

CIS 505: Software Systems Lecture Note on Naming: Part 2

Insup Lee 
Department of Computer and Information Science
University of Pennsylvania

Spring 2007

1

Domain Name Service (DNS)

- A planetary name service that translates Internet domain names, which are the basis for the Web's global URL space.
 - maps <node name> to <IP address>
 - names for autonomous *naming domains*, e.g., *cis.upenn.edu*
 - names for specific nodes, e.g., *eniac.cis.upenn.edu*
 - names for service aliases (e.g., *www*, *mail* servers)
 - (mostly) independent of location, routing, etc.
- Use of hierarchy for scalability
 - decentralized administration of the name space
 - hierarchical authority and trust

CIS 505, Spring 2007

Naming

2

Top-Level Domains (TLDs)

- Mostly controlled by Network Solutions, Inc. today
 - *.com*: commercial
 - *.edu*: educational institution
 - *.gov*: US government
 - *.mil*: US military
 - *.net*: networks and ISPs (now also a number of other things)
 - *.org*: other organizations
 - 244, 2-letter country suffixes, e.g., *.us*, *.uk*, *.cz*, *.tv*, ...
 - and a bunch of new suffixes that are not very common, e.g., *.biz*, *.name*, *.pro*, ...
- Finding the Root
 - 13 "root servers" store entries for all **top level domains** (TLDs)
 - DNS servers have a hard-coded mapping to root servers so they can "get started"

[Z. Ives]

CIS 505, Spring 2007

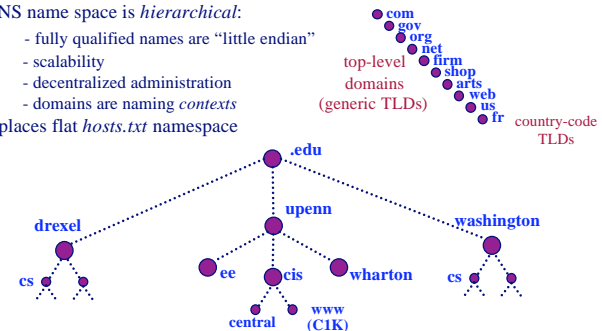
Naming

3

Domain Name Hierarchy

DNS name space is *hierarchical*:

- fully qualified names are "little endian"
 - scalability
 - decentralized administration
 - domains are naming *contexts*
- replaces flat *hosts.txt* namespace

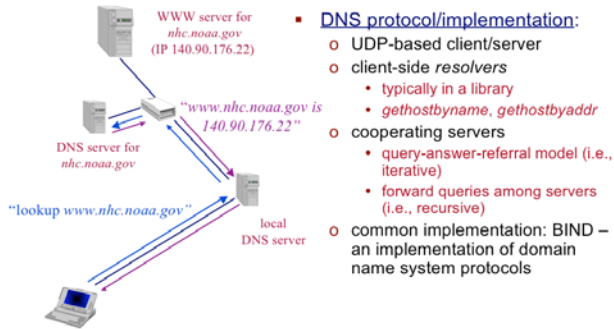


CIS 505, Spring 2007

Naming

4

DNS Implementation



- **DNS protocol/implementation:**
 - UDP-based client/server
 - client-side *resolvers*
 - typically in a library
 - *gethostbyname, gethostbyaddr*
 - cooperating servers
 - query-answer-referral model (i.e., iterative)
 - forward queries among servers (i.e., recursive)
 - common implementation: BIND – an implementation of domain name system protocols

CIS 505, Spring 2007


Naming

5

DNS Name Server Hierarchy

DNS servers are organized into a hierarchy that mirrors the name space.

Specific servers are designated as *authoritative* for portions of the name space.



CIS 505, Spring 2007

Naming

6

The DNS Name Space

Type of record	Associated entity	Description
SOA	Zone	Holds information on the represented zone
A	Host	Contains an IP address of the host this node represents
MX	Domain	Refers to a mail server to handle mail addressed to this node
SRV	Domain	Refers to a server handling a specific service
NS	Zone	Refers to a name server that implements the represented zone
CNAME	Node	Symbolic link with the primary name of the represented node
PTR	Host	Contains the canonical name of a host
HINFO	Host	Holds information on the host this node represents
TXT	Any kind	Contains any entity-specific information considered useful

- The most important types of resource records forming the contents of nodes in the DNS name space.

CIS 505, Spring 2007

Naming

7

DNS Table

- An excerpt from the DNS database for the zone *cs.vu.nl*.

Name	Record type	Record value
cs.vu.nl	SOA	star (1999121502,7200,3600,2419200,86400)
cs.vu.nl	NS	star.cs.vu.nl
cs.vu.nl	NS	top.cs.vu.nl
cs.vu.nl	NS	solo.cs.vu.nl
cs.vu.nl	TXT	"Vrije Universiteit - Math. & Comp. Sc."
cs.vu.nl	MX	1 zephyr.cs.vu.nl
cs.vu.nl	MX	2 tornado.cs.vu.nl
cs.vu.nl	MX	3 star.cs.vu.nl
star.cs.vu.nl	HINFO	Sun Unix
star.cs.vu.nl	MX	1 star.cs.vu.nl
star.cs.vu.nl	MX	10 zephyr.cs.vu.nl
star.cs.vu.nl	A	130.37.24.5
star.cs.vu.nl	A	192.31.231.42
zephyr.cs.vu.nl	HINFO	Sun Unix
zephyr.cs.vu.nl	MX	1 zephyr.cs.vu.nl
zephyr.cs.vu.nl	MX	2 tornado.cs.vu.nl
zephyr.cs.vu.nl	A	192.31.231.66
www.cs.vu.nl	CNAME	soling.cs.vu.nl
ftp.cs.vu.nl	CNAME	soling.cs.vu.nl
soling.cs.vu.nl	HINFO	Sun Unix
soling.cs.vu.nl	MX	1 soling.cs.vu.nl
soling.cs.vu.nl	MX	10 zephyr.cs.vu.nl
soling.cs.vu.nl	A	130.37.24.11
laser.cs.vu.nl	HINFO	PC MS-DOS
laser.cs.vu.nl	A	130.37.30.32
vucs-das.cs.vu.nl	PTR	0.26.37.130.in-addr.arpa
vucs-das.cs.vu.nl	A	130.37.26.0

CIS 505, Spring 2007

DNS: some issues

- Naming contexts
 - I want to use short, unqualified names like *central* instead of *central.cis.upenn.edu* when I'm in the *cis.upenn.edu* domain.
- What about trust? How can we know if a server is authoritative, or just an impostor?
 - What happens if a server lies or behaves erratically? What denial-of-service attacks are possible? What about privacy?
- What if an "upstream" server fails?
- Is the hierarchical structure sufficient for scalability?
 - more names vs. higher request rates
- DNS caching

Hierarchical Implementations: LDAP (1)

Attribute	Abbr.	Value
Country	C	NL
Locality	L	Amsterdam
Organization	O	Vrije Universiteit
OrganizationalUnit	OU	Comp. Sc.
CommonName	CN	Main server
Mail_Servers	—	137.37.20.3, 130.37.24.6, 137.37.20.10
FTP_Server	—	130.37.20.20
WWW_Server	—	130.37.20.20

Hierarchical Implementations: LDAP (2)




Figure 5-23. (a) Part of a directory information tree.

Hierarchical Implementations: LDAP (3)

Attribute	Value	Attribute	Value
Country	NL	Country	NL
Locality	Amsterdam	Locality	Amsterdam
Organization	Vrije Universiteit	Organization	Vrije Universiteit
OrganizationalUnit	Comp. Sc.	OrganizationalUnit	Comp. Sc.
CommonName	Main server	CommonName	Main server
Host_Name	star	Host_Name	zephyr
Host_Address	192.31.231.42	Host_Address	137.37.20.10


(b)

Figure 5-23. (b) Two directory entries having *Host_Name* as RDN.

Distributed Hash Tables

General Mechanism

- Figure 5-4. Resolving key 26 from node 1 and key 12 from node 28 in a Chord system.




CIS 505, Spring 2007

Naming

13

Mapping to Distributed Hash Tables (1)

```
description {
  type = book
  description {
    author = Tolkien
    title = LOTR
  }
  genre = fantasy
}
```



(a)

(b)

- Figure 5-24. (a) A general description of a resource. (b) Its representation as an AVTree.


CIS 505, Spring 2007

Naming

14

Mapping to Distributed Hash Tables (2)

```
description {
  type = book
  description {
    author = Tolkien
    title = *
  }
  genre = *
}
```



(a)

(b)


- Figure 5-25. (a) The resource description of a query. (b) Its representation as an AVTree.

CIS 505, Spring 2007

Naming

15

Semantic Overlay Networks



- Figure 5-26. Maintaining a semantic overlay through gossiping.

CIS 505, Spring 2007

Naming

16