

Chapter 5 Naming

- Properties of a true 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

Names, Identifiers, And Addresses

Tanenbaum & Van Steen, Distributed Systems: Principles and Paradigms, 2e. (c) 2007

Forwarding Pointers (1)

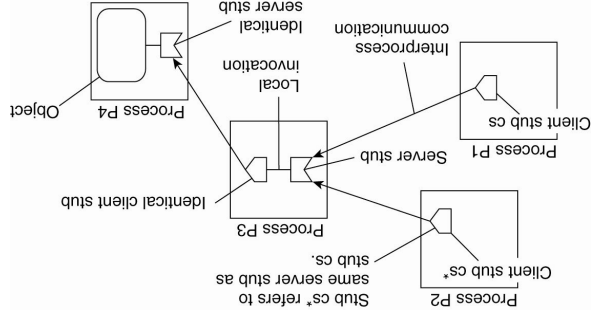


Figure 5-1. The principle of forwarding pointers using (client stub, server stub) pairs.

Tanenbaum & Van Steen, Distributed Systems: Principles and Paradigms, 2e. (c) 2007

Forwarding Pointers (3)

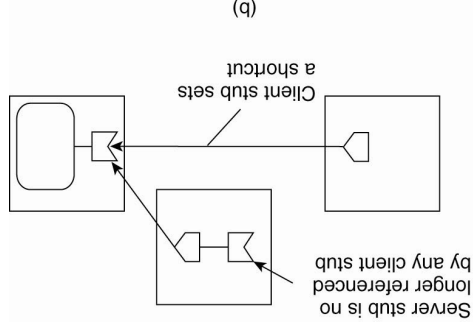


Figure 5-2. Redirecting a forwarding pointer by storing a shortcut in a client stub.

Tanenbaum & Van Steen, Distributed Systems: Principles and Paradigms, 2e. (c) 2007

Forwarding Pointers (2)




Figure 5-2. Redirecting a forwarding pointer by storing a shortcut in a client stub.

Tanenbaum & Van Steen, Distributed Systems: Principles and Paradigms, 2e. (c) 2007

Home-Based Approaches

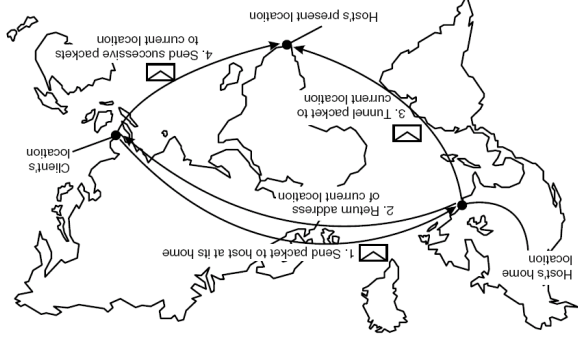


Figure 5-3. The principle of Mobile IP.

Tanenbaum & Van Steen, Distributed Systems: Principles and Paradigms, 2e. (c) 2007

Distributed Hash Tables

General Mechanism

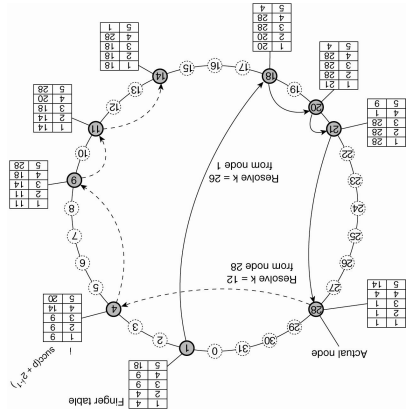


Figure 5-4. Resolving key 26 from node 1 and key 12 from node 28 in a Chord system.

Tanenbaum & Van Steen, Distributed Systems: Principles and Paradoms, 2e. (c) 2007

Hierarchical Approaches (1)




Figure 5-5. Hierarchical organization of a location service into domains, each having an associated directory node.

Tanenbaum & Van Steen, Distributed Systems: Principles and Paradoms, 2e. (c) 2007

Hierarchical Approaches (2)

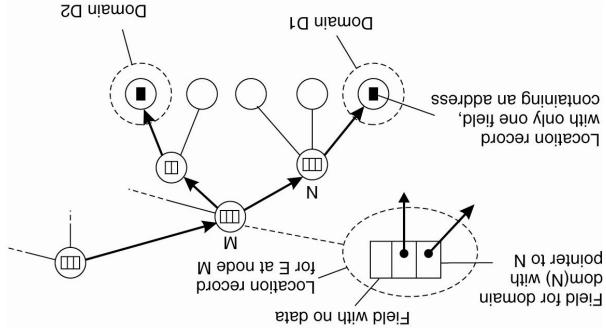


Figure 5-6. An example of storing information of an entity having two addresses in different leaf domains.

Tanenbaum & Van Steen, Distributed Systems: Principles and Paradoms, 2e. (c) 2007

Hierarchical Approaches (3)

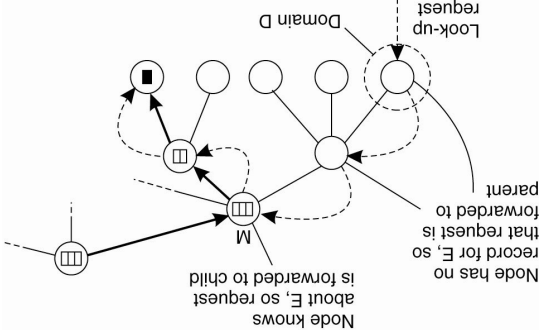


Figure 5-7. Looking up a location in a hierarchically organized location service.

Tanenbaum & Van Steen, Distributed Systems: Principles and Paradoms, 2e. (c) 2007

Hierarchical Approaches (4)




Figure 5-8. (a) An insert request is forwarded to the first node that knows about entity E.

Tanenbaum & Van Steen, Distributed Systems: Principles and Paradoms, 2e. (c) 2007

Hierarchical Approaches (5)





Figure 5-8. (b) A chain of forwarding pointers to the leaf node is created.

Tanenbaum & Van Steen, Distributed Systems: Principles and Paradoms, 2e. (c) 2007

Figure 5-10. The general organization of the UNIX file system implementation on a logical disk of contiguous disk blocks.



Name Spaces (2)

- Information required to mount a foreign name space in a distributed system
- The name of an access protocol.
- The name of the server.
- The name of the mounting point in the foreign name space.

Tanenbaum & Van Steen, Distributed Systems: Principles and Paradoms, 2e. (c) 2007

Name Space Distribution (1)





Figure 5-13. An example partitioning of the DNS name space, including Internet-accessible files, into three layers.

Tanenbaum & Van Steen, Distributed Systems: Principles and Paradoms, 2e. (c) 2007

Figure 5-9. A general naming graph with a single root node.



Name Spaces (1)

Tanenbaum & Van Steen, Distributed Systems: Principles and Paradoms, 2e. (c) 2007

Linking and Mounting (1)

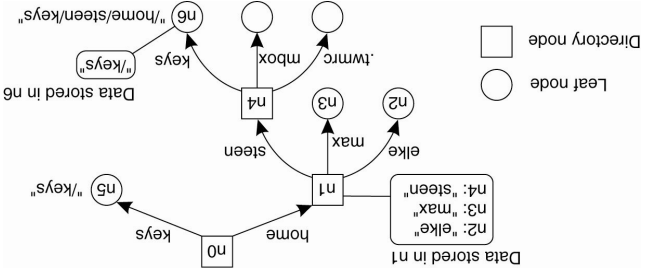


Figure 5-11. The concept of a symbolic link explained in a naming graph.

Tanenbaum & Van Steen, Distributed Systems: Principles and Paradoms, 2e. (c) 2007

Linking and Mounting (3)

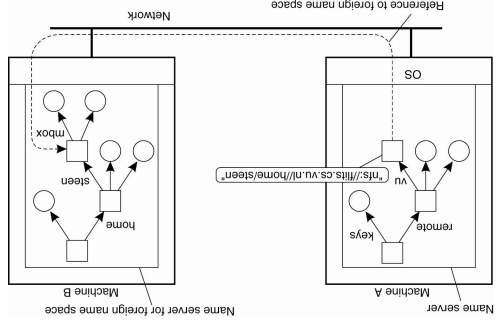


Figure 5-12. Mounting remote name spaces through a specific access protocol.

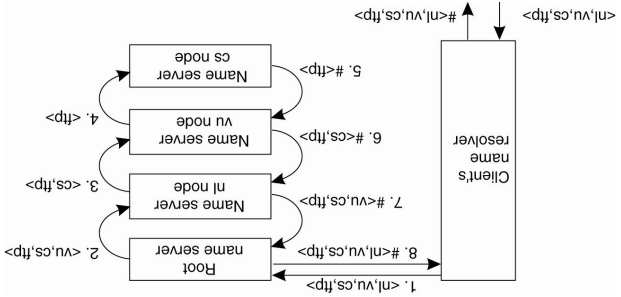
Tanenbaum & Van Steen, Distributed Systems: Principles and Paradoms, 2e. (c) 2007

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

Figure 5-14. A comparison between name servers for implementing nodes from a large scale name space partitioned into a global layer, an administrative layer, and a managerial layer.

Tanenbaum & Van Steen, Distributed Systems: Principles and Paradoms, 2e. (c) 2007


Implementation of Name Resolution (2)



Tanenbaum & Van Steen, Distributed Systems: Principles and Paradoms, 2e. (c) 2007

Figure 5-16. The principle of recursive name resolution.

Example: The Domain Name System



Tanenbaum & Van Steen, Distributed Systems: Principles and Paradoms, 2e. (c) 2007

Figure 5-18. The comparison between recursive and iterative name resolution with respect to communication costs.

Implementation of Name Resolution (1)

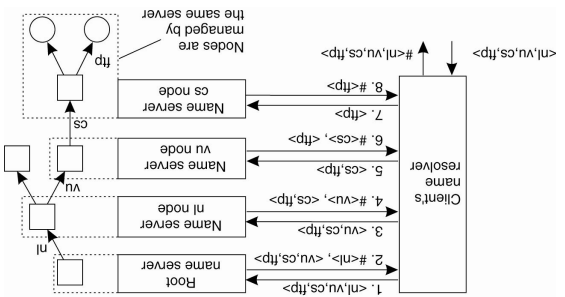


Figure 5-15. The principle of iterative name resolution.

Tanenbaum & Van Steen, Distributed Systems: Principles and Paradoms, 2e. (c) 2007

Implementation of Name Resolution (3)

Server for node	Should resolve	Looks up	Passes to child	Receives and caches	Returns
cs	<ftp>	<ftp>	—	—	<ftp>
vu	<cs,ftp>	<cs,ftp>	<ftp>	<cs,ftp>	<cs,ftp>
nl	<vu,cs,ftp>	<vu,cs,ftp>	<cs,ftp>	<vu,cs,ftp>	<vu,cs,ftp>
root	<n1,vu,cs,ftp>	<n1,vu,cs,ftp>	<vu,cs,ftp>	<n1,vu,cs,ftp>	<n1,vu,cs,ftp>

Figure 5-17. Recursive name resolution of <n1, vu, cs, ftp>. Name servers cache intermediate results for subsequent lookups.

Tanenbaum & Van Steen, Distributed Systems: Principles and Paradoms, 2e. (c) 2007


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

Figure 5-19. The most important types of resource records forming the contents of nodes in the DNS name space.


Tanenbaum & Van Steen, Distributed Systems: Principles and Paradoms, 2e. (c) 2007

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



Hierarchical Implementations: LDAP (2)

Figure 5-21. Part of the description for the `vu.nl` domain which contains the `cs.vu.nl` domain.



DNS Implementation (3)

Figure 5-20. An excerpt from the DNS database for the zone `cs.vu.nl`.

Name	Record type	Record value
<code>cs.vu.nl.</code>	SOA	<code>star.cs.vu.nl. hostmaster.cs.vu.nl. 2005092900 7200 3600 2419200 3600</code>
<code>cs.vu.nl.</code>	TXT	<code>"Vrije Universiteit - Math. & Comp. Sc."</code>
<code>cs.vu.nl.</code>	MX	<code>1 mail.few.vu.nl.</code>
<code>cs.vu.nl.</code>	NS	<code>ns.vu.nl.</code>
<code>cs.vu.nl.</code>	NS	<code>top.cs.vu.nl.</code>
<code>cs.vu.nl.</code>	NS	<code>solo.cs.vu.nl.</code>
<code>cs.vu.nl.</code>	NS	<code>star.cs.vu.nl.</code>
<code>star.cs.vu.nl.</code>	A	<code>130.37.24.6</code>
<code>star.cs.vu.nl.</code>	A	<code>192.31.231.42</code>
<code>star.cs.vu.nl.</code>	A	<code>130.37.20.10</code>
<code>zephyr.cs.vu.nl.</code>	A	<code>130.37.20.10</code>
<code>zephyr.cs.vu.nl.</code>	MX	<code>1 zephyr.cs.vu.nl.</code>
<code>zephyr.cs.vu.nl.</code>	MX	<code>2 tomado.cs.vu.nl.</code>
<code>zephyr.cs.vu.nl.</code>	HINFO	<code>"Sun" "Unix"</code>
<code>star.cs.vu.nl.</code>	HINFO	<code>"Sun" "Unix"</code>
<code>star.cs.vu.nl.</code>	MX	<code>666 zephyr.cs.vu.nl.</code>
<code>star.cs.vu.nl.</code>	MX	<code>1 star.cs.vu.nl.</code>
<code>star.cs.vu.nl.</code>	A	<code>192.31.231.42</code>
<code>star.cs.vu.nl.</code>	A	<code>130.37.24.6</code>
<code>star.cs.vu.nl.</code>	NS	<code>star.cs.vu.nl.</code>

DNS Implementation (1)

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

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

Hierarchical Implementations: LDAP (3)

Figure 5-22. A simple example of an LDAP directory entry using LDAP naming conventions.

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 (1)

Figure 5-20. An excerpt from the DNS database for the zone `cs.vu.nl`.

Record type	Record value
<code>ftp.cs.vu.nl.</code>	<code>soling.cs.vu.nl.</code>
<code>www.cs.vu.nl.</code>	<code>soling.cs.vu.nl.</code>
<code>soling.cs.vu.nl.</code>	<code>A</code>
<code>soling.cs.vu.nl.</code>	<code>130.37.20.20</code>
<code>soling.cs.vu.nl.</code>	<code>MX</code>
<code>soling.cs.vu.nl.</code>	<code>1 soling.cs.vu.nl.</code>
<code>soling.cs.vu.nl.</code>	<code>MX</code>
<code>soling.cs.vu.nl.</code>	<code>666 zephyr.cs.vu.nl.</code>
<code>soling.cs.vu.nl.</code>	<code>HINFO</code>
<code>soling.cs.vu.nl.</code>	<code>"Sun" "Unix"</code>
<code>soling.cs.vu.nl.</code>	<code>HINFO</code>
<code>soling.cs.vu.nl.</code>	<code>192.168.4.3</code>
<code>soling.cs.vu.nl.</code>	<code>"OCÉ" "Proprietary"</code>
<code>soling.cs.vu.nl.</code>	<code>A</code>
<code>soling.cs.vu.nl.</code>	<code>192.168.4.2</code>
<code>soling.cs.vu.nl.</code>	<code>A</code>
<code>soling.cs.vu.nl.</code>	<code>127.0.0.1</code>

DNS Implementation (2)

Mapping to Distributed Hash Tables (1)




Figure 5-24. (a) A general description of a resource. (b) Its representation as an AVTree. Tanenbaum & Van Steen, Distributed Systems: Principles and Paradoms, 2e. (c) 2007

Semantic Overlay Networks

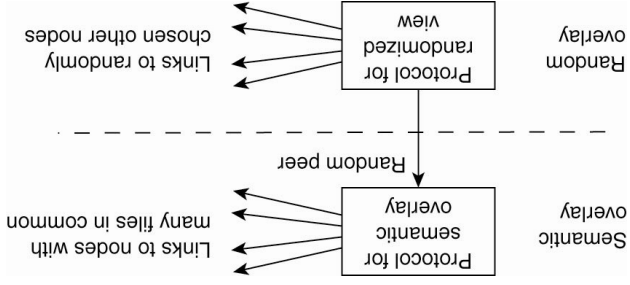


Figure 5-26. Maintaining a semantic overlay through gossiping. Tanenbaum & Van Steen, Distributed Systems: Principles and Paradoms, 2e. (c) 2007

Mapping to Distributed Hash Tables (2)

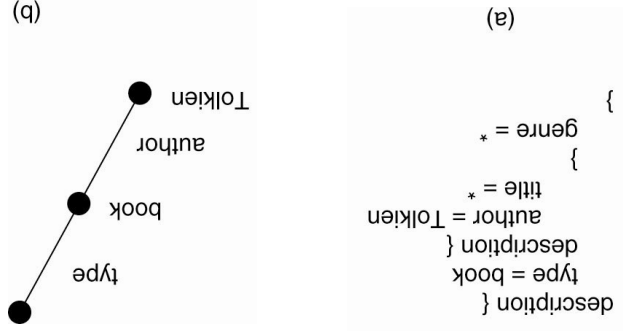


Figure 5-25. (a) The resource description of a query. (b) Its representation as an AVTree. Tanenbaum & Van Steen, Distributed Systems: Principles and Paradoms, 2e. (c) 2007