

Today's Objectives

- Naming Challenges
- Domain Name System

Oct 9, 2017

Sprenkle - CSCI325

1

Review

- What is RPC?
 - What is its purpose?
 - How does it work?
 - What implementation of RPC are we using in Project 2?
- What is RMI? How does it relate to RPC?

Oct 9, 2017

Sprenkle - CSCI325

2

Terminology

- **Pure names** - uninterpreted bit patterns; must be “looked up” before they are used
 - Example: Mac address
- **Non-pure names** - contain information about the object they name (such as location or address); inadequate for object identification over time
 - Example: IP address
- **Resolution** - translate a name into useful data about an object
- **Binding** - association between name and object
- **Attribute** - value of a property associated with an object (names are often bound to attributes)
- **Contexts** - sets of bindings between names and attributes for objects

Oct 9, 2017

Sprenkle - CSCI325

3

Naming Services Design Goals

- Handle an arbitrary number of names
- Have a long (infinite?) lifetime
- Provide high availability
- Hide (or isolate) faults
- Tolerate mistrust

Oct 9, 2017

Sprenkle - CSCI325

4

Types of Naming Systems

- Flat
- Structured
- Attribute-based

Oct 9, 2017

Sprenkle - CSCI325

5

Flat Systems

- Tend to be for LAN
- No location or other information
- Example: RPC or RMI naming

Oct 9, 2017

Sprenkle - CSCI325

6

Structured naming

DOMAIN NAME SYSTEM

Oct 9, 2017

Sprenkle - CSCI325

7

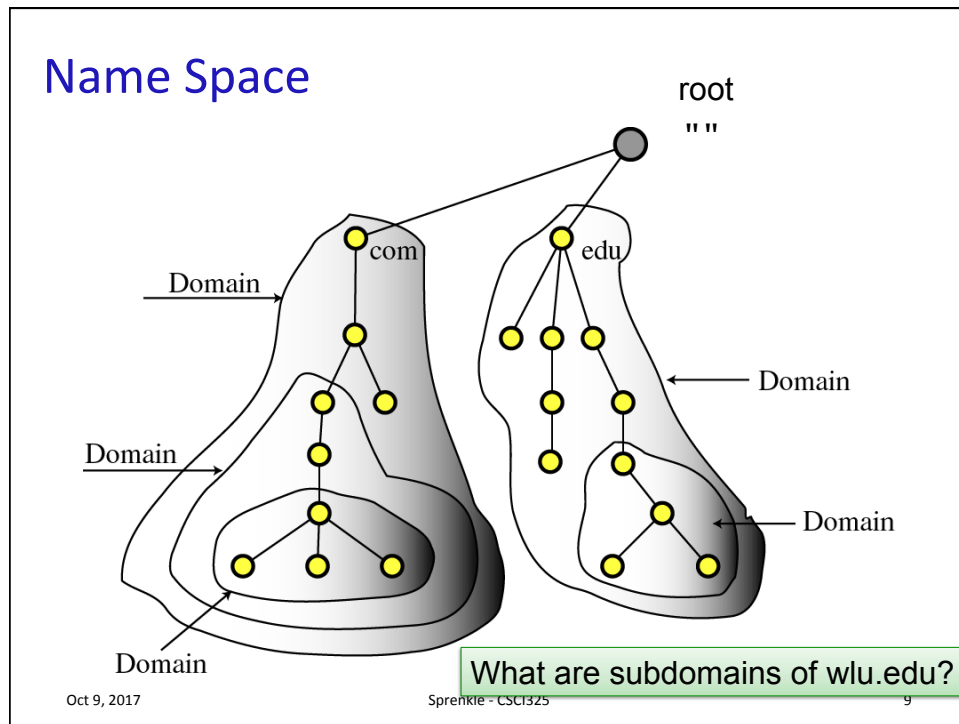
DNS Components

- A globally distributed, scalable, reliable database
- Name Space:
 - Specifications for a structured name space and data associated with the names
- Resolvers:
 - Client programs that extract information from Name Servers.
- Name Servers:
 - Server programs which hold information about the structure and the names.

Oct 9, 2017

Sprenkle - CSCI325

8



Global Distribution

- Data is maintained locally, but retrievable globally
 - No single computer has all DNS data
- DNS lookups can be performed by any device
- Remote DNS data is locally cachable to improve performance

Scalability

- No limit to the size of the database
- No limit to the number of queries
 - Tens of thousands of queries handled easily every second
- Queries distributed among primaries, secondaries, and caches

Oct 9, 2017

Sprenkle - CSCI325

11

Reliability

- Data is replicated
 - Data from primary is copied to multiple secondaries
- Clients can query
 - Primary server
 - Any of the copies at secondary servers
- Clients will typically query local caches
- DNS protocols can use either UDP or TCP
 - If UDP, DNS protocol handles retransmission, sequencing, etc.

Oct 9, 2017

Sprenkle - CSCI325

12

Dynamicity

- Primary database can be updated dynamically
 - Add/delete/modify of any record
- Modification of the primary database triggers replication

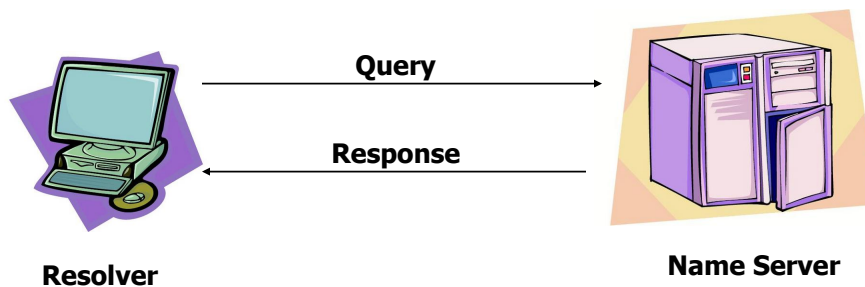
Oct 9, 2017

Sprenkle - CSCI325

13

Resolvers

- Resolver maps a name to an address and vice versa.

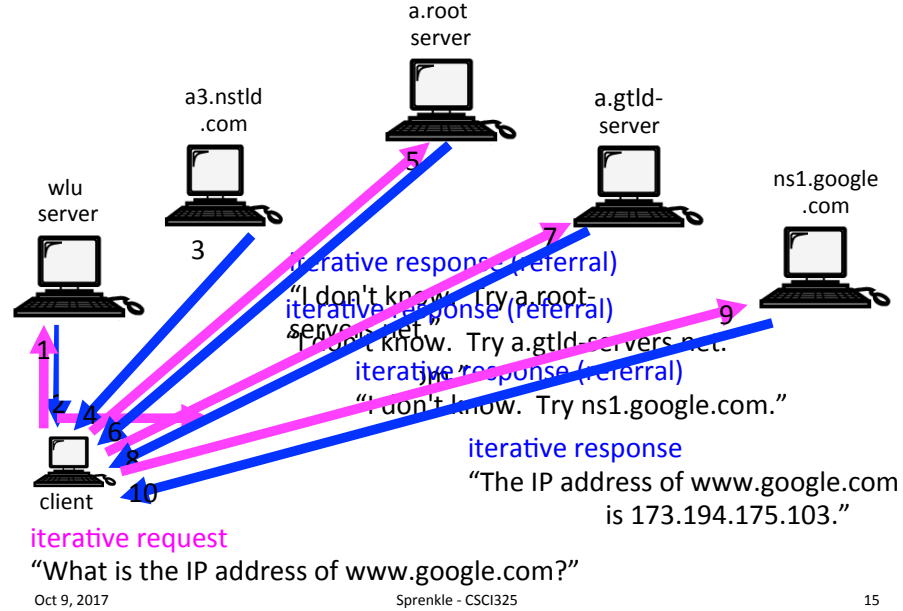


Oct 9, 2017

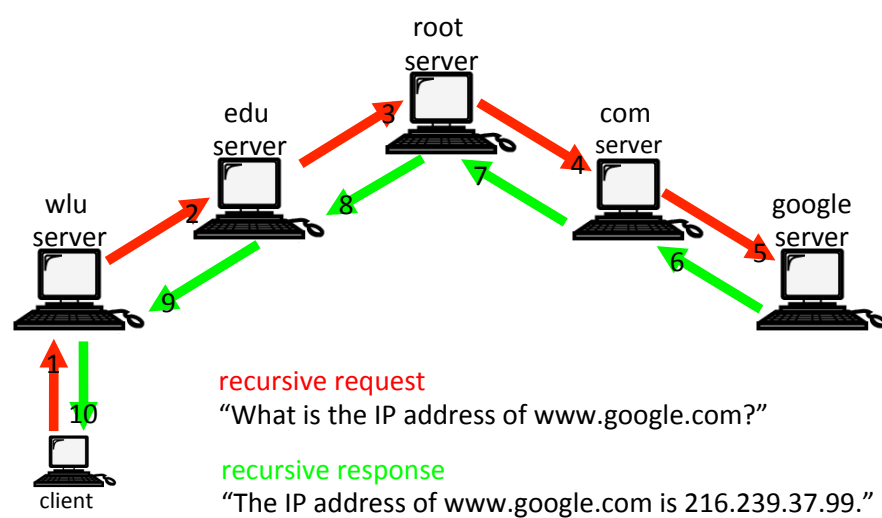
Sprenkle - CSCI325

14

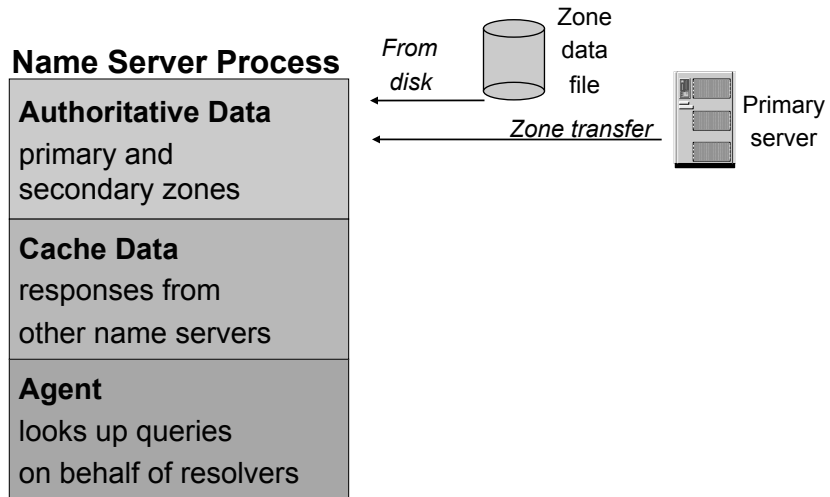
Iterative Resolution: Referrals



Recursive Resolution



Name Server: Architecture

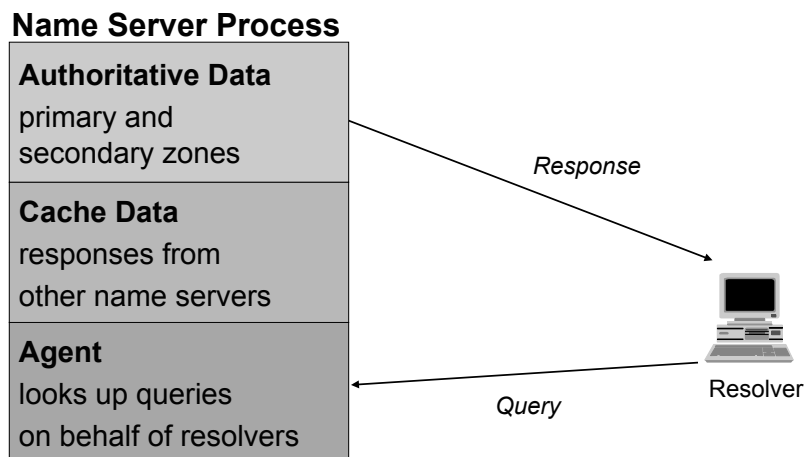


Oct 9, 2017

Sprenkle - CSCI325

17

Name Server: Authoritative Data



Oct 9, 2017

Sprenkle - CSCI325

18

Name Server: Using Other Name Servers

Name Server Process

Authoritative Data


Primary and
Secondary zones

Cache Data

responses from
other name servers

Agent

looks up queries
on behalf of resolvers



Oct 9, 2017

Sprenkle - CSCI325

19

Name Server: Cached Data

Name Server Process

Authoritative Data


Primary and
Secondary zones

Cache Data

responses from
other name servers

Agent

looks up queries
on behalf of resolvers

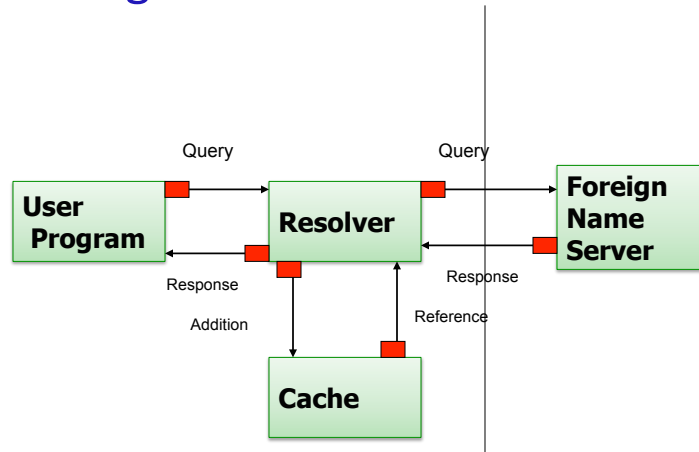


Oct 9, 2017

Sprenkle - CSCI325

20

Block Diagram

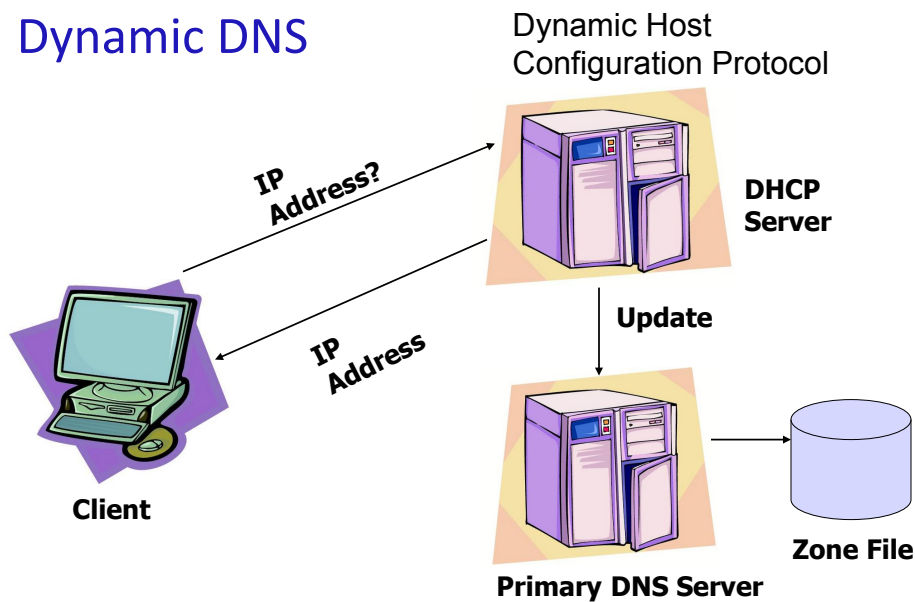


Oct 9, 2017

Sprengle - CSCI325

21

Dynamic DNS



Oct 9, 2017

Sprengle - CSCI325

22

Looking Ahead

- COD paper – due tonight
- Preliminary Bookstore deadline next Monday