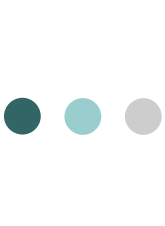

CS514: Intermediate Course in Computer Systems

Lecture 4: January 27, 2003
"Introduction to Naming"

CS514

- **"Any problem in computer science can be solved with another layer of indirection"**

  David Wheeler

Naming is a layer of indirection

CS514

- What problems does it solve?
  - Makes objects human readable
  - Hides complexity and dynamics
    - Multiple lower-layer objects can have one name
    - Changes in lower-layer objects hidden
  - Allows an object to be found in different ways
    - One object can have multiple names

Names map to objects through a resolution service

CS514

Identifiers and Locators

CS514

- A name is always an *identifier* to a greater or lesser extent
  - Can be persistent or non-persistent
  - Can be globally unique, locally unique, or even non-unique
- If a name has structure that helps the resolution service, then the name is also a *locator*

Naming in networks

CS514

DNS names map into addresses

CS514

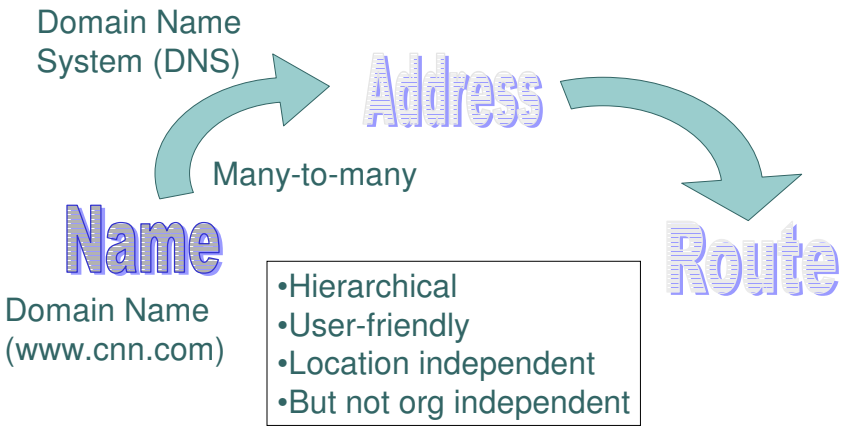

- Hierarchical
- User-friendly
- Location independent
- But not org independent

Addresses map into routes

CS514

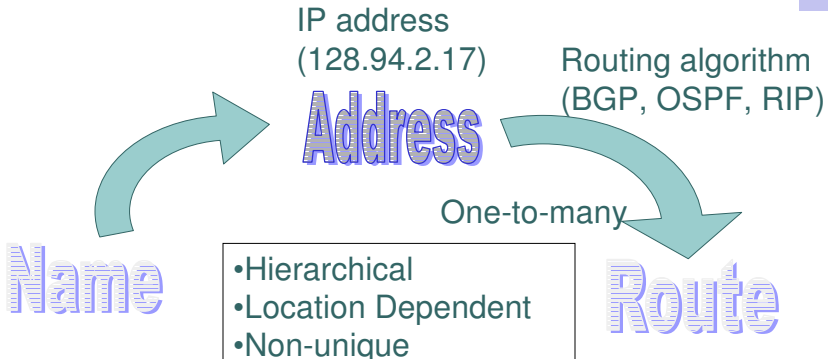

- Hierarchical
- Location Dependent
- Non-unique
- Can change often
- Refers to an interface, not a host

Routes get packets to interfaces

CS514

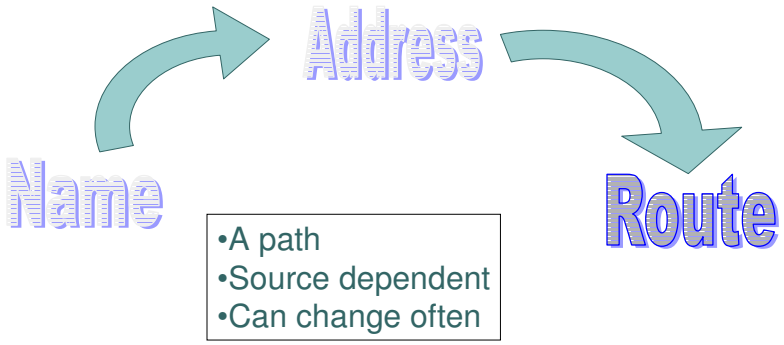

- A path
- Source dependent
- Can change often

DNS names and IP addresses are identifiers and locators

CS514

- Both are typically non-persistent
- Private IP addresses identify only in the context of an IP realm
- Domain names are good identifiers
  - woodstock.cs.cornell.edu identifies a host
  - www.cnn.com identifies a service
- URLs are good identifiers

Domain Name System (DNS)

CS514

- Distributed directory service
- Hierarchical name space
- Each level separated by '.'
  - Analogous to '/' separator in file systems
- One global root
  - Replicated across <20 root servers!
  - There have been Denial of Service (DoS) attacks on these root servers, none real successful
  - Because of caching, queries to root servers relatively rare
- DNS maybe only global directory service???

DNS is simple but powerful

CS514

- Only one type of query
  - Query(domain name, RR type)
    - Resource Record (RR) type is like an attribute type
  - Answer(values, additional RRs)
- Limited number of RR types
- Hard to make new RR types
  - Not for technical reasons…
  - Rather because each requires global agreement

DNS is the core of the Internet

CS514

- Global name space
  - Can be the core of a naming or identifying scheme
- Global directory service
  - Can resolve a name to nearly every computer on the planet

Important DNS RR types

CS514

- **NS**: Points to next Name Server down the tree
- **A**: Contains the IP address
  - **AAAA** for IPv6
- **MX**: Contains the name of the mail server
- Service-oriented RR types
  - **SRV**: Contains addresses and ports of services on servers
    - One way to learn what port number to use
  - **NAPTR**: Essentially a generalized mapping from one name space (i.e. phone numbers) to another (i.e. SIP URL)

DNS tree structure

CS514

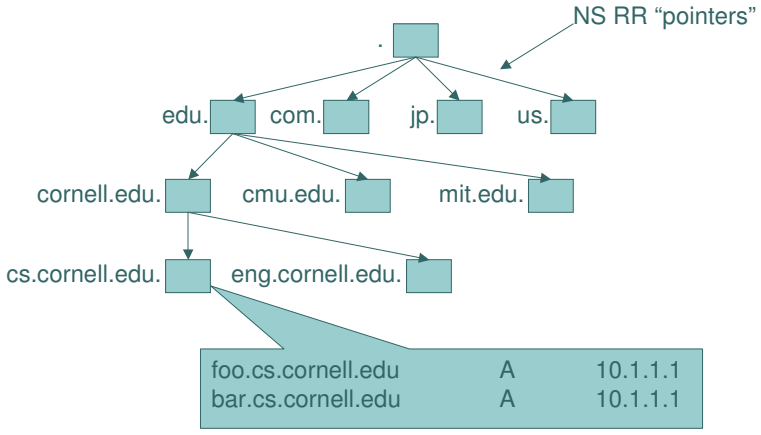

Primary and secondary servers

CS514

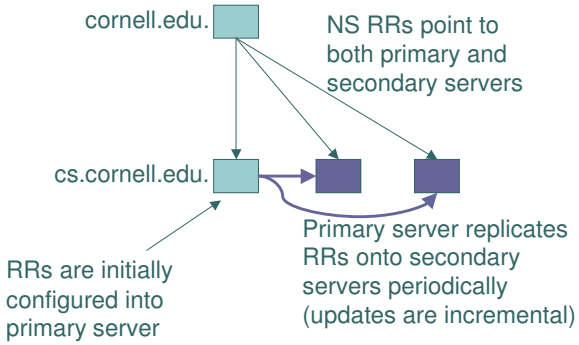

Resolver structure and configuration

CS514

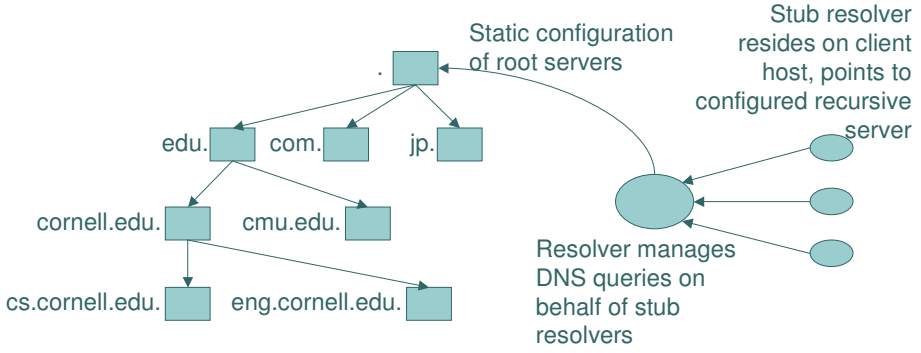

Resolver structure and configuration

CS514

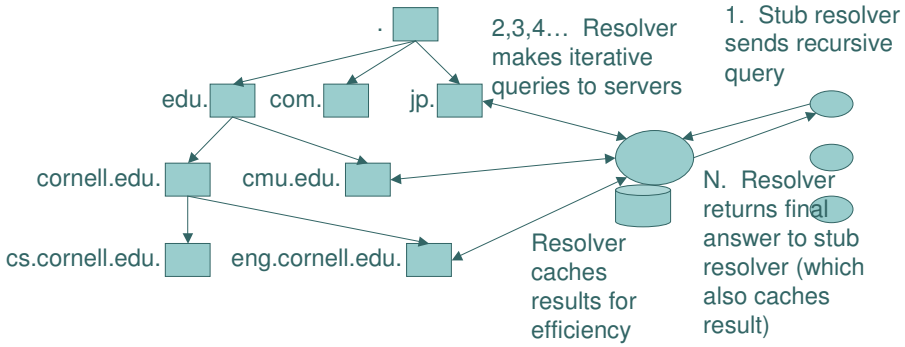

DNS cache management

CS514

- All RRs have Time-to-live (TTL) values
- When TTL expires, cache entries are removed
- NS RRs tend to have long TTLs
  - Cached for a long time
  - Reduces load on higher level servers
- A RRs may have very short TTLs
  - Order one minute for some web services
  - Order one day for typical hosts

Why is DNS iterative and not recursive?

CS514

- AT/MvS teach that recursive is more efficient
  - Better caching characteristics
    - Caches in servers, not just resolvers
  - Shorter paths
- However, high-performance recursive server much harder to implement
  - Maintain state for thousands of concurrent queries
  - Manage cache
- Recursive server prone to DoS attacks

* AT/MvS = Andrew Tanenbaum/Martin van Steen text

LDAP is another popular distributed directory service

CS514

- Richer and more general than DNS
  - Has generalized attribute/value scheme
  - Can search on attribute, not just name
- Simpler and more efficient than a full relational database
- Not a global directory service, though namespace is global
  - Its predecessor, X.500, was meant to be
  - But “local” LDAP services can point to each other
- Commonly used for personnel RR databases, subscriber databases

URLs, URNs, and URIs

CS514

- Uniform Resource <Locator, Name, Identifier>
- URL tells a computer where and how to reach a resource
  - These came first
- URN is a true identifier
  - Unique, persistent
- URI refers to both URLs and URNs
  - Defines syntax for current and future URLs and URNs
- *For now we only really care about URLs*

URL

CS514

- Consists of:

<scheme>:<scheme-specific-part>

URL

CS514

- Consists of:

<scheme>:<scheme-specific-part>

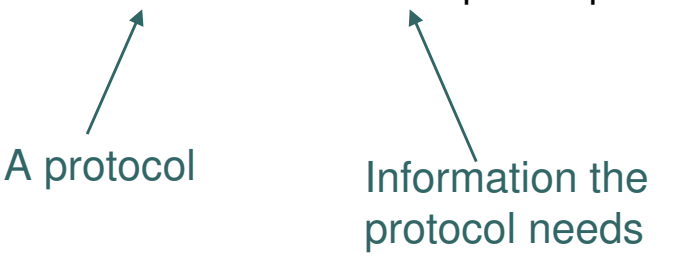

URL examples

CS514

- HTTP (web)
  - http://www.cnn.com/news/story.html
- Email
  - mailto://francis@cs.cornell.edu
- Newsgroups
  - news:cornell/class/cs514
- SIP (Session Initiation Protocol)
  - sip://service@phone.verizon.com

Note the central role of DNS

CS514

- HTTP (web)
  - http://***www.cnn.com***/news/story.html
- Email
  - mailto://francis@***cs.cornell.edu***
- Newsgroups
  - news:cornell/class/cs514
- SIP (Session Initiation Protocol)
  - sip://service@***phone.verizon.com***

Locating mobile entities (section 4.2, AT/MvS)

CS514

- What is a mobile entity?
- From naming perspective, it is an entity whose address changes often
- This doesn't require physical mobility!
  - Every time you dial up, you may get a new address
- So, "mobility" existed well before laptops became common
  - Though laptops create more mobility

Is mobility a problem for DNS?

CS514

- Not really
  - Even though DNS was designed with relatively stable IP addresses in mind
- Because mobility only effects leaf DNS servers
  - Recall: A RR TTL is short, but NS RR TTL is long
- Note: *non-mobile* web server's A RRs often have very short TTLs
  - To allow quick failover to another web server

Is mobility a problem at all?

CS514

- Less than you'd think
- Most mobile systems are clients; servers are rarely mobile
  - Clients are initiators of connections, not recipients
  - Therefore, there is not a client locating problem
- What about email, instant messaging, and VoIP (Voice over IP)?
  - Clients receive emails, instant messages, and phone calls

Application specific registration as a mobility solution

CS514

- To receive email, client connects to an email server
- To do instant messaging, client registers with an IM server
- To do VoIP, client registers with a SIP server

This is an adequate solution to 90% of mobility issues

- This is why Mobile IP hasn't gotten traction (i.e. Microsoft has not implemented it)

Mobile IP uses an IP-level registration

CS514

Mobile Node has a stable home address at its home network

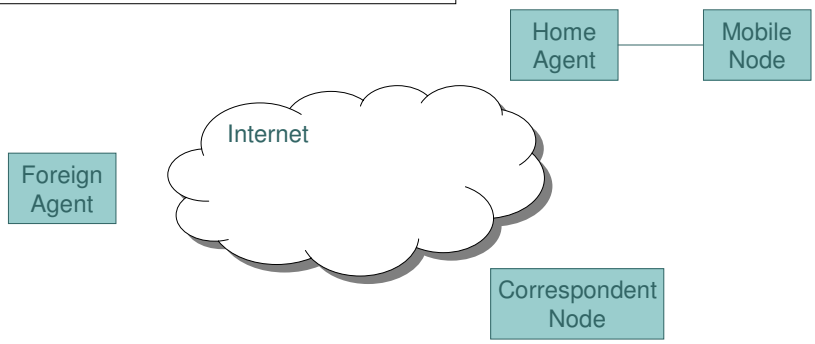

Mobile IP uses an IP-level registration

CS514

Mobile Node moves to foreign network, gets a Care-of Address

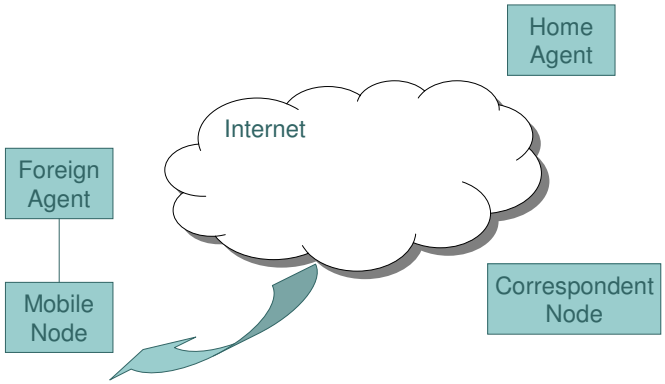

Mobile IP uses an IP-level registration

CS514

Mobile Node registers with Home Agent, creates IP tunnel

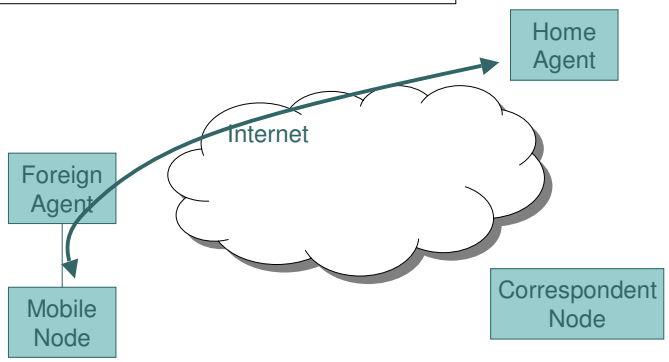

Mobile IP uses an IP-level registration

CS514

Connection initiated by Correspondent Node will be tunneled to Mobile Node

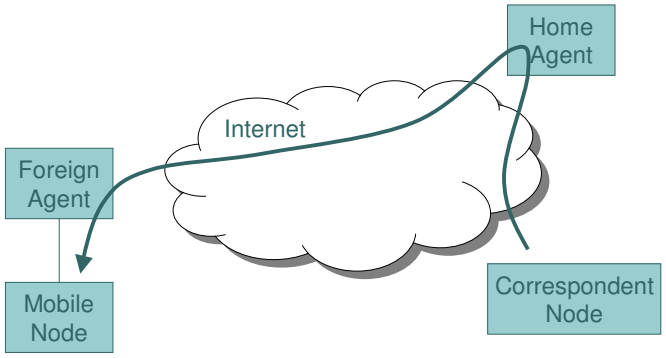

Mobile IP adds a layer of indirection

CS514

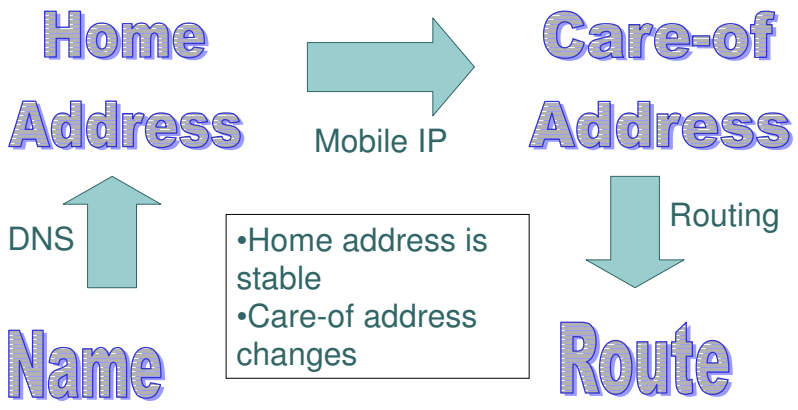

Client identification

CS514

- Servers cannot locate clients, but often must be able to identify them
- HTTP cookies serve this role
- HTTP cookies also contain many attributes about the client or session
- They also typically contain some kind of signature
  - To prevent tampering

Identifiers must be made hard to spoof

CS514

- That is why driver's licenses have pictures and credit cards have signatures
- In networking, two ways:
  1. Identifier is also a locator
     - Reverse routability
  2. Some kind of secret-protected signature

Reverse routability: DoS and Mobile IP

CS514

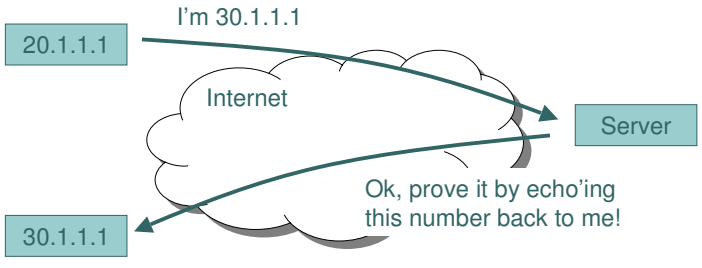

Since challenge doesn't go back to 20.1.1.1 (i.e. is not reverse routable), 20.1.1.1 cannot spoof 30.1.1.1

Summary of Lecture

CS514

Introduction to Naming

- Naming basics:
  - Names, Addresses, Routes
  - Identifiers and Locators
- DNS is *the* global directory service
  - LDAP is a popular local directory service
- URLs build on DNS (and also URIs and URNs)
- Mobility is not much of a problem
  - Chapter 4 got this wrong
- Identifiers must be hard to spoof
  - Reverse routability, cryptographic signatures