

# Network Service Management & Virtualization

DNS

# Outline

- DNS Overview
- Introduction to DNS Design

# DNS – starter activity (4 mins)

- What is it?
- When is it used?
- Is it important?
- What if it were not to work? Would the Internet still function like we know it today?

# Naming

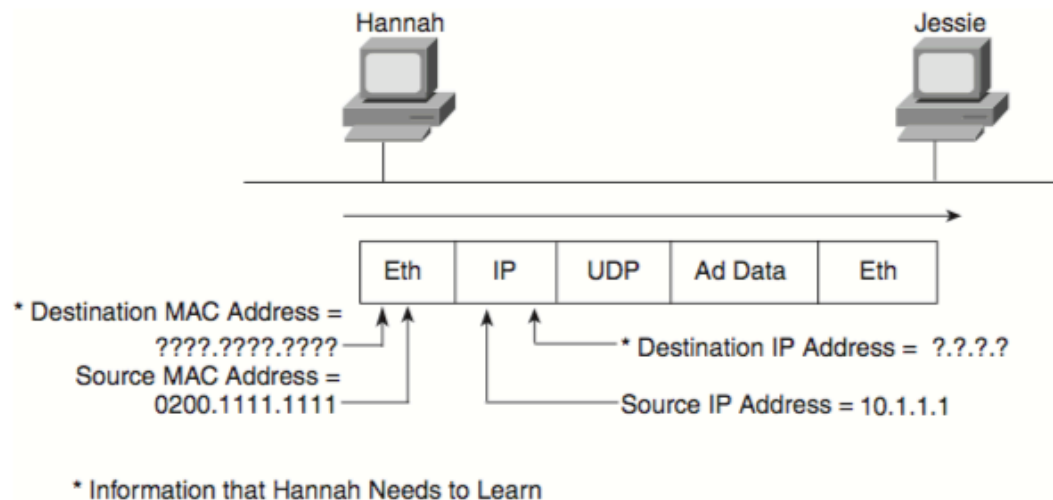
- The Domain Name System (DNS) is one of the most essential parts of the Internet's infrastructure.
- By using DNS, you can connect to a website like `www.google.com` without having to know the website's IP address. (After all, who wants to remember every website's IP address?)
- How do we efficiently locate resources?
  - DNS: name → IP address
- Challenge
  - How do we scale these to the wide area?

# DNS, ARP & recap of DHCP! (1)

- Network designers should try to make using the network as simple as possible. At most, users might want to remember the name of another computer with which they want to communicate, such as remembering the name of a website.
- They certainly do not want to remember the IP address, nor do they want to try to remember any MAC addresses!
- So, TCP/IP needs protocols that dynamically discover all the necessary information to allow communications, without the user knowing more than a name.
- User typically identifies a remote computer by using a name.

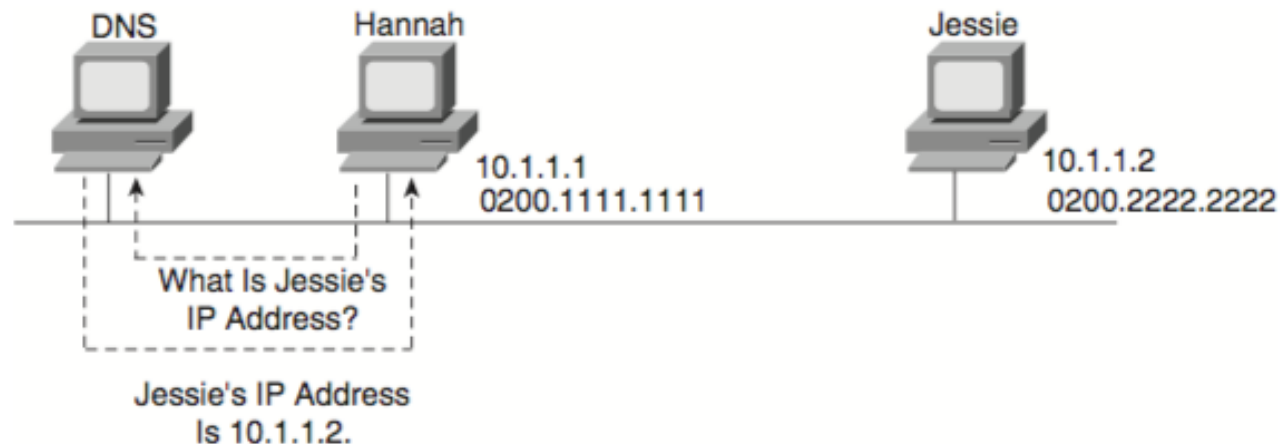
# DNS, ARP & recap of DHCP! (2)

- TCP/IP needs a way to let a computer find the IP address of another computer based on its name.
- TCP/IP also needs a way to find MAC addresses associated with other computers on the same LAN subnet.
- Hannah Knows Jessie's Name, Needs IP Address and MAC Address




# DNS, ARP & recap of DHCP! (3)

- Hannah knows the IP address of a DNS server because the address was either preconfigured on Hannah's machine or was learned with DHCP
- As soon as Hannah somehow identifies the name of the other computer (for example, jessie.example.com), she sends a DNS request to the DNS, asking for Jessie's IP address. The DNS replies with the address, 10.1.1.2.



# DNS, ARP & recap of DHCP! (4)

- Effectively, the same thing happens when you surf the Internet and connect to any website. Your PC sends a request, just like Hannah's request for Jessie, asking the DNS to resolve the name into an IP address. After that happens, your PC can start requesting that the web page be sent.

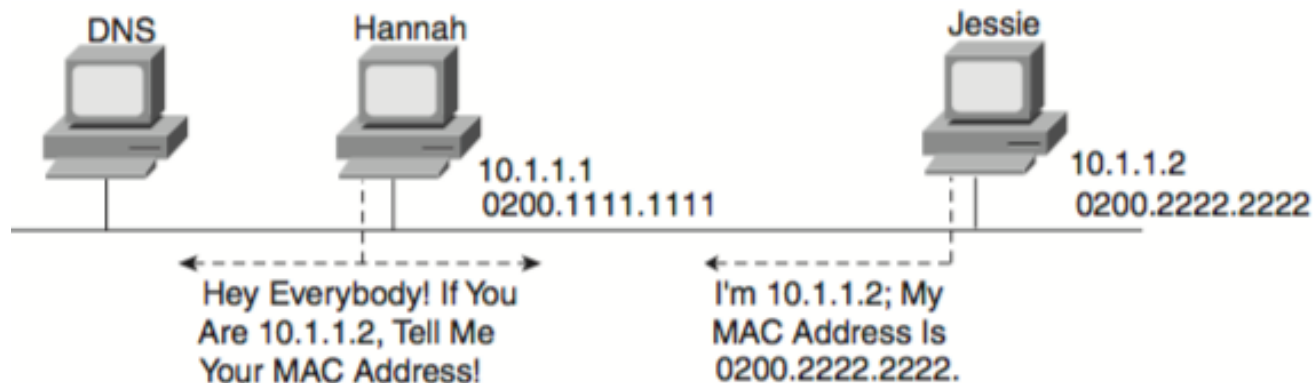




# DNS, ARP & recap of DHCP! (5)

## The ARP Process

- As soon as a host knows the IP address of the other host, the sending host may need to know the MAC address used by the other computer. For example, Hannah still needs to know the Ethernet MAC address used by 10.1.1.2, so Hannah issues something called an ARP broadcast. An ARP broadcast is sent to a broadcast Ethernet address, so everyone on the LAN receives it. Because Jessie is on the same LAN, she receives the ARP broadcast. Because Jessie's IP address is 10.1.1.2 and the ARP broadcast is looking for the MAC address associated with 10.1.1.2, Jessie replies with her own MAC address.



# DNS, ARP & recap of DHCP! (6)

## The ARP Process

- Now Hannah knows the destination IP and Ethernet addresses that she should use when sending frames to Jessie, and the packet can be sent successfully.
- Hosts may or may not need to ARP to find the destination host's MAC address based on the two-step routing logic used by a host.
- **If the destination host is on the same subnet, the sending host sends an ARP looking for the destination host's MAC address**

# DNS, ARP & recap of DHCP! (7)

## The ARP Process

- **However, if the sending host is on a different subnet than the destination host, the sending host's routing logic results in the sending host needing to forward the packet to its default gateway.** For example, if Hannah and Jessie had been in different subnets, Hannah's routing logic would have caused Hannah to want to send the packet to Hannah's default gateway (router).
- In that case, Hannah would have used ARP to find the router's MAC address instead of Jessie's MAC address.
- Additionally, hosts need to use ARP to find MAC addresses only once in a while. Any device that uses IP should retain, or cache, the information learned with ARP, placing the information in its ARP cache. Each time a host needs to send a packet encapsulated in an Ethernet frame, it first checks its ARP cache and uses the MAC address found there. If the correct information is not listed in the ARP cache, the host then can use ARP to discover the MAC address used by a particular IP address. Also, a host learns ARP information when receiving an ARP as well. For example, the ARP process shown in results in both Hannah and Jessie learning the other host's MAC address. (use `arp -a` on PC to see arp cache)

# Potential Solution (1) – Local file

## Why not use /etc/hosts?

- Original Name to Address Mapping
  - Flat namespace
  - SRI kept main copy (The Stanford Research Institute (now SRI International) maintained a text file named HOSTS.TXT that mapped host names to the numerical addresses of computers on the ARPANET)
  - Downloaded regularly
- Count of hosts was increasing: machine per domain → machine per user
  - Many more downloads
  - Many more updates

# Potential Solution (2) - Centralize Server

Why not centralize DNS?

- One place where all mappings are stored
- All queries go to the central server

Many practical problems:

- Single point of failure (SPOF)
- High Traffic volume
- Distant centralized database
- Single point of update
  
- Doesn't *scale!*
- Need a distributed, hierarchical collection of servers!

# Domain Name System

- Properties of DNS
  - Hierarchical name space divided into zones
  - Distributed over a collection of DNS servers
- Hierarchy of DNS servers
  - Root DNS servers
  - Top-Level Domain (TLD) DNS servers
  - Authoritative DNS servers
- For performing translations
  - Local DNS servers
  - Resolver software

# Domain Name System Goals

- Basically a wide-area distributed database
  - The sheer size of the database and frequency of updates suggest that it must be maintained in a distributed manner, with local caching to improve performance.
- Scalability
- Decentralized maintenance
- Robustness
- Global scope
  - Names mean the same thing everywhere
- Don't need
  - Atomicity
  - Strong consistency

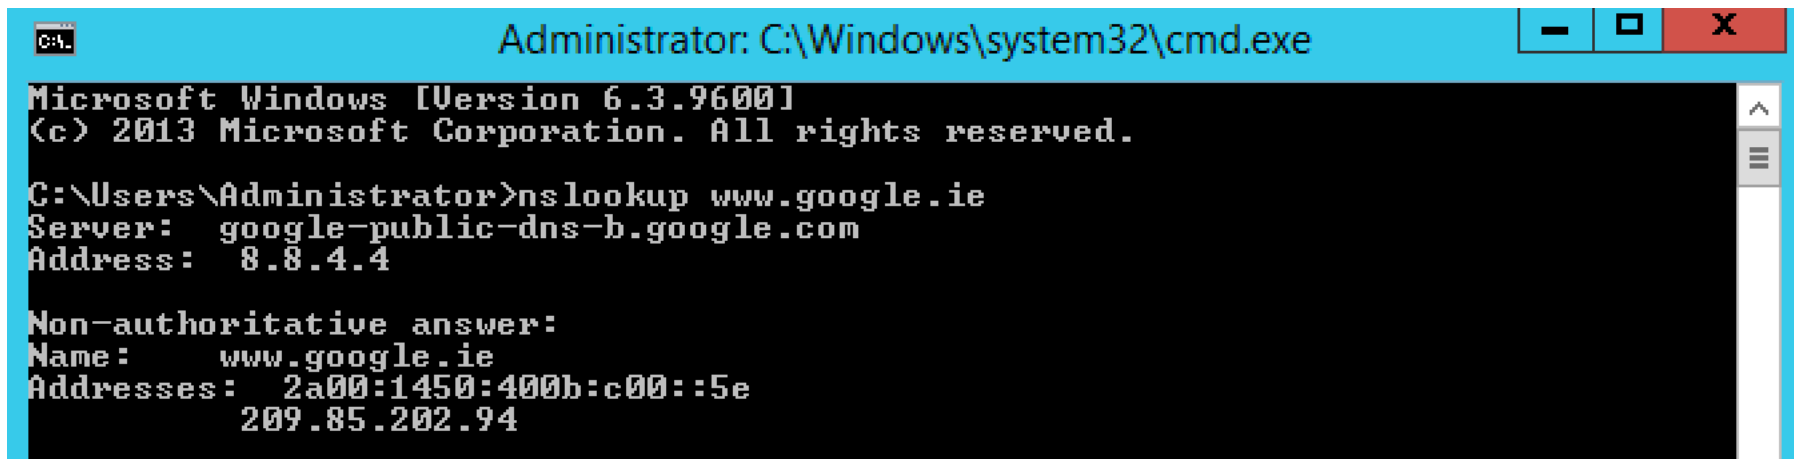
# Prep for activity – nslookup & netstat

- Nslookup which stands for name server lookup is a network administration command-line tool available for many computer operating systems for querying the Domain Name System (DNS) to obtain domain name or IP address mapping or for any other specific DNS record.
- netstat (network statistics) is a command-line network utility tool that displays network connections for the Transmission Control Protocol (both incoming and outgoing), routing tables, and a number of network interface (network interface controller or software-defined network interface) and network protocol statistics.



# Activity 1

- Open a command prompt (Windows) /terminal window (Mac)
- Type `nslookup www.google.ie`
- What information do you get back?



```
Administrator: C:\Windows\system32\cmd.exe
Microsoft Windows [Version 6.3.9600]
(c) 2013 Microsoft Corporation. All rights reserved.

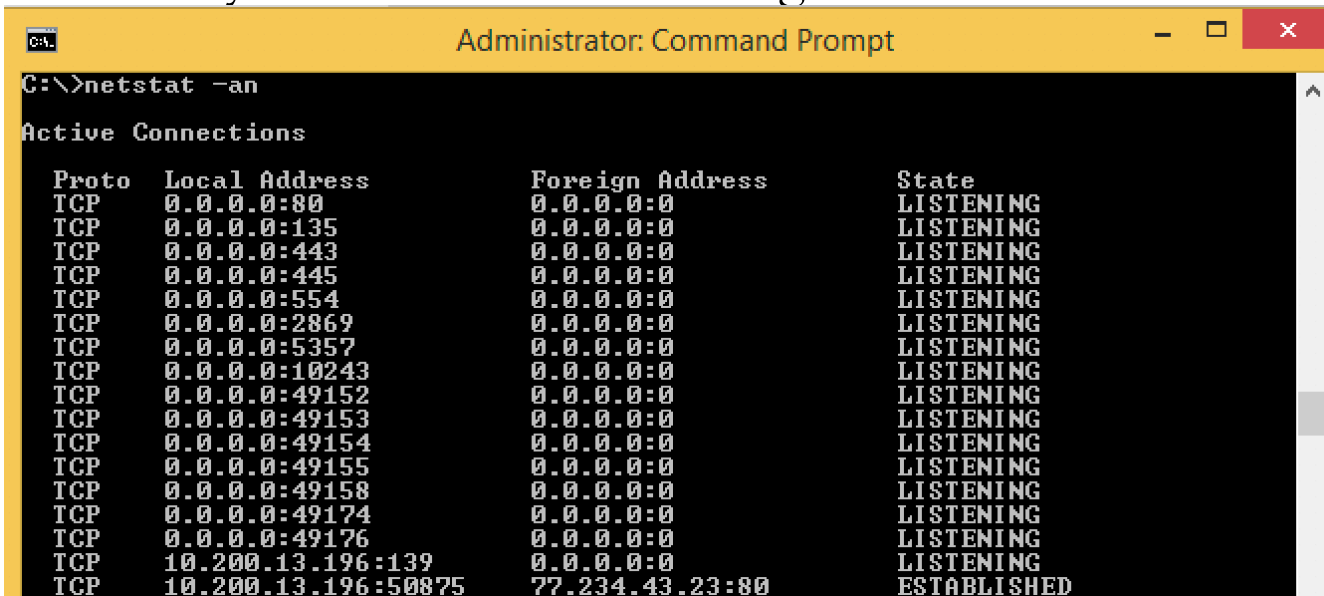
C:\Users\Administrator>nslookup www.google.ie
Server:  google-public-dns-b.google.com
Address:  8.8.4.4

Non-authoritative answer:
Name:     www.google.ie
Addresses: 2a00:1450:400b:c00::5e
          209.85.202.94
```

- Nameserver Lookup
- In DNS, so-called "non-authoritative answers" refer to DNS records kept on third-party DNS servers, which they obtained from the "authoritative" servers that provide the original source of the data.

# Activity 2

- Type netstat -an (may be different on MAC)
  - -a Displays all active connections and the TCP and UDP ports on which the computer is listening.
  - -n Displays active TCP connections, however, addresses and port numbers are expressed numerically and no attempt is made to determine names.
- From the list of IP addresses perform nslookups
- Where is your machine connecting to?



```
Administrator: Command Prompt
C:\>netstat -an
Active Connections

Proto Local Address          Foreign Address        State
TCP   0.0.0.0:80              0.0.0.0:0              LISTENING
TCP   0.0.0.0:135            0.0.0.0:0              LISTENING
TCP   0.0.0.0:443            0.0.0.0:0              LISTENING
TCP   0.0.0.0:445            0.0.0.0:0              LISTENING
TCP   0.0.0.0:554            0.0.0.0:0              LISTENING
TCP   0.0.0.0:2869           0.0.0.0:0              LISTENING
TCP   0.0.0.0:5357           0.0.0.0:0              LISTENING
TCP   0.0.0.0:10243          0.0.0.0:0              LISTENING
TCP   0.0.0.0:49152          0.0.0.0:0              LISTENING
TCP   0.0.0.0:49153          0.0.0.0:0              LISTENING
TCP   0.0.0.0:49154          0.0.0.0:0              LISTENING
TCP   0.0.0.0:49155          0.0.0.0:0              LISTENING
TCP   0.0.0.0:49158          0.0.0.0:0              LISTENING
TCP   0.0.0.0:49174          0.0.0.0:0              LISTENING
TCP   0.0.0.0:49176          0.0.0.0:0              LISTENING
TCP   10.200.13.196:139      0.0.0.0:0              LISTENING
TCP   10.200.13.196:50875    77.234.43.23:80        ESTABLISHED
```

# What is DNS?

- DNS is a globally distributed, scalable, hierarchical, and dynamic database that provides a mapping between hostnames, IP addresses (both IPv4 and IPv6), text records, mail exchange information (MX records), name server information (NS records), and security key information defined in Resource Records (RRs).
- The information defined in RRs is grouped into zones and maintained locally on a DNS server so it can be retrieved globally through the distributed DNS architecture.
- DNS can use either the User Datagram Protocol (UDP) or Transmission Control Protocol (TCP) and historically uses a destination port of 53.

Source: Cisco.com

# DNS explained

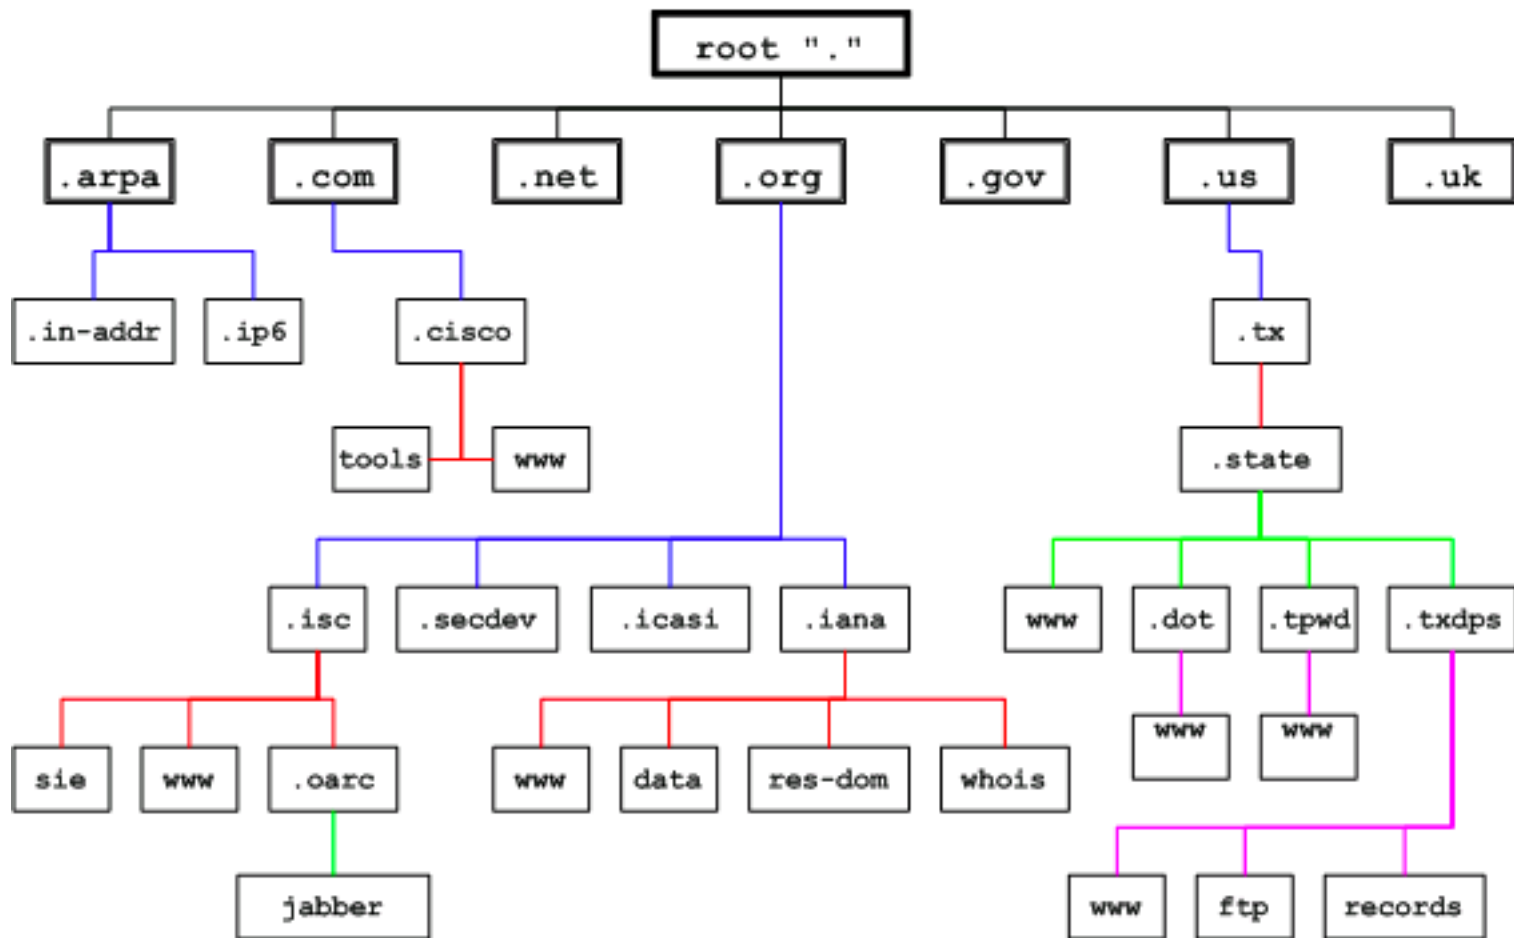
- At the top of the tree are the root domain nameservers. Every domain has an implied/hidden “.” at the end that designates the DNS root nameservers at the top of the hierarchy.
- Root domain nameservers know the IP addresses of the authoritative nameservers that handle DNS queries for the Top Level Domains (TLD) like “.com”, “.edu” or “.gov”. It first asks the root domain nameservers for the IP address of the TLD server, e.g. , “.com” (e.g. for google.com).

# DNS explained

- DNS is composed of a hierarchical domain name space that contains a tree-like data structure of linked domain names (nodes).
- Domain name space uses Resource Records (RRs) that may or may not exist to store information about the domain.
- The tree-like data structure for the domain name space starts at the root zone ".", which is the top most level of the DNS hierarchy. Although it is not typically displayed in user applications, the DNS root is represented as a trailing dot in a fully qualified domain name (FQDN).
- For example, the right-most dot in "www.cisco.com." represents the root zone.

# DNS explained

- From the root zone, the DNS hierarchy is then split into sub-domain (branches) zones. Each domain name is composed of one or more labels. Labels are separated with "." and may contain a maximum of 63 characters. A FQDN may contain a maximum of 255 characters, including the ".". Labels are constructed from right to left, where the label at the far right is the top level domain (TLD) for the domain name.
- The following example shows how to identify the TLD for a domain name: com is the TLD for www.cisco.com as it is the label furthest to the right.
- The following diagram illustrates a sample of the Domain Name System hierarchy starting from the root ".". Everything below the ".org" domain name space is in the org domain and everything below ".cisco.com" domain name space is in the cisco.com domain.



`.arpa`: primarily used for address to host mappings

`.com`, `.net`, `.org`, `.org`: are generic TLDs (gTLD)

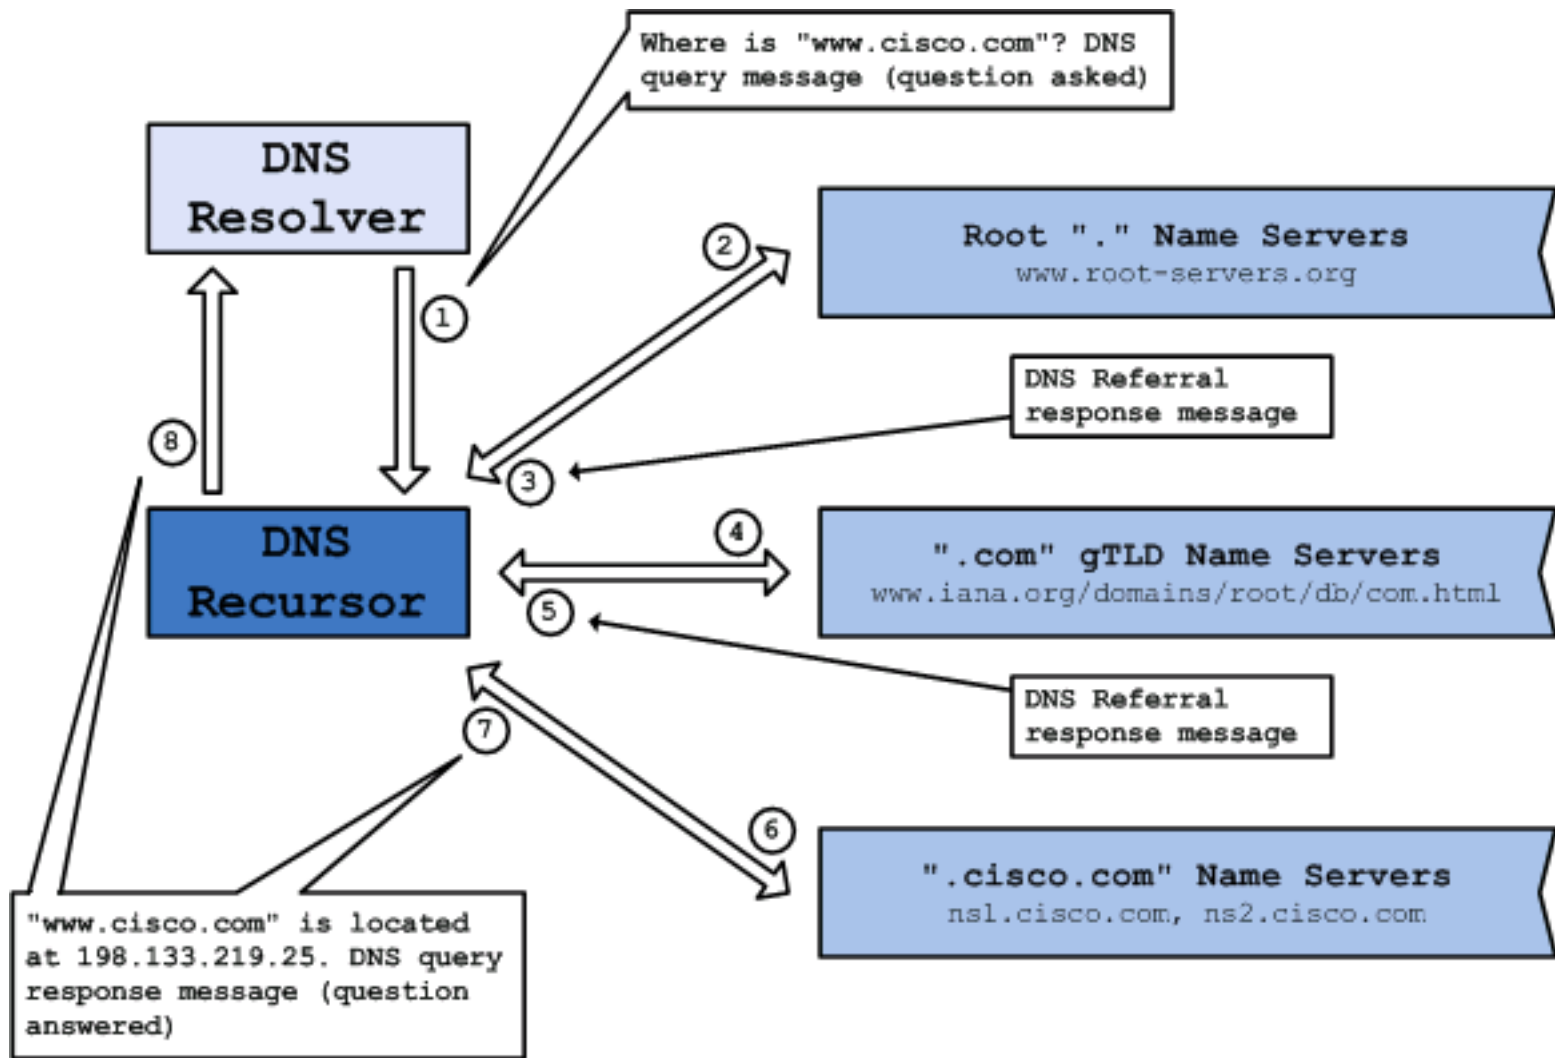
`.us`, `.uk`: are country code TLDs (ccTLD)

# Important DNS Terminology

- **Resolver:** A DNS client that sends DNS messages to obtain information about the requested domain name space.
- **Recursion:** The action taken when a DNS server is asked to query on behalf of a DNS resolver.
- **Authoritative Server:** A DNS server that responds to query messages with information stored in RRs for a domain name space stored on the server.
- **Recursive Resolver:** A DNS server that recursively queries for the information asked in the DNS query.
- **FQDN:** A Fully Qualified Domain Name is the absolute name of a device within the distributed DNS database.
- **RR:** A Resource Record is a format used in DNS messages that is composed of the following fields: NAME, TYPE, CLASS, TTL, RDLENGTH, and RDATA.
- **Zone:** A database that contains information about the domain name space stored on an authoritative server.







# Recursive query

1. The DNS resolver sends a query message to the recursive resolver asking for the address of [www.cisco.com](http://www.cisco.com).
2. The DNS recursor sends a query message to the root name servers looking for the *.com* domain name space.
3. The root name servers send a DNS referral response message to the DNS recursor informing it to ask the gTLD name servers for the *.com* domain name space.
4. The DNS recursor sends a query message to the gTLD name servers looking for the *.cisco.com* domain name space.
5. The gTLD name servers send a DNS referral response message to the DNS recursor informing it to ask the *.cisco.com* name servers, *ns1.cisco.com* or *ns2.cisco.com*, about this domain name space.
6. The DNS recursor sends a query to *ns1.cisco.com* or *ns2.cisco.com* asking for [www.cisco.com](http://www.cisco.com).
7. The *.cisco.com* name servers, *ns1.cisco.com* or *ns2.cisco.com*, send an authoritative DNS query response message to the DNS recursor with the A (address) RR information for [www.cisco.com](http://www.cisco.com).
8. The DNS recursor sends a DNS query response message to the DNS resolver with the A (address) RR information for [www.cisco.com](http://www.cisco.com).

Questions?