

CSE 265:


System & Network Administration

- DNS – The Domain Name System
 - What does DNS do?
 - The DNS namespace
 - BIND software
 - How DNS works
 - DNS database
 - Testing and debugging (tools)

What does DNS do?

- Provides hostname – IP lookup services
 - `www.lehigh.edu` = `128.180.2.57`
- DNS defines
 - A hierarchical namespace for hosts and IP addresses
 - A distributed database of hostname and address info
 - A “resolver” – library routines that query this database
 - Improved routing for email
 - A mechanism for finding services on a network
 - A protocol for exchanging naming information
- DNS is essential for any org using the Internet

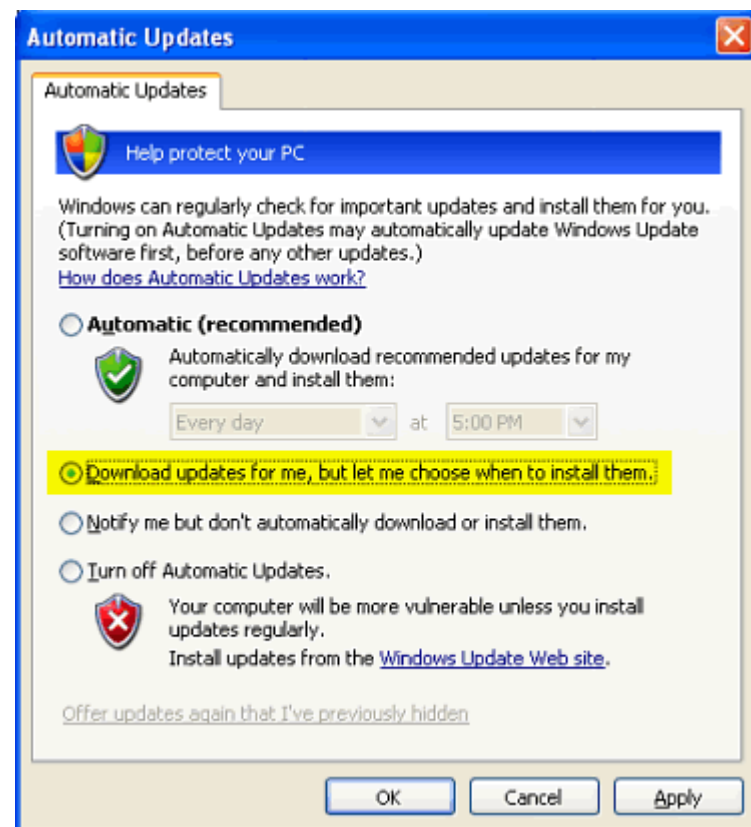

Internet Domain Survey Host Count



Source: Internet Systems Consortium (www.isc.org)


What uses DNS?

- Any application that operates over the Internet
- Such as
 - email
 - Spam filters
 - WWW
 - FTP
 - IRC, IM
 - Windows update
 - telnet, ssh




The DNS namespace

- A tree of “domains”
- Root is “.” (dot), followed by top-level (root-level) domains
- Two branches of tree
 - One maps hostnames to IP addresses
 - Other maps IP address back to hostnames
- Two types of top-level domain names used today
 - gTLDs: generic top-level domains
 - ccTLDs: country code top-level domains



Some illustrations from
O'Reilly's *DNS & Bind*



Generic top-level domains

Domain Purpose		Domain Purpose	
com	Companies	aero	Air transport industry
edu	Educational institutions	biz	Businesses
gov	(US) government agencies	coop	Cooperatives
mil	(US) military agencies	info	Unrestricted
net	Network providers	jobs	Human resources folks
org	Nonprofit organizations	museum	Museums
int	International organizations	name	Individuals
arpa	IP address lookup	pro	Professionals (attorneys, etc.)

- But today there are an abundance of top-level domains
 - .black, .blue, .airforce, .agency, .audio, etc.
- See <http://www.iana.org/domains/root/db/>

Common country codes

Code	Country	Code	Country
au	Australia	hu	Hungary
br	Brazil	jp	Japan
ca	Canada	md	Moldovia
cc	Cocos Islands	mx	Mexico
ch	Switzerland	nu	Niue
de	Germany	se	Sweden
fi	Finland	tm	Turkmenistan
fr	France	tv	Tuvalu
hk	Hong Kong	us	United States

- See <http://www.iana.org/domains/root/db/>

Domain name management

- Network Solutions (now VeriSign) used to manage .com, .org, .net, and .edu directly
- VeriSign now manages infrastructure for .com, .net, .tv, .name and .cc
 - Dozens of others manage country codes and other top-level domains
- Organizations can now register with many different registrars (even when VeriSign manages the underlying database)
- Domain holders must have two name servers authoritative for the domain



Selecting a domain name



- Most good (short) names in .com and other old gTLDs are already in use
- Domain names are up to 63 characters per segment (but a 12 character length limit is recommended), and up to 255 chars overall
- Identify two authoritative name servers
- Select a registrar, and pay ~\$1-\$35/year for registration

BIND software


- Berkeley Internet Name Domain system
 - By far, the most popular nameserver
[[Measurement Factory 2010 study](#)]
- Three components
 - a daemon called named that answers queries
 - library routines that resolve host queries by contacting DNS servers
 - command-line utilities (nslookup, dig, host)

How DNS works

- A client calls `gethostbyname()`, which is part of the resolver library
- The resolver library sends a lookup request to the first nameserver that it knows about (from `/etc/resolv.conf`)
- If the nameserver knows the answer, it sends it back to the client
- If the nameserver doesn't know, it either
 - asks the next server, or
 - returns a failure, and suggests that the client contact the next server

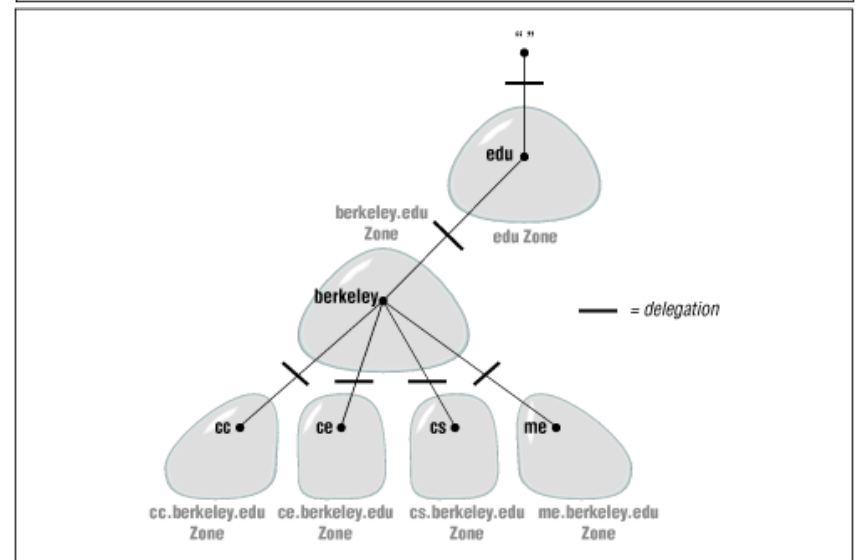

Resolving process

- 1** Name server *A* receives a query from the resolver.
- 2** *A* queries *B*.
- 3** *B* refers *A* to other name servers, including *C*.
- 4** *A* queries *C*.
- 5** *C* refers *A* to other name servers, including *D*.
- 6** *A* queries *D*.
- 7** *D* answers.
- 8** *A* returns answer to resolver.



Delegation

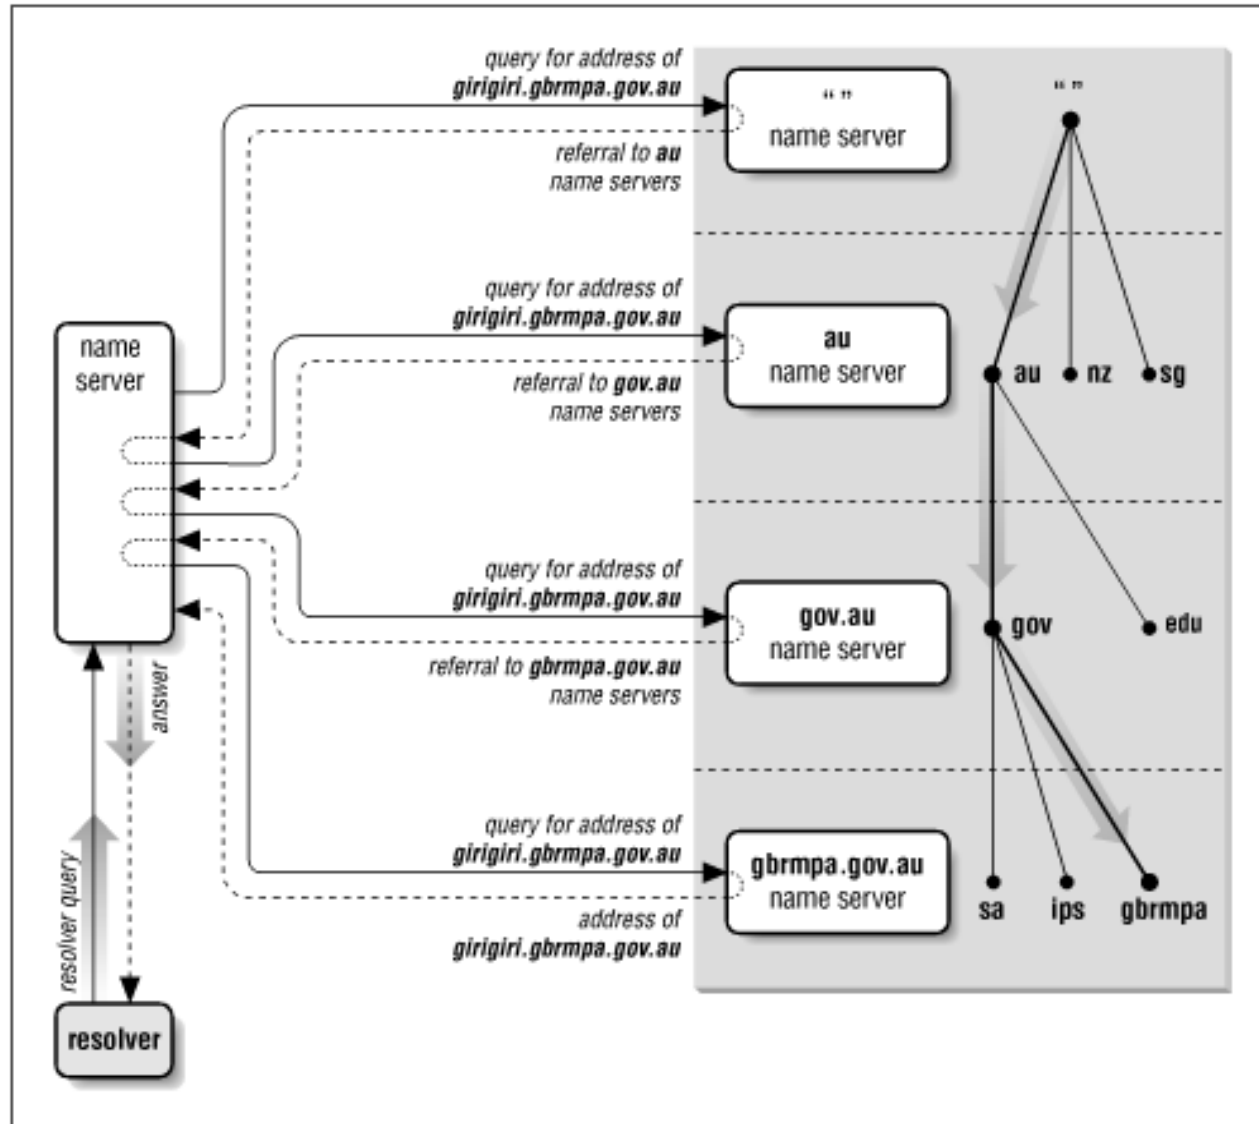
- Impractical for high-level servers to know about all hosts (or even sub-domains) below
- Servers delegate specific zones to other servers
- Names and addresses of authoritative servers for the relevant zone are returned in referrals



What servers know

- All servers know about the 13 root servers
 - hardcoded (rarely changes!), or in hint file
 - a.root-servers.net ... m.root-servers.net
- Each root server knows about servers for every top-level domain (.com, .net, .uk, etc.)
- Each top-level domain knows the servers for each second-level domain within the top-level domain
- Authoritative servers know about their hosts

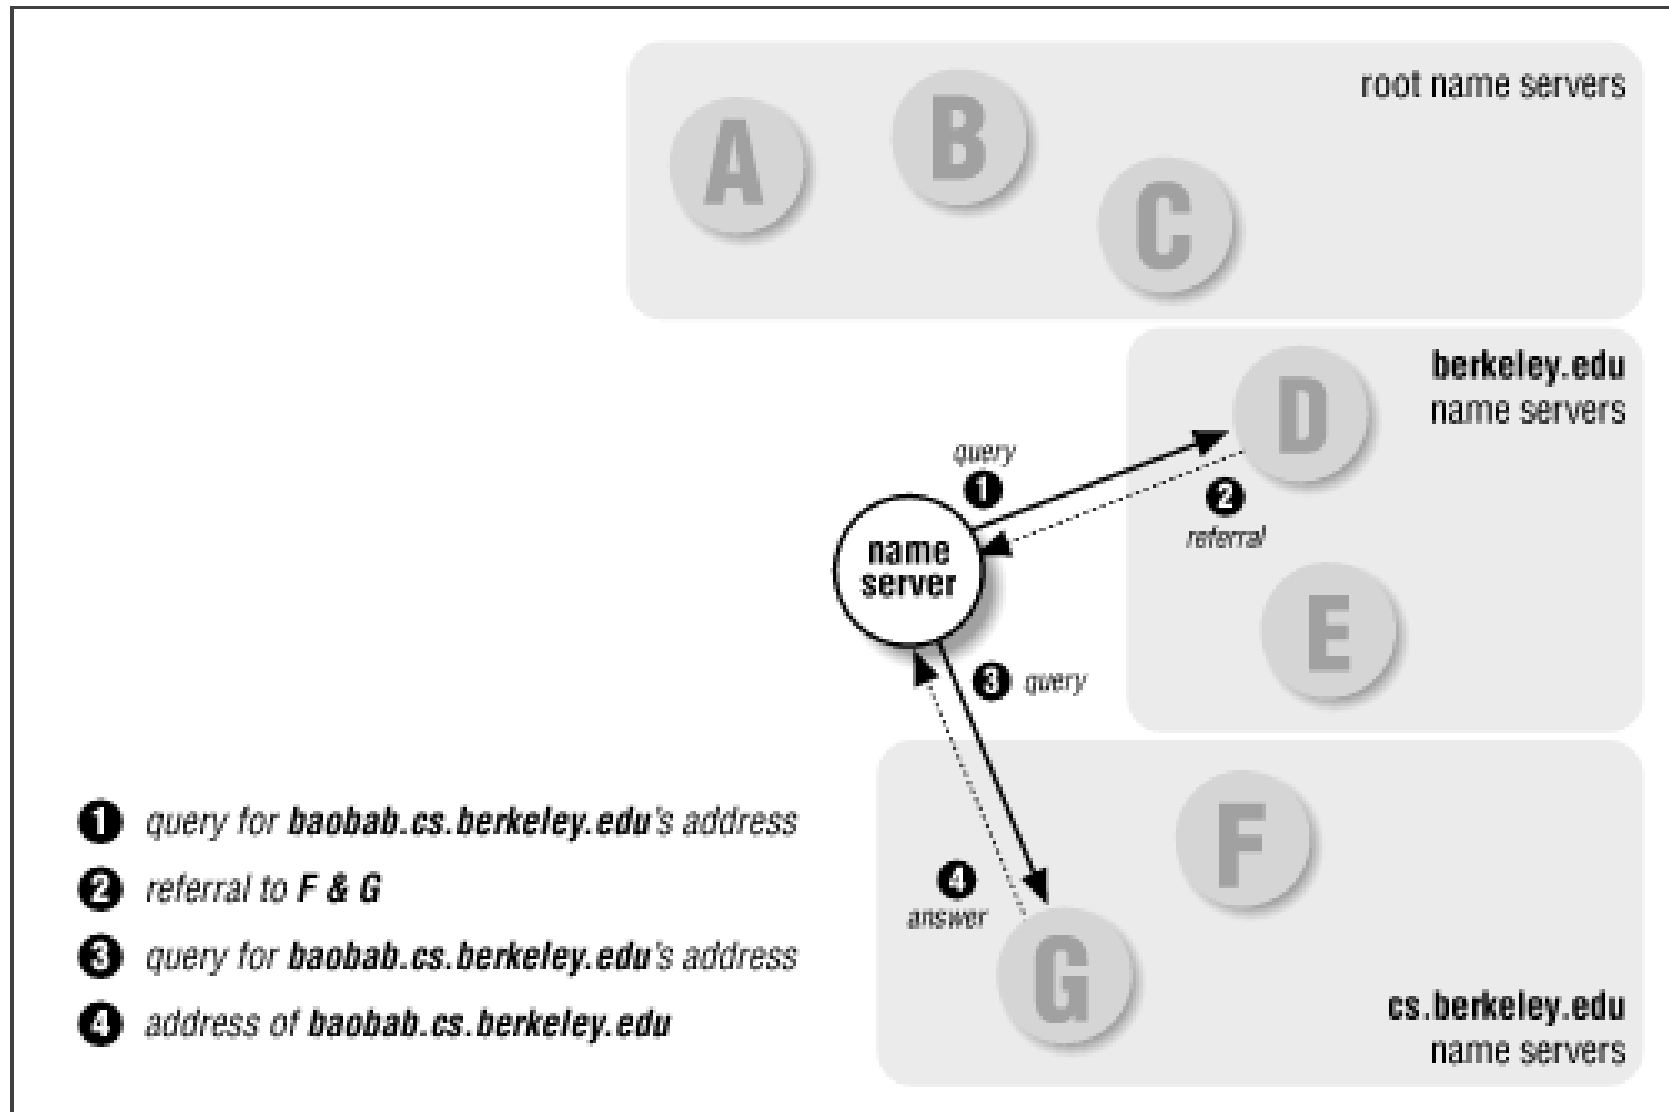
Example resolution



Types of name servers


- Recursive vs. nonrecursive servers
 - Servers that allow recursive queries will do all the work
 - Nonrecursive servers will only return referrals or answers
- Authoritative vs. caching-only servers
 - Authoritative servers have the original data
 - Caching servers retain data previously seen for future use

Caching reduces DNS load



IP-to-hostname resolution

- IP resolution works essentially the same as hostname resolution
- Query for 15.16.192.152
 - Rendered as query for 152.192.16.15.in-addr.arpa
- Each layer can delegate to the next



BIND client configuration

- Each host has `/etc/resolv.conf` which lists DNS servers
 - Can be set manually, or via DHCP
 - Example from sunlab:

```
search cse.lehigh.edu eeecs.lehigh.edu
nameserver 128.180.120.6
nameserver 128.180.120.4
nameserver 128.180.2.9
```
- Servers must be recursive, and should have a cache
- Servers are contacted in order, only after timing out previous attempt

BIND server issues

- named is typically started at boot time
- Configured using /etc/named.conf
- Can decide between
 - caching vs. authoritative
 - slave vs. master (per zone)
 - answering recursive or only iterative queries
- Lots more options
 - Who can access, what port, etc.

DNS on Linux

- Linux uses `/etc/nsswitch.conf` to determine what sources to use for name lookups

```
# /etc/nsswitch.conf
#
passwd: files nisplus
shadow: files nisplus
group: files nisplus
hosts: files dns
```

- Configuration is in `/etc/named.conf`
- Other files in `/var/named`

DNS database

- Exactly what data is stored?
- Resource records
 - Specify nameservers
 - Name to address translation
 - Address to name translation
 - Host aliases
 - Mail routing
 - Free text, location, etc.
- Format
 - [name] [ttl] [class] type data

Resource record: name

[name] [ttl] [class] type data

- name is host or domain for the record
- Absolute names must end with a dot
- Relative names do not – the current domain is added (sometimes causing mistakes!)
 - `www.cse.lehigh.edu.cse.lehigh.edu`

Resource record: ttl

[name] **[ttl]** [class] type data

- The time to live (ttl) field specifies in seconds the time that the data item may still be cached
- Increasing the ttl (say to a week) decreases traffic and DNS load substantially
- Setting a value too low can hurt web site performance
- Typical values are in days or weeks

Resource record: class

[name] [ttl] **class** type data

- Three values of class are supported
 - IN: Internet
 - default (and only one modern systems care about)
 - CH: ChaosNet
 - obsolete protocol used by obsolete machines
 - HS: Hesiod
 - database service built on top of BIND (from MIT)

Resource record: type

[name] [ttl] [class] **type** data

- Many DNS record types

- Zone

- SOA: Start of authority (define a zone)
- NS: Name server

- Basic

- A: IPv4 address (name to address translation)
- AAAA: IPv6 address (name to address translation)
- PTR: address-to-name translation
- MX: Mail exchanger

- Other

- CNAME: Canonical name (implements aliases)

SOA record

```
cs.colorado.edu 86400 IN SOA ns.cs.colorado.edu. hostmaster.cs.colorado.edu.  
(  
    2001111300 ; serial number  
    7200      ; refresh (2 hours)  
    1800     ; retry (30 minutes)  
    604800  ; expire (1 week)  
    7200 )   ; minimum (2 hours)
```

- refresh = how often slave servers must check master
- retry = when the slave will try again after failure
- expire = how long data can be considered valid without master
- minimum = TTL for cached negative answers

NS record

lehigh.edu.	86400	IN	NS	cerberus.CC.lehigh.edu.
lehigh.edu.	86400	IN	NS	spot.CC.lehigh.edu.
lehigh.edu.	86400	IN	NS	rover.CC.lehigh.edu.
cse.lehigh.edu.	86400	IN	NS	kato.eecs.lehigh.edu.
cse.lehigh.edu.	86400	IN	NS	rosie.eecs.lehigh.edu.
cse.lehigh.edu.	86400	IN	NS	cerberus.cc.lehigh.edu.
cse.lehigh.edu.	86400	IN	NS	spot.cc.lehigh.edu.
cse.lehigh.edu.	86400	IN	NS	rover.cc.lehigh.edu.

- Can't tell whether the nameserver is master or slave (but it is definitely authoritative, not caching)

A and PTR records

rover.cc.lehigh.edu.	45355	IN	A	128.180.2.9
spot.cc.lehigh.edu.	45355	IN	A	128.180.1.3
cerberus.cc.lehigh.edu.	45355	IN	A	69.7.224.17
kato.eecs.lehigh.edu.	86400	IN	A	128.180.120.6
rosie.eecs.lehigh.edu.	86400	IN	A	128.180.120.4
6.120.180.128.in-addr.arpa.	7200	IN	PTR	kato.eecs.lehigh.edu.
4.120.180.128.in-addr.arpa.	7200	IN	PTR	rosie.eecs.lehigh.edu.

- lehigh.edu and 180.128.in-addr.arpa are different zones
 - each has own SOA and resource records
- Some apps require that A and PTR records match (for authentication)

MX and CNAME records

```
piper          IN  MX  10 piper
               IN  MX  20 mailhub
               IN  MX  50 boulder.colorado.edu.
xterm1         IN  MX  10 mailhub
```

```
ftp           IN  CNAME anchor
www           IN  CNAME anchor
```

```
www.cse.lehigh.edu. 6754 IN  CNAME telstar.eecs.lehigh.edu.
```

- Every host should have MX records
- Machines that accept mail for others need to be configured to do so (e.g., mailhub)
- CNAMEs can nest eight deep in BIND

Dynamic updates to DNS

- DNS was originally designed for an environment in which hostnames (and other DNS info) changed slowly, if at all
- DHCP breaks this assumption
- Recent versions of BIND allow DHCP to notify BIND of address assignments

Testing and debugging (tools)

- named supports lots of logging options
- typical BIND tools
 - nslookup (old, possibly deprecated)
 - host
 - dig
- whois – find domain and network registration info

Other Issues

- Many aspects of DNS haven't been covered in lecture
 - Lots of details!
 - Security issues
 - IPv6
 - Internationalization – now supported!
- DNS is generally case-insensitive
- VeriSign Site Finder product
 - See <http://cyber.law.harvard.edu/tlds/sitefinder/>