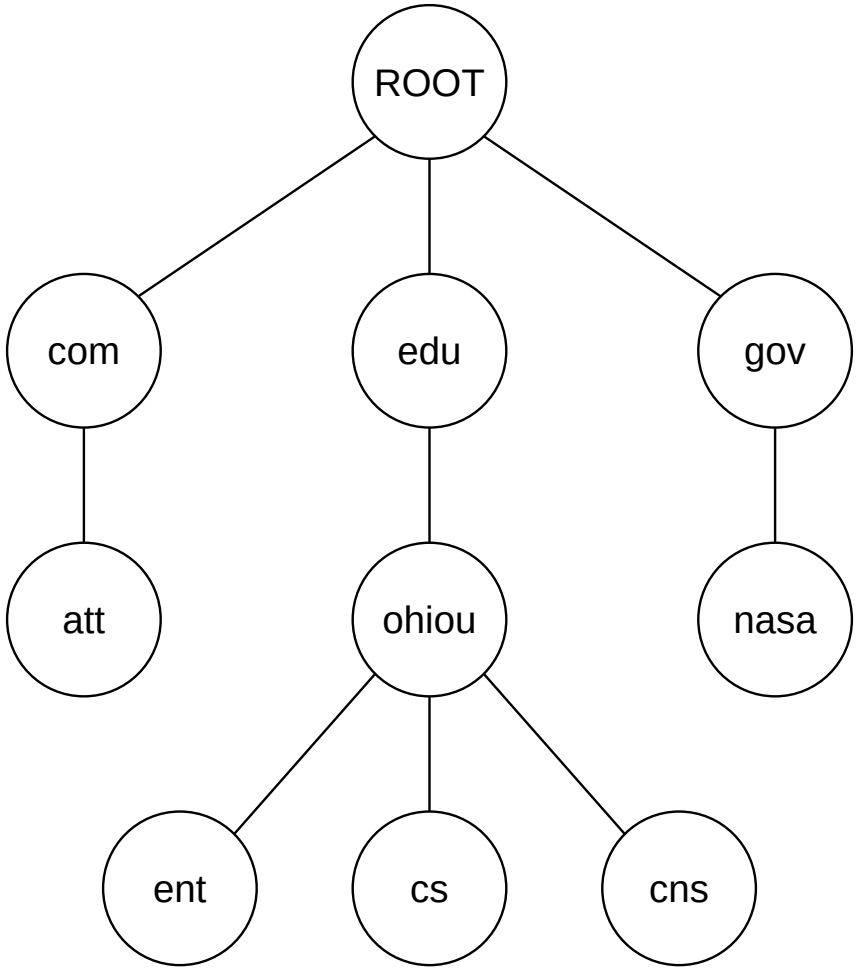

Root Name Servers

- There are 13 "Root" name servers around the world:
  - Designated A-M
  - letter.root-servers.net

Domain Name Example

- Ohio University agrees to operate a domain name server
- Ohio University obtains authority for domain `ohiou.edu`
- Computer Science Department at Ohiou obtains authority for domain `cs.ohiou.edu`
  - Allocates machines `ace` and `boss` as *nameservers*
- To add a new machine, *jarok*, owner obtains permission for the name from the local department
- Authority operating the domain name server installs name `jarok.cs.ohiou.edu` along with appropriate Internet address

Domain Query Processing

- Database contains several types of entries, for example:
  - Machine name and Internet address pairs
  - Mail destination name and Internet address pairs
- Each query specifies the type of lookup desired
- Given name can map to different address depending on query type
- Query may specify iterative (one step at a time) or recursive (complete) resolution

Caching In Domain Servers

- Each server caches pairs during lookup
- Hosts may also cache pairs
- Answer to query tells source of information as well as how to reach the source if information obtained from a cache
- Each answer specifies the length of time that the information can be cached
  - Consider the following "unusual" names:
    - "`mirror.archive.umich.edu`"
    - "`ohiou.edu`"
    - "`www.cs.ohiou.edu`"

Name Server Information
(Example from Ace /etc/named.db/hosts/cs)

```
;   local aliases so we can change stuff easy
gopher IN       CNAME   oucsace
www    IN       CNAME   boss
news   IN       CNAME   boss
ftp    IN       CNAME   oucsace
time   IN       CNAME   bigbird
boss   IN       A       132.235.1.1
       IN       HINFO   "Sun Sparc 20" "Solaris 2.4"
       IN       MX      0       boss
       IN       MX      100     oucsace
oucsboss        IN       CNAME   boss
oucsace         IN       A       132.235.1.2
ace    IN       CNAME   oucsace
oucsace2        IN       CNAME   oucsace
view   IN       A       132.235.1.3
duce   IN       A       132.235.1.6
prime  IN       A       132.235.1.7
```

Shorthand

- The name of a host using all of the fields is called a *Fully Qualified Domain Name*, or *FQDN*
- FQDN's are hard to type
- Most of the time, you can use shorthand
  - More correctly performed by the host
  - Sometimes supported by the DNS server
- Most machines used to use an automated search strategy
- Now, most just tack your FQDN suffix onto the end
  - Sometimes you can get it to look in other domains too (/etc/resolv.conf)

Naming Politics

- The growth of "browsers" and shorthand created a strange situation that hadn't been anticipated by the designers of the DNS system
  - Everybody wanted X.com
  - Hierarchies because pointless because they weren't in use
- "Now there's money and that changes everything" – Comer
  - Domain registration became very lucrative
- Legalities problems
  - The IETF had no interest in making policy about `billgatesstinks.com`
- Formed a group known as ICANN to resolve the dispute
  - The Internet Corporation for Assigned Names and Numbers

Current Top Level Domains

- There are 3 kinds of Top Level Domains (TLDs):
  - 2 letter domains that refer to countries
  - Generic TLDs (gTLDs) with 3 or more letters
    - Sponsored
    - Unsponsored
  - Special TLDs - ".arpa"
- TLD History
  - There were originally 7 TLDs
    - .com .edu .gov .int .mil .net and .org
      - 3 of them could be used without restriction (.com .net and .org)

Current Top Level Domains (continued)

- 7 new TLDs were added in 2001/2002
  - aero - air transport industry (sponsored)
  - biz - businesses (unsponsored)
  - coop - Cooperatives (sponsored)
  - info - unrestricted (unsponsored)
  - museum - museums (sponsored)
  - name - for registration by individuals (unsponsored)
  - pro - Accountants, lawyers, physicians, and other professionals (unsponsored)

Current Top Level Domains (continued 2)

- 6 were added in 2004
  - asia
  - cat
  - jobs
  - mobi
  - tel
  - travel
- There continue to be numerous requests for new TLDs
  - kids
  - xxx
  - law
  - ads
  - ...

TLD .arpa

- The .arpa domain is used for other mapping purposes
  - .in-addr.arpa is used to map from IP addresses back to names
  - .ip6.arpa is the same, but for IPv6
  - .e164.arpa is used to map phone numbers into URIs
- To look up the name of 132.235.1.1, you look up

  1.1.235.132.in-addr.arpa

Tuvalu - Just for Fun

- Tuvalu is a small island nation in the south pacific
  - Former british colony
  - Became in independant nation in 1978
  - Consists of 9 coral atolls
  - Population of 11,000 (2001)
  - 26 square km total
- The official United Nations abbreviation for Tuvalu is "tv"
- In 2000, control of ".tv" was sold for $50M in royalties
  - Verisign bought it for $45 in cash in 2002
- mtv.tv public.tv mlb.tv ufc.tv hollywood.tv we.tc etc...

DNS Details

- Legal characters
  - 26 ASCII letters
  - 9 digits
  - the dash
- Cannot begin or end with a dash
- Capitalization is ignored (everything treated as lower case)
- "`www.`" doesn't mean ANYTHING!!!