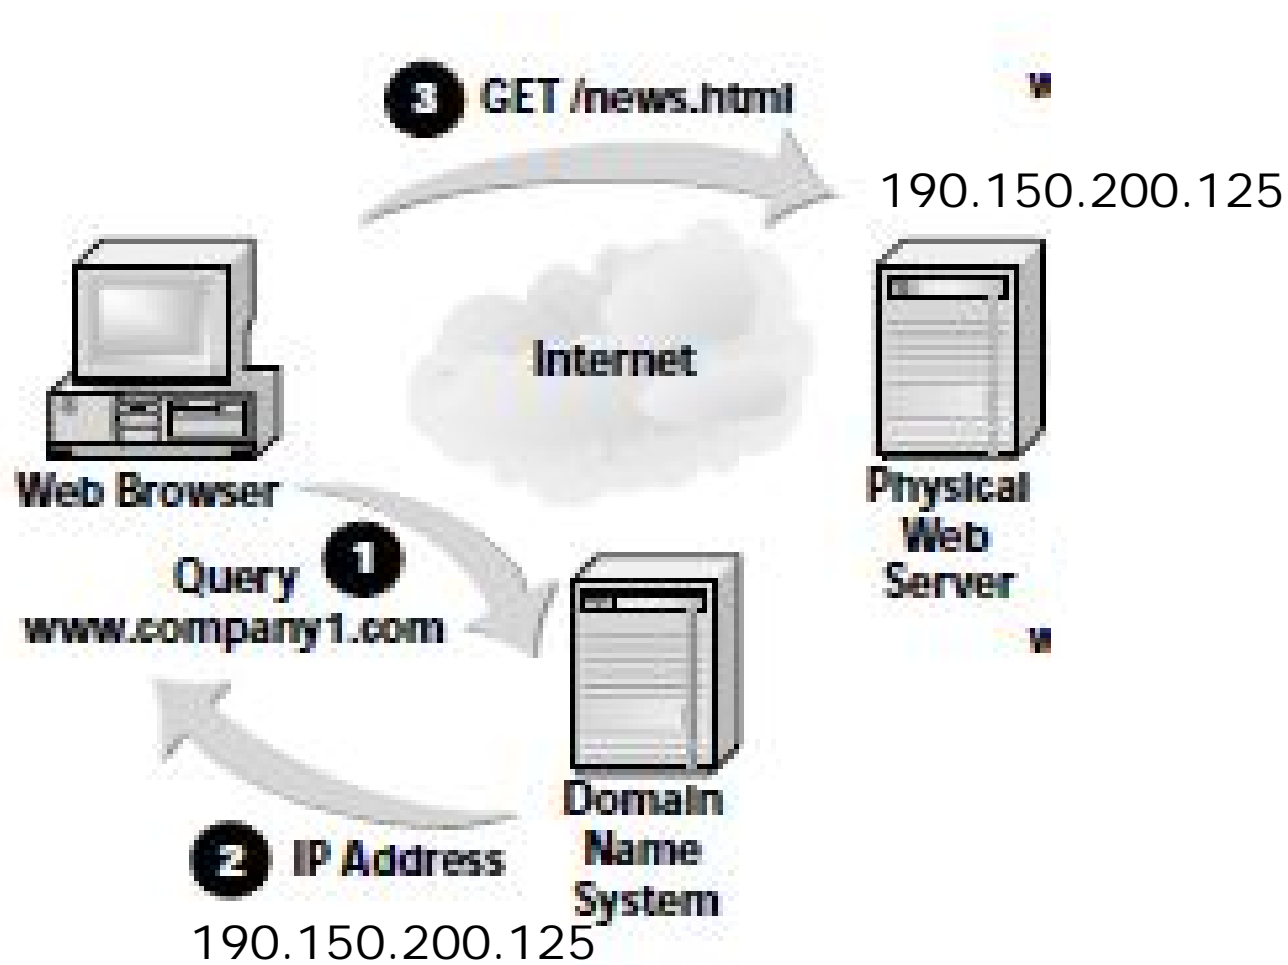


# (DNS) Domain Name System


Mahesh.Chahar@gmail.com

# How it Works




# History

- ▶ Late 1960s, the ARPA (later DARPA), began funding the ARPAnet, an experimental wide area computer network that connected important research organizations in the United States.
- ▶ The TCP/IP (1980s) quickly became the standard host-networking protocol on the ARPAnet. Its inclusion in Berkeley's (University of California) popular BSD Unix operating system was instrumental in democratizing internetworking. BSD Unix was virtually free to universities.



# History

- ▶ The network grew from a handful of hosts to tens of thousands of hosts. The original ARPANet became the backbone of a confederation of local and regional networks based on TCP/IP, called the **Internet**.
- ▶ In 1988, DARPA experiment was over. Another network, funded by the National Science Foundation called the NSFNET, replaced the ARPANet as the backbone of the Internet.




# History

- ▶ In 1995, the Internet made a transition from publicly funded NSFNET (as a backbone) to multiple commercial backbones, run by long-distance carriers like **MCI** and **Sprint**, and long-time commercial internetworking players like **PSINet** and **UUNET**.
- ▶ Through the 1970s, the ARPAnet was a small, friendly community of a few hundred hosts. A single file, *HOSTS.TXT*, contained all the information needed to know about those hosts: it held name-to-address mappings for every host connected to the ARPAnet.
- ▶ *HOSTS.TXT* was maintained by SRI's *Network Information Center* (dubbed "the NIC") and distributed from a single host, SRI-NIC.




# History

- ▶ ARPAnet administrators typically emailed their changes to the NIC, and periodically FTPed to SRI-NIC and grabbed the current *HOSTS.TXT*.
- ▶ Their changes were compiled into a new *HOSTS.TXT* once or twice a week. As the ARPAnet grew, however, this scheme became unworkable.




# Challenges

- ▶ Moreover, the traffic generated by the update process increased even faster: every additional host meant not only another line in HOSTS.TXT, but potentially another host updating from SRI-NIC
- ▶ *Traffic and load*
- ▶ *Name collisions*
- ▶ *Consistency*

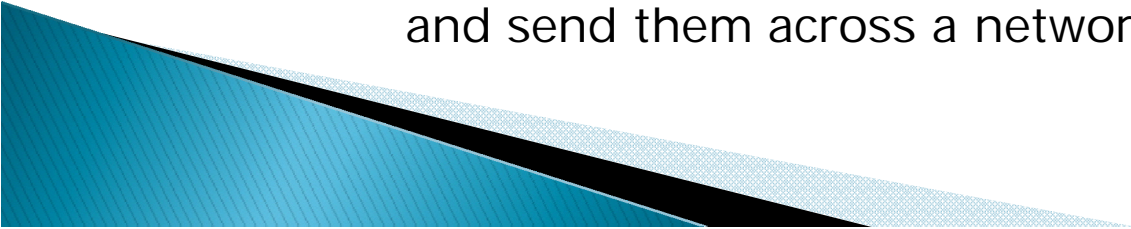


# Challenges

- ▶ The essential problem was that the *HOSTS.TXT* mechanism didn't scale well. Ironically, the success of the ARPAnet as an experiment led to the failure and obsolescence of *HOSTS.TXT*.
  - ▶ The ARPAnet's governing bodies chartered an investigation into a successor for *HOSTS.TXT*.
  - ▶ Paul Mockapetris, was given responsibility for designing the architecture of the new system.
- 




# DNS

- ▶ It is a *distributed database*.
    - Allows local control of the segments of the overall database, yet the data in each segment is available across the entire network through a client/server scheme.
  - ▶ Robustness and adequate performance
    - are achieved through replication and caching.
  - ▶ *Nameservers*
    - It's a program that contains information about some segments of the database and make it available to clients, called *resolvers*.
  - ▶ *Resolvers*
    - *Resolvers are often just library routines that create queries and send them across a network to a nameserver.*
- 


# DNS Database structure

- Similar to Unix Filesystem

Unix Filesystem

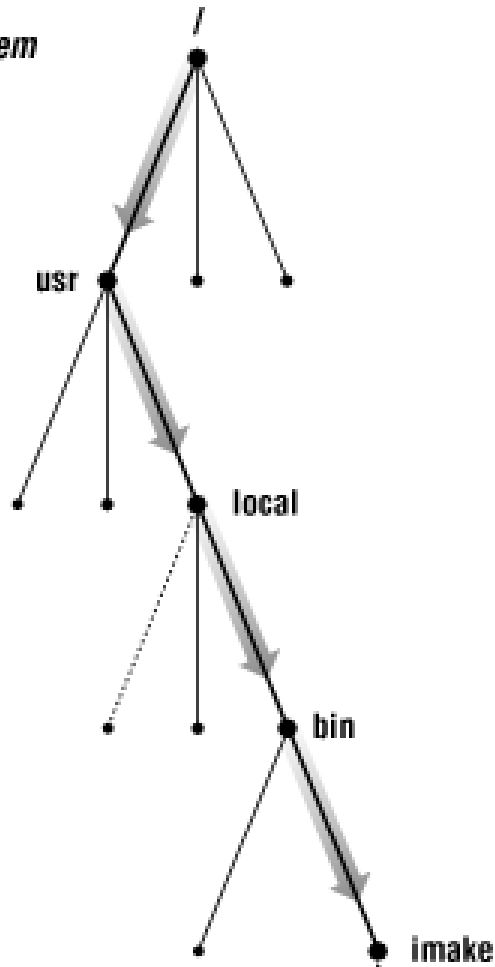


DNS Database



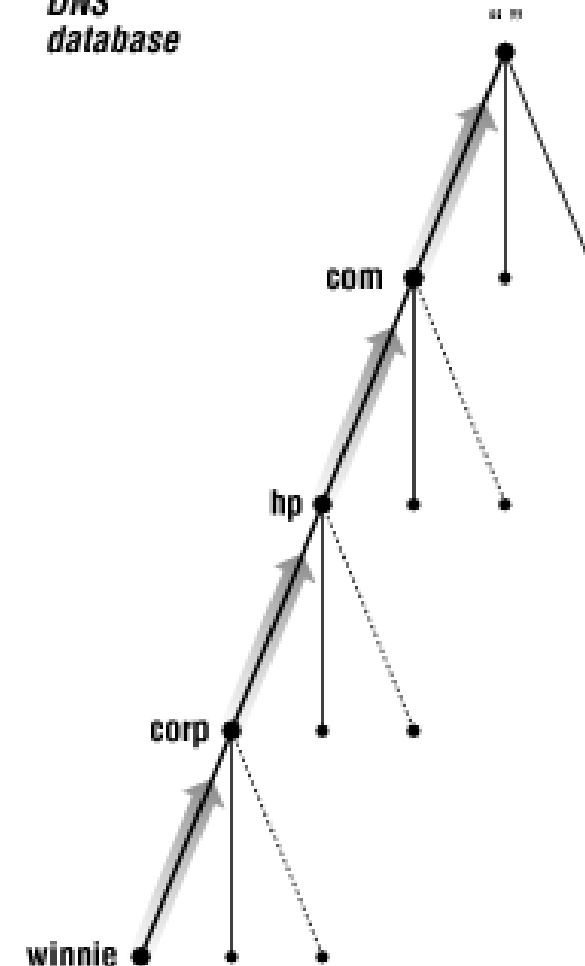
# Reading Name

*Unix  
filesystem*



*/usr/local/bin/imake*

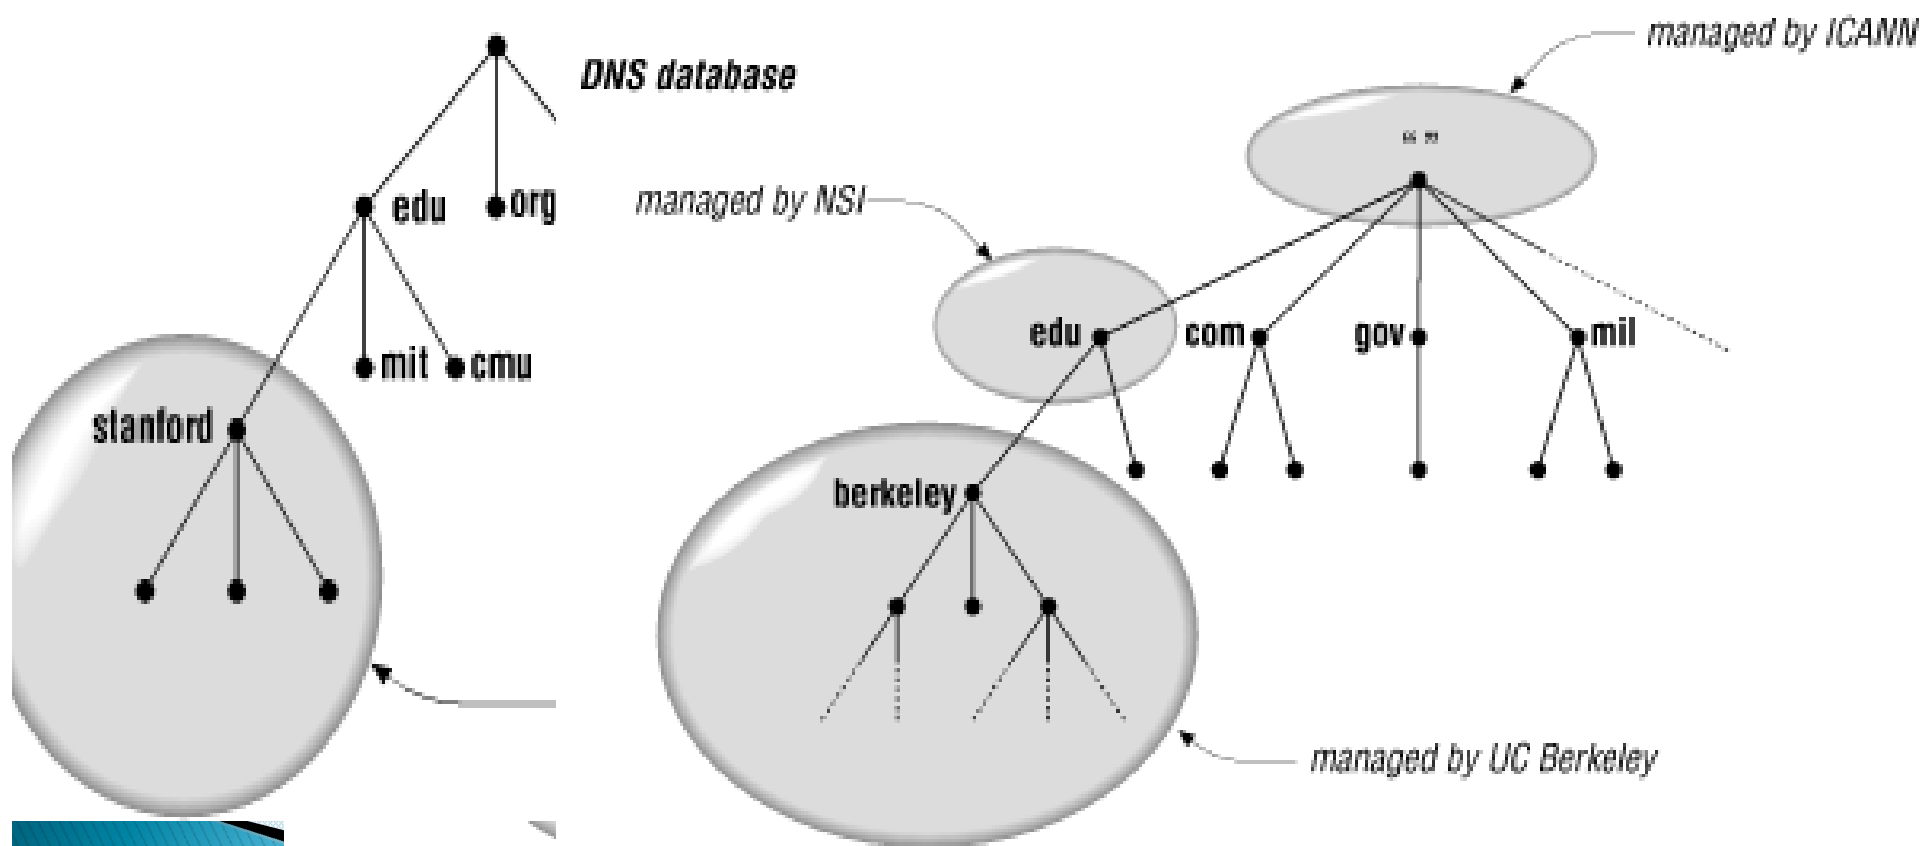
*DNS  
database*



*winnie.corp.hp.com*


# DNS Delegation

Delegating like remotely mounting a filesystem



# Domain Name

- ▶ Domains are often referred to by *level* ; These terms simply refer to a domain's position in the domain name space:
  - *top-level domain or second-level domain*
  - A top-level domain is a child of the root.
  - A first-level domain is a child of the root (i.e., a top-level domain).
  - A second-level domain is a child of a first-level domain, and so on.




# Top-Level Domains (TLDs)

- ▶ divides the Internet domain name space organizationally into seven domains:  
**com, edu, gov, mil, net, org, int**
- These original domains are called *generic top-level domains, or gTLDs*. maintained by the IANA
- ▶ IANA distinguishes the following groups of top-level domains:
  - infrastructure top-level domain(arpa)
  - country top-level domains (ccTLD)
  - internationalized top-level domains (IDNs)
  - generic top-level domains (gTLD)
    - [com](#), [info](#), [net](#), [org](#), and [biz](#)<sup>\*</sup>, [name](#)<sup>\*</sup>, [pro](#)<sup>\*</sup>



# Domains

- ▶ Domains [edu](#), [gov](#), [int](#), and [mil](#) are now considered sponsored top-level domains, much like the many newly created *themed* domain names (e.g., [jobs](#)).
- ▶ *GeoTLD* invokes an association with a geographical, geopolitical, ethnic, linguistic or cultural community. [.cat](#), [.asia](#), [.kiwi](#), [.paris](#), [.gal](#)




# *Name Servers and Zones*

- ▶ NSs have complete information about some part of the domain name space (a *zone*), which they load from a file / name server.
- ▶ The name server is then said to have *authority for that zone*.
- ▶ *Name servers can be authoritative for multiple zones, too.*
- ▶ Domains (*berkeley.edu, hp.com*) are broken into smaller, more manageable units by delegation. *These units are called zones.*







# Zones



# Zones



# Domain Vs Zone

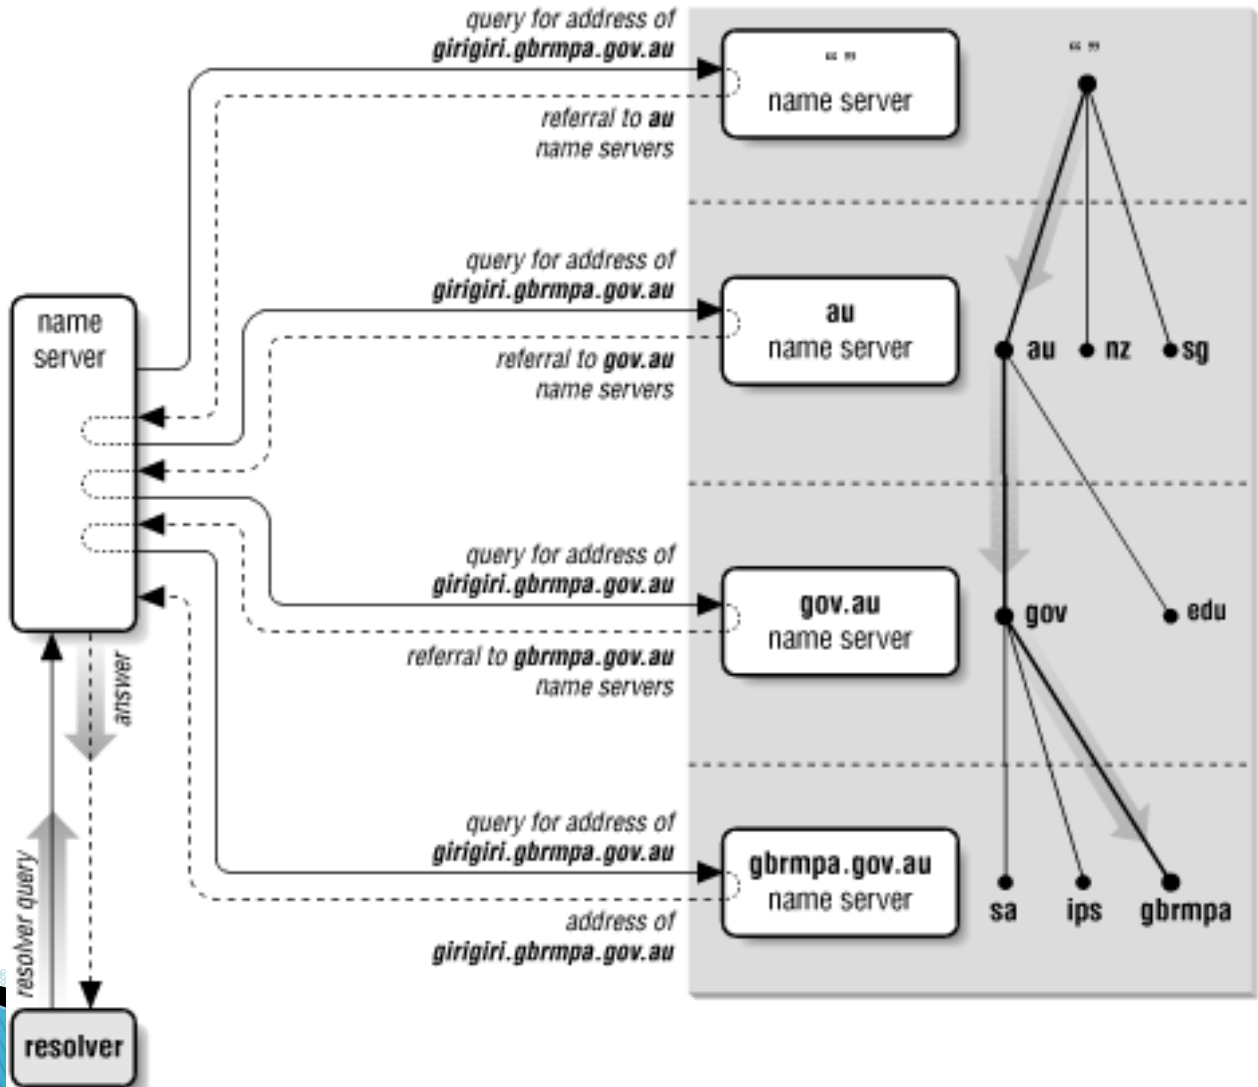


# Root servers (Root Nameservers)

- ▶ 13 Root Nameservers

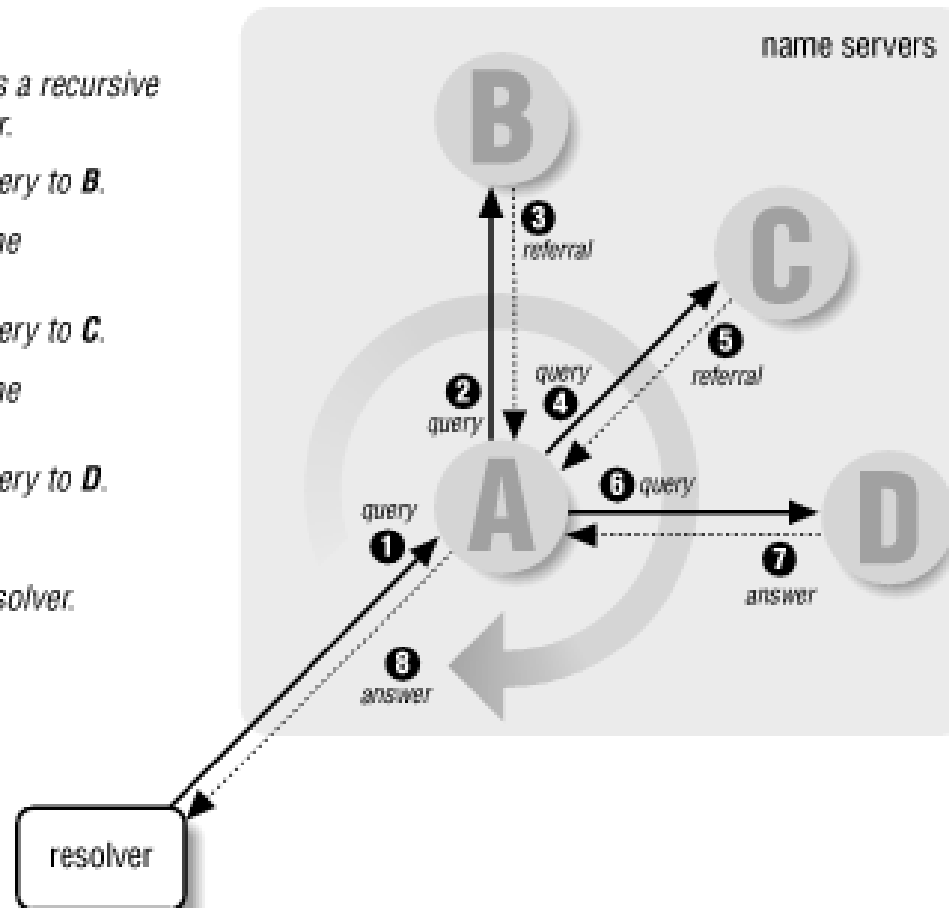


# Resolution




# Resolution Process

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



# Address to Name Mapping



most specific

least specific

winnie . corp . hp . com  
152 . 192 . 16 . 15

hostname *winnie.corp.hp.com*

# Caching

