Chapter 24
The Domain Name System (DNS)

Raj Jain

Raj Jain is now at
Washington University in Saint Louis
Jain@cse.wustl.edu
http://www.cse.wustl.edu/~jain/

- Naming hierarchy
- Server hierarchy
- Name resolution
- Other information in name servers

Why Names?

- Computers use addresses
- Humans cannot remember IP addresses
  $\Rightarrow$ Need names
  Example, Liberia for 164.107.51.28
- 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: Liberia.cis.ohio-state.edu

Name Hierarchy

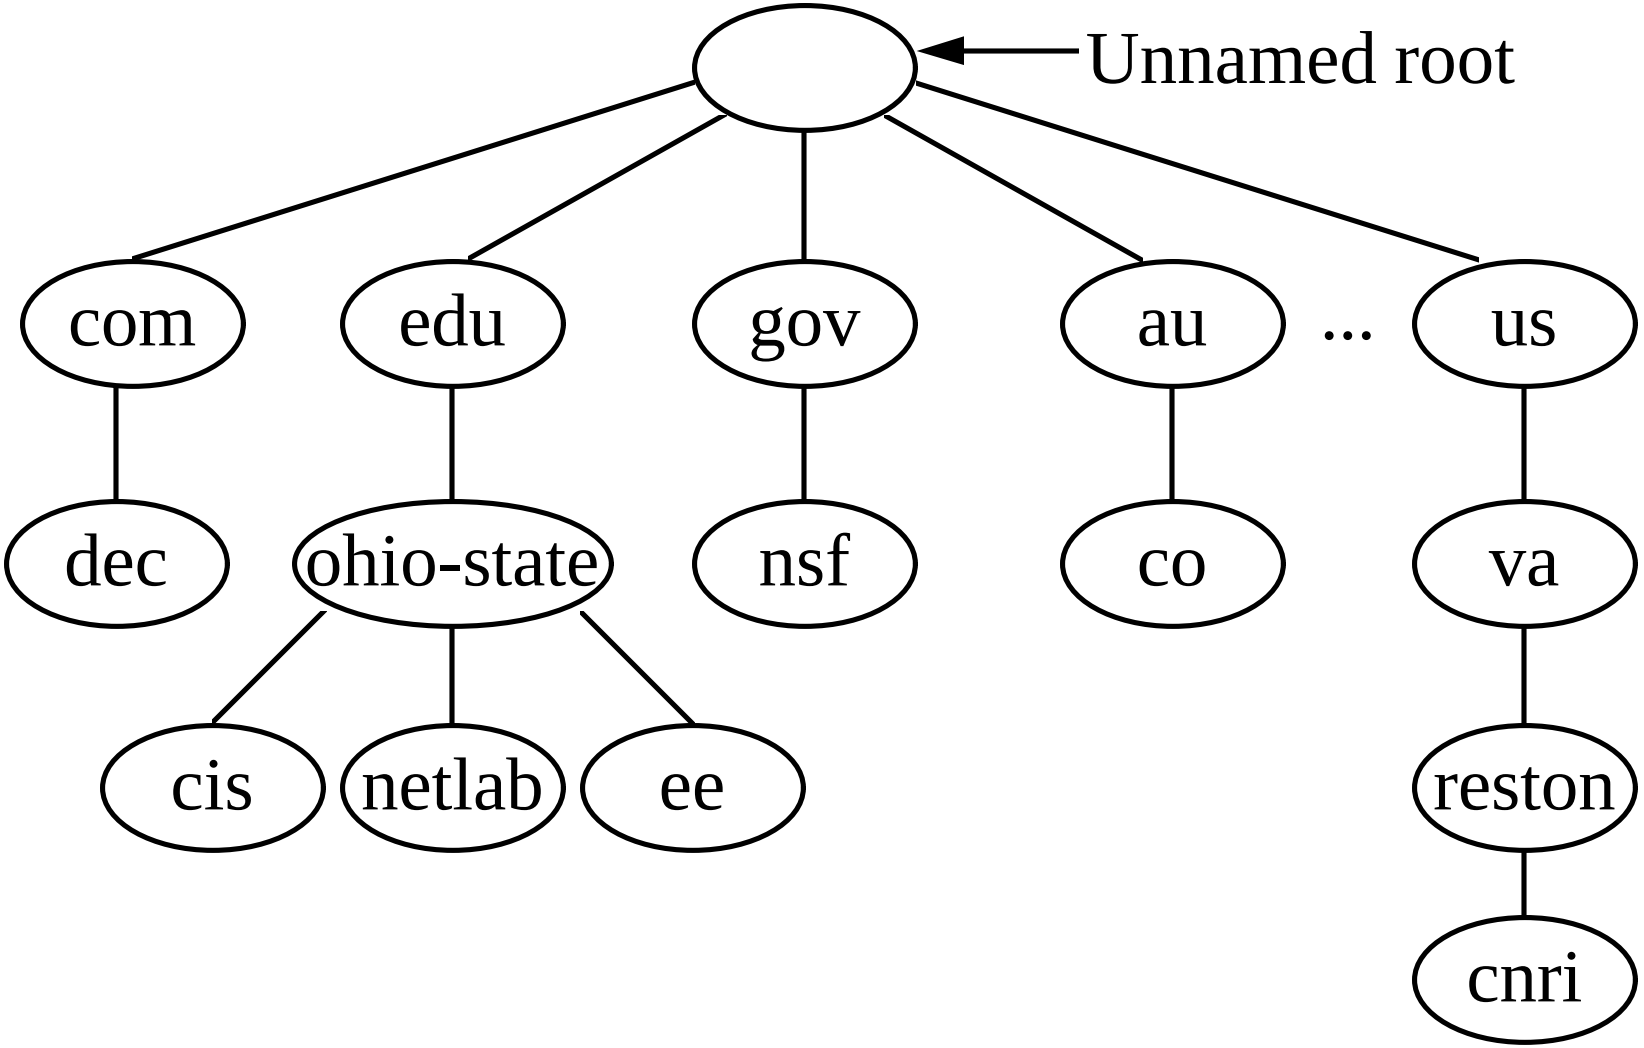

Name Hierarchy

- 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., math.ohio-state and cis.ohio-state 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 or an individual object

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

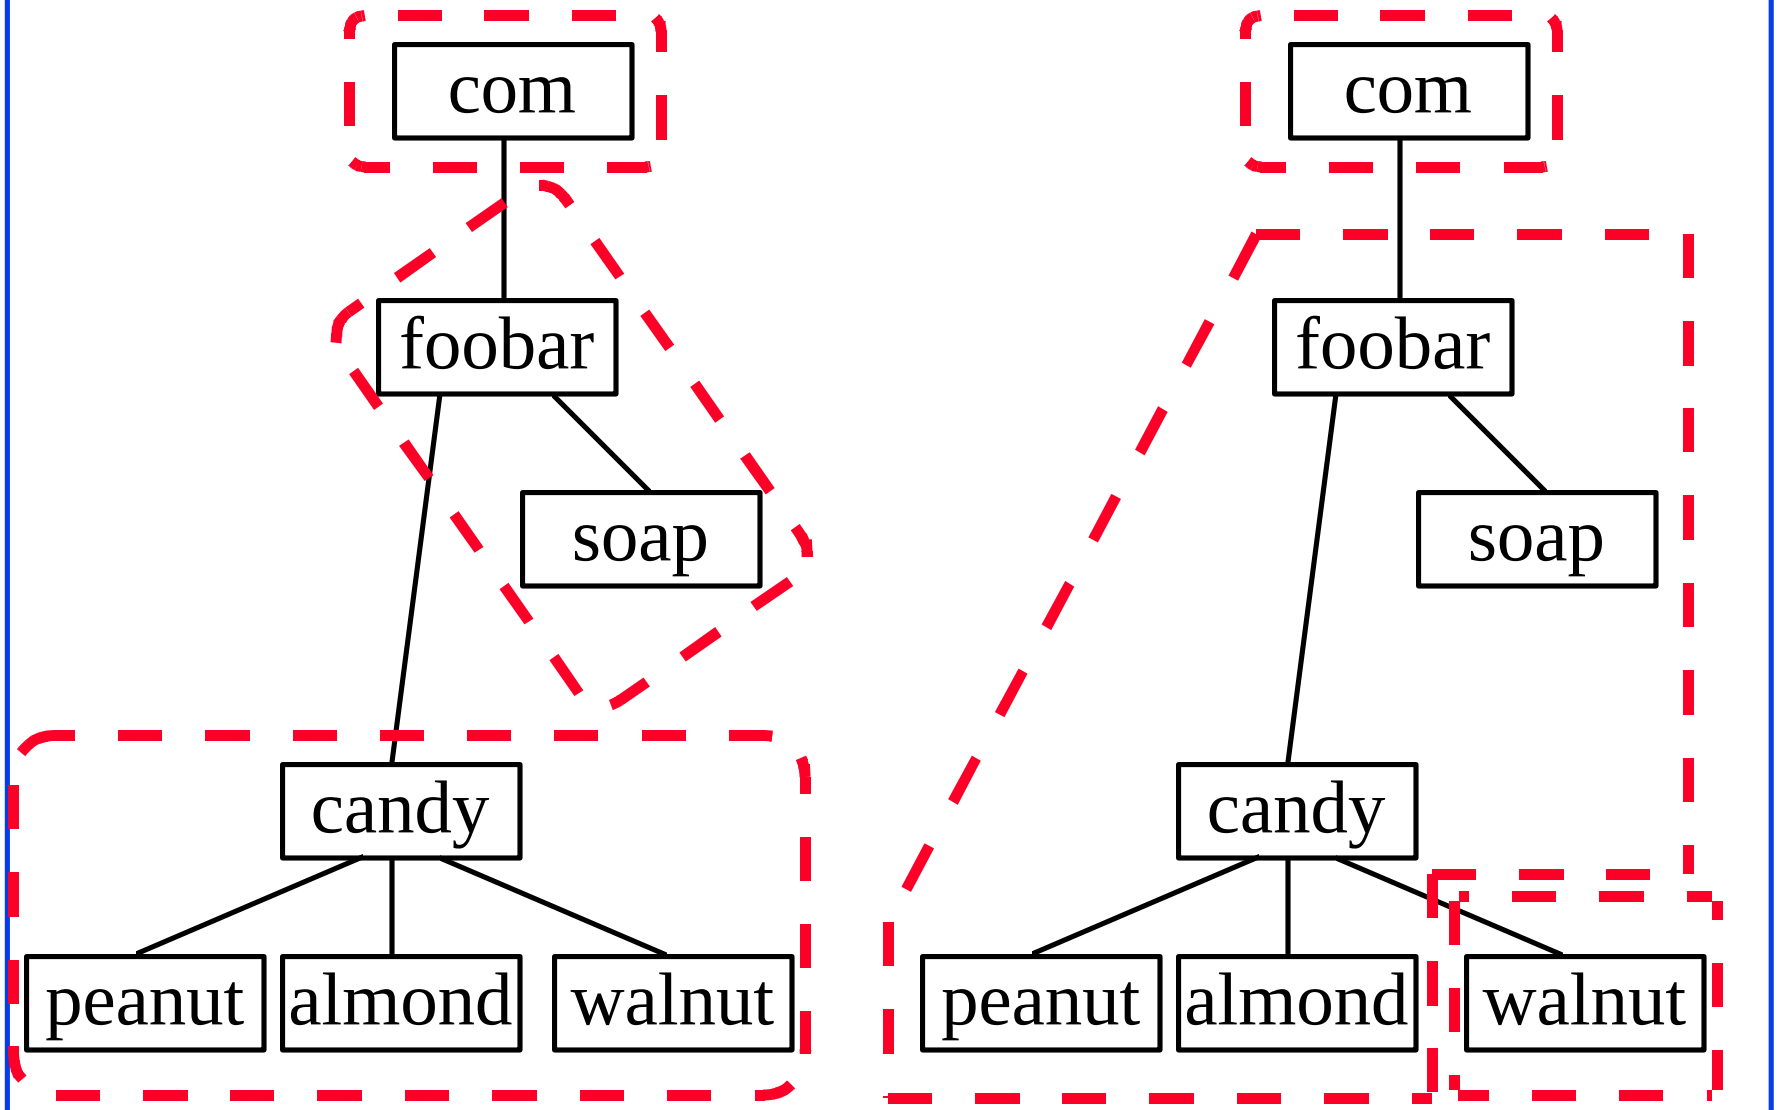

Server Hierarchy (Cont)

- 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
- Contiguous space ⇒ A single node in the naming tree cannot be split
- A given level of hierarchy can be partitioned into multiple servers

Server Hierarchy (Cont)

- Authority ⇒ has the name to address translation table
- Responsible ⇒ Either has the name to address translation table or knows the server who has
- 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

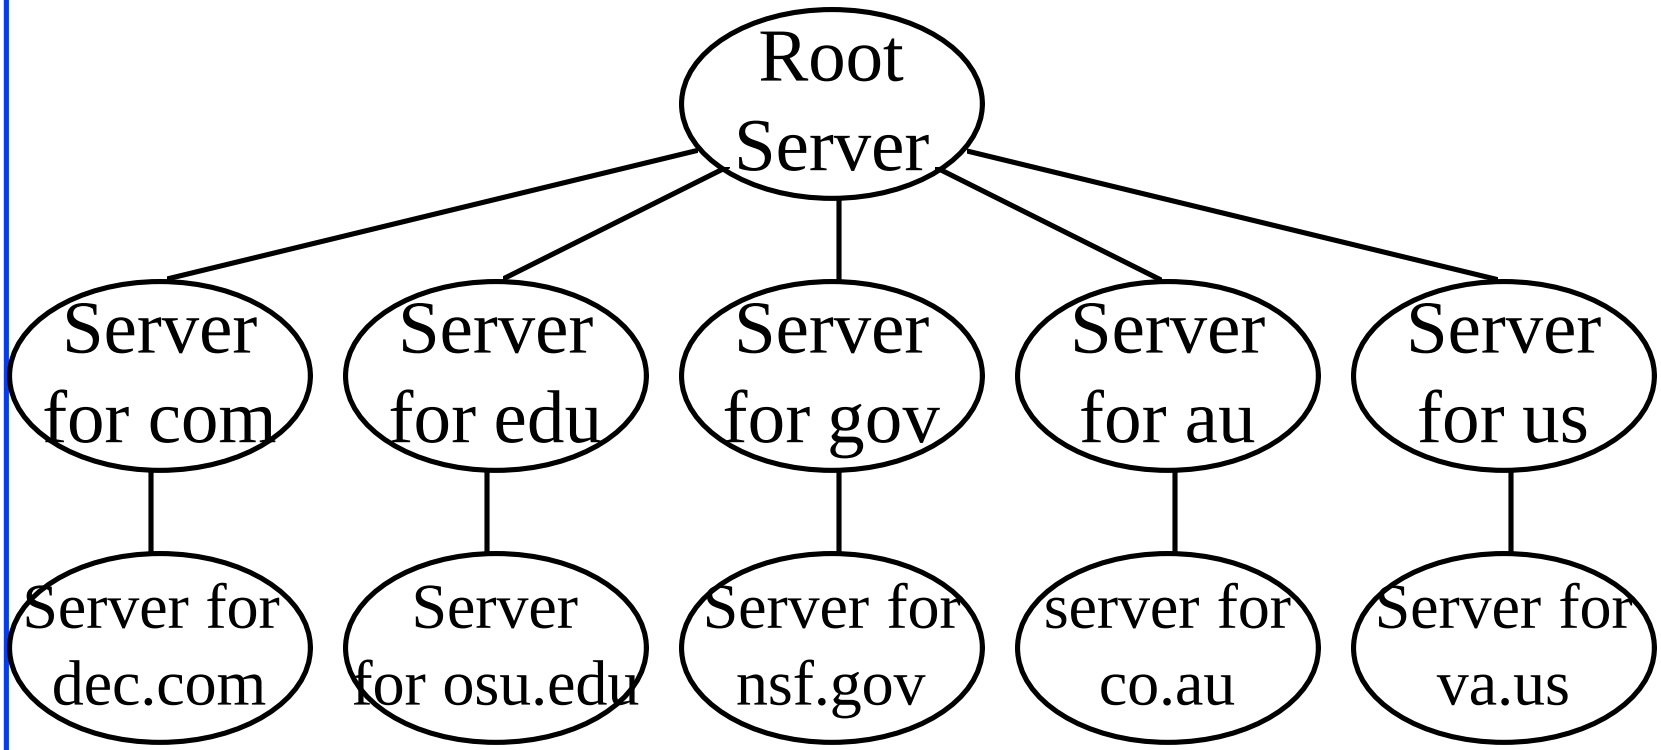

Server Hierarchy: Better

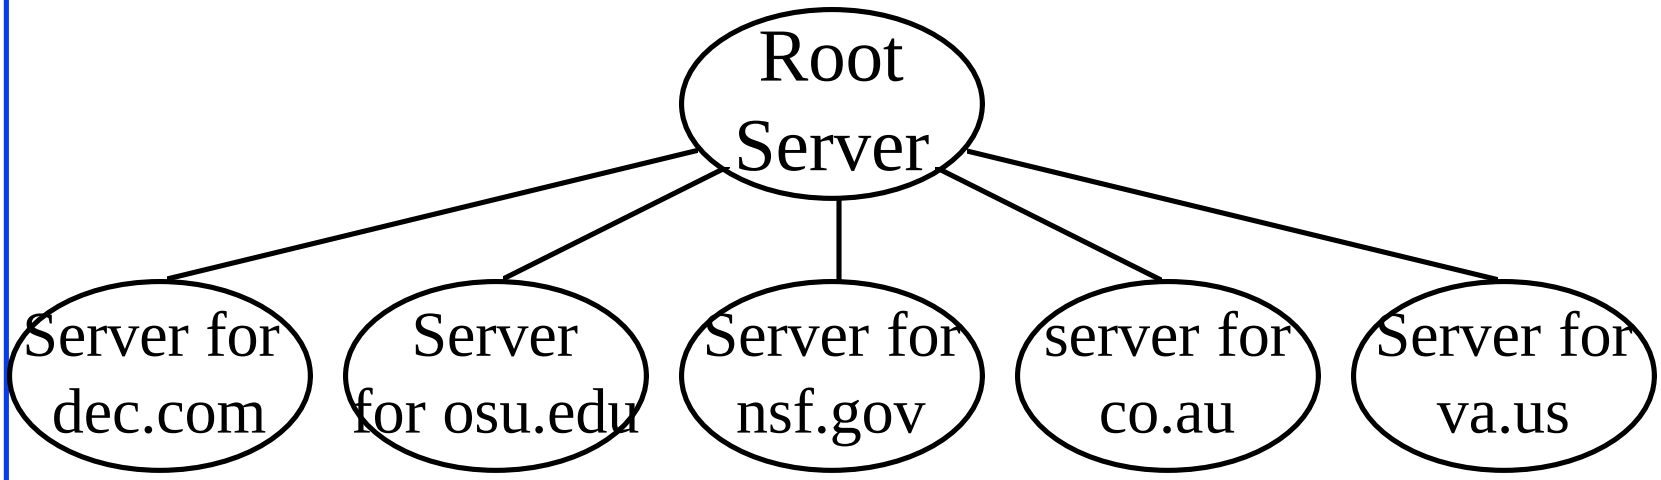

- Fewer servers

Name Resolution

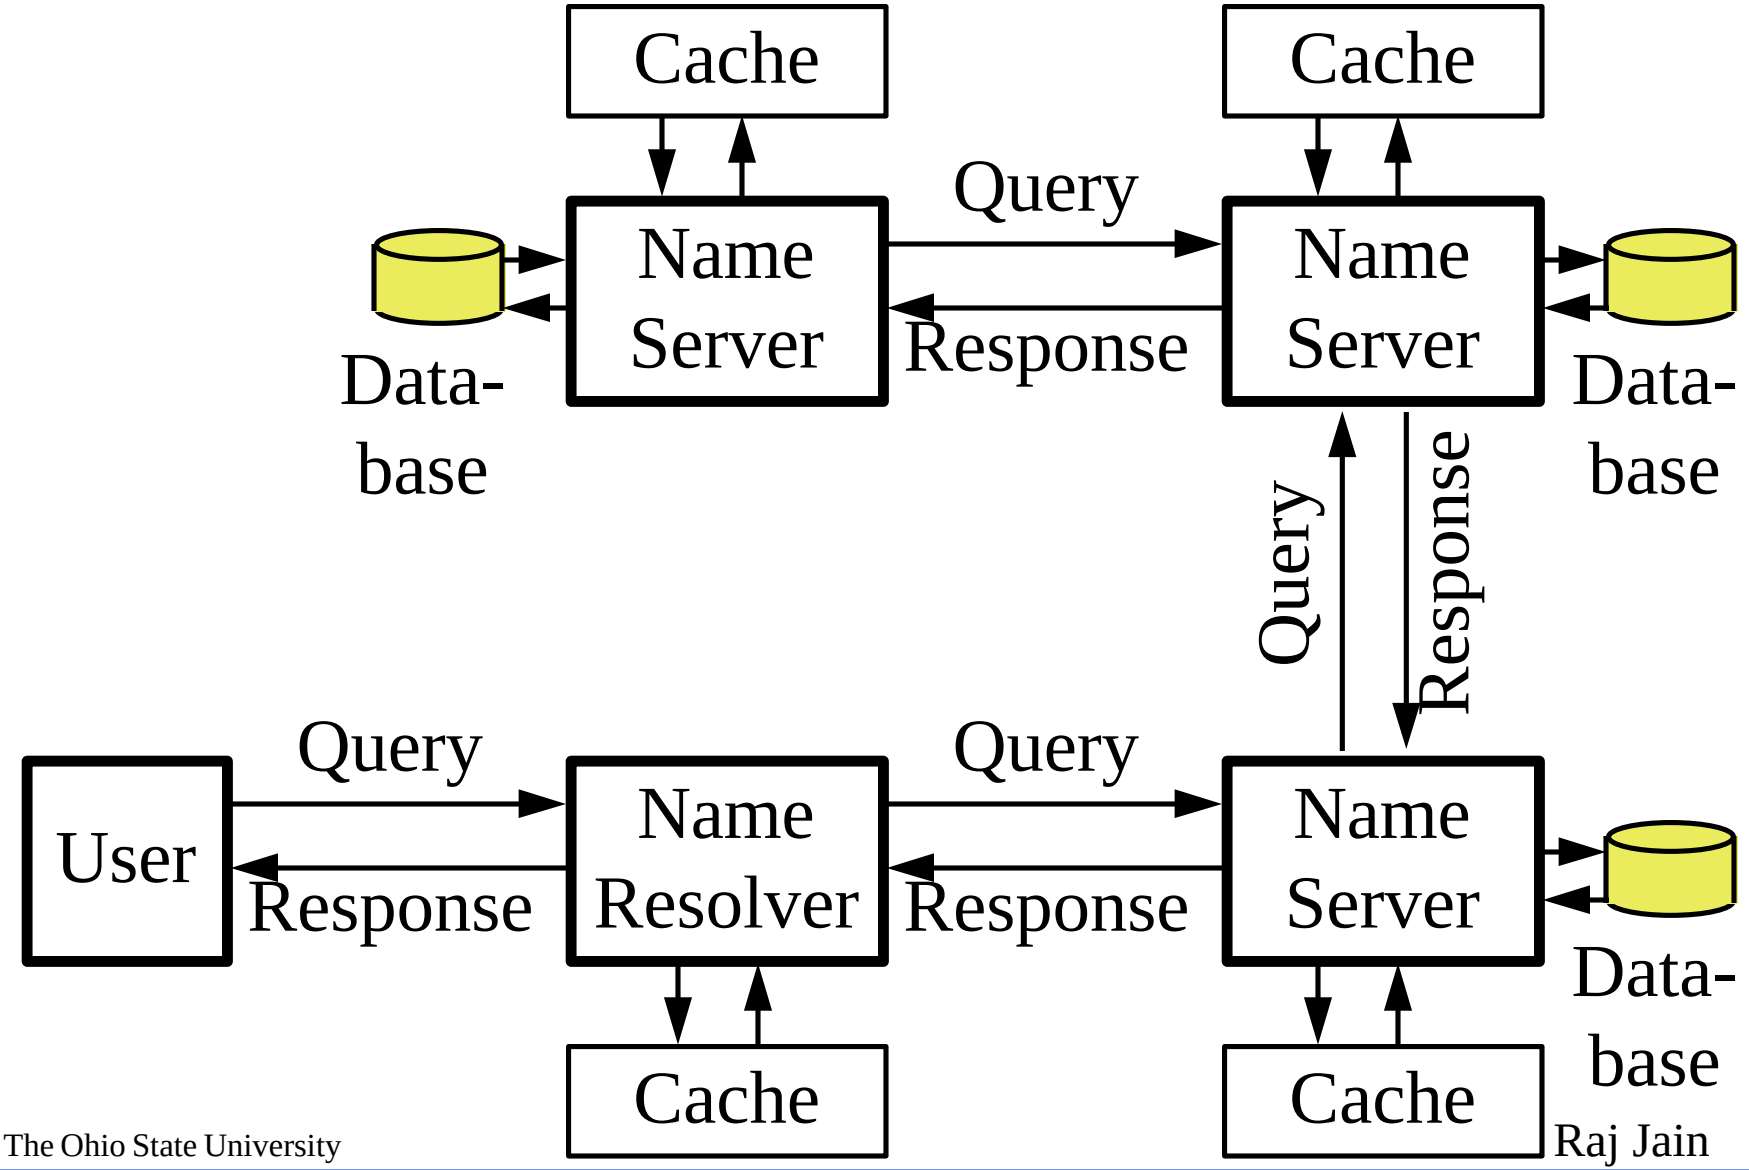

Name Resolution (Cont)

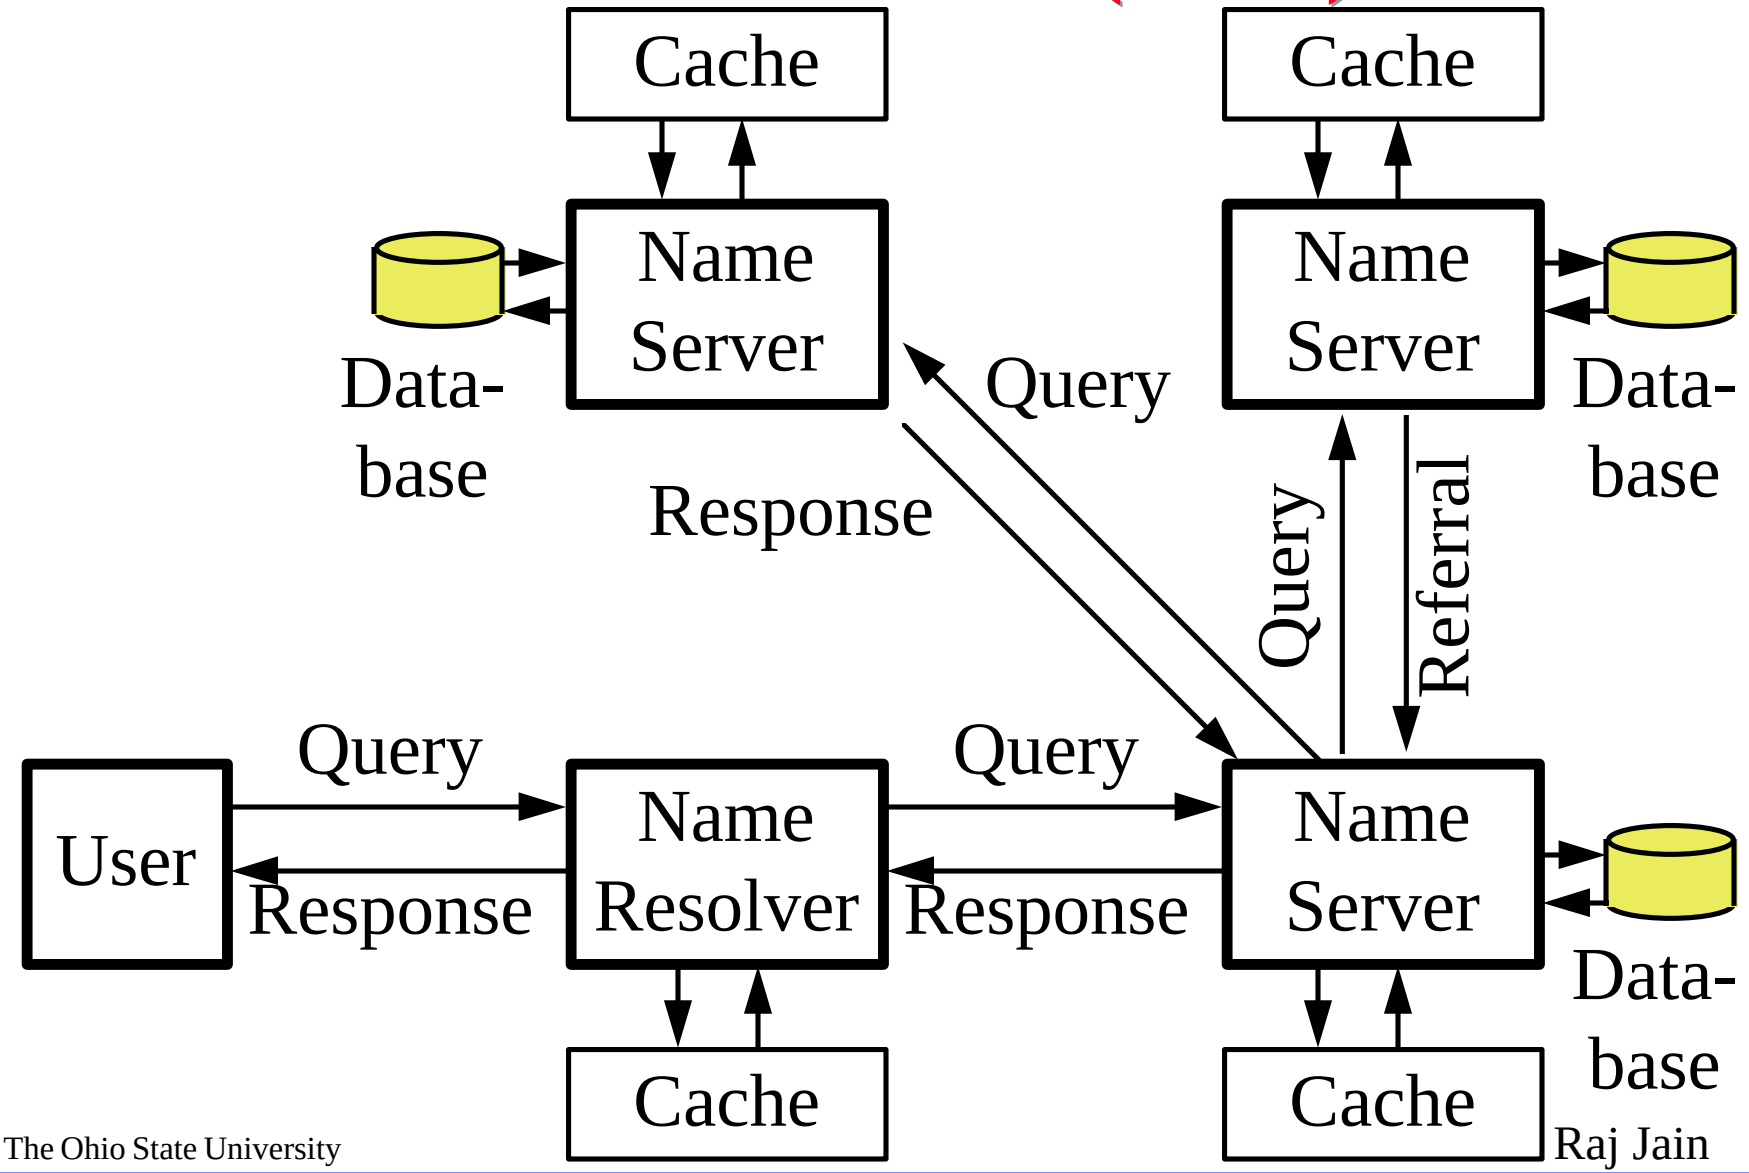

Name Resolution (Cont)

- Each computer has a name resolver routine, e.g., gethostbyname 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 (Cont)

- Resolvers use UDP (single name)
  or TCP (whole group of names)
- Knowing the address of the root server is sufficient
- Recursive Query:
  Give me an answer (Don't give me a referral)
- Iterative Query:
  Give me an answer or a referral to the next server
- Resolvers use recursive query.
- Servers use iterative query.

DNS Optimization

- Spatial Locality: Local computers referenced more often than remote
- Temporal Locality: Same set of domains referenced repeatedly ⇒ 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., Liberia.cis = Liberia.cis.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

Inverse Mapping

- Given an address, what is the name?
- nnn.nnn.nnn.nnn.in-addr.arpa

Types of DNS Entries

- DNS is used not just for name to address resolution
- But 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
- CNAME entry = Alias name (like a file link), "see name"
- www.foobar.com = hobbes.foobar.com

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

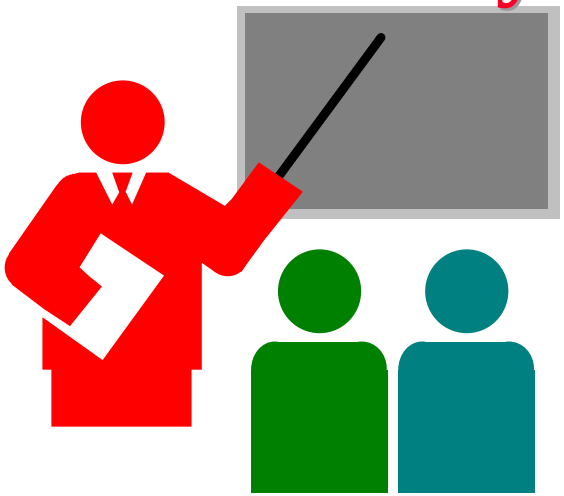

- DNS: Maps names to addresses
- Names are hierarchical. Administration is also hierarchical.
- No standard for number of levels
- Replication and caching is used for performance optimization.

Homework

- Read Chapter 24 and RFC 1034 and 1035
- Submit answer to exercise 24.1

DNS: References

- D. Comer, "Computer Networks and Internets," Prentice Hall, NJ, 1997, Chapter 24.
- D. Comer, "Internetworking with TCP/IP," Vol 1, 3rd Ed, Prentice Hall, 1995, Chapter 22.
- U. Black, "TCP/IP and Related Protocols," 2nd Ed, McGraw-Hill, Chapter 4.

DNS: RFCs

- [RFC2065] D. Eastlake, C. Kaufman, "Domain Name System Security Extensions", 01/03/1997, 41 pages. (Updates RFC1034)
- [RFC2052] A. Gulbrandsen, P. Vixie, "A DNS RR for specifying the location of services (DNS SRV)", 10/31/1996, 10 pages.
- [RFC2010] B. Manning, P. Vixie, "Operational Criteria for Root Name Servers", 10/14/1996, 7 pages.
- [RFC1995] M. Ohta, "Incremental Zone Transfer in DNS", 08/28/1996, 8 pages. (Updates RFC1035)
- [RFC1956] D. Engebretson, R. Plzak, "Registration in the MIL Domain", 06/06/1996, 2 pages.
- [RFC1912] D. Barr, "Common DNS Operational and Configuration Errors", 02/28/1996, 16 pages.
- [RFC1877] S. Cobb, "PPP Internet Protocol Control Protocol Extensions for Name Server Addresses", 12/26/1995, 6 pages.

- [RFC1876] C. Davis, P. Vixie, T. Goodwin, I. Dickinson, "A Means for Expressing Location Information in the Domain Name System", 01/15/1996, 18 pages. (Updates RFC1034)
- [RFC1816] F. Networking Council (FNC), "U.S. Government Internet Domain Names", 08/03/1995, 8 pages.
- [RFC1811] F. Networking Council, "U.S. Government Internet Domain Names", 06/21/1995, 3 pages.
- [RFC1788] W. Simpson, "ICMP Domain Name Messages", 04/14/1995, 7 pages.
- [RFC1779] S. Kille, "A String Representation of Distinguished Names", 03/28/1995, 8 pages.
- [RFC1737] K. Sollins, L. Masinter, "Functional Requirements for Uniform Resource Names", 12/20/1994, 7 pages.
- [RFC1713] A. Romao, "Tools for DNS debugging", 11/03/1994, 13 pages. (FYI 27)
- [RFC1685] H. Alvestrand, "Writing X.400 O/R Names", 08/11/1994, 11 pages. (RTR 12)

- [RFC1664] C. Allocchio, A. Bonito, B. Cole, S. Giordano, R. Hagens, "Using the Internet DNS to Distribute RFC1327 Mail Address Mapping Tables", 08/11/1994, 23 pages.
- [RFC1649] R. Hagens, A. Hansen, "Operational Requirements for X.400 Management Domains in the GO-MHS Community", 07/18/1994, 14 pages.
- [RFC1591] J. Postel, "Domain Name System Structure and Delegation", 03/03/1994, 7 pages.
- [RFC1537] P. Beertema, "Common DNS Data File Configuration Error", 10/06/1993, 9 pages.
- [RFC1536] A. Kumar, J. Postel, C. Neuman, P. Danzig, S. Miller, "Common DNS Implementation Errors and Suggested Fixes.", 10/06/1993, 12 pages.
- [RFC1535] E. Gavron, "A Security Problem and Proposed Correction With Widely Deployed DNS Software", 10/06/1993, 5 pages.
- [RFC1485] S. Hardcastle-Kille, "A String Representation of Distinguished Names (OSI-DS 23 (v5))", 07/28/1993, 7 pages.

- [RFC1480] A. Cooper, J. Postel, "The US Domain", 06/28/1993, 47 pages.
- [RFC1464] R. Rosenbaum, "Using the Domain Name System To Store Arbitrary String Attributes", 05/27/1993, 4 pages.
- [RFC1401] Internet Architecture Board, L. Chapin, "Correspondence between the IAB and DISA on the use of DNS throughout the Internet", 01/13/1993, 8 pages.
- [RFC1394] P. Robinson, "Relationship of Telex Answerback Codes to Internet Domains", 01/08/1993, 15 pages.
- [RFC1386] A. Cooper, J. Postel, "The US Domain", 12/28/1992, 31 pages.
- [RFC1383] C. Huitema, "An Experiment in DNS Based IP Routing", 12/28/1992, 14 pages.
- [RFC1279] S. Kille, "X.500 and Domains", 11/27/1991, 13 pages.
- [RFC1183] R. Ullman, P. Mockapetris, L. Mamakos, C. Everhart, "New DNS RR Definitions", 10/08/1990, 11 pages.
- [RFC1178] D. Libes, "Choosing a Name for Your Computer", 09/04/1990, 8 pages. (FYI 5)

- [RFC1101] P. Mockapetris, "DNS encoding of network names and other types", 04/01/1989, 14 pages. (Updates RFC1034)
- [RFC1035] P. Mockapetris, "Domain names - implementation and specification", 11/01/1987, 55 pages. (STD 13) (Updated by RFC1348, RFC1995, RFC1996)
- [RFC1034] P. Mockapetris, "Domain names - concepts and facilities", 11/01/1987, 55 pages. (STD 13) (Updated by RFC1982, RFC2065, RFC1876, RFC1101)
- [RFC1033] M. Lottor, "Domain administrators operations guide", 11/01/1987, 22 pages.
- [RFC1032] M. Stahl, "Domain administrators guide", 11/01/1987, 14 pages.
- [RFC1031] W. Lazear, "MILNET name domain transition", 11/01/1987, 10 pages.
- [RFC0982] H. Braun, "Guidelines for the specification of the structure of the Domain Specific Part DSP of the ISO standard NSAP address", 04/01/1986, 11 pages.

- [RFC0974] C. Partridge, "Mail routing and the domain system", 01/01/1986, 7 pages. (STD 14)
- [RFC0973] P. Mockapetris, "Domain system changes and observations", 01/01/1986, 10 pages. (Updates RFC0882)
- [RFC0921] J. Postel, "Domain name system implementation schedule - revised", 10/01/1984, 13 pages. (Updates RFC0897)
- [RFC0920] J. Postel, J. Reynolds, "Domain requirements", 10/01/1984, 14 pages.
- [RFC0897] J. Postel, "Domain name system implementation schedule", 02/01/1984, 8 pages. (Updates RFC0881) (Updated by RFC0921)
- [RFC0883] P. Mockapetris, "Domain names: Implementation specification", 11/01/1983, 73 pages.
- [RFC0882] P. Mockapetris, "Domain names: Concepts and facilities", 11/01/1983, 31 pages. (Updated by RFC0973)
- [RFC0881] J. Postel, "Domain names plan and schedule", 11/01/1983, 10 pages. (Updated by RFC0897)

- [RFC0830] Z. Su, "Distributed system for Internet name service", 10/01/1982, 16 pages.
- [RFC0819] Z. Su, J. Postel, "Domain naming convention for Internet user applications", 08/01/1982, 18 pages.
- [RFC0799] D. Mills, "Internet name domains", 09/01/1981, 6 pages.