

# DNS Root NameServers

## An Overview

Dr. Farid Farahmand

Updated: 9/24/12

# Who-is-Who!

- Over half million networks are connected to the Internet – 5 billion users by 2015!
- Network numbers are managed by ICANN (Internet Corporation for Assigned Names and Numbers) - <http://www.icann.org/>
  - Delegates part of address assignments to regional authorities called **registrars**
    - Registrars are authorized by ICANN to assign blocks of addresses
    - IP address blocks are given to ISPs and companies
    - ISPs **distribute** individual addresses to users and organizations

# ICANN Organization

- The Internet Corporation for Assigned Names and Numbers (ICANN)
  - ICANN is a non-profit organization
  - It is under a contract with DoC (U.S. department of commerce)
    - The United States Department of Commerce who must approve all changes requested to addressing (Zone files) by ICANN.
  - Responsible for coordinating the Internet's systems of unique identifiers, including the systems of domain names and numeric addresses that are used to reach computers on the Internet
- ICANN assigns address blocks to regional Internet registries (RIR)
  - There are five RIR (e.g., Africa or US-Canada)
  - In U.S. RIR is called **The American Registry for Internet Numbers (ARIN)**

# IANA Function

- The ICANN is under contract (since 1998) with the United States Department of Commerce to perform the **IANA function**
  - Internet Assigned Numbers Authority – IANA
- The IANA functions includes
  - Internet Protocol (IP) address space allocation,
  - protocol identifier assignment
  - **generic** (gTLD) and **country code** (ccTLD) Top-Level Domain name system management
  - root server system management functions

# ARIN & AS

- In U.S. Regional Internet Registries is called **The American Registry for Internet Numbers** (ARIN)
- ARIN manages the distribution of Internet number resources, including IPv4 and IPv6 address space and AS numbers
  - **Autonomous System** (AS) is a collection of connected Internet Protocol (IP) routing prefixes under the control of one or more network operators
  - Example: AT&T has AS# 7018
  - Border Gateway Protocol (BGP) uses the AS# for routing purposes

AS #	Provider
701	UUnet (U.S. domestic) (AS 701-705)
1239	Sprintlink U.S. Domestic
3356	Level 3
7018	AT&T WorldNet
209	Qwest
3561	Cable and Wireless (aq'd by SAVVIS)
3549	Global Crossing
2914	Verio
6461	AboveNet
702	UUnet (International)
1299	TeliaNet
5511	OpenTransit
5459	LINX
16631	Cogent
6453	Teleglobe

# Nameserver

- The entire Internet is managed through special **hierarchical addressing** system
- In order to reach a destination, each request must find out about the IP address of the domain (destination's physical location) it is trying to reach
- Thus, before sending a request, the source must perform a query to learn about the IP address of the destination node
  - The queries (questions) are sent to authoritative **nameservers**
- An **authoritative nameserver** is a name server that gives answers in response to questions asked about names in a zones
  - Authoritative only
    - Only answer to queries about a zone
  - Caching name server
    - They are configured to give authoritative answers to queries for some zones and act as a caching name server for all other zones.
- **DNS zones** may consist of only one domain, or may comprise many domains and sub-domains
  - Each Zone is defined by a Zone File
- A **Zone File** contains specification for host addressing, name aliasing, electronic mail routing, backup server systems, geographic location, administrative contacts, and many other pieces of information
  - Each entry has a DNS record types (e.g., A=address record; MX=Mail exchange record)
- The **Root Zone** is controlled by the United States Department of Commerce who must approve all changes to the root zone file requested by ICANN.

# A fully qualified domain name (FQDN)

- A fully qualified domain name (FQDN) is a domain name that specifies its **exact location** in the tree hierarchy of the Domain Name System (DNS)
  - It is an **authoritative name server**
  - It specifies all domain levels
  - For example, given a device with a local hostname myhost and a parent domain name example.com, the fully qualified domain name is **myhost.example.com**
  - The FQDN therefore uniquely identifies the device —while there may be many hosts in the world called myhost, there can only be one myhost.example.com.
  - In DNS zone files, a fully qualified domain name is specified with a trailing dot. For example, **myhost.example.com.**

# BIND Software

- The obvious question is how does DNS operation actually take place?
  - Using DNS software
- **Berkeley Internet Name Domain**(BIND) is the de facto standard for running DNS on Unix-like OS
  - Developed by four graduate students at the Computer Systems Research Group at Berkeley
- A new version of BIND (**BIND 9**) was written by the ISC (**Internet Systems Consortium, Inc.**, ) from scratch
  - Included new features: IPv6, remote name daemon control, etc.
- All Zone-files, thus follow BIND-style

- 1.1 BIND
- 1.2 Microsoft DNS
- 1.3 Dnsmasq
- 1.4 djbdns
- 1.5 Simple DNS Plus
- 1.6 NSD
- 1.7 PowerDNS
- 1.8 MaraDNS
- 1.9 Nominum ANS
- 1.10 Nominum Vantio
- 1.11 Posadis
- 1.12 Unbound
- 1.13 pdnsd
- 1.14 Cisco Network Registrar
- 1.15 Domain Name Relay Daemon (dnrd)
- 1.16 Geographic DNS daemon (gdnsd)



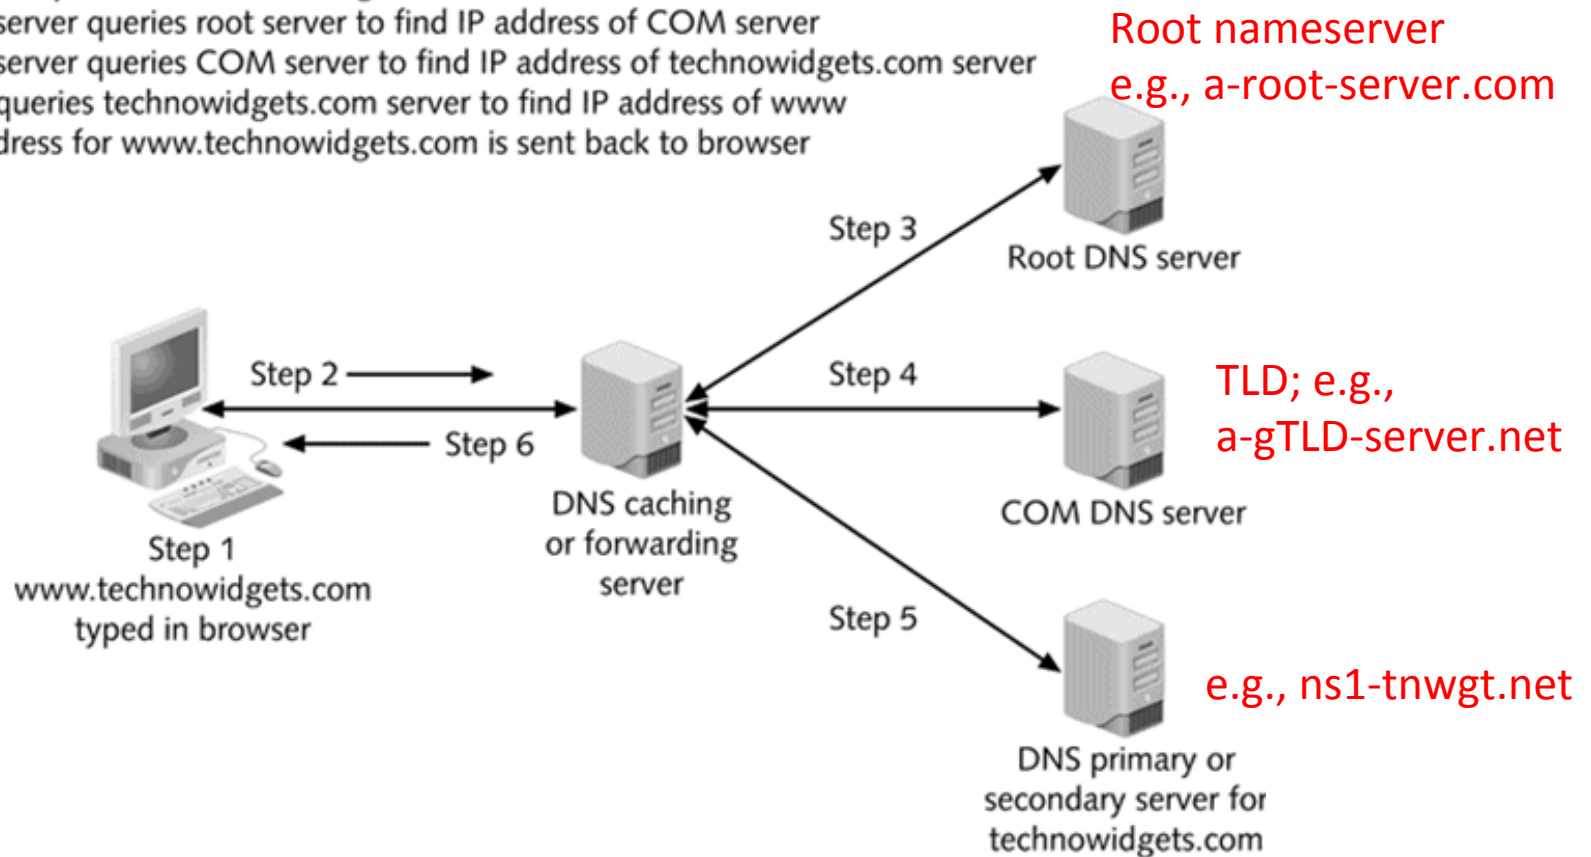
# NSD Software

- Another notable software is NSD for **name server daemon**
  - Daemon is a background process that handles requests for service
- NSD is an open-source server program for the Domain Name System
  - Developed by NLnet Labs of Amsterdam
  - Uses the standard TCP/UDP port 53
  - Latest version is 3
  - Main advantage is **more efficient memory** usage: e.g., for serving domains, NSD can save significant RAM space (PROJECT IDEA)
  - Remember: It is all about cache!
- Three root nameservers have switched from BIND to NSD
  - k.root-servers.net
  - h.root-servers.net (there are three H1, H2, H3)
  - l.root-servers.net

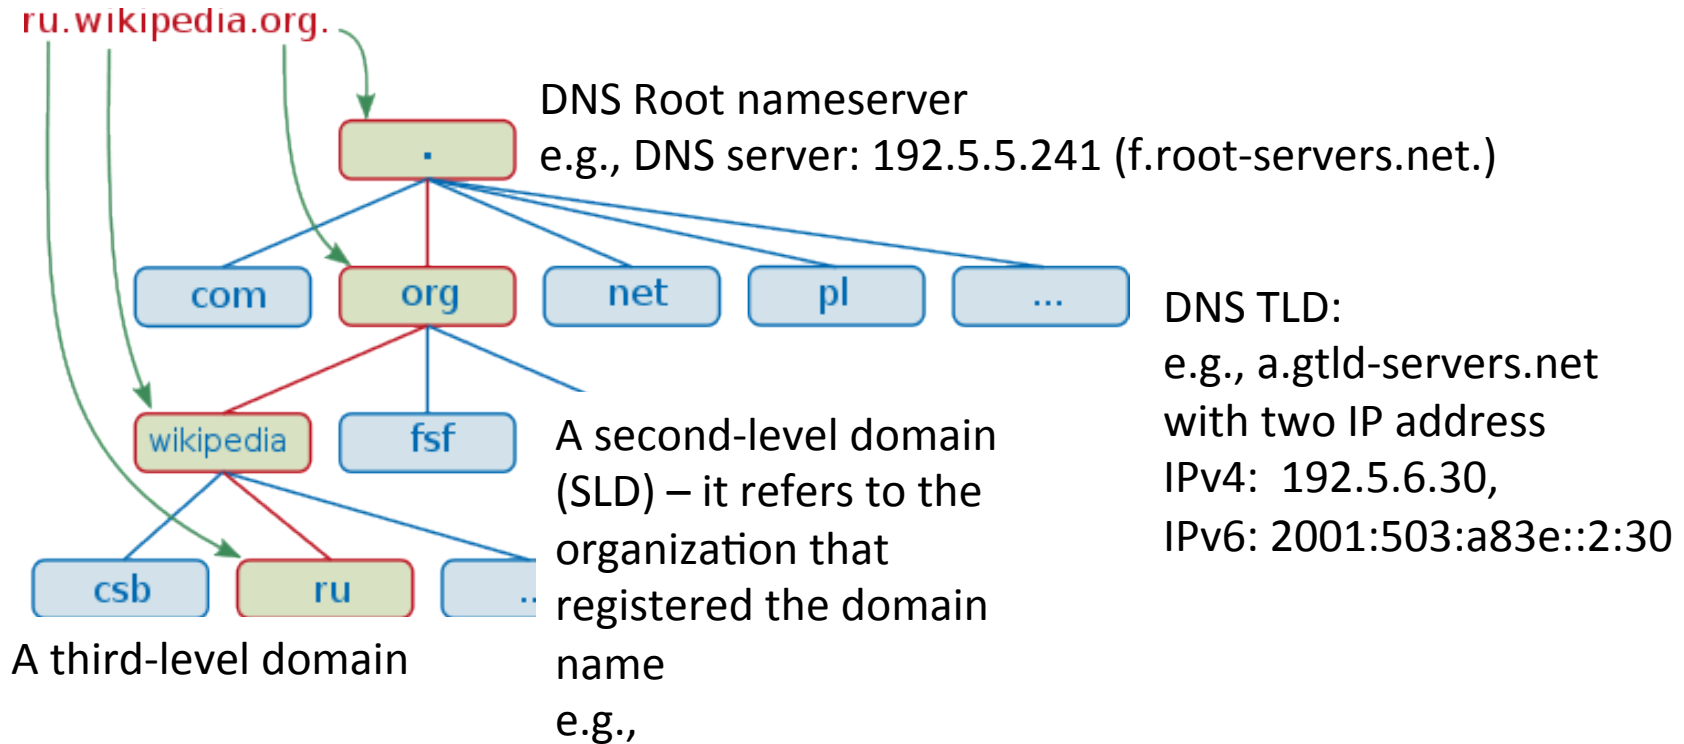
- 1.1 BIND
- 1.2 Microsoft DNS
- 1.3 Dnsmasq
- 1.4 djbdns
- 1.5 Simple DNS Plus
- 1.6 NSD
- 1.7 PowerDNS
- 1.8 MaraDNS
- 1.9 Nominum ANS
- 1.10 Nominum Vantio
- 1.11 Posadis
- 1.12 Unbound
- 1.13 pdnsd
- 1.14 Cisco Network Registrar
- 1.15 Domain Name Relay Daemon (dnrd)
- 1.16 Geographic DNS daemon (gdnsd)

# Finding the IP Address for a Domain (Name Resolution)

1. User types `www.technowidgets.com` in browser
2. Browser queries DNS server to get IP address
3. DNS server queries root server to find IP address of COM server
4. DNS server queries COM server to find IP address of `technowidgets.com` server
5. DNS queries `technowidgets.com` server to find IP address of `www`
6. IP address for `www.technowidgets.com` is sent back to browser



# Example of Hierarchical Naming



# Example of Hierarchical Naming

ru.wikipedia.org.



DNS Root nameserver

e.g., DNS server: 192.5.5.241 (f.root-servers.net.)

```
village-158-231:~ farid11$ nslookup 192.5.5.241
Server:          130.157.27.50
Address:         130.157.27.50#53

Non-authoritative answer:
241.5.5.192.in-addr.arpa      name = f.root-servers.net.

Authoritative answers can be found from:
5.5.192.in-addr.arpa        nameserver = sfba.sns-pb.isc.org.
5.5.192.in-addr.arpa        nameserver = ams.sns-pb.isc.org.
5.5.192.in-addr.arpa        nameserver = ord.sns-pb.isc.org.
sfba.sns-pb.isc.org         internet address = 149.20.64.3
sfba.sns-pb.isc.org         has AAAA address 2001:4f8:0:2::19
ord.sns-pb.isc.org          internet address = 199.6.0.30
ord.sns-pb.isc.org          has AAAA address 2001:500:71::30
ams.sns-pb.isc.org          internet address = 199.6.1.30
ams.sns-pb.isc.org          has AAAA address 2001:500:60::30
```

DNS TLD:

e.g., a.gtld-servers.net  
with two IP address

IPv4: 192.5.6.30,

IPv6: 2001:503:a83e::2:30

# Root nameservers

There are currently **13 root name servers** specified, with names in the form *letter.root-servers.net*, where *letter* ranges from A to M.

```
> server 213.199.161.77
(root) nameserver = j.root-servers.net
(root) nameserver = k.root-servers.net
(root) nameserver = l.root-servers.net
(root) nameserver = m.root-servers.net
(root) nameserver = a.root-servers.net
(root) nameserver = b.root-servers.net
(root) nameserver = c.root-servers.net
(root) nameserver = d.root-servers.net
(root) nameserver = e.root-servers.net
(root) nameserver = f.root-servers.net
(root) nameserver = g.root-servers.net
(root) nameserver = h.root-servers.net
(root) nameserver = i.root-servers.net
Default Server: [213.199.161.77]
Address: 213.199.161.77

> status
Server: [213.199.161.77]
Address: 213.199.161.77

Name: status
Addresses: 65.55.39.12
           207.46.31.61
```














There are currently **13 root name servers** specified, with names in the form ***letter.root-servers.net***, where *letter* ranges from A to M.


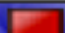
Letter	IPv4 address	IPv6 address	AS-number <sup>[3]</sup>	Old name	Operator	Location #sites (global/local) <sup>[4]</sup>	Software
A	198.41.0.4	2001:503:ba3e::2:30	AS19836 <a href="#">↗</a>	ns.internic.net	Verisign	Distributed using <i>anycast</i> 6/0	BIND
B	192.228.79.201 (since January 2004; originally was 128.9.0.107) <sup>[5]</sup>	2001:478:65::53 (not in root zone yet)	none	ns1.isi.edu	USC-ISI	Marina Del Rey, California, U.S. 0/1	BIND
C <a href="#">↗</a>	192.33.4.12	2001:500:2::c (not in root zone yet)	AS2149 <a href="#">↗</a>	c.psi.net	Cogent Communications	Distributed using <i>anycast</i> 6/0	BIND
D <a href="#">↗</a>	128.8.10.90	2001:500:2d::d	AS27 <a href="#">↗</a>	terp.umd.edu	University of Maryland	College Park, Maryland, U.S. 1/0	BIND
E	192.203.230.10	N/A	AS297 <a href="#">↗</a>	ns.nasa.gov	NASA	Mountain View, California, U.S. 1/0	BIND
F <a href="#">↗</a>	192.5.5.241	2001:500:2f:f	AS3557 <a href="#">↗</a>	ns.isc.org	Internet Systems Consortium	Distributed using <i>anycast</i> 2/47	BIND <sup>g[6]</sup>
G <a href="#">↗</a>	192.112.36.4	N/A	AS5927 <a href="#">↗</a>	ns.nic.ddn.mil	Defense Information Systems Agency	Distributed using <i>anycast</i> 6/0	BIND
H <a href="#">↗</a>	128.63.2.53	2001:500:1::803f:235	AS13 <a href="#">↗</a>	aos.arl.army.mil	U.S. Army Research Lab	Aberdeen Proving Ground, Maryland, U.S. 2/0	NSD
I <a href="#">↗</a>	192.36.148.17	2001:7fe::53	AS29216 <a href="#">↗</a>	nic.nordu.net	Autonomica	Distributed using <i>anycast</i> 36	BIND
J	192.58.128.30 (since November 2002; originally was 198.41.0.10)	2001:503:c27::2:30	AS26415 <a href="#">↗</a>		Verisign	Distributed using <i>anycast</i> 63/7	BIND
K <a href="#">↗</a>	193.0.14.129	2001:7fd::1	AS25152 <a href="#">↗</a>		RIPE NCC	Distributed using <i>anycast</i> 5/13	NSD <sup>[7]</sup>
L <a href="#">↗</a>	199.7.83.42 (since November 2007; originally was 198.32.64.12) <sup>[8]</sup>	2001:500:3::42	AS20144 <a href="#">↗</a>		ICANN	Distributed using <i>anycast</i> 37/1	NSD <sup>[9]</sup>
M <a href="#">↗</a>	202.12.27.33	2001:dc3::35	AS7500 <a href="#">↗</a>		WIDE Project	distributed using <i>anycast</i> 5/1	BIND

IN\_MY\_MAC:~ farid11\$ dig



There are currently **13 root name servers** specified, with names in the form ***letter.root-servers.net***, where ***letter*** ranges from A to M.

No.	Root Server	Location	IP Address	Bind	Icann-Root	Public-Root
1	a.public-root.net	Amsterdam, Netherlands	84.22.100.2		Resolving	Resolving
2	b.public-root.net	Budapest, Hungary	79.172.201.120		Resolving	Resolving
3	c.public-root.net	Paris, France	91.121.45.127		<b>Failed</b>	<b>Failed</b>
4	d.public-root.net	London, UK	80.252.121.2		Resolving	Resolving
5	e.public-root.net	Kelowna, British Columbia, Canada	209.97.202.107		<b>Time Out</b>	<b>Time Out</b>
6	f.public-root.net	Melbourne, Victoria, Australia	84.22.100.250		<b>Time Out</b>	<b>Time Out</b>
7	g.public-root.net	Chicago, Illinois, USA	199.5.157.131		Resolving	Resolving
8	h.public-root.net	Des Moines, Iowa, USA	208.71.35.137		Resolving	Resolving
9	i.public-root.net	Chennai, Tamilnadu, India	122.183.133.220		<b>Time Out</b>	<b>Time Out</b>
10	j.public-root.net	Singapore	84.22.100.89		Resolving	Resolving
11	k.public-root.net	Moscow, Russia	82.146.40.113		<b>Time Out</b>	<b>Time Out</b>
12	l.public-root.net	Tehran, Iran	91.186.213.6		Resolving	Resolving
13	m.public-root.net	Chimbote, Peru	200.37.61.62		<b>Time Out</b>	<b>Time Out</b>

	=	DNS is Up
	=	DNS is Down

Resolving	=	Resolving Root
<b>Time Out</b>	=	Unable to Resolve Root

# Who Controls the Nameserver

- As an example **Internet Systems Consortium** (ISC) operates one of the 13 global authoritative DNS root servers, **F-root**
- This server uses BIND software
- It has two addresses (IPv4 and IPv6)

<a href="#">F</a>	192.5.5.241	2001:500:2f:f	<a href="#">AS3557</a>	<a href="#">ns.isc.org</a>	<a href="#">Internet Systems Consortium</a>	Distributed using anycast 2/47	<a href="#">BIND</a>
-------------------	-------------	---------------	------------------------	----------------------------	---	-----------------------------------	----------------------




# Who is ISC

## Internet Systems Consortium

- ISC, is a non-profit corporation (in Delaware) supporting the infrastructure of the Internet
- It develops and maintains software, protocols, and operations
  - For example they developed BIND 9.0 and Dynamic Host Configuration Protocol (DHCP)

# How many gTLD (generic Top-Level Domain)?



# gTLD

[http://en.wikipedia.org/wiki/Generic\\_top-level\\_domain](http://en.wikipedia.org/wiki/Generic_top-level_domain)

Name <sup>[k]</sup>	Entity <sup>[k]</sup>	Notes
.aero	air-transport industry	Must verify eligibility for registration; only those in various categories of air-travel-related entities may register.
.asia	Asia-Pacific region	This is a TLD for companies, organizations, and individuals based in the region of Asia, Australia, and the Pacific.
.biz	business	This is an open TLD; any person or entity is permitted to register; however, registrations may be challenged later if they are not held by commercial entities in accordance with the domain's charter. This TLD was created to provide relief for the wildly-popular .com TLD.
.cat	Catalan	This is a TLD for Web sites in the <a href="#">Catalan language</a> or related to Catalan culture.
.com	commercial	This is an open TLD; any person or entity is permitted to register. Though originally intended for-profit business entities, for a number of reasons it became the "main" TLD for domain names and is currently used by all types of entities including nonprofits, schools and private individuals. Domain name registrations may be challenged if the holder cannot prove an outside relation justifying reservation of the name, to prevent "squatting".
.coop	cooperatives	The .coop TLD is limited to cooperatives as defined by the <a href="#">Rochdale Principles</a> .
.edu	educational	The .edu TLD is limited to specific educational institutions such as, but not limited to, primary schools, middle schools, secondary schools, colleges, and universities. In the US, its usability was limited in 2001 to post-secondary institutions accredited by an agency on the list of <a href="#">nationally recognized accrediting agencies</a> maintained by the <a href="#">United States Department of Education</a> . This domain is therefore almost exclusively used by U.S. colleges and universities. Some institutions that do not meet the current registration criteria have <a href="#">grandfathered</a> domain names.
.gov	governmental	The .gov TLD is limited to U.S. governmental entities and agencies.
.info	information	This is an open TLD; any person or entity is permitted to register.
.int	international organizations	The .int TLD is strictly limited to organizations, offices, and programs which are <a href="#">endorsed by a treaty between two or more nations</a> .
.jobs	companies	The .jobs TLD is designed to be added after the names of established companies with jobs to advertise. At this time, owners of a "company.jobs" domain are not permitted to post jobs of third party employers.
.mil	U.S. military	The .mil TLD is limited to use by the U.S. military.
.mobi	mobile devices	Must be used for mobile-compatible sites in accordance with standards.
.museum	museums	Must be verified as a legitimate museum.
.name	individuals, by name	This is an open TLD; any person or entity is permitted to register; however, registrations may be challenged later if they are not by individuals (or the owners of fictional characters) in accordance with the domain's charter
.net	network	This is an open TLD; any person or entity is permitted to register. Originally intended for use by domains pointing to a distributed network of computers, or "umbrella" sites that act as the portal to a set of smaller websites
.org	organization	This is an open TLD; any person or entity is permitted to register. Originally intended for use by non-profit organizations, and still primarily used by same.
.pro	professions	Currently, .pro is reserved for licensed or certified lawyers, accountants, physicians and engineers in France, Canada, NL, UK and the U.S. A professional seeking to register a .pro domain must provide their registrar with appropriate credentials.
.tel	Internet communication services	A contact directory housing all types of contact information directly in the Domain Name System.
.travel	travel and tourism industry related sites	Must be verified as a legitimate travel-related entity.
.xxx	adult entertainment	For sites providing sexually-explicit content, such as pornography.



# Testing the Public Root Servers

- Go to <http://public-root.com>
- Do Root-Server-Check and examine which Root Servers are up
- Do Root-Server-Location and see where they are located at
- Get information about E.Root.Server – Where is it?

# A Practical Example!

- When you visit a Web site, you need the DNS server to resolve your requested domain name.
- The DNS server of your workstation queries for name resolution and it is typically run by your ISP
- If you find out that the DNS server is too slow, you can change your DNS!!

# A Practical Example! – cont.

- Using my MacBook Pro I did:
  - cat /etc/resolv.conf

```
#  
# Mac OS X Notice  
#  
# This file is not used by the host name and address resolution  
# or the DNS query routing mechanisms used by most processes on  
# this Mac OS X system.  
#  
# This file is automatically generated.  
#  
domain sbx02888.rohneca.wayport.net  
nameserver 192.168.5.1
```

I am connected to Wayport.net machines  
The dynamic DNS that I have received is 192.168.5.1 –  
This is where my machine goes and make query

Let's say the DNS ended up being very slow. So, I want to change it to another machine which is faster so I can brows quicker!

I decided to use Google Public DNS, instead (8.8.8.8)

[http://www.itistimed.com/?DATA=8.8.8.8&ACTION\\_TYPE=Resolve](http://www.itistimed.com/?DATA=8.8.8.8&ACTION_TYPE=Resolve)

Go to [http://www.plus.net/support/software/dns/changing\\_dns\\_mac.shtml](http://www.plus.net/support/software/dns/changing_dns_mac.shtml)

To learn how to change your DNS in your MAC.

# A Practical Example! – cont.

```
village-158-231:etc farid11$ time nslookup bbc.com 8.8.8.8
Server:          8.8.8.8
Address:         8.8.8.8#53
```

```
Non-authoritative answer:
Name:   bbc.com
Address: 212.58.241.131
```

```
real    0m0.078s
```

```
user    0m0.001s
```

```
sys     0m0.004s
```

```
village-158-231:etc farid11$ time nslookup bbc.com
Server:          130.157.27.50
Address:         130.157.27.50#53
```

```
Non-authoritative answer:
Name:   bbc.com
Address: 212.58.241.131
```

```
real    0m0.159s
```

```
user    0m0.001s
```

```
sys     0m0.004s
```

This is using the default DNS!

Let's say the DNS ended up being very slow. So, I want to change it to another machine which is faster so I can brows quicker!

I decided to use Google Public DNS, instead (8.8.8.8)

[http://www.itistimed.com/?DATA=8.8.8.8&ACTION\\_TYPE=Resolve](http://www.itistimed.com/?DATA=8.8.8.8&ACTION_TYPE=Resolve)

Go to

[http://www.plus.net/support/software/dns/changing\\_dns\\_mac.shtml](http://www.plus.net/support/software/dns/changing_dns_mac.shtml)

To learn how to change your DNS in your MAC.



# Commands

```
47229q-shz2010a:~ farid11$ time nslookup www.google.com 125.22.47.125
Server:          125.22.47.125
Address:         125.22.47.125#53
```

```
Non-authoritative answer:
www.google.com canonical name = www.l.google.com.
Name:   www.l.google.com
Address: 74.125.113.104
Name:   www.l.google.com
Address: 74.125.113.105
Name:   www.l.google.com
Address: 74.125.113.147
Name:   www.l.google.com
Address: 74.125.113.99
Name:   www.l.google.com
Address: 74.125.113.103
Name:   www.l.google.com
Address: 74.125.113.106
```

```
real    0m0.015s
user    0m0.001s
sys     0m0.004s
```

- Here are a series of command I used on my MAC to measure the Address Resolution using different DNS servers:
  - *dscacheutil -flushcache // flush the cache*
  - *time nslookup www.google.com 125.22.47.125*
  - *time nslookup www.google.com 208.67.222.222*
- I used these to compare the performance of the two DNS servers

Thus, for DNS server 125.22.47.125, it took 15 millisecond to resolve my Google query!!

Here is the information about the DNS server: [http://www.itistimed.com/?DATA=125.22.27.125&ACTION\\_TYPE=Resolve](http://www.itistimed.com/?DATA=125.22.27.125&ACTION_TYPE=Resolve)

# DNS Cache Poisoning

- DNS cache poisoning is a data integrity compromise in the Domain Name System (DNS)
- Read:  
[http://adventuresinsecurity.com/Papers/DNS\\_Cache\\_Poisoning.pdf](http://adventuresinsecurity.com/Papers/DNS_Cache_Poisoning.pdf)
- Short Video:  
<http://www.youtube.com/watch?v=1d1tUefYn4U>
- This is a nice demo if you can follow it:
  - <http://www.videosurf.com/video/dns-cache-poisoning-demo-1240529251>

# References

- Learn about Google DNS  
<http://code.google.com/speed/public-dns/>
- Free DNS servers
- <http://theos.in/windows-xp/free-fast-public-dns-server-list/>