

Domain Name System (DNS)



Shivkumar Kalyanaraman
Rensselaer Polytechnic Institute
shivkuma@ecse.rpi.edu

<http://www.ecse.rpi.edu/Homepages/shivkuma>




- Naming hierarchy
- Server hierarchy
- Name resolution
- Other information in name servers
- Ref: Chap 14

Why Names?

- ❑ Computers use addresses
- ❑ Humans cannot remember IP addresses
 - ⇒ Need names
 - Example, “shiva” for 128.113.50.56
- ❑ Simplest Solution: Each computer has a unique name and has a built in table of name to address translation
- ❑ Problem: Not scalable
- ❑ Solution: DNS (Adopted in 1983)
- ❑ Hierarchical Names: shiva.ecse.rpi.edu

Name Hierarchy



Name Hierarchy

- Unique domain suffix is assigned by Internet Authority
- The domain administrator has complete control over the domain
- No limit on number of subdomains or number of levels
- computer.site.division.company.com
- computer.site.subdivision.division.company.com
- Domains within an organization do not have to be uniform in number of subdomains or levels

Name Hierarchy (Cont)

- Name space is not related to physical interconnection, e.g., ecse.rpi.edu and ipl.rpi.edu could be on the same floor or in different cities
- Geographical hierarchy is also allowed, e.g., cnri.reston.va.us
- A name could be a *subdomain* (eg: ecse.rpi.edu) or an individual *object* (eg: cortez.rpi.edu)

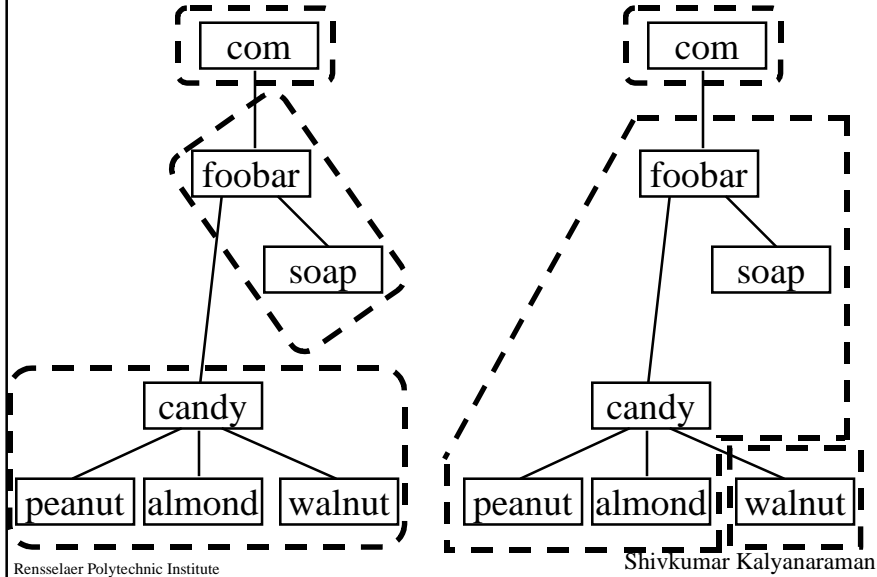
Top Level Domains

Domain Name/Assignment	
com	Commercial
edu	Educational
gov	Government
mil	Military
net	Network
org	Other organizations
arpa	Advanced Research Project Agency
country code	au, uk, ca

Server Hierarchy

- Servers are organized in a hierarchy
- Each server has an authority over a part of the naming hierarchy
- The server does not need to keep all names.
- It needs to know other servers who are responsible for other subdomains
- A single node in the naming tree cannot be split among multiple servers
- A given level of hierarchy can be partitioned into multiple servers

Server Hierarchy (example)




9


Server Hierarchy (Cont)

- Authority \Rightarrow has the name to address translation table
- Responsible \Rightarrow Either has the name to address translation table or knows the server who has
 - But such a reply is called “non-authoritative” reply
- A single server can serve multiple domains, e.g., purdue.edu and laf.in.us
- Root server knows about servers for top-level domains, e.g., com
- Each server knows the root server

Server Hierarchy: Example



Server Hierarchy: Better

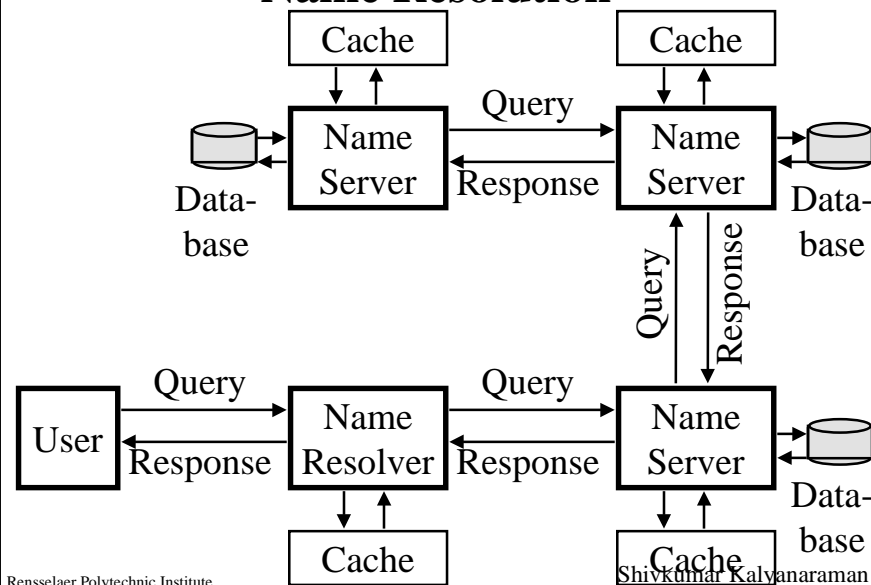


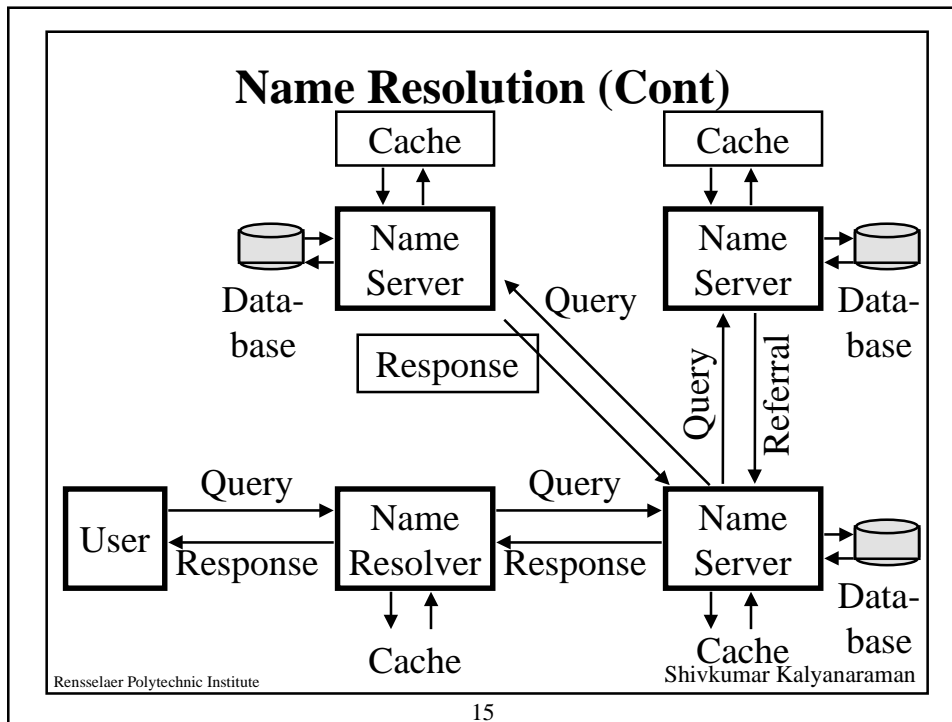
- Fewer servers
- More entries/links per server
- Fewer levels to traverse before resolving a name

Name Resolution (Cont)

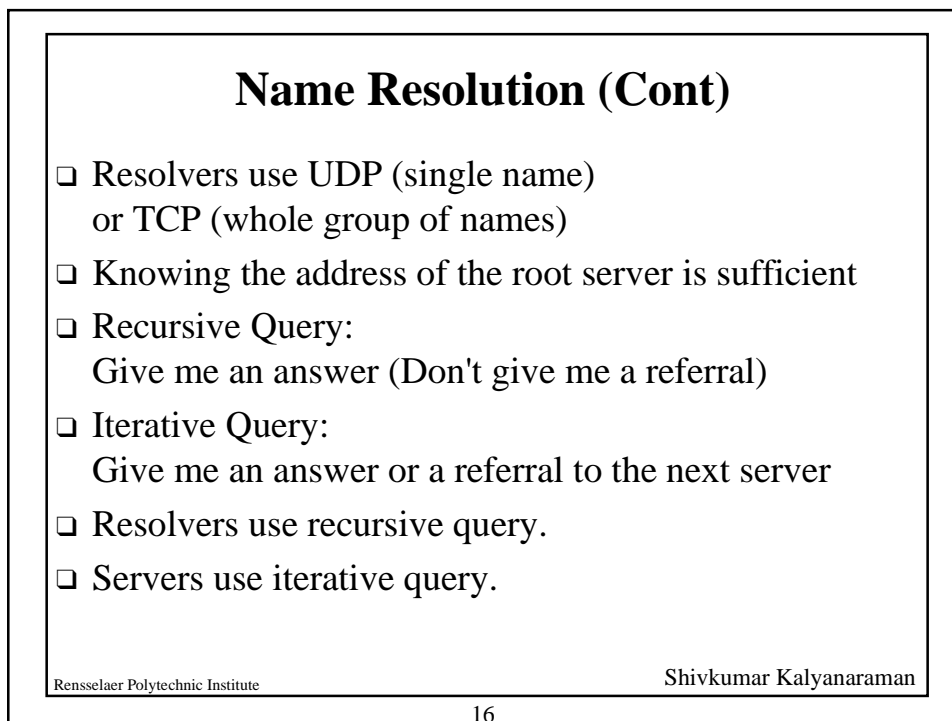
- ❑ Each computer has a name resolver routine, e.g., gethostbyname & gethostbyaddr in UNIX
- ❑ Each resolver knows the name of a local DNS server
- ❑ Resolver sends a DNS request to the server
- ❑ DNS server either gives the answer, forwards the request to another server, or gives a referral
- ❑ Referral = Next server to whom request should be sent

Name Resolution





15



16

DNS Optimization

- Spatial Locality: Local computers referenced more often than remote
- Temporal Locality: Same set of domains referenced repeatedly \Rightarrow Caching
- Each entry has a time to live (TTL)
- Replication: Multiple servers. Multiple roots. Ask the geographically closest server.

Abbreviations

- Servers respond to a full name only
- However, humans may specify only a partial name
- Resolvers may fill in the rest of the suffix, e.g., shiv = shiv.ecse.ohio-state.edu
- Each resolver has a list of suffixes to try

DNS Message Format

Identification	Parameter
Number of Questions	Number of Answers
Number of Authority	Number of Additional
Question Section	
...	
Answer Section	
...	
Authority Section	
...	
Additional Information Section	
...	

Format (Cont)

- Format of the query section entries:

Query Domain Name	
...	
Query Type	Query Class

- Format of other section entries:

Resource Domain Name	
Type	Class
Time to live	
Resource Data Length	Resource Data

DNS Message Format

- Length = 0 ⇒ End of names. Length < 64
Two msbs (most significant bits) = 11 ⇒ Pointer
- Resource data contains serial (version) number of the zone, refresh interval, retry interval, expiry interval, mailbox of the responsible person, etc.

DNS Message Format (Cont)

Bit	Meaning
0	Operation: 0=Query, 1=Response
1-4	Query type: 0=Standard, 1=Inverse, 2,3 obsolete
5	Set if answer authoritative
6	Set if message truncated
7	Set if recursion desired
8	Set if recursion available
9-11	Reserved
12-15	Response type: 0=No error, 1=Format error, 2=Server Failure, 3=Name does not exist


Types of DNS Entries

- ❑ DNS used other types of resolution
- ❑ Eg: also for finding mail server, pop server, responsible person, etc for a computer
- ❑ DNS database has multiple “*types*”
 - ❑ Record type A ⇒ Address of X
 - ❑ Record type MX ⇒ Mail exchanger of X
- ❑ DNS database may also have multiple “*classes*”
 - ❑ Can support name resolution for multiple protocols eg: IP, SNA, DECbit etc
- ❑ Pointer queries: given IP address find name

Resource Record Types

Type	Meaning
A	Host Address
CNAME	Canonical Name (alias)
HINFO	CPU and O/S
MINFO	Mailbox Info
MX	Mail Exchanger
NS	Authoritative name server for a domain
PTR	Pointer to a domain name (link)
RP	Responsible person
SOA	Start of zone authority (Which part of naming hierarchy implemented)
TXT	Arbitrary Text

Summary



- ❑ DNS: Maps names to addresses
- ❑ Names are hierarchical. Administration is also hierarchical.
- ❑ No standard for number of levels
- ❑ Replication and caching is used for performance optimization.

Informal Exercises

- ❑ /etc/hosts is a table of name-to-IP-address mappings
 - ❑ Find out approximately how many hosts, subnets and domains are there in the RPI campus net
 - ❑ Does this table give some addresses of root servers as well ?
 - ❑ Why does the /etc/hosts in ECSE have the entire table for the campus net.
- ❑ Lookup the file /etc/resolv.conf which gives the domain name and addresses of nameservers.
 - ❑ Why are multiple nameservers listed ?
 - ❑ Lookup the name of an IP address using nslookup. This generates a pointer query - you can watch it using tcpdump. What is unusual about pointer queries on the wire ?