

CS 550: **Advanced Operating Systems**

Naming

Wei Tang
Computer Science Department
Illinois Institute of Technology

CS 550
Advanced Operating Systems
February 24th, 2011

Outline

- Definition
 - Names, Identifiers and Addresses
- Name spaces
- Name resolution
- Example: The Domain Name System
- Example: X.500, LDAP

Naming

- A name in a distributed system is a string of bits or characters that is refer to an entity
 - Example of entity?

Hosts, printers, disks, files,

Processes, users, mailboxes, news groups,

Web pages, graphical windows, message, network connections,

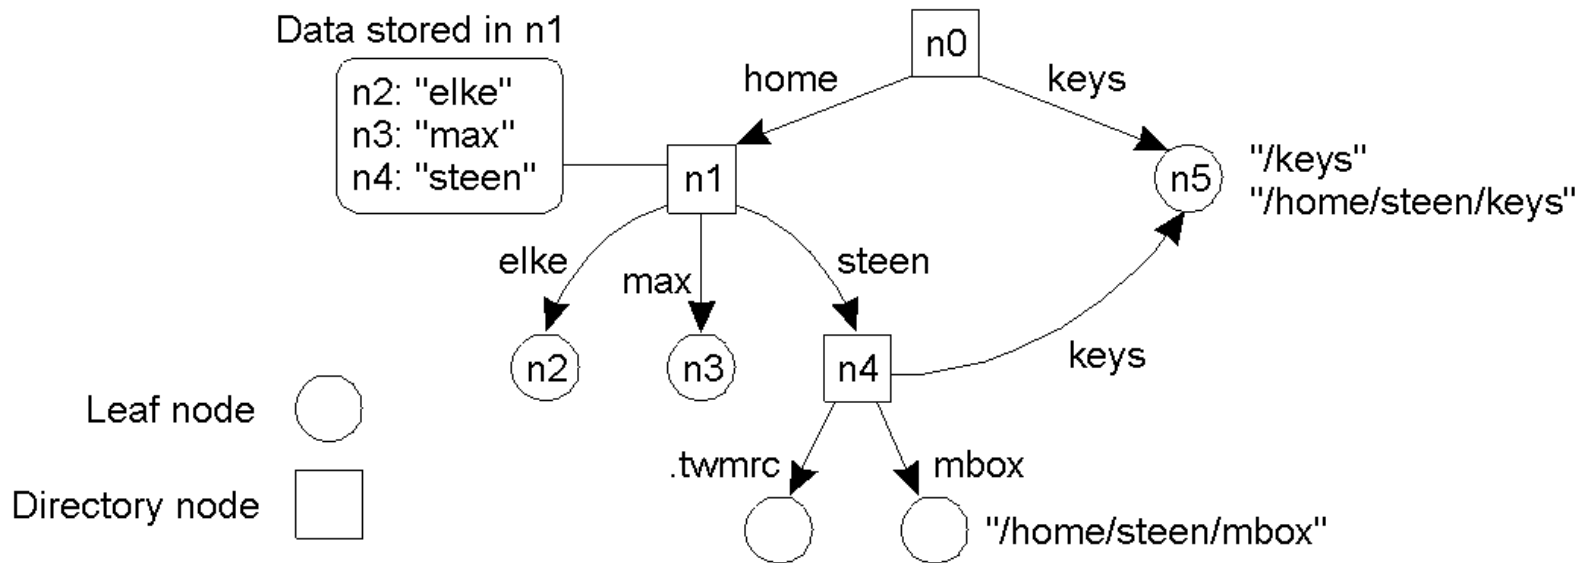
And so on.....

Naming

- Three types of names:
 - **Address:** the name of an access point to an entity
 - **Identifier:**
 - An identifier refers to at most one entity
 - Each entity is referred to by at most one identifier
 - An identifier always refers to the same entity
 - **Human-friendly name:**
 - E.g. unix file name, DNS names
- *Names are always organized in a name space*
 - *A **name space** is an organization mechanism for a group of names.*

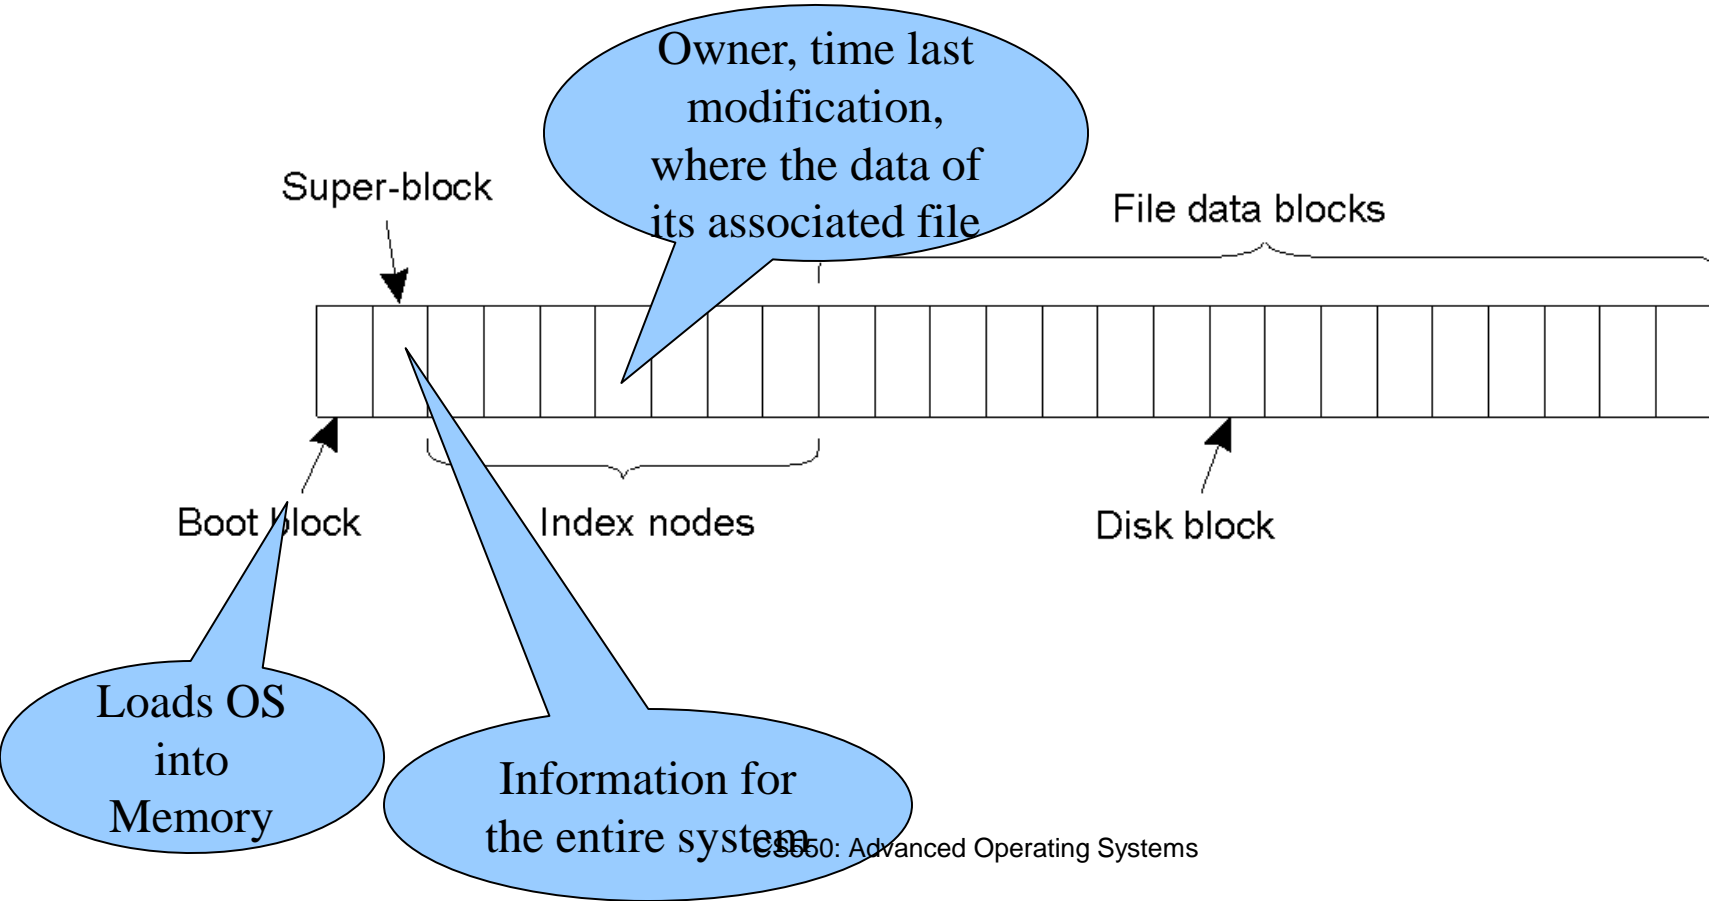
Name Spaces

- Directed acyclic graph (DAG)
 - The graph does not have a cycle

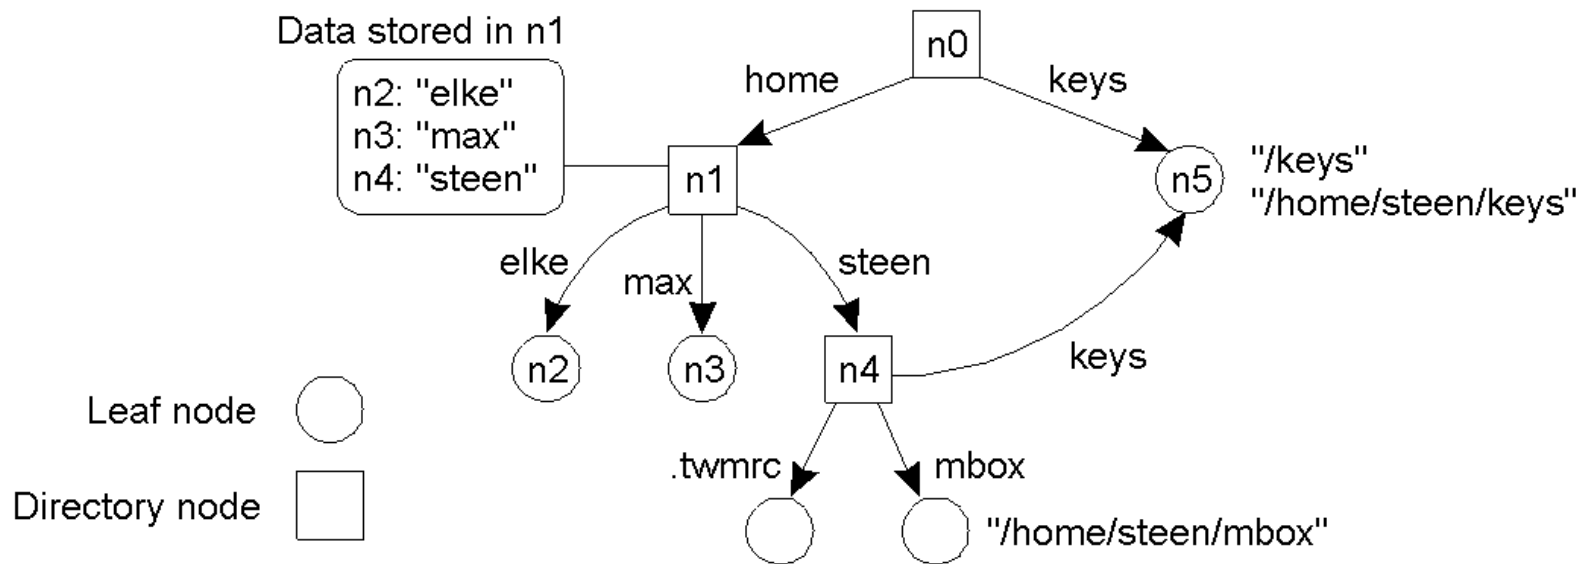


Name Spaces

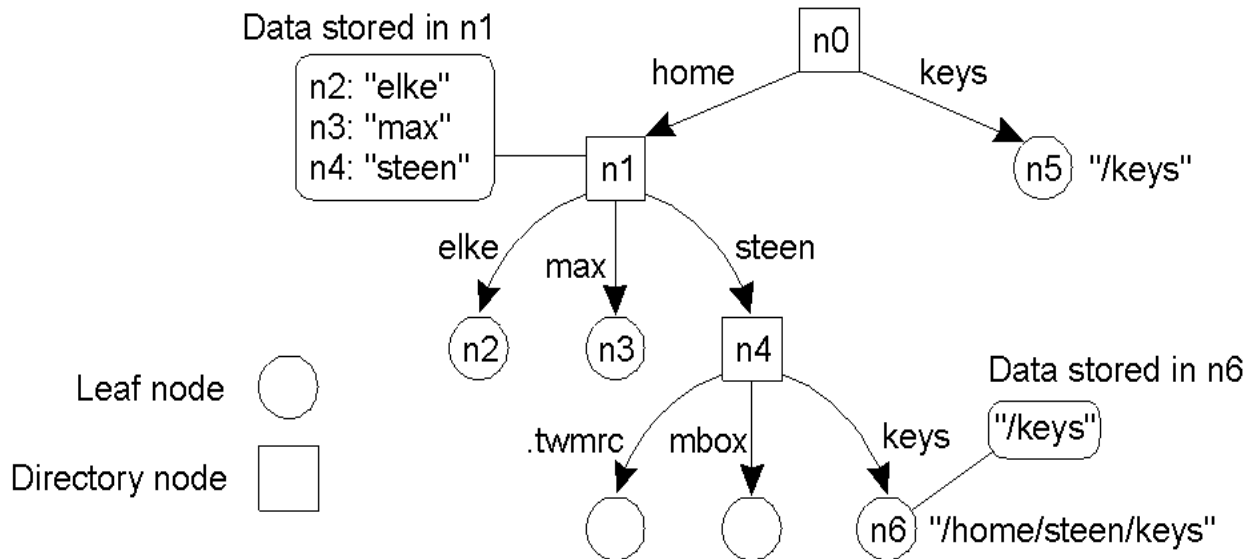
- The general organization of the UNIX file system implementation on a logical disk of contiguous disk blocks.



Hard Link

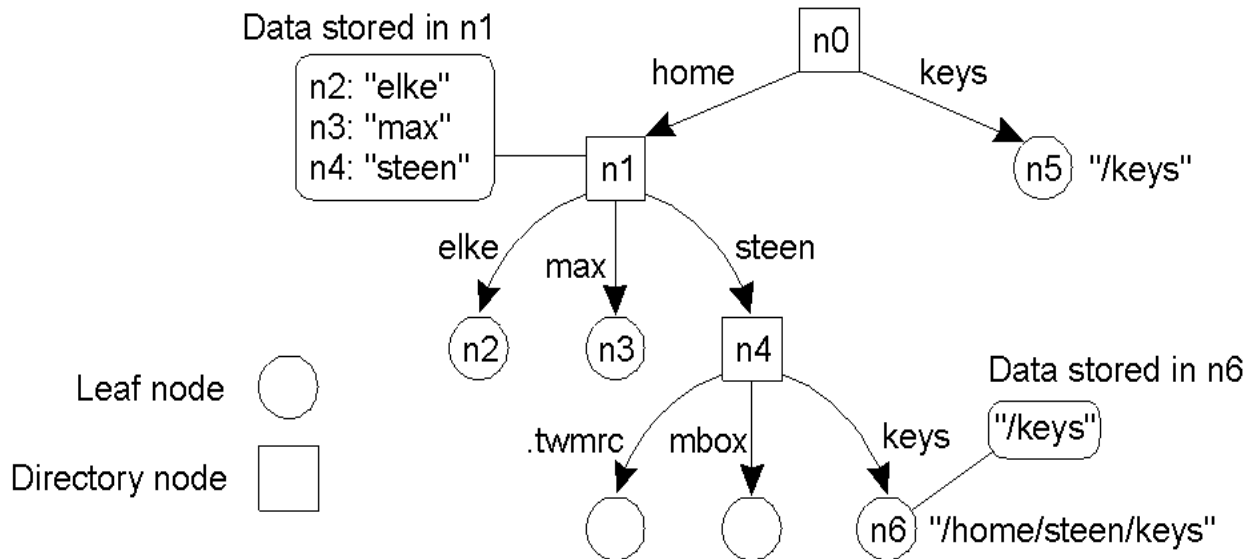


Symbolic Link



Name Resolution

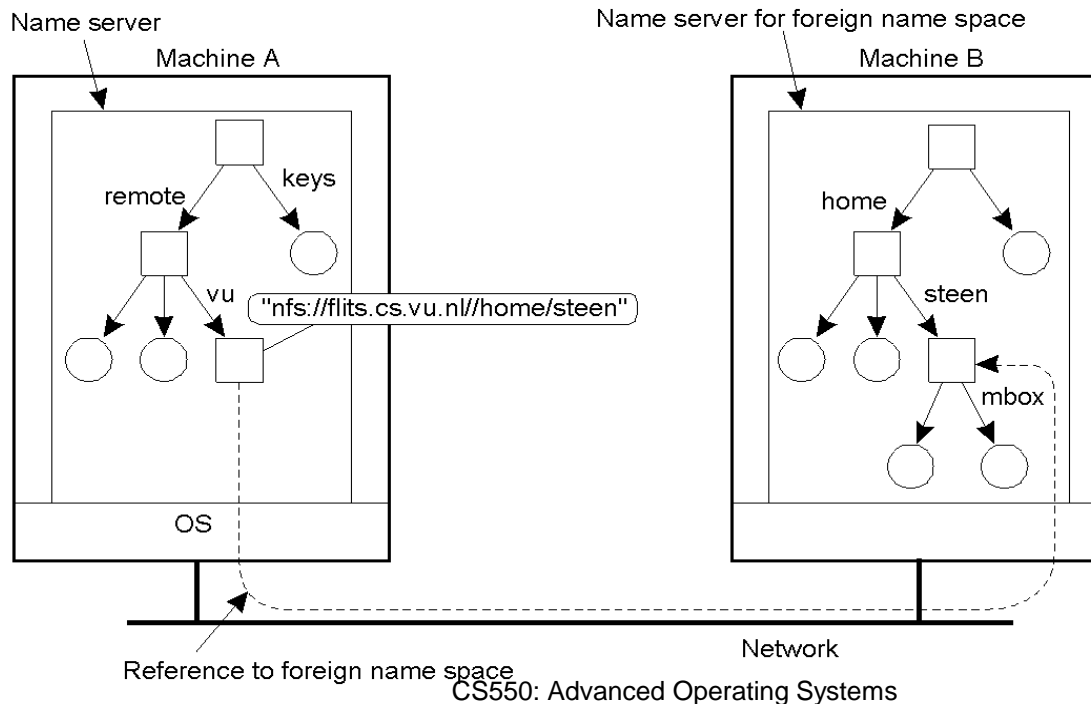
- Name resolution: the process of looking up a name



- Ways to merge different names spaces
 - Mounting
 - DEC's Global Name Service (GNS)

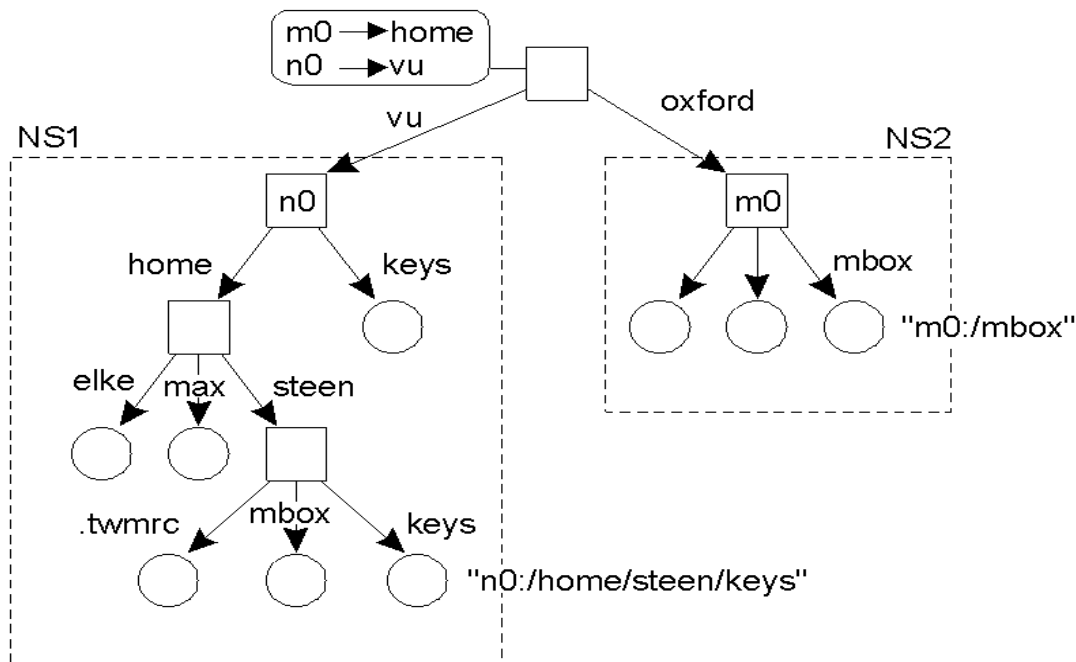
Name Resolution: Mounting

- NFS mount protocol: map a remote node onto local DAG
 - Remote files are accessed using local names! (*location independence*)
 - OS maintains a mount table with the mappings



Name Resolution: GNS

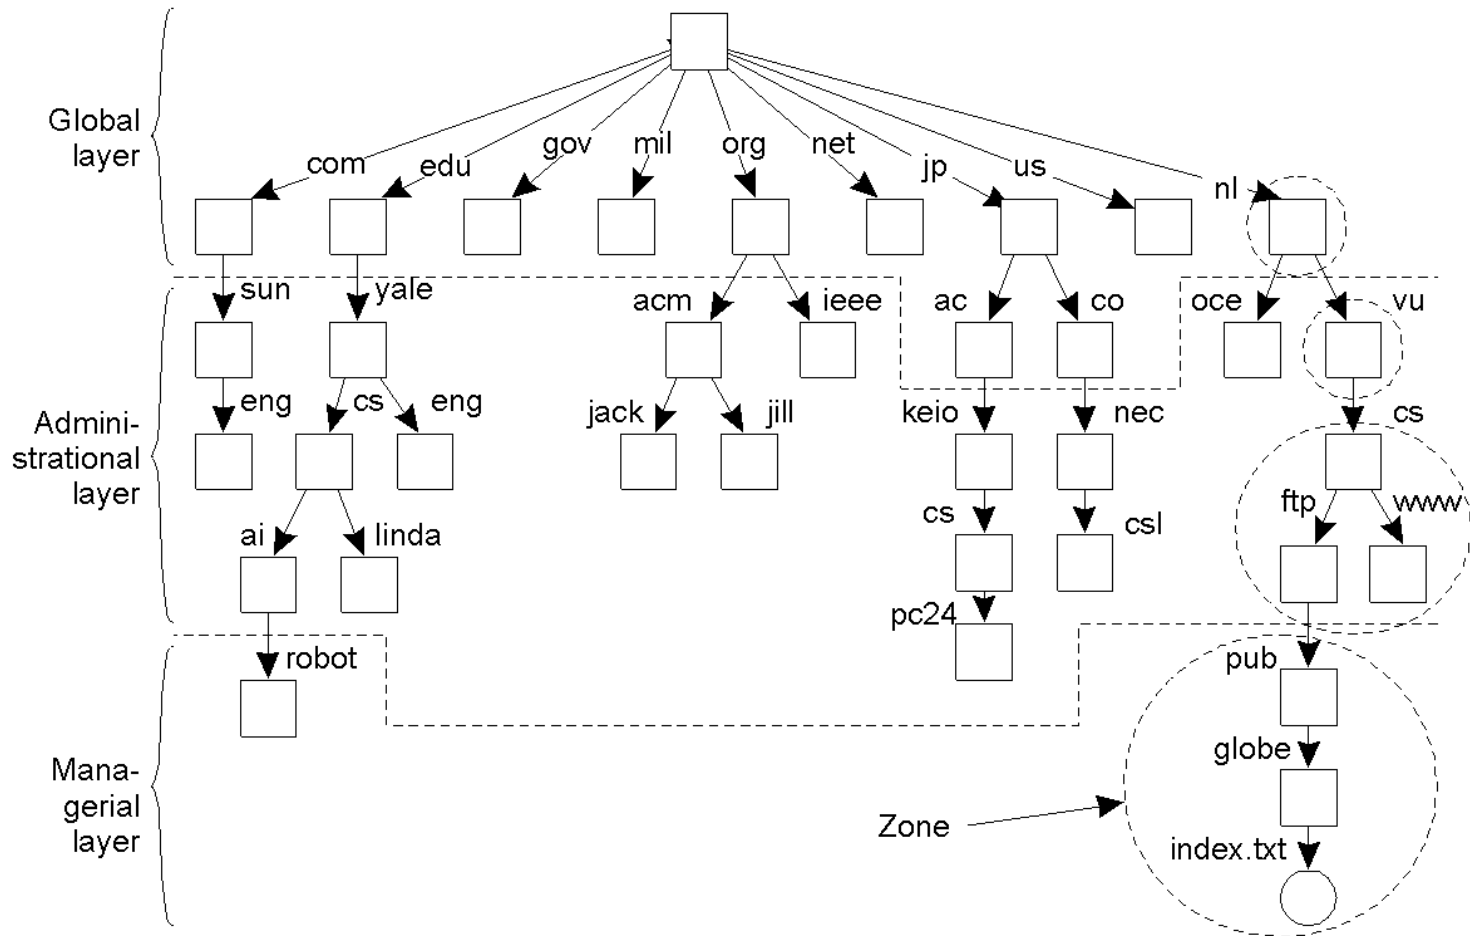
- Another way to merge different name spaces, DEC's Global Name Service (GNS)
 - To add a new root node and to make the existing root nodes its children



Implementing Name Spaces

- Naming service: a service that allows users and processes to add, remove and lookup names
- Name spaces for large-scale widely distributed systems are typically organized hierarchically
- Name space is distributed and has three logical layers
 - Global layer : formed by highest-level nodes, rarely changed
 - Administrational layer: directory nodes within a single organization
 - Managerial layer: nodes may change frequently

Name Space Distribution Example



Example DNS name space

Name Space Distribution

Item	Global	Administrational	Managerial
Geographical scale of network	Worldwide	Organization	Department
Total number of nodes	Few	Many	Vast numbers
Responsiveness to lookups	Seconds	Milliseconds	Immediate
Update propagation	Lazy	Immediate	Immediate
Number of replicas	Many	None or few	None
Is client-side caching applied?	Yes	Yes	Sometimes

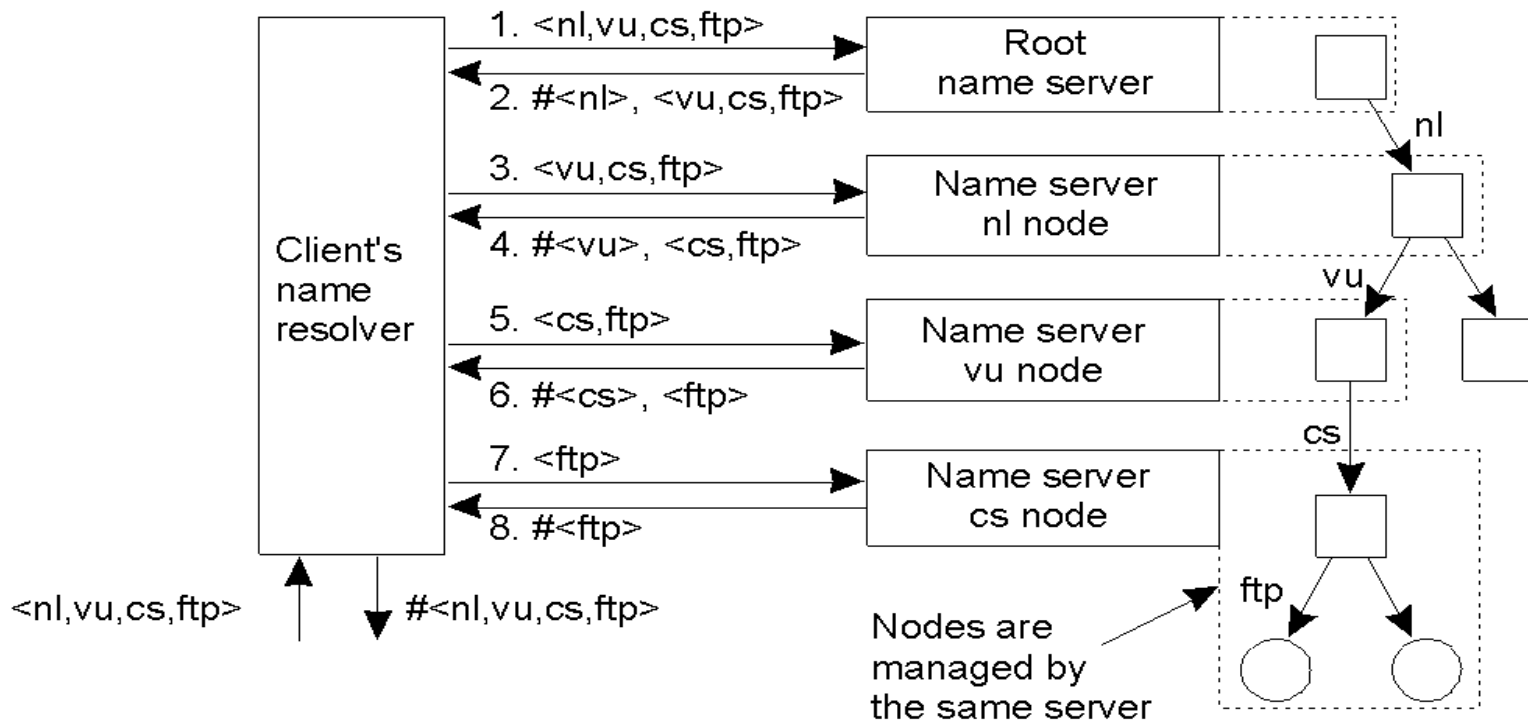
- A comparison between name servers for implementing nodes from a large-scale name space partitioned into a global layer, as an administrative layer, and a managerial layer.
- The more stable a layer, the longer are the lookups valid (and can be cached longer)

Name resolution

- Name resolution: The process of looking up a name
- Two techniques: iterative and recursive

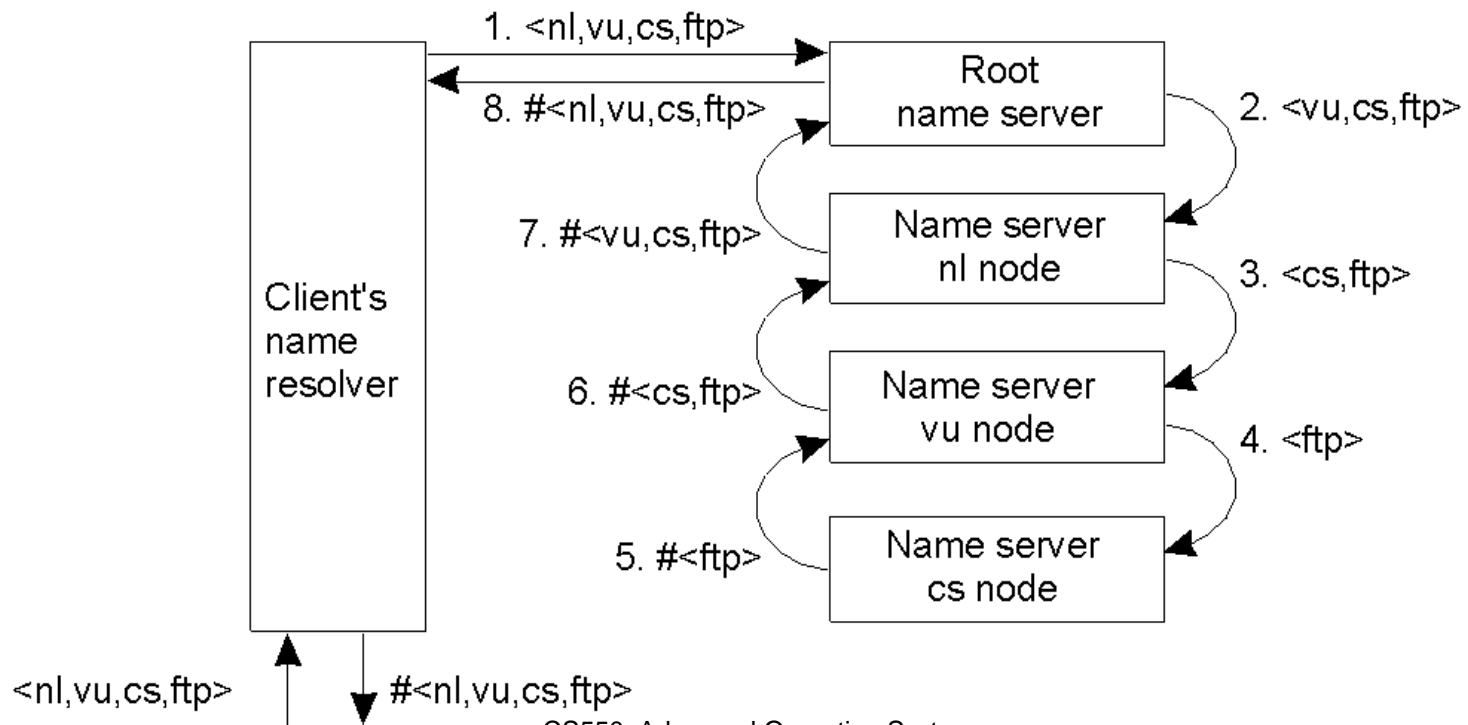
Implementing Name Resolution

- Iterative name resolution
 - Start with the root
 - Each layer resolves as much as it can and returns address of next name server



Implementing Name Resolution

- Recursive name resolution
 - Start at the root
 - Each layer resolves as much as it can and hands the rest to the next layer



Which is better?

- Disadvantage of recursive name resolution
 - Higher performance demand on each name server
- Advantages of recursive name resolution
 - Caching possible at name servers, why?
 - Communication cost reduced

Case Studies

- So far, we have discussed:
 - Definition of name
 - Name space
 - Name resolution
- Now, let's look at some naming service systems (aka directory service systems)
 - DNS
 - LDAP

Domain Name Space

- One of the largest distributed naming services
- Primarily used for looking up host addresses and mail servers
- Comparable to a telephone book
- Its name space is hierarchically organized as a rooted tree
- DNS also maintains an inverse mapping of IP addresses to host names

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

DNS Implementation

- DNS name space is divided into a global and an administrative layer
- Each zone is implemented by a name server, which is always replicated
- Example:

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	130.37.24.6
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

Name Servers

Mail Servers

Attribute-based Naming

- Flat and structured names generally provide a unique and location-independent way of referring to entities.
- How to identify an entity merely by a description?
- Describe an entity in terms of (*attribute, value*) pairs, generally referred to as **attribute-based** naming.
- Structured Naming -> naming systems
- Attribute-based Naming -> directory service

X.500 Directory Service

- OSI Standard
- Directory service: special kind of naming service where:
 - Clients can lookup entities based on attributes instead of full name
 - Real-world example: Yellow pages: look for a plumber

LDAP

- Lightweight Directory Access Protocol (LDAP)
 - LDAP: Simplified version of X.500
 - Widely used for Internet services
 - Application-level protocol, uses TCP
 - Lookups and updates can use strings instead of OSI encoding
 - Use master servers and replicas servers for performance improvements
 - Example LDAP implementations:
 - Active Directory (Windows 2000)
 - Novell Directory services
 - iPlanet directory services (Netscape)
 - Typical uses: user profiles, access privileges, network resources

LDAP

- LDAP originated at the University of Michigan.
- LDAP can be used as a front-end to X.500 or stand-alone.
- LDAP is now available commercially from a number of sources (including Netscape)
- Refs:
 - Netscape LDAP server docs
 - U. of Michigan LDAP docs
 - www.openldap.org docs
 - RFCs: 1777, 1773, 1823, ...

LDAP Name Space


- Portions of namespace are contained on each LDAP server
 - These portions are sometimes called partitions
 - Some servers cooperate to provide the same portion
- Each entry on an LDAP server is uniquely identified by its *distinguished name* (*DN* for short)
 - Each *DN* is a sequence of attribute=value pairs
- In X.500, the theory is that each *DN* is globally unique; in LDAP, the theory is that the combination of the LDAP host plus the *DN* are globally unique

Attributes

- The attributes can be anything, but there is a standard hierarchy used (for a *global* LDAP namespace):

C	<i>country name</i>
O	<i>organization name</i>
OU	<i>organizational unit</i>
CN	<i>common name</i>
L	<i>locality name</i>
ST	<i>state or province</i>
STREET	<i>street address</i>

Sample DN & Hierarchy



dn: CN=Ioan Raicu, OU=CS, O=IIT, C=US

LDAP Services

- LDAP ObjectClass
 - LDAP uses the concept of object classes to define which attributes are allowed for objects of any given type
- LDAP replication
 - LDAP servers can be set to replicate some or all of their data, on a push or a pull basis, using simple authentication or certificate-based authentication
- Security and access control
 - LDAP provides for a complex level of access control instances, or ACIs

Summary

- Names
- Name spaces
- Name resolution
- Example: The Domain Name System
- Example: X.500, LDAP
- Readings:
 - Chpt 5 of AST

Questions

