96 Overheads

Part 1: Networking basic concepts, DNS

More about this course about Internet application protocols can be found at URL:

http://www.dsv.su.se/~jpalme/internet-course/Int-app-prot-kurs.html

Last update: 2005-09-24 19:09

Important information about this course segment

Lectures are not mandatory

But there can be questions in the exam on what is said during the lecture. Reading the written material carefully, and trying to understand or find out the ideas behind the overhead slides in the compendiums, will give you the necessary information to pass the exam.

Lectures may not exactly follow the lecture schedule, and I may skip some things in the end.

Requirements

Exam: Some of the compendiums are allowed during the exam, others are not. This is marked on the front page of the compendium. Even though some exams may be marked for KTH or for SU, any student can go to any exam, provided that you notify in advance.

Work task: Prepare an XML DTD and an XML code using this DTD. Check them against an XML validator. See course description for more info.

Prio: When an item in the course schedule is marked Prio, this means that certain computer rooms are booked for work on the work task. You can go to computer rooms at other times, if there are seats available. No supervisor will help you with the work task during these periods.

Mailing list

Either participate in the First Class conference for this course, or subscribe to the mailing list.

Network layering

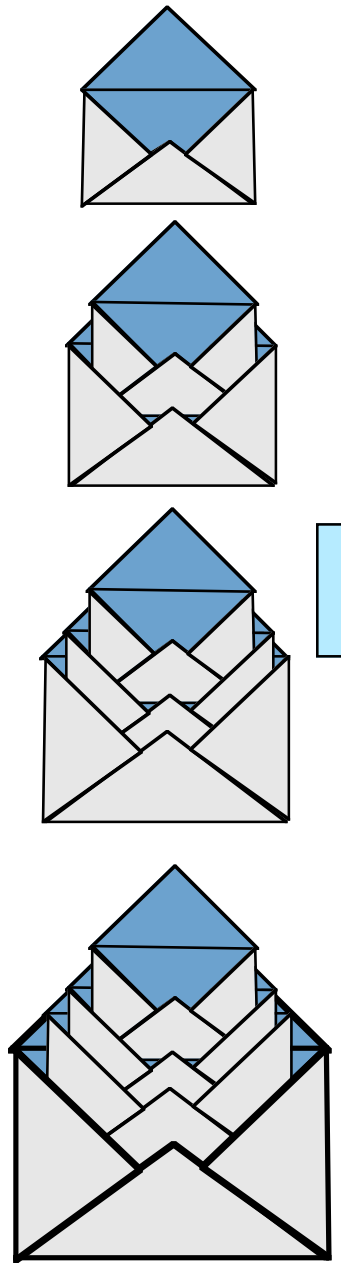

Application screen

Draw polygons

Draw lines and characters

Position cursor, colour pixel

Mail content format

Send, route e-mail

Connections, multiple packets

Packet routing

Send packets

Send bits

Mail content format

Send, route e-mail

Connections, multiple packets

Packet routing

Send packets

Send bits

Understanding layering

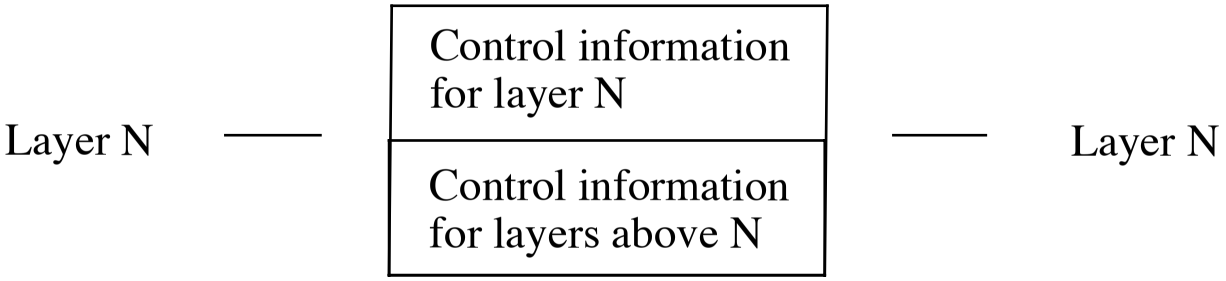

Structured nondistributed application

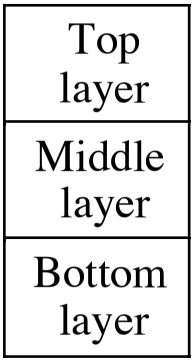

Structured and distributed application

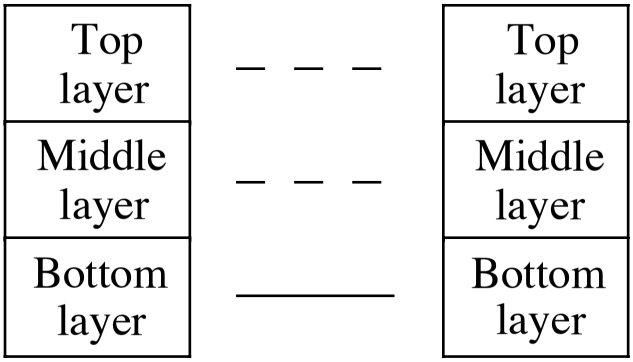

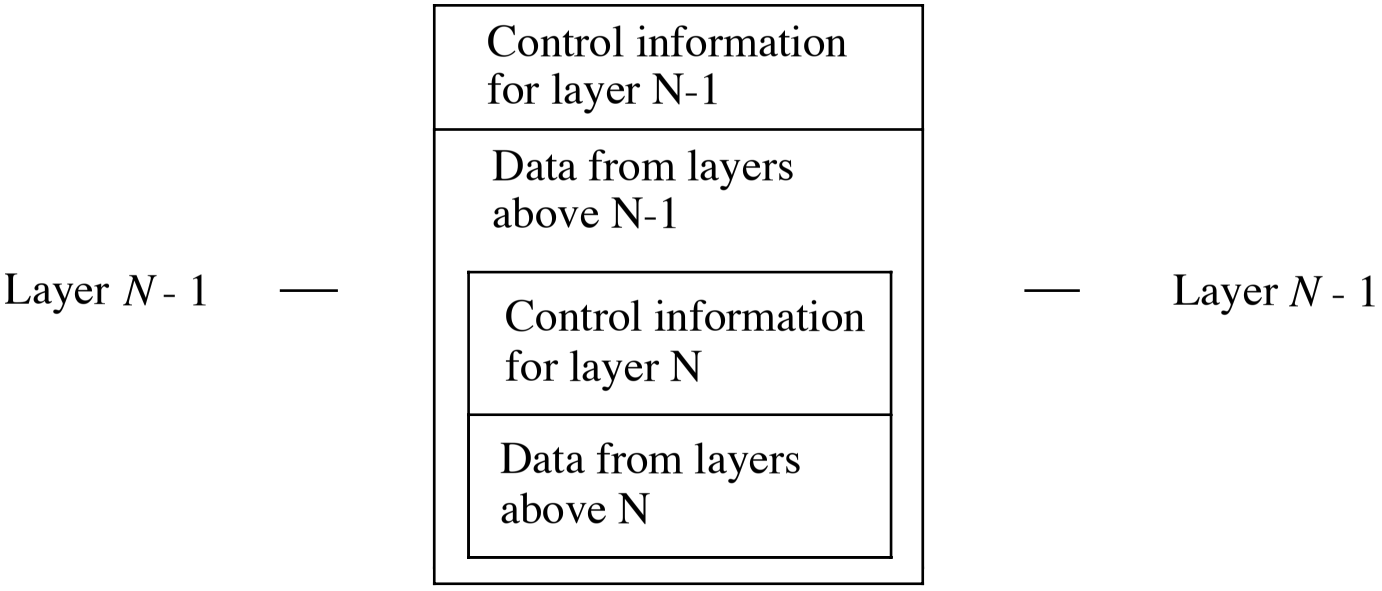

Overview of Internet protocols and services

Protocol name	Main usage	Clients	Servers
DNS	Translating domain names to numerical host addresses	All kinds of clients and name servers	Name servers
HTTP (and HTML)	Downloading web pages in the WWW. Can also be used to send in filled in forms and to send in files. Also used for many specialized protocols based on HTTP.	Web browsers	HTTP servers
SMTP (and RFC822 and MIME)	Sending and forwarding of e-mail to and between MTAs (Message Transfer Agents)	Mail clients and SMTP servers	SMTP servers
POP and IMAP	Downloading of e-mail to the mail clients of their recipients	Mail clients	POP or IMAP servers
NNTP	Downloading and forwarding of Usenet News articles.	News clients and news servers	News servers
FTP	Anonymous downloading of files, non-anonymous transfer of files between logged in directories.	FTP clients, Web browsers	FTP servers
Gopher	An old, nowadays not much used protocols, which can be seen as a limited subset of HTTP.	Web browsers, Gopher clients	Gopher servers
PICS	”Protection” of children from material on the net regarded as unsuitable for them.	All kinds of clients	PICS servers
LDAP	Searching in directories.	LDAP clients, often built into e-mail clients.	LDAP servers

Layering example

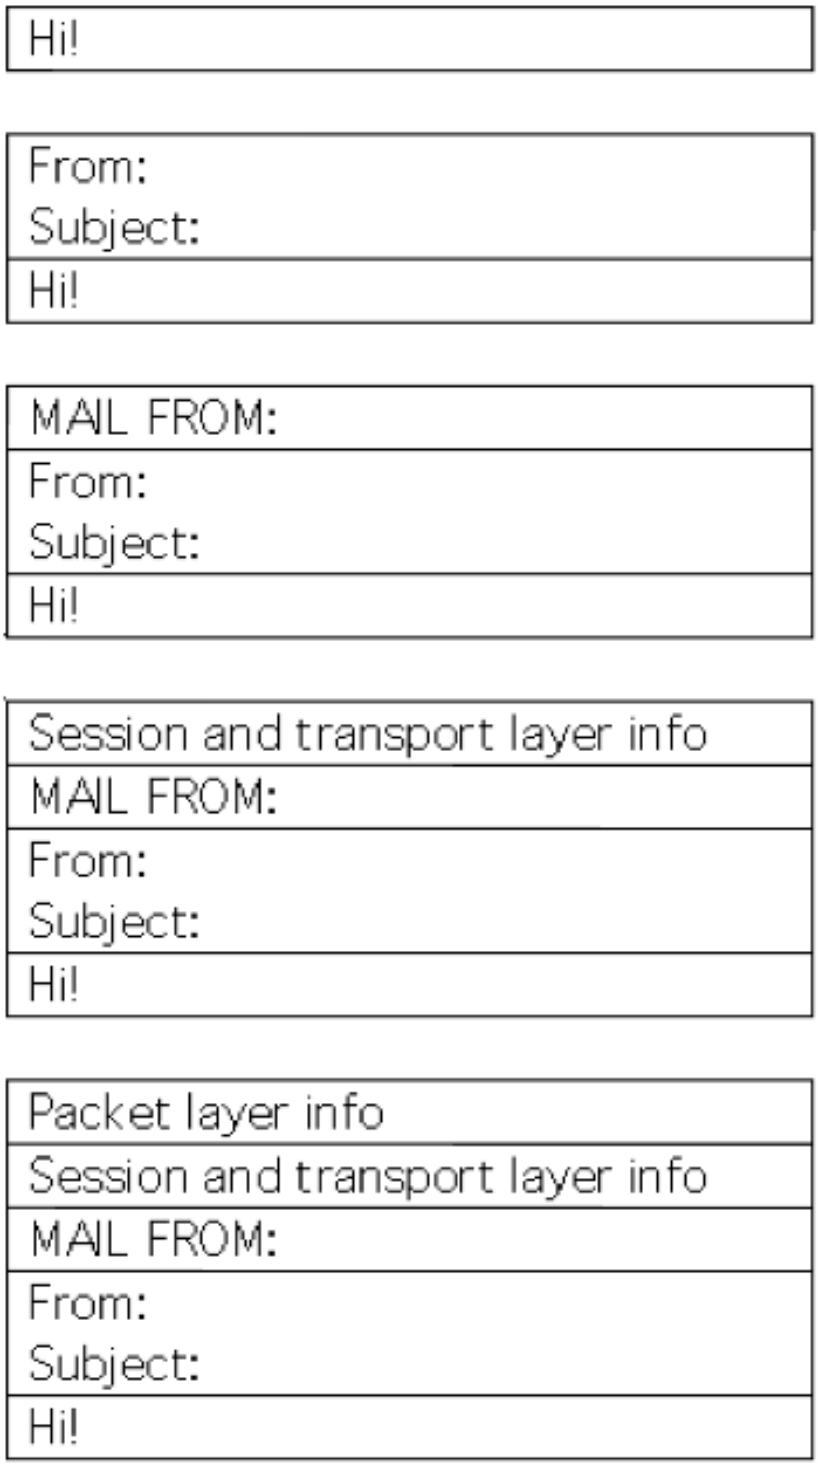

Computers, applications, ports, packets

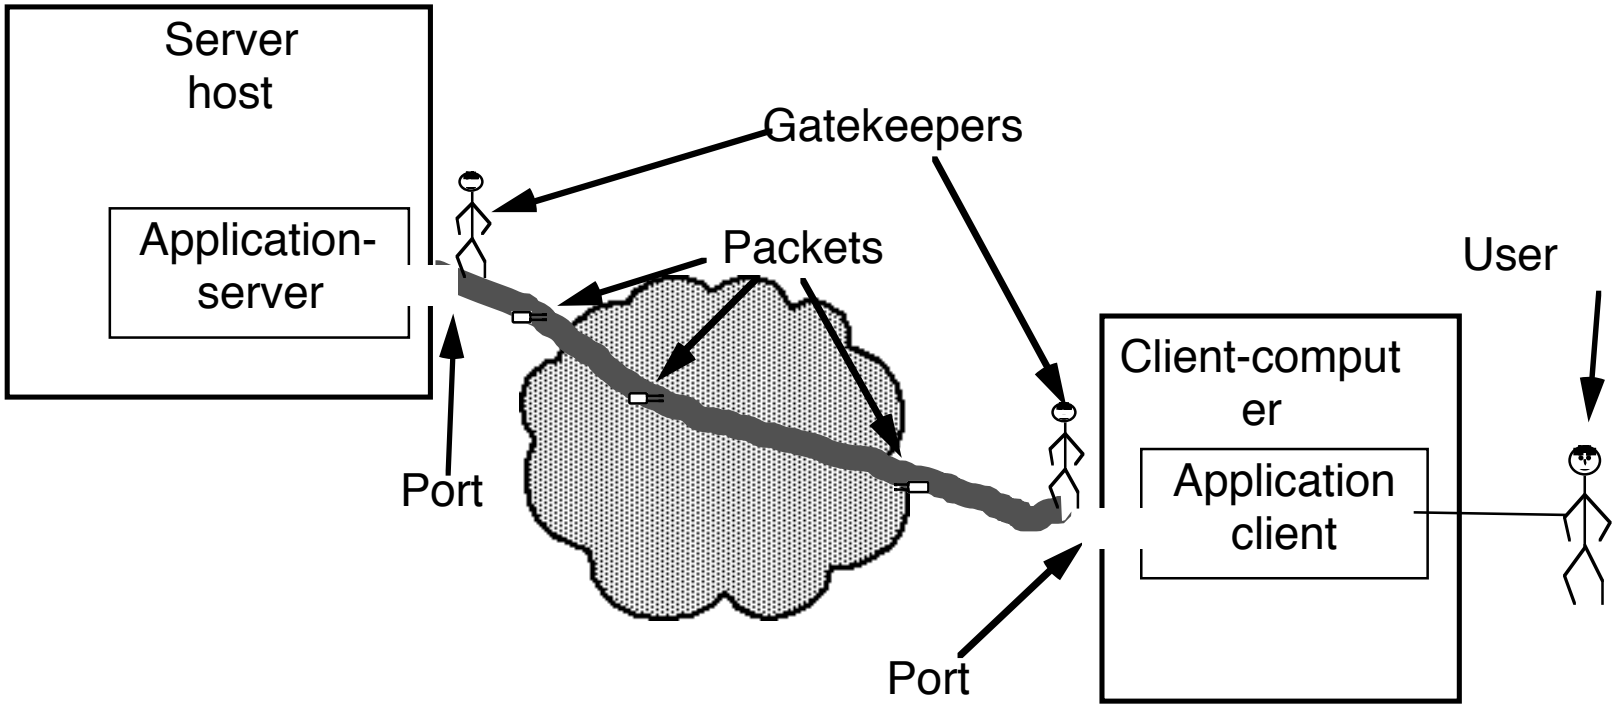

One host can have many different ports fpr different applications.
Exampels of ports: E-mail, file transfer, World Wide Web.
All communication to one particular port uses one particular language.

Registered port numbers

Port numbers are registered with IANA (Internet Assigned Numbers Authority)

The port numbers are divided into three ranges: the Well Known Ports, the Registered Ports, and the Dynamic and/or Private Ports.
The Well Known Ports are those from 0 through 1023.
The Registered Ports are those from 1024 through 49151.
The Dynamic and/or Private Ports are those from 49152 through 65535.

```
ftp-data     20  File Transfer [Default Data]
ftp          21  File Transfer [Control]
telnet       23  Telnet
smtp         25  Simple Mail Transfer
nameserver   42  Host Name Server
nicname      43  Who Is
domain       53  Domain Name Server
whois++      63  whois++
gopher       70  Gopher
finger       79  Finger
http         80  World Wide Web HTTP
www-http     80  World Wide Web HTTP
kerberos     88  Kerberos
hostname    101  NIC Host Name Server
pop2        109  Post Office Protocol - Version 2
pop3        110  Post Office Protocol - Version 3
sunrpc      111  SUN Remote Procedure Call
auth        113  Authentication Service
uucp-path   117  UUCP Path Service
nntp        119  Network News Transfer Protocol
imap2       143  Interim Mail Access Protocol v2
imap3       220  Interactive Mail Access Prot. v3
```

Architectures

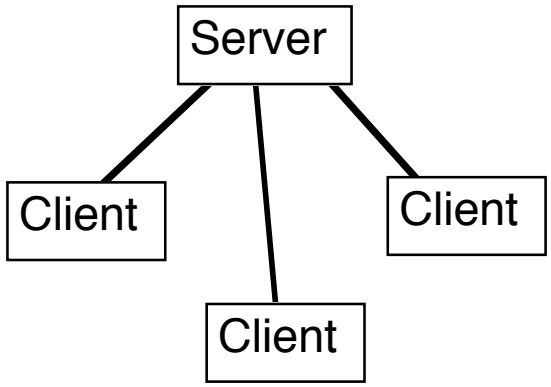

Example: LAN data base

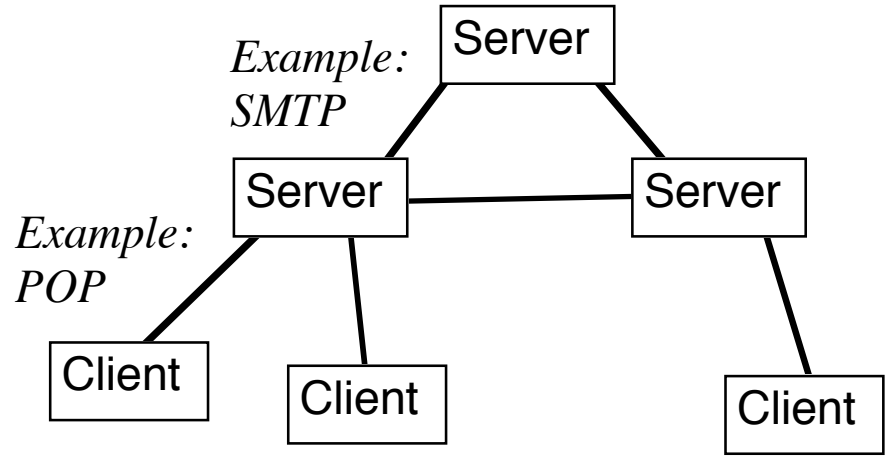

Example: E-mail, Usenet News

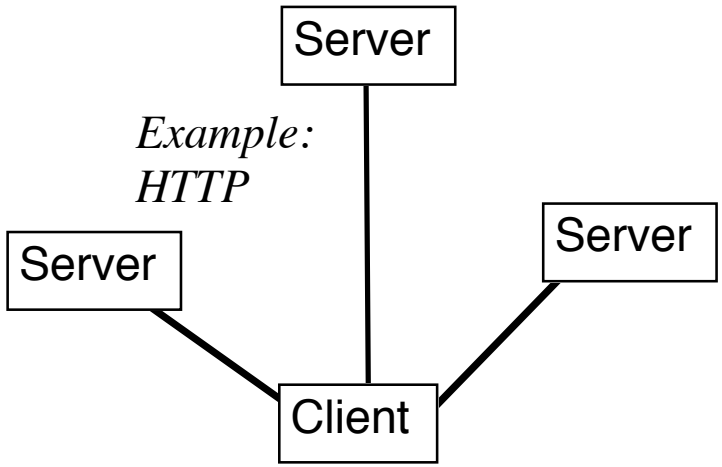

Example: WWW

Symmetric and asymmetric protocols

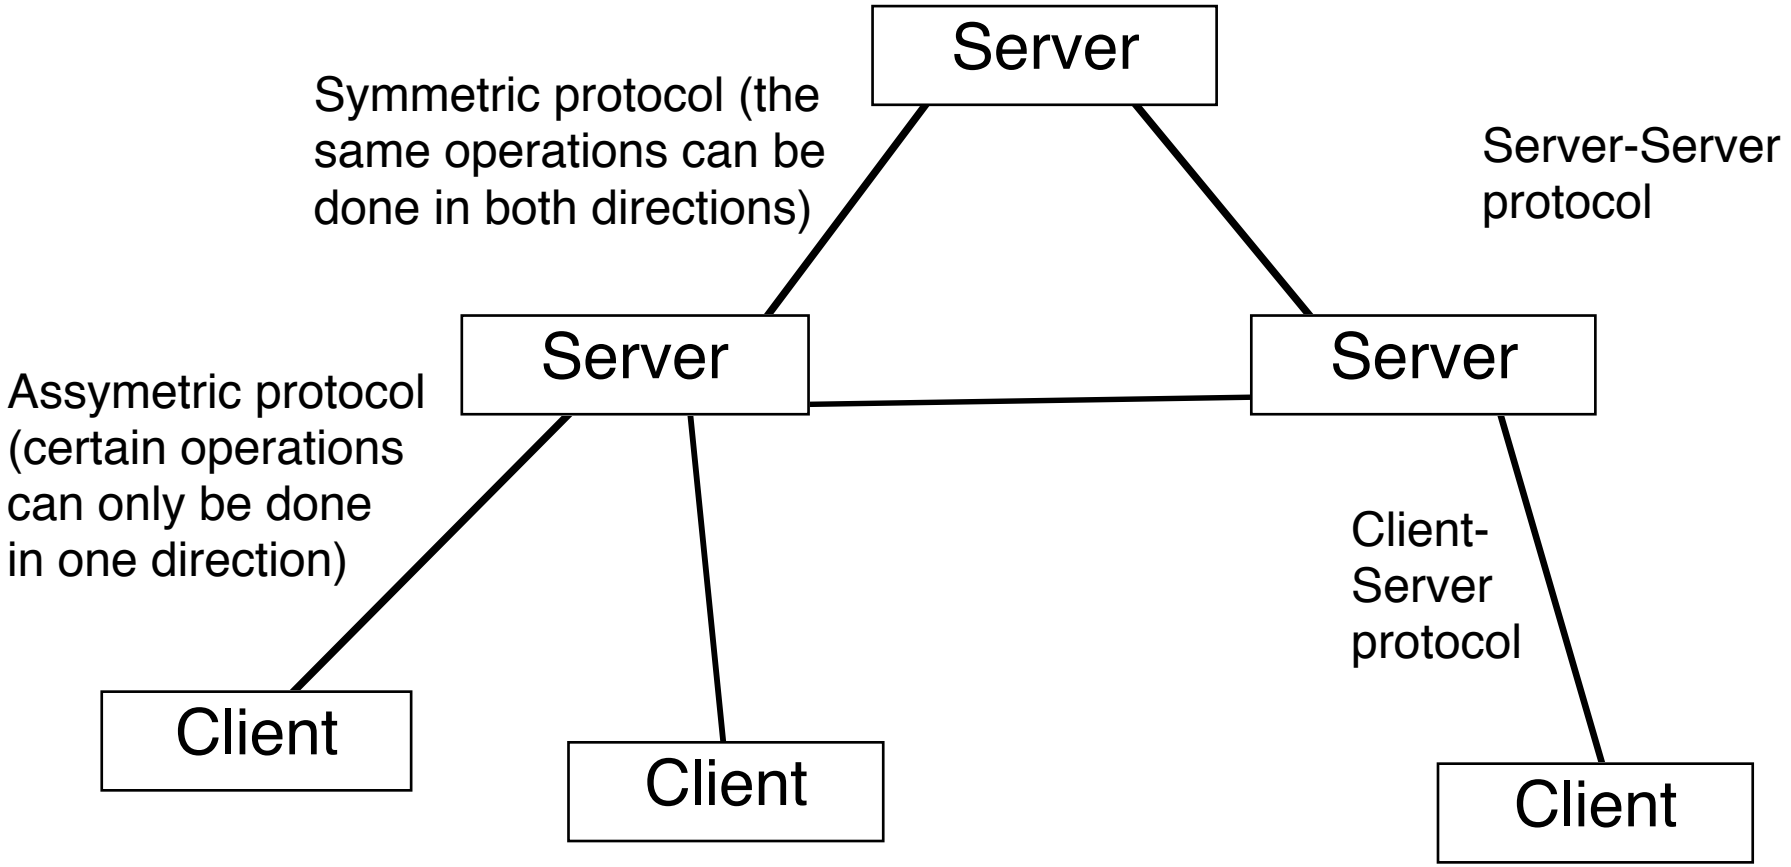

Protocols

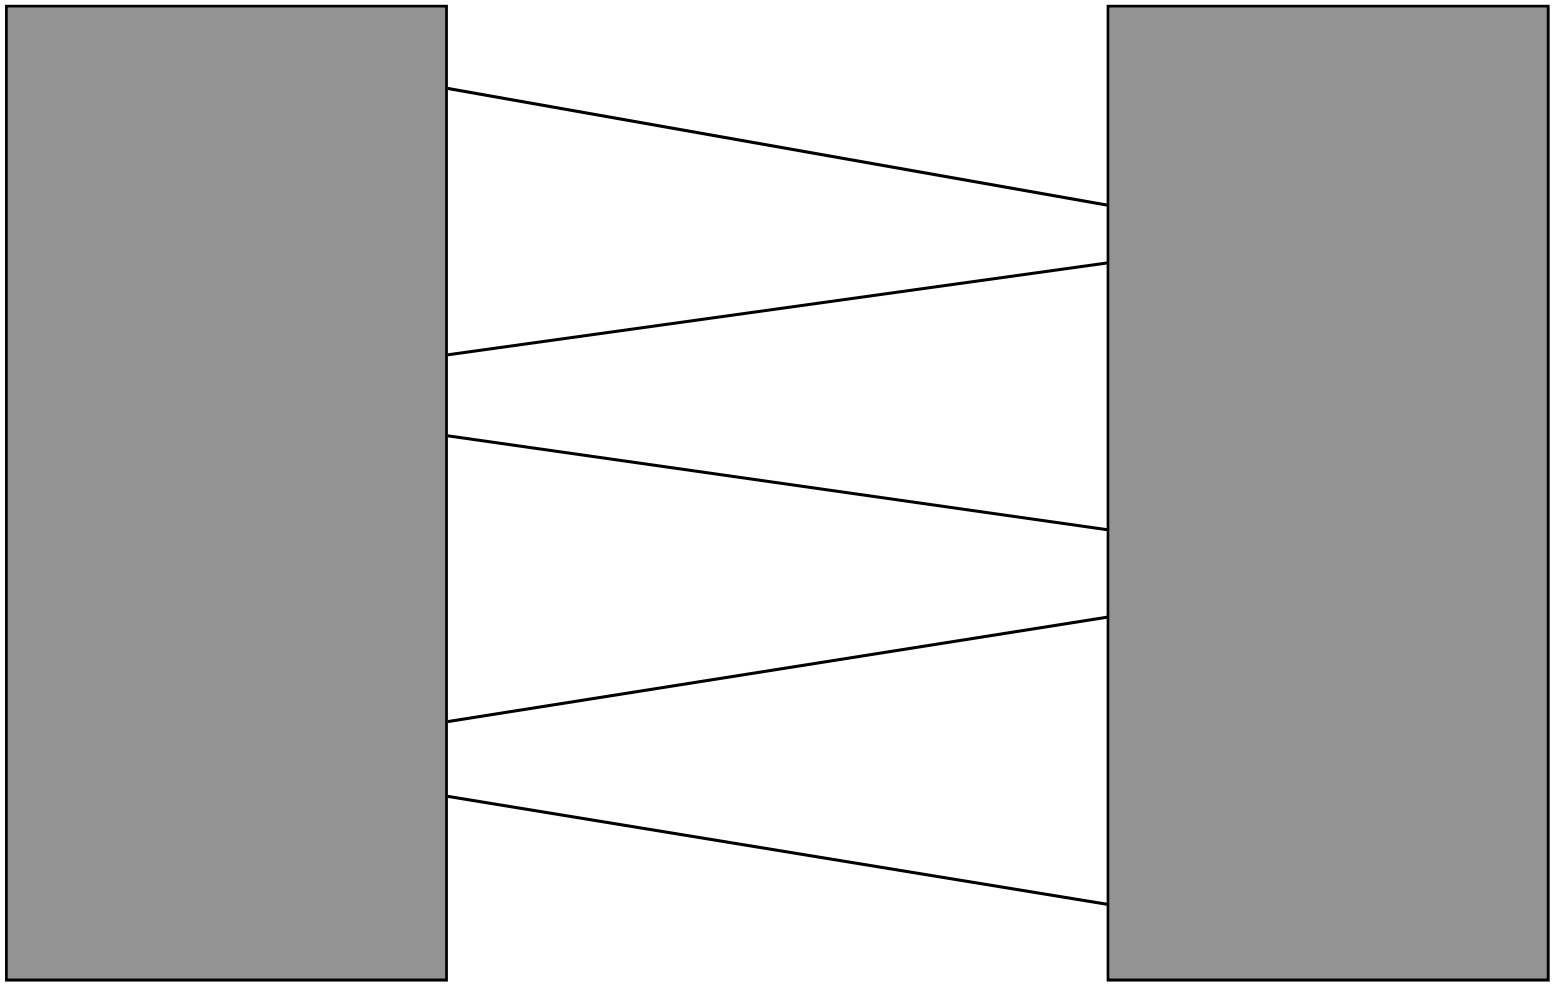

Confirmations, error codes, responses

Transfer of responsibility

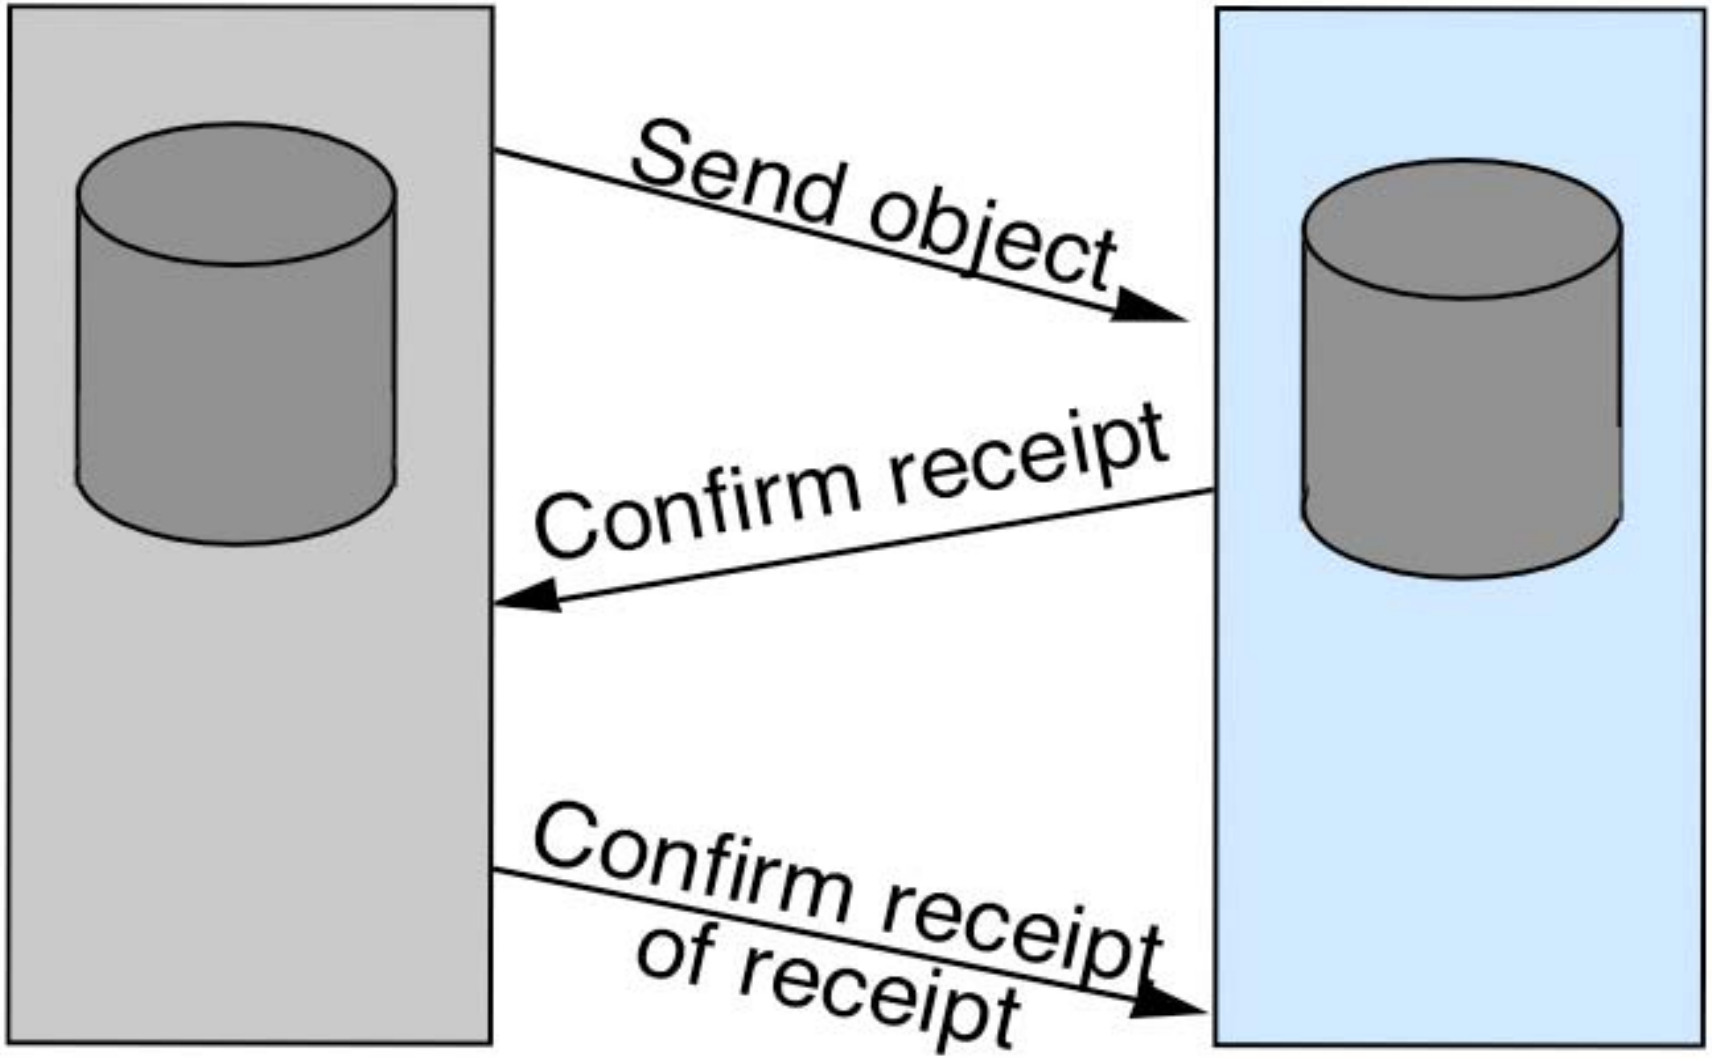

Identification

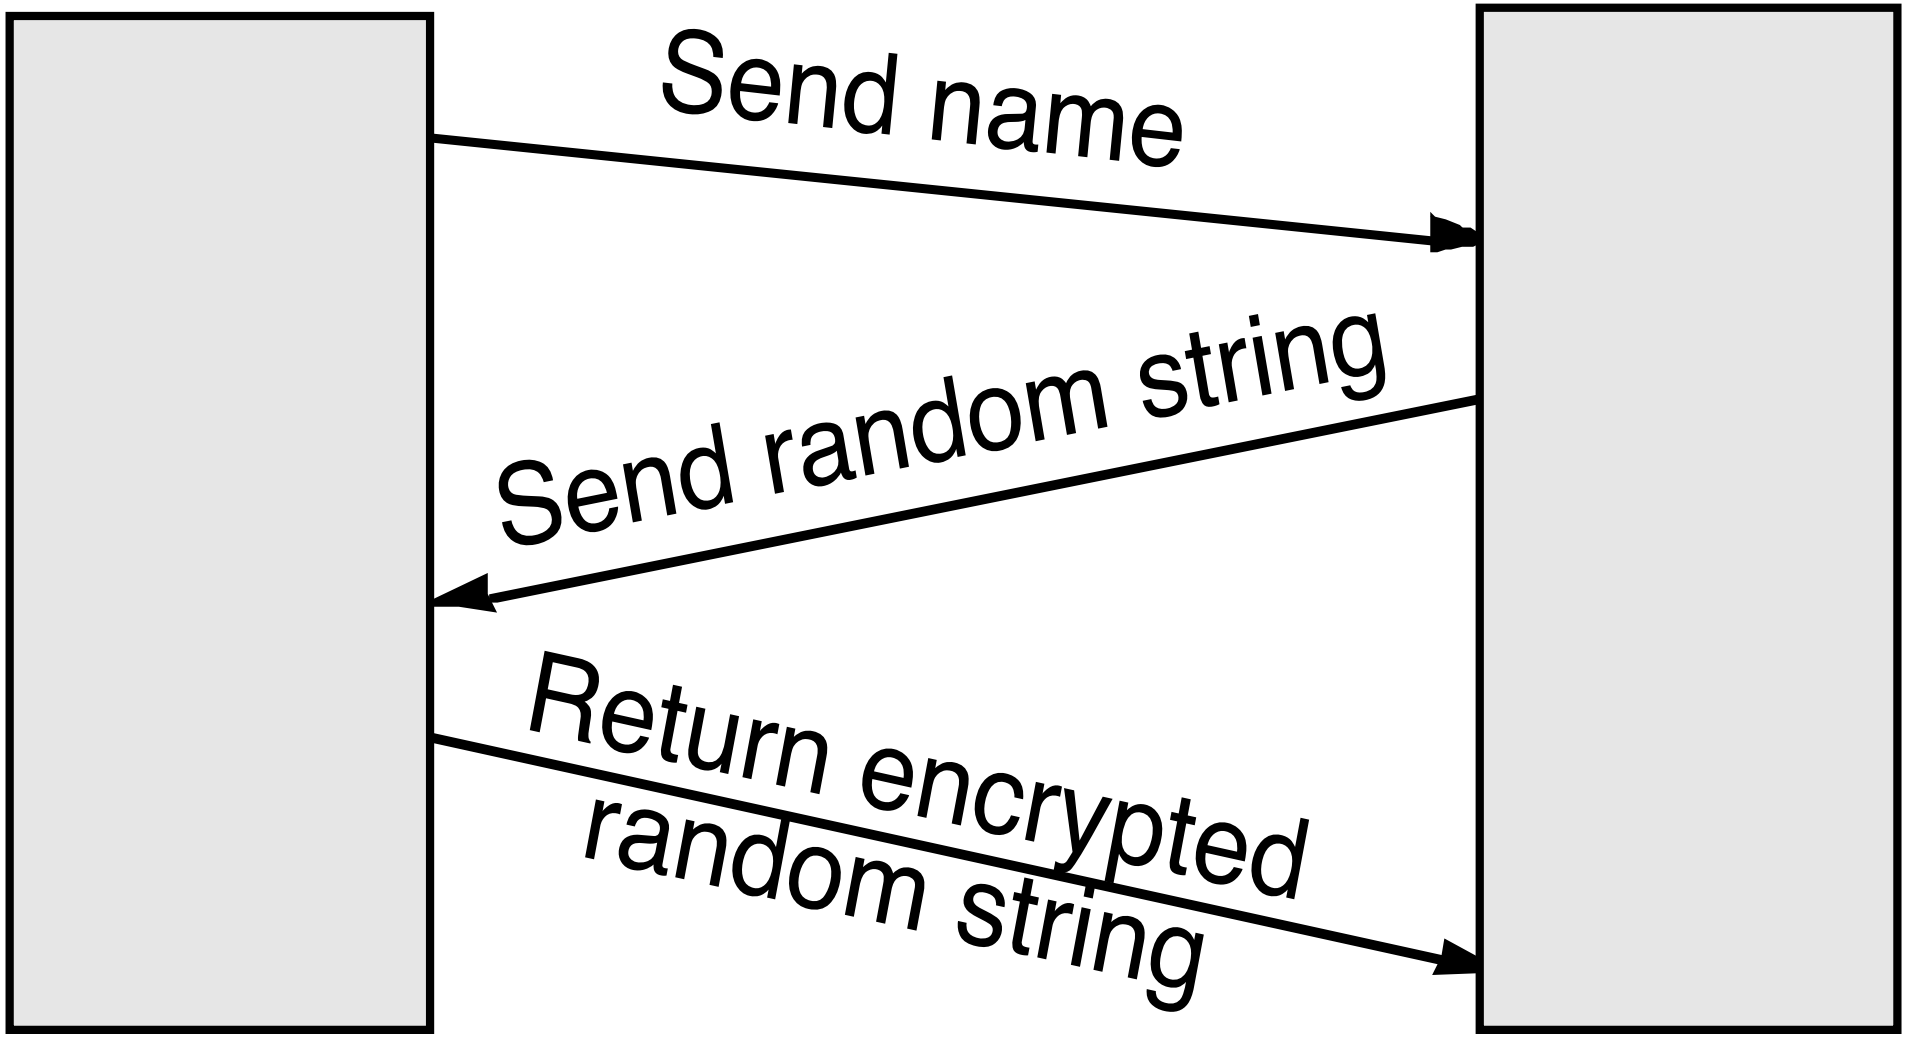

Public/secret key encryption

encrypted text = f_1(original text)

original text = f_2(encrypted text)

Can f_2 be derived from f_1?

Pros and cons of public key encryption

+ Solves partly key transportation problem

– More CPU-time consuming

Authentication, authorization

- To verify the sender of a message
- Payments, agreements
- UA-UA or MTA-MTA

Authentication methods

(a) Passwords

(b) Specially designed networks

(c) Public key cryptography

Three levels of protection of message transmission:

(1) The agents identify each other using noninvertible forms of ordinary passwords. This is called *weak authentication*.

(2) The agents identify each other using public key encryption algorithms. This is called *strong authentication*.

(3) Strong authentication is combined with encryption of all messages during the whole transmission.

Digital Signatures and Digital Seals

Methods: Secret key encryption of signature or checksum, which anyone can decrypt with public key

- Number of interactions
- Need of a neutral third party
- Bilateral or open to groups

Certificate Authorities

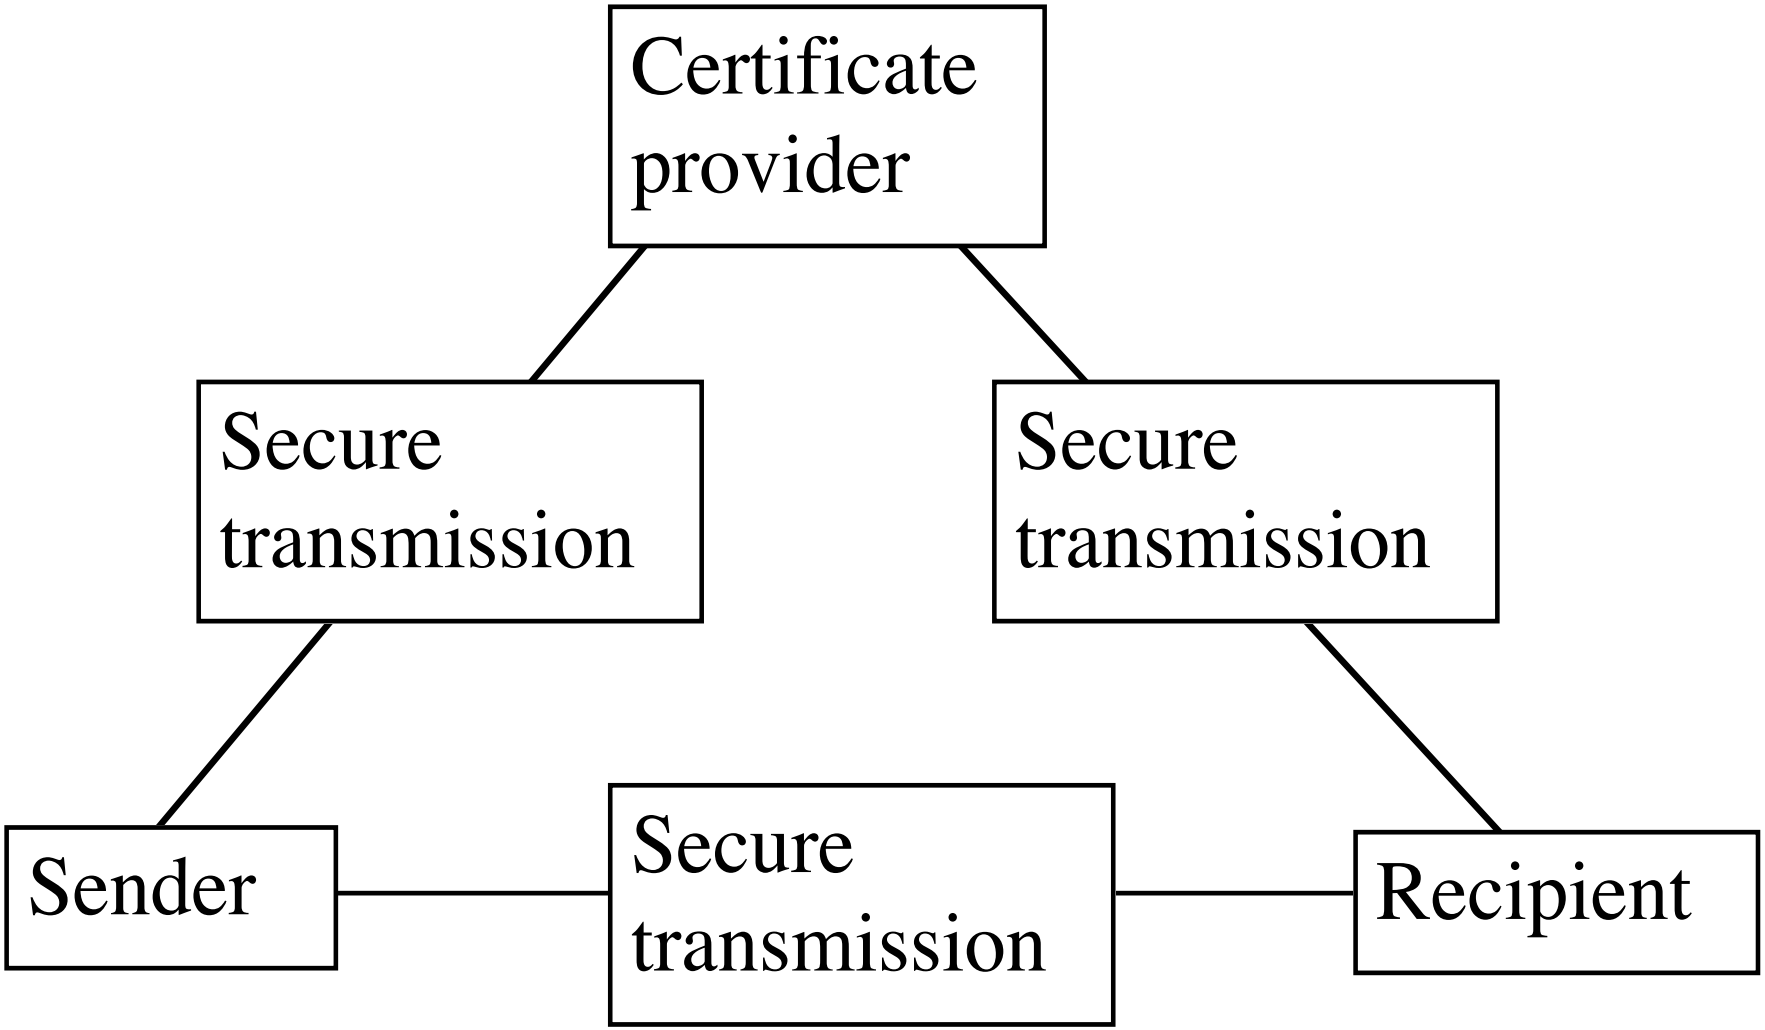

Store-and-forward transmission

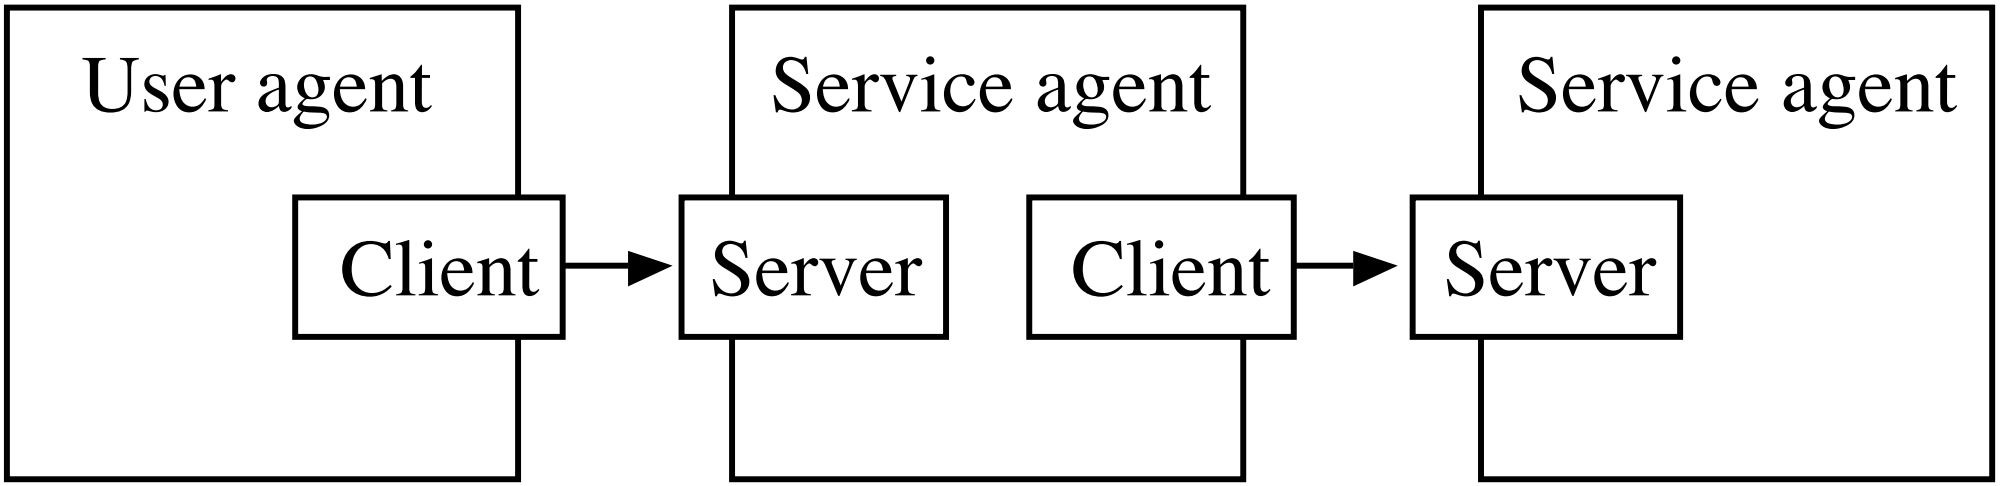

Ending a connection 1: notation to be used

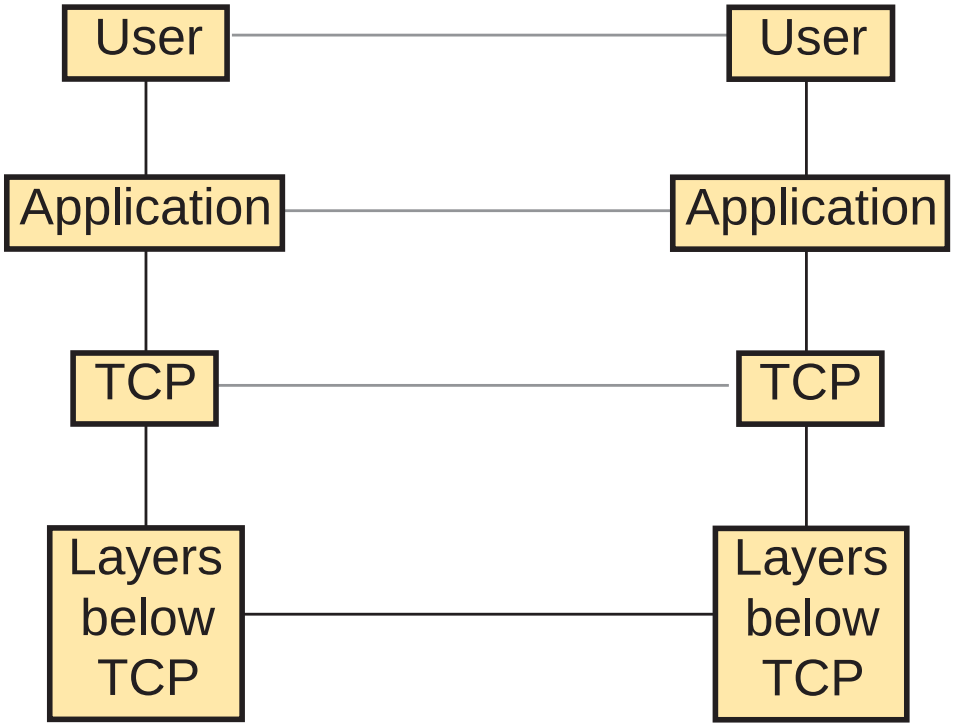

HTTP GET Operation

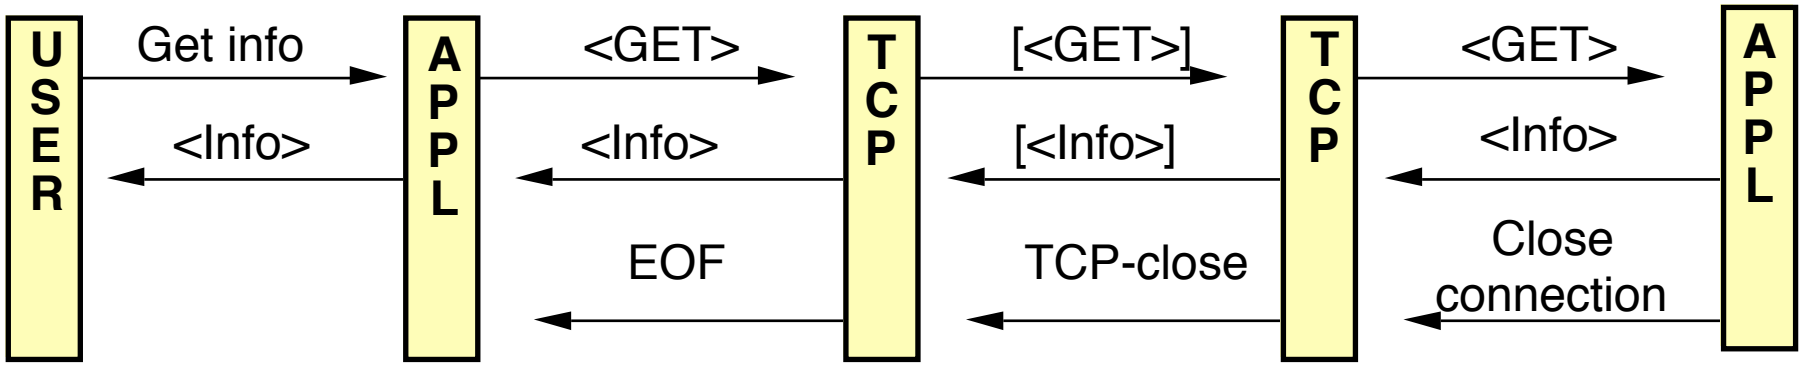

SMTP Sending a message

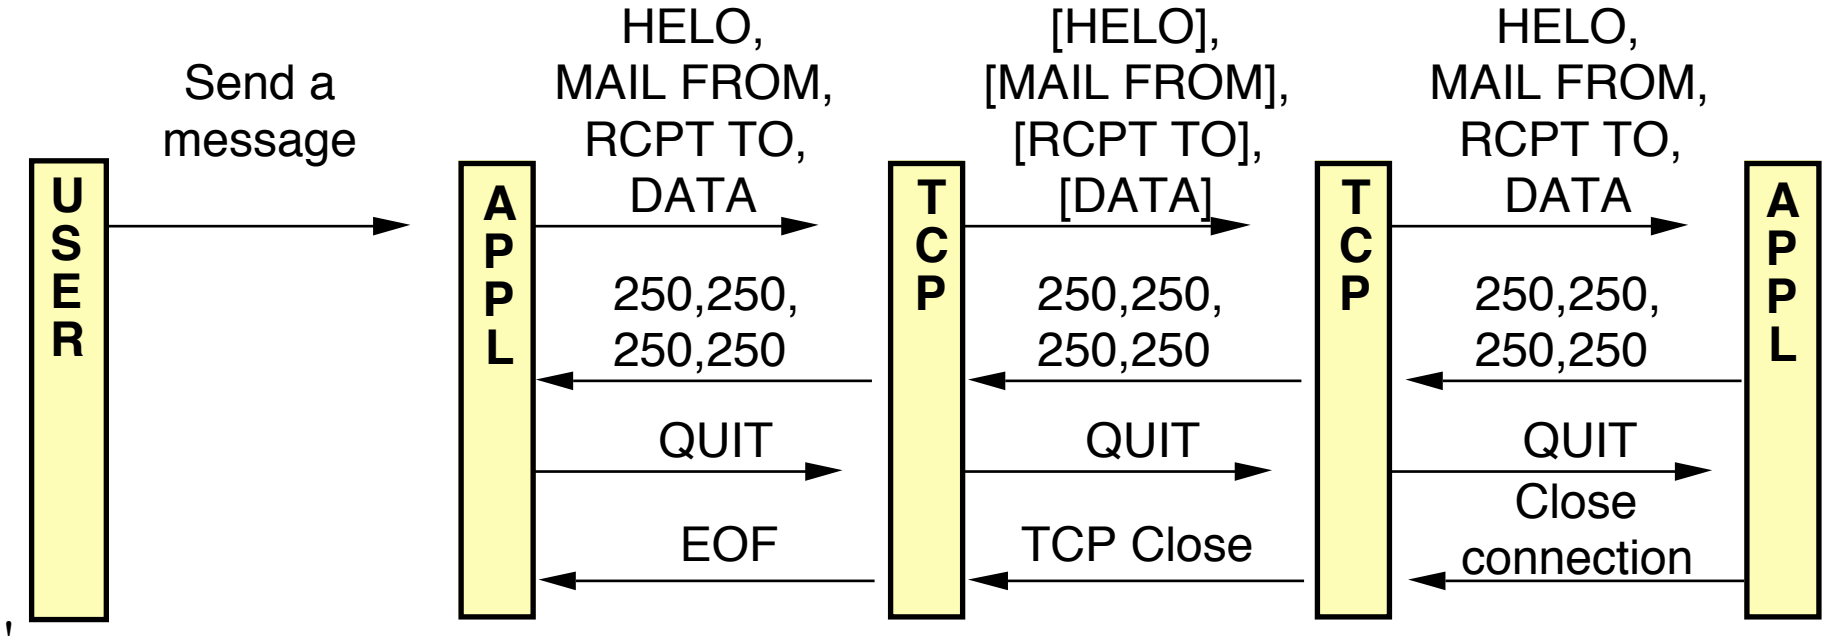

Connection retention

Transaction processing versus connection-oriented protocols

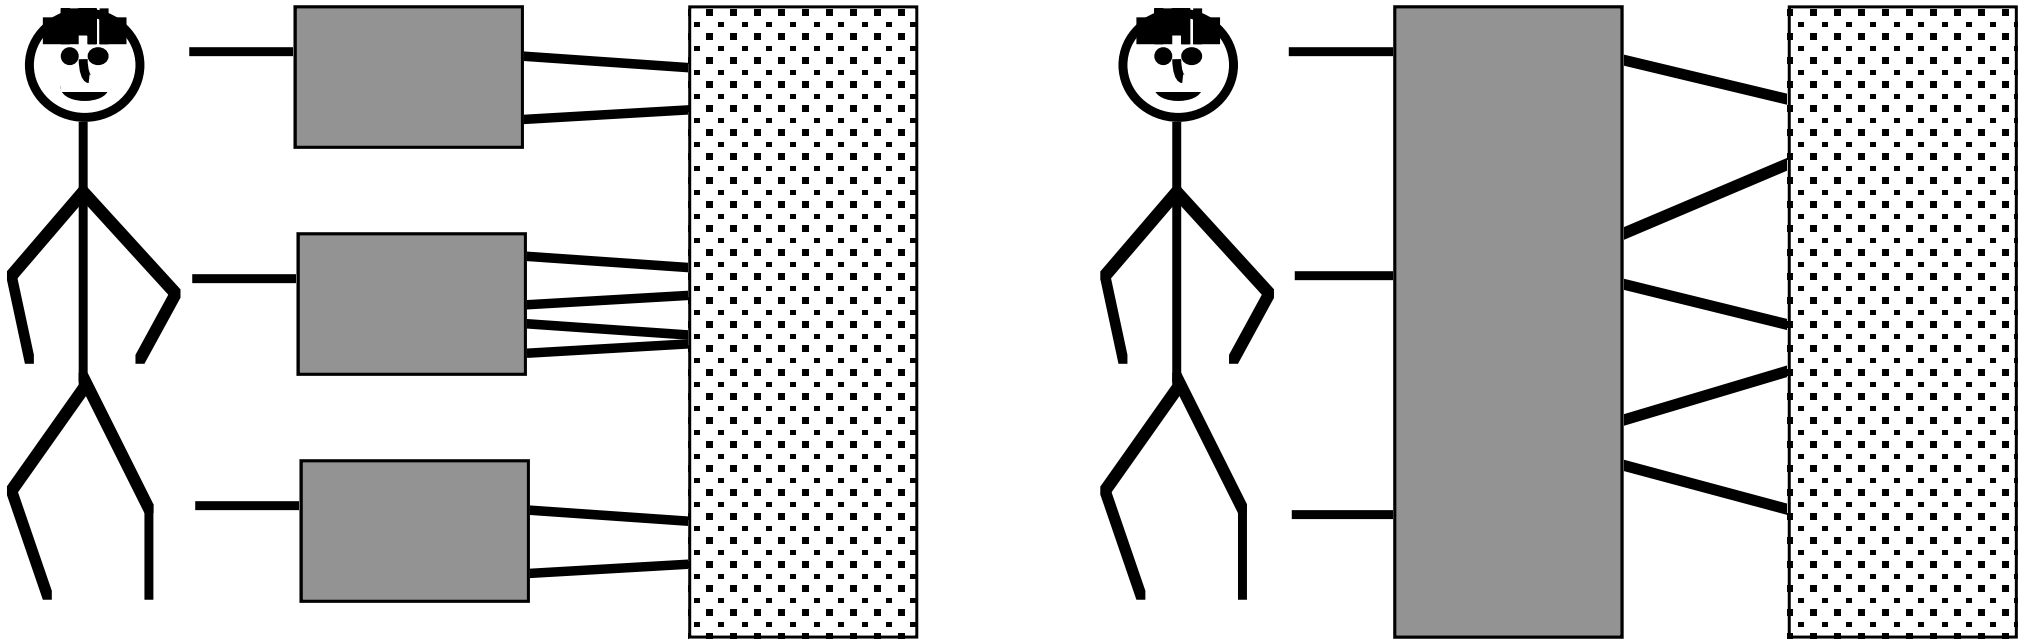

Stateful and stateless sessions

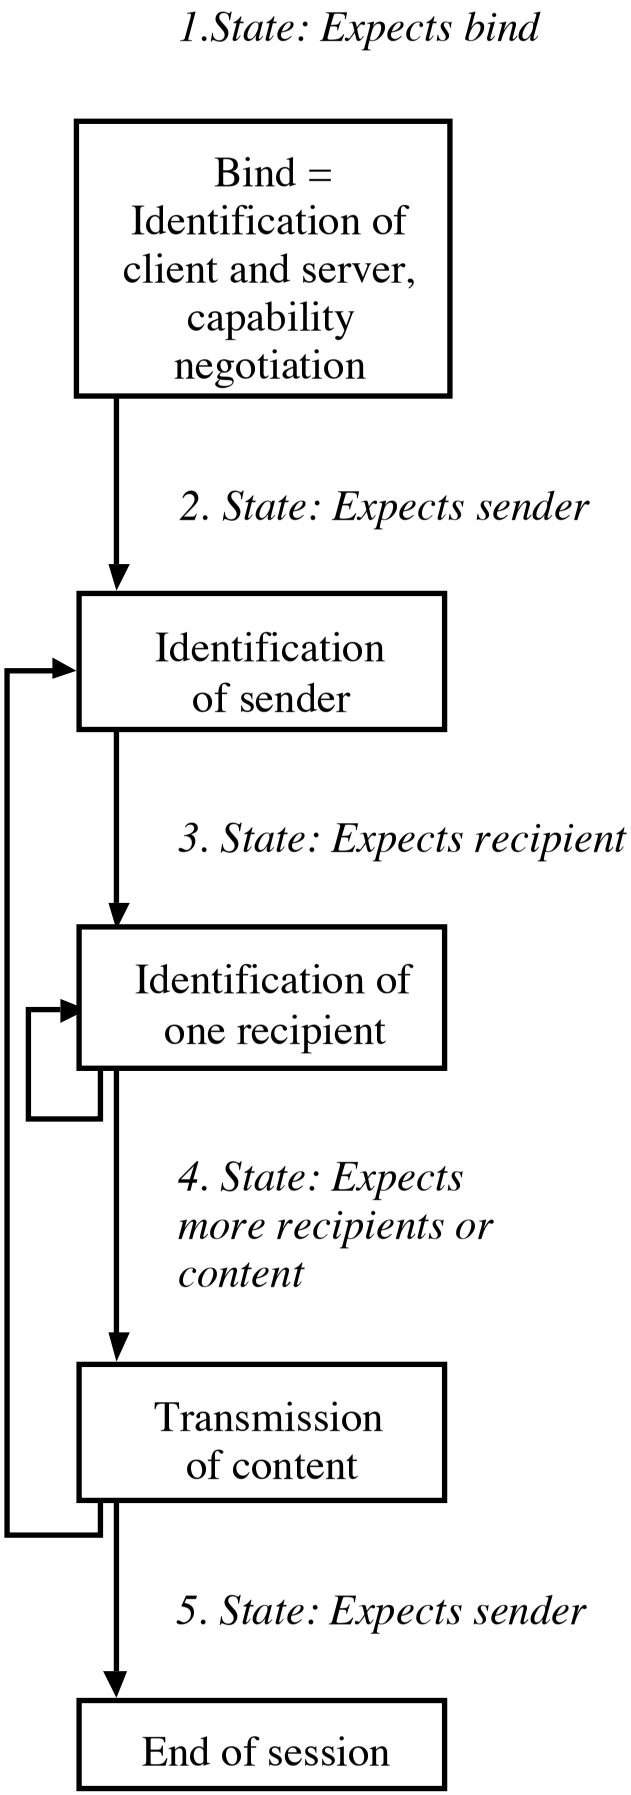

Reducing turn-around time

1. *Specify more powerful commands*, where more is done in one command, so that fewer interactions are needed.
2. *Open several parallel connections.* HTTP clients (web browsers) often keep four parallel connections for downloading the different parts of a web page (text, pictures, applets). Too many parallel connections is costly in resources for both the client and server, but with too few connections, dead time may occur when the client is waiting for data from all the connections.
3. In protoocols which use many small interactions, such as SMTP and NNTP, the delay can be used with *pipelining*, see next overhead.

Pipelining

Wait for response before sending the next command

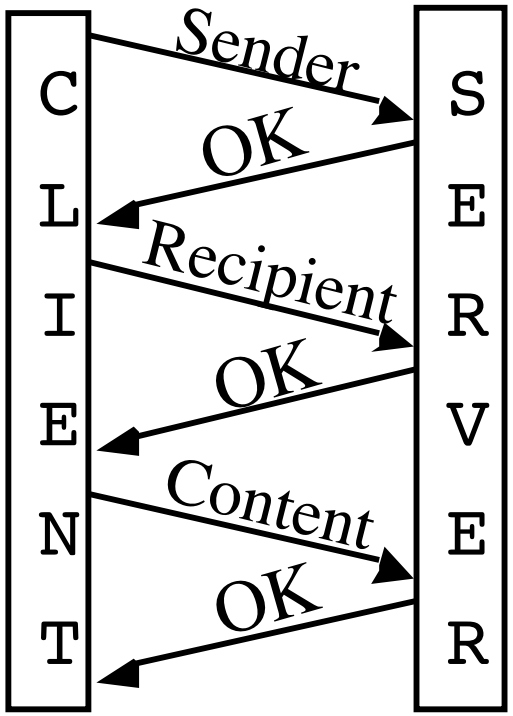

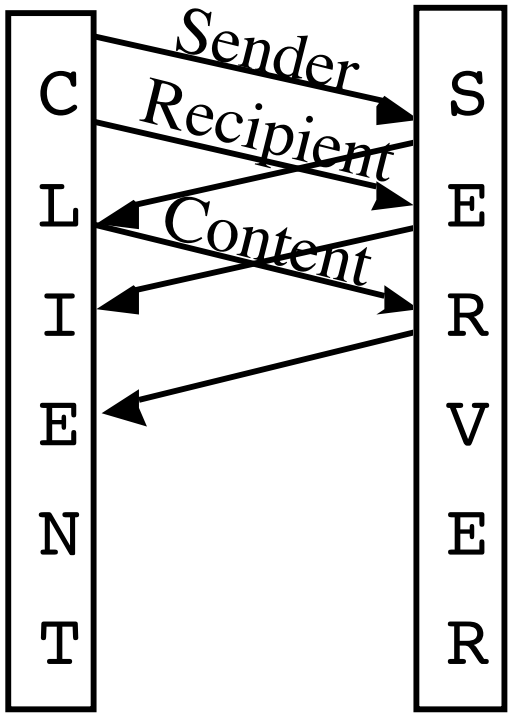

Pipelining: Send commands without waiting for response

Chaining, referral, multicasting

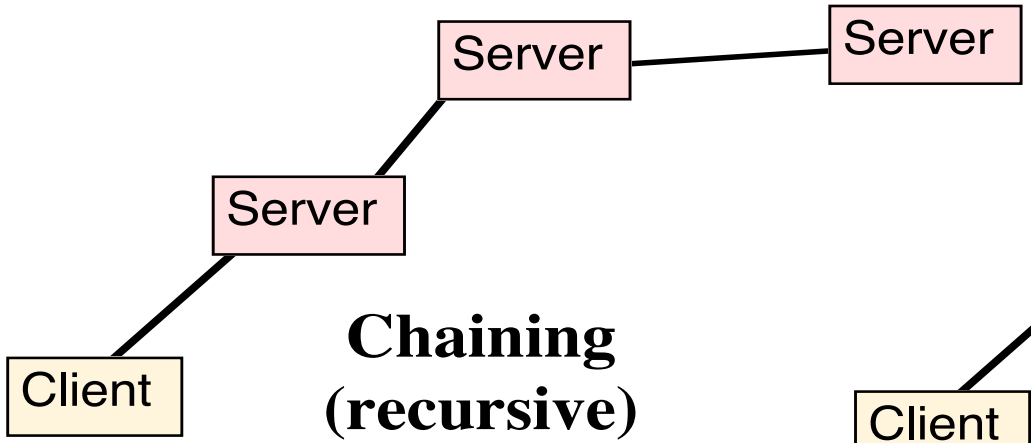

Chaining (recursive)

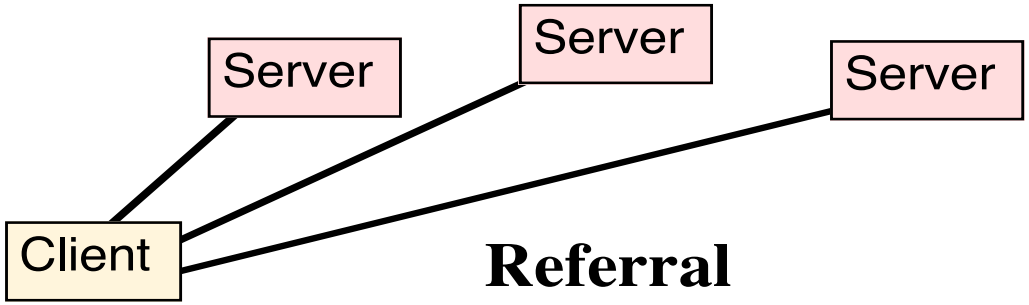

Referral (iterative)

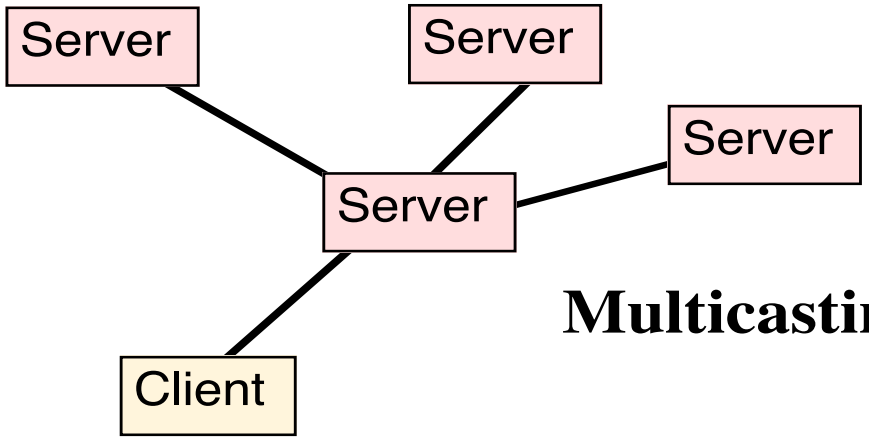

Multicasting

Extension problem

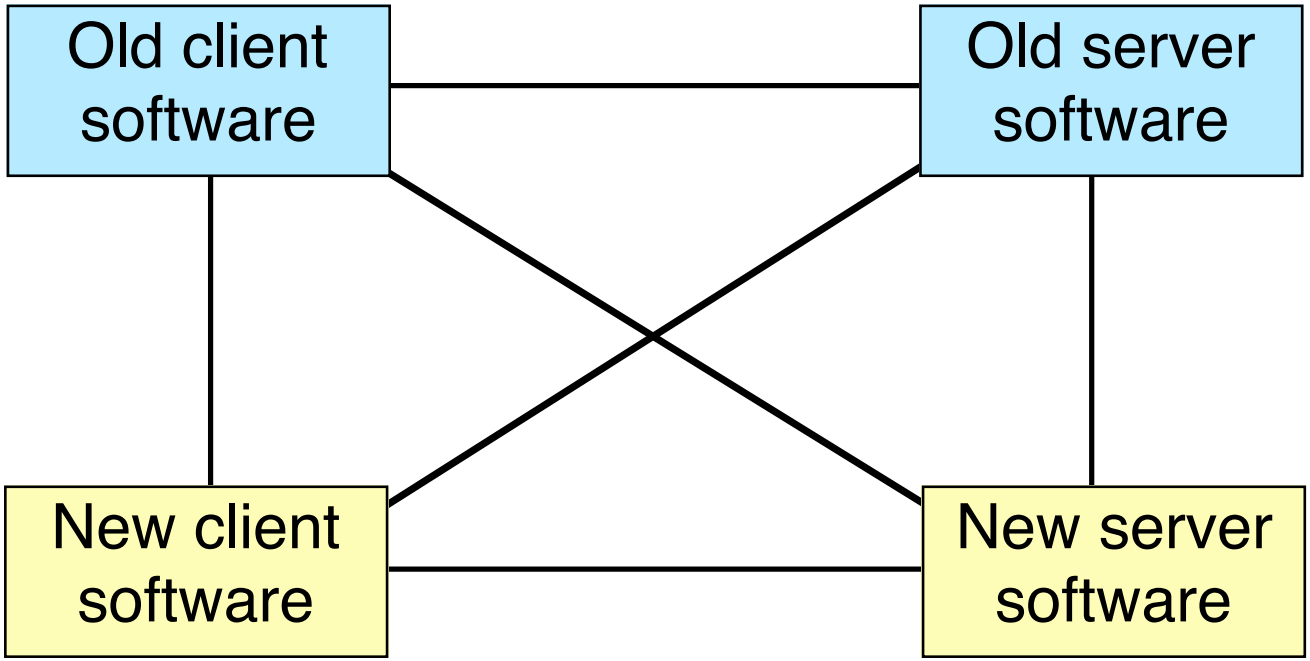

Horror example: Binary files through 7-bit e-mail

Extension by levels: for example HTML 1.0, HTML 2.0

Extension by feature selection

Built-in extension points

Registration facility vers. X-headers

Version number method

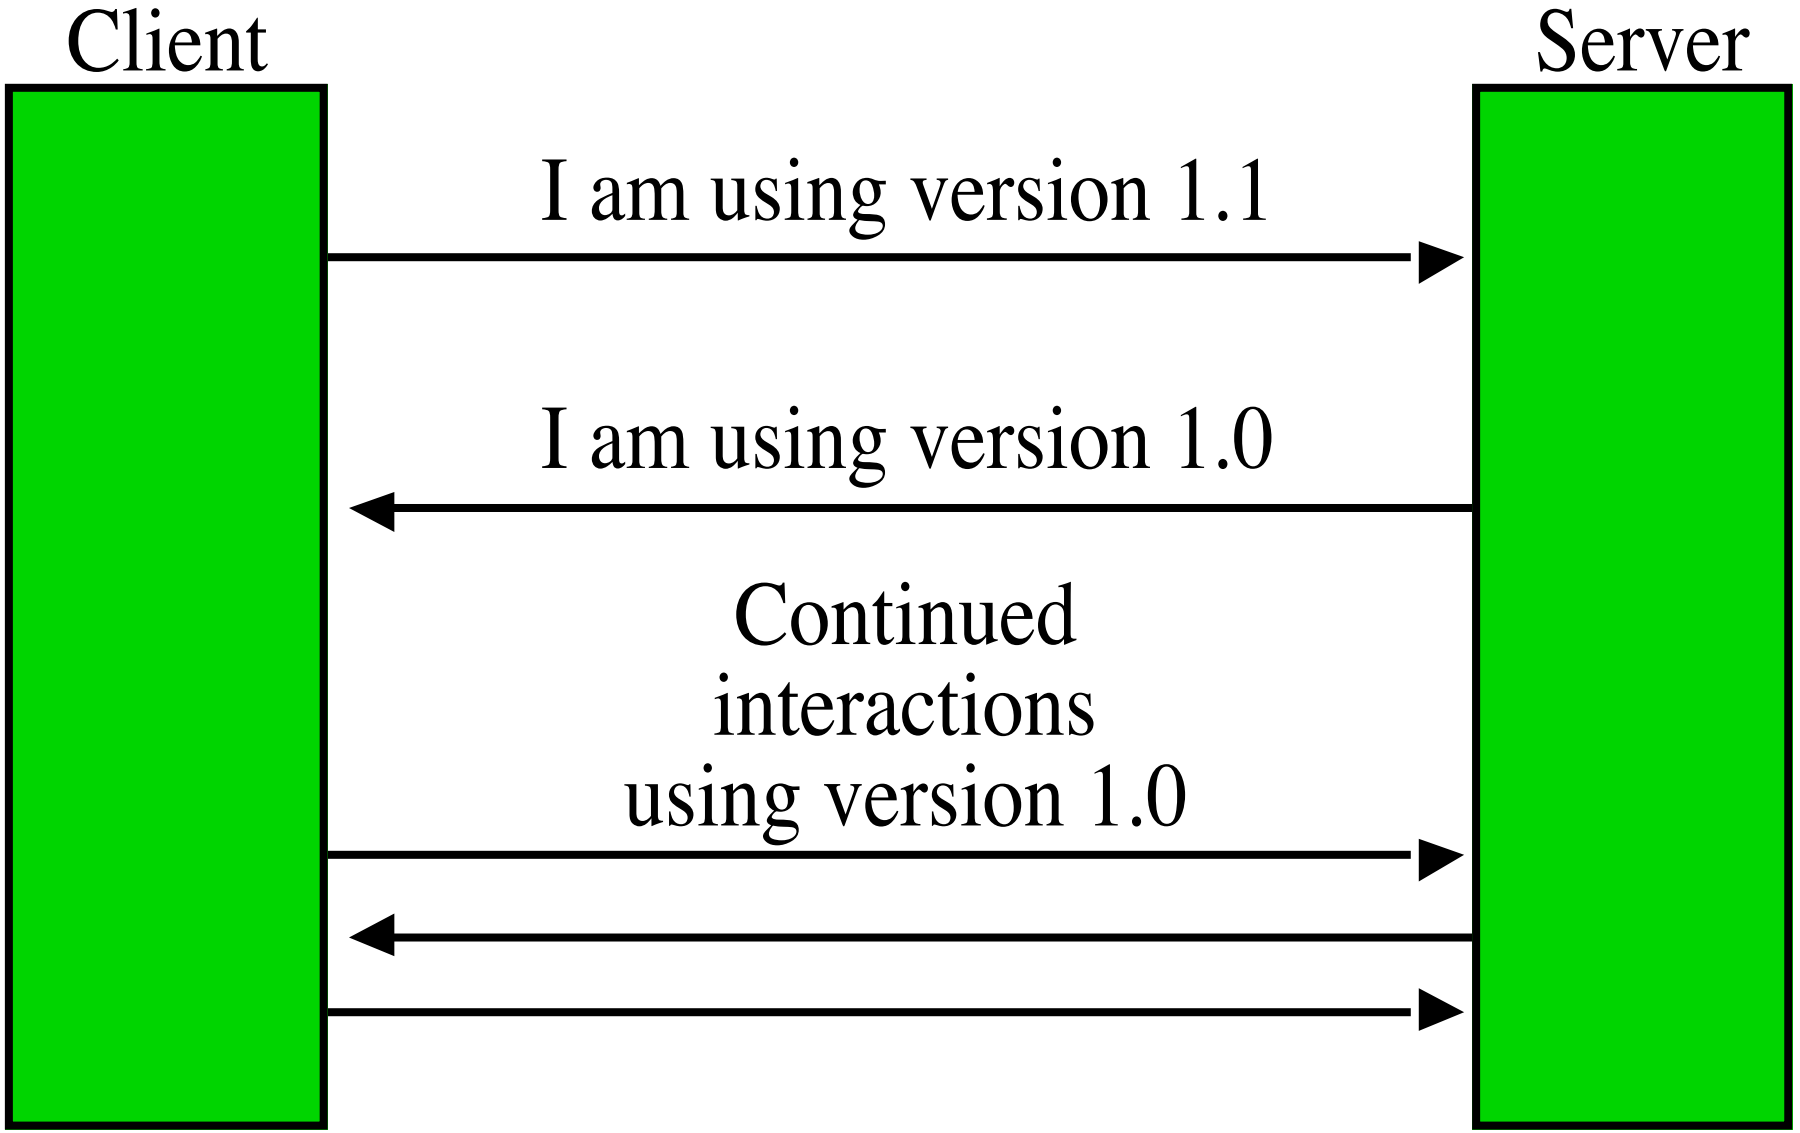

Feature selection method

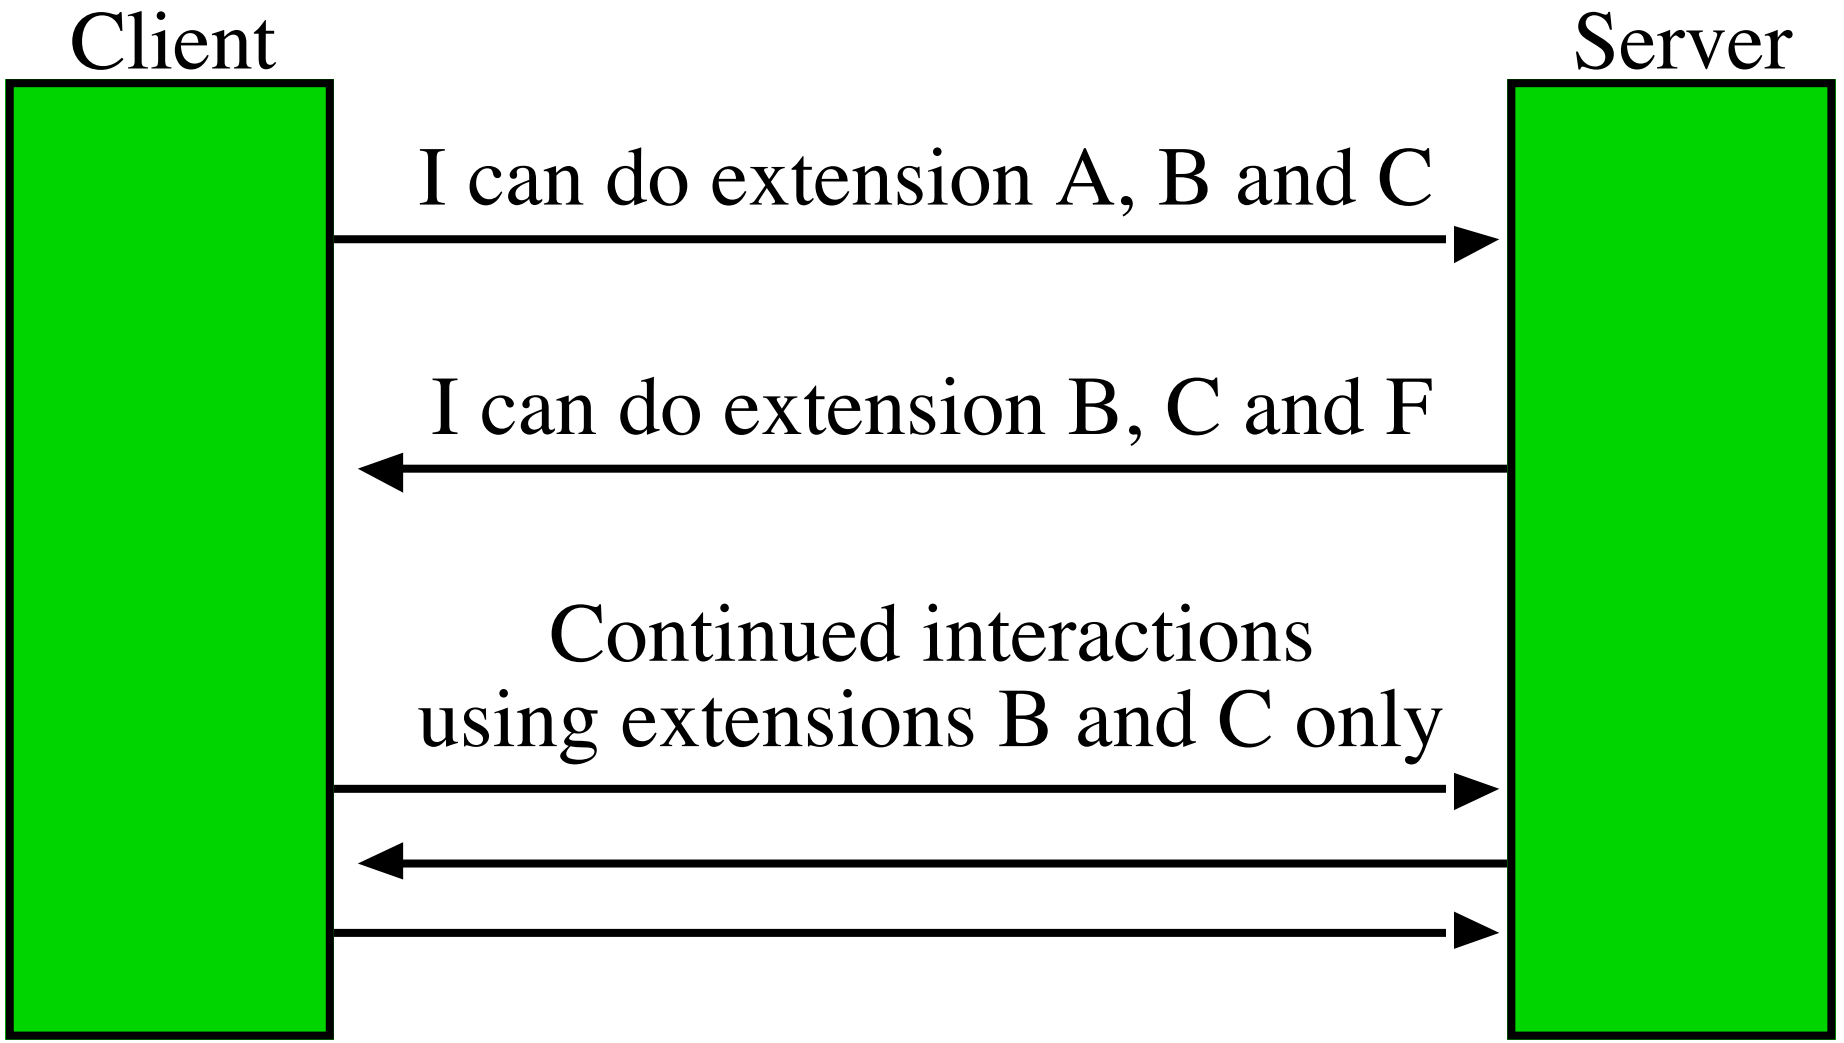

Intermediaries

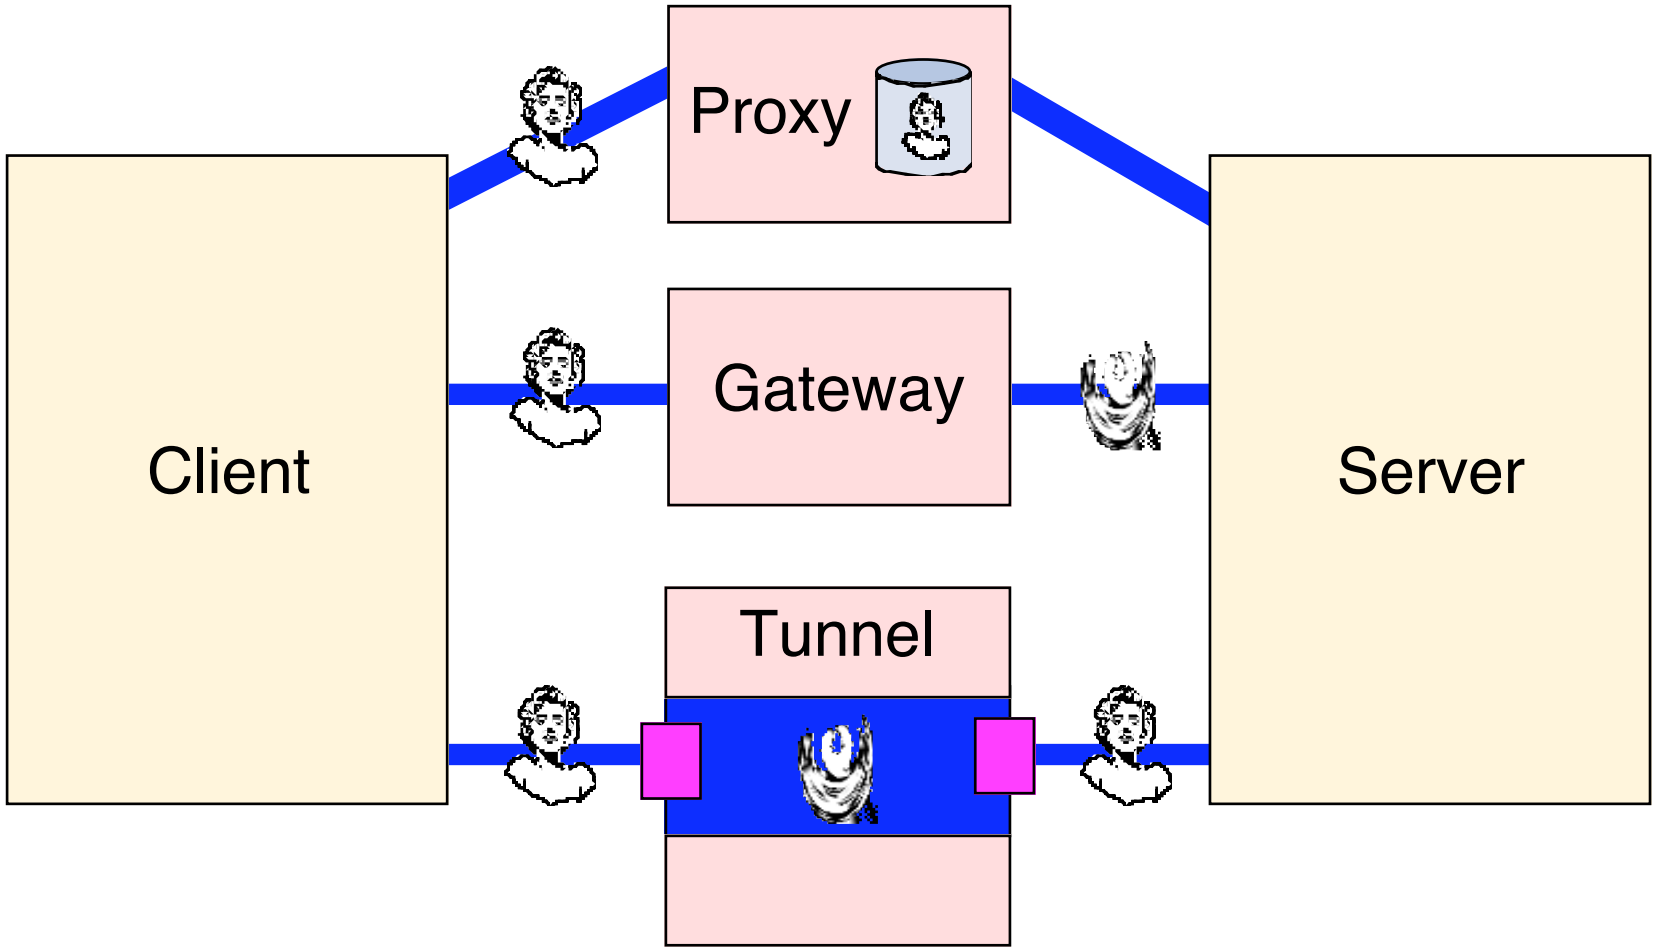

Caching: Saving a copy

Shadowing (push-caching): Agreed, controlled replication

Mirroring: Duplicated data base

Prefetching: Guess in advance

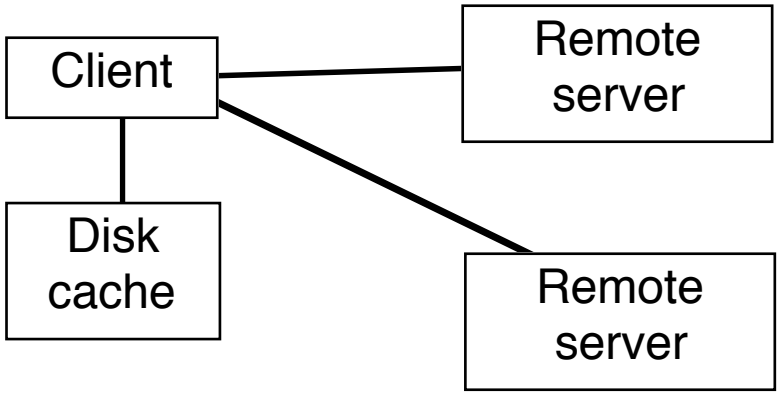

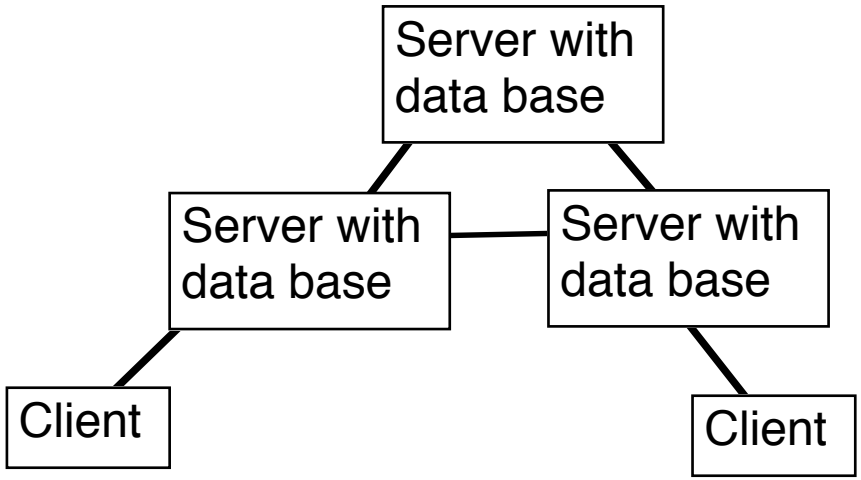

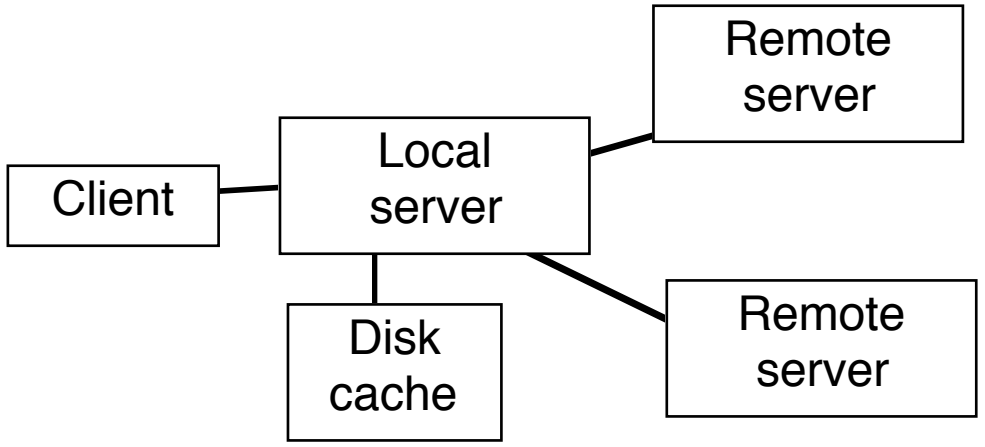

Replication

Why replication

Reduce network loade

Reduce load on very popular servers

- Example: Popular home page required nine dedicated workstations and rotating DNS server to distribute load

Faster response times

Master versus equalitarian replication

One master copy: All other instances are copies of the master

Replication can be pre-ordered, automatic, initiated by the master or the slave

Example: Most shareware data bases

No master copy: All instances are equal

Example: Usenet News

Pros and cons

Master copy gives simpler and safer updating

Master copy gives central control - *note virus control!*

Master copy depends on the master server

Replicate much or little?

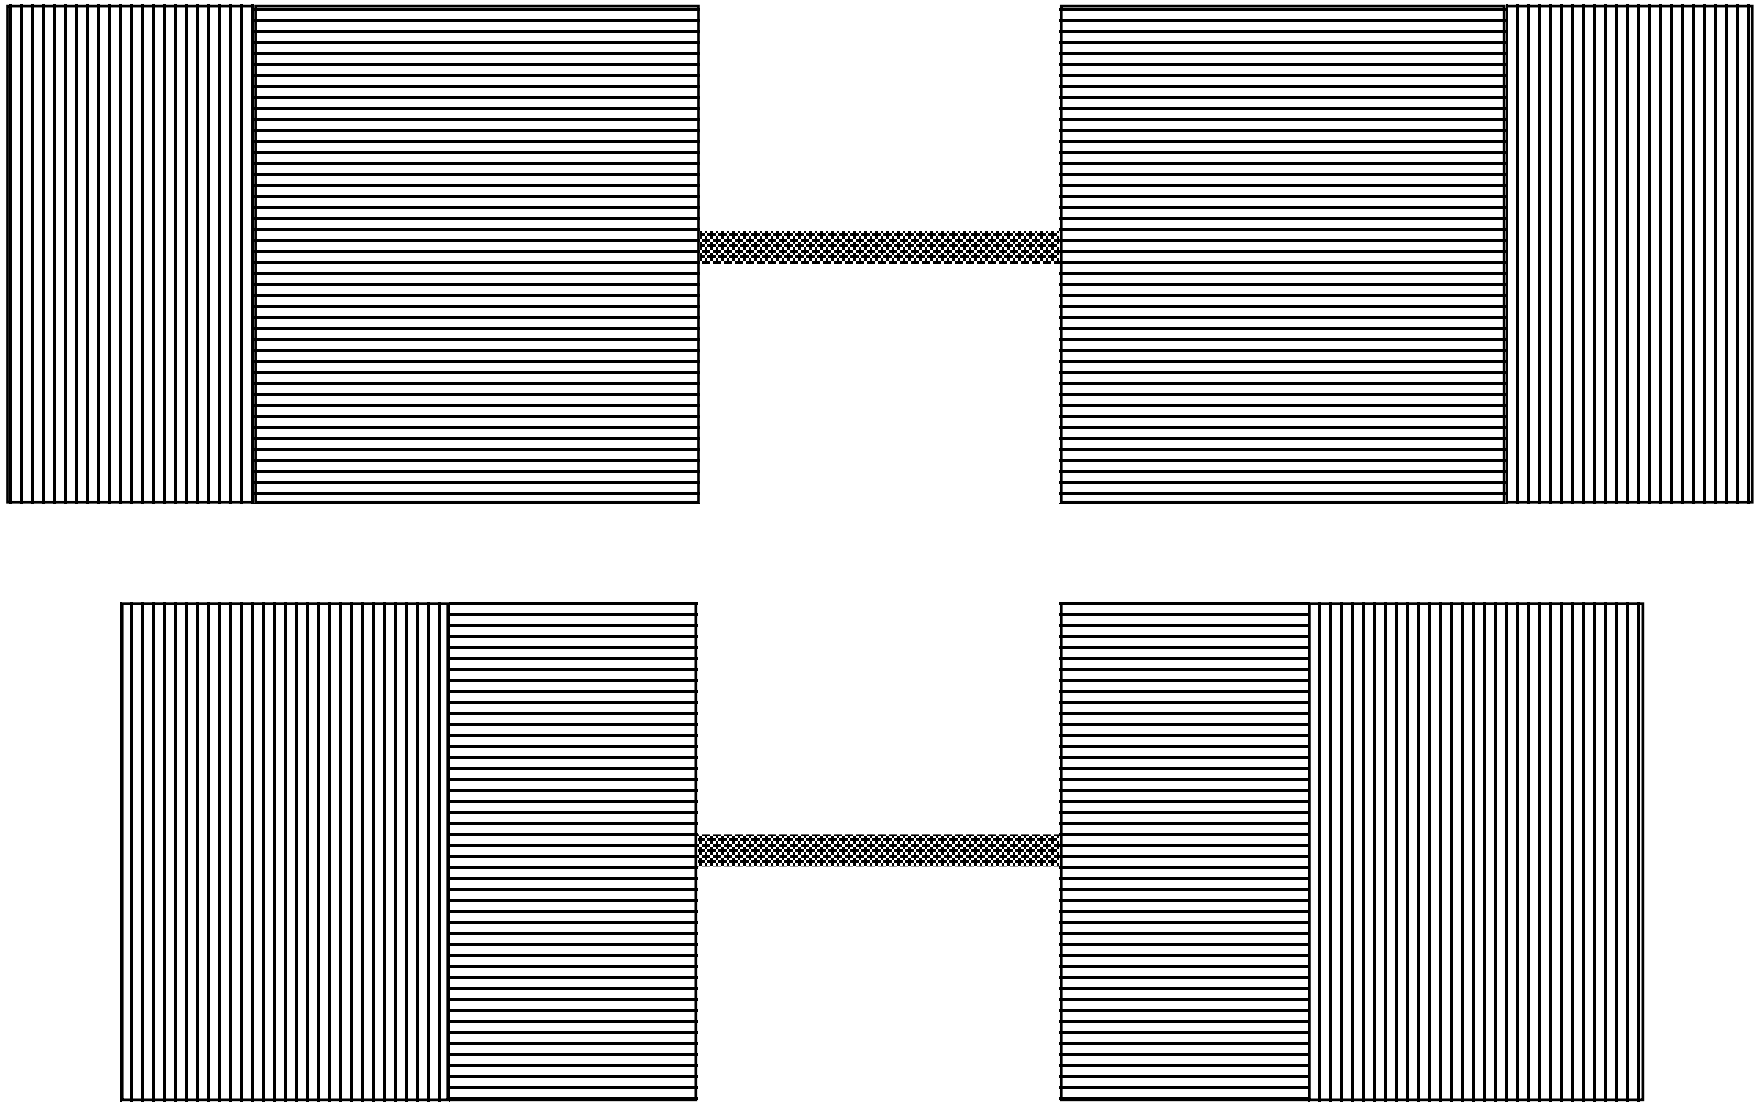

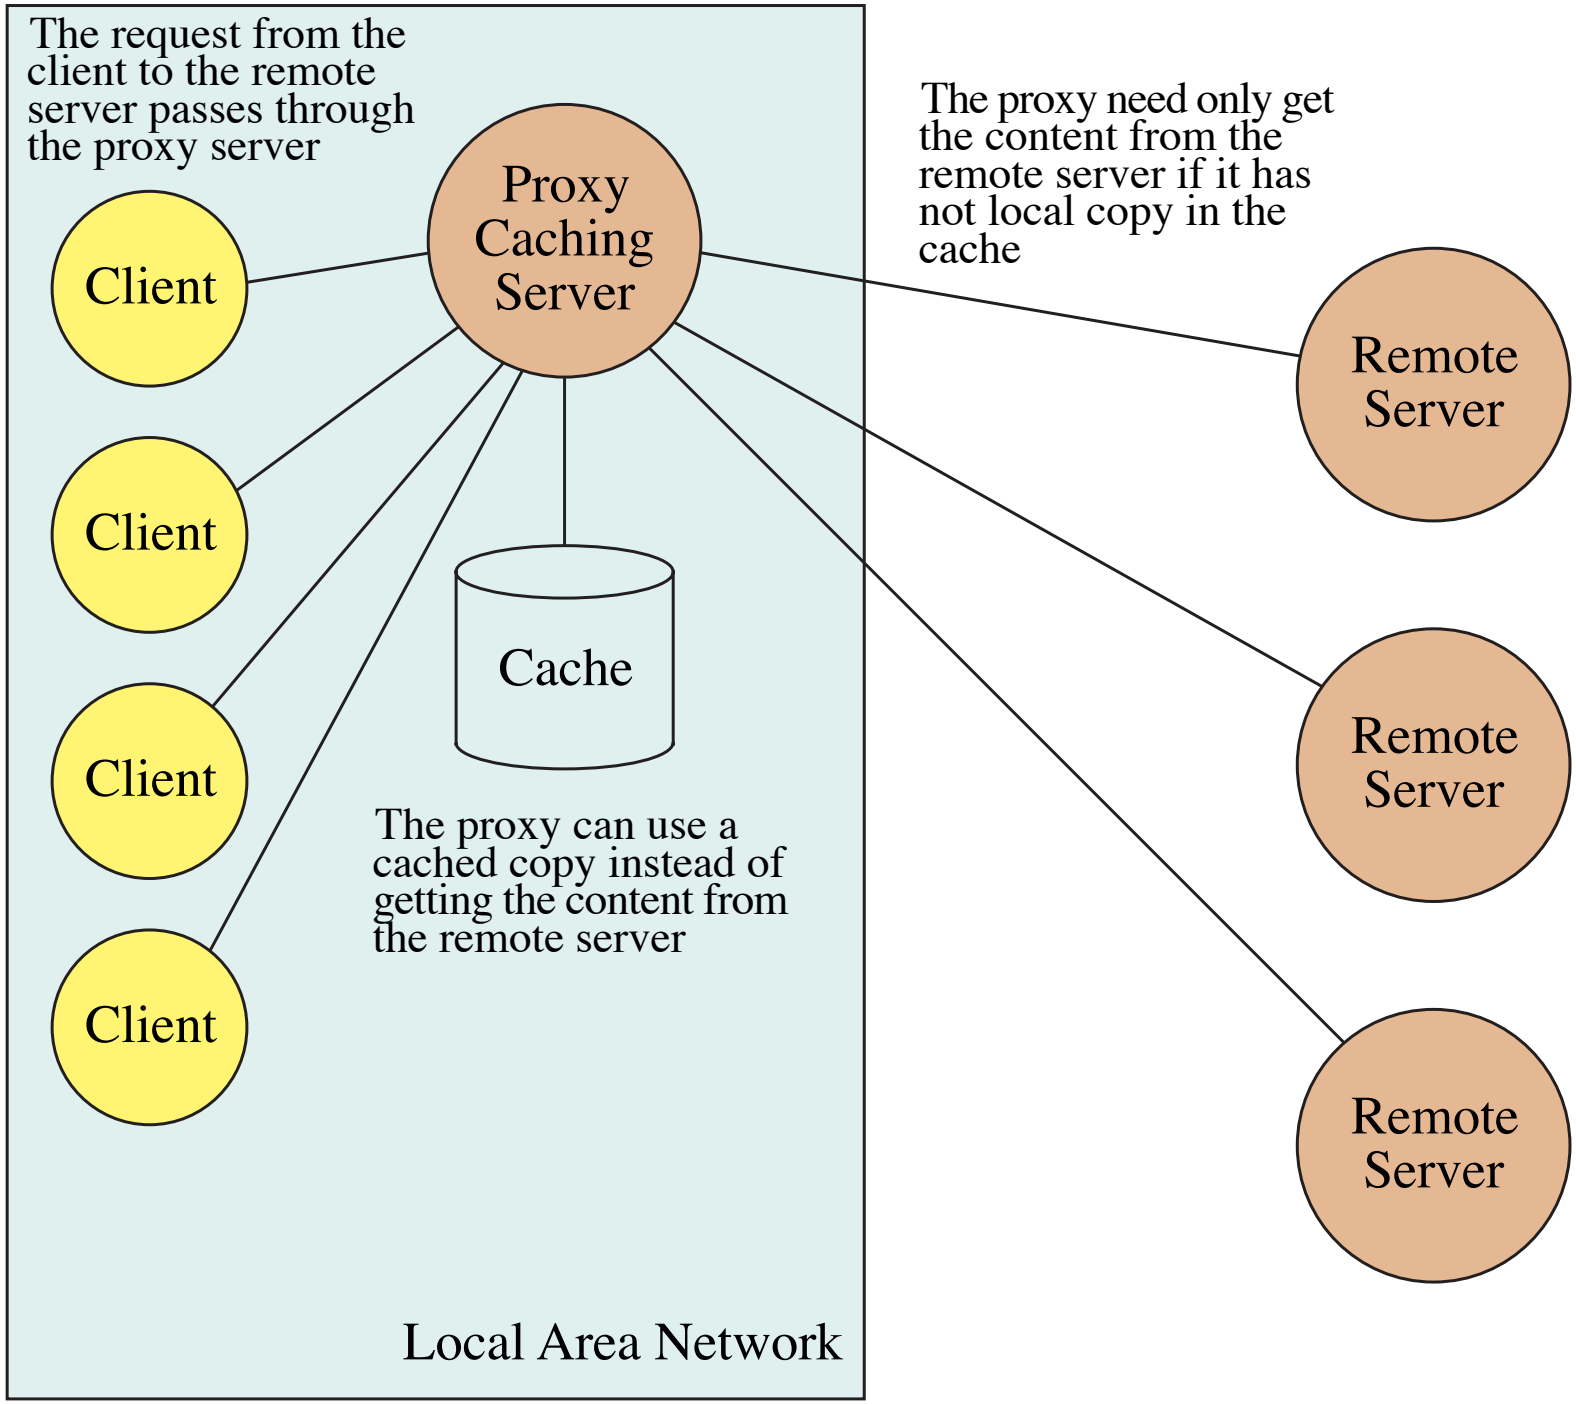
The request from the client to the remote server passes through the proxy server
Proxy Caching Server
Client
Client
Client
Client
Cache
The proxy can use a cached copy instead of getting the content from the remote server
Local Area Network
The proxy need only get the content from the remote server if it has not local copy in the cache
Remote Server
Remote Server
Remote Server

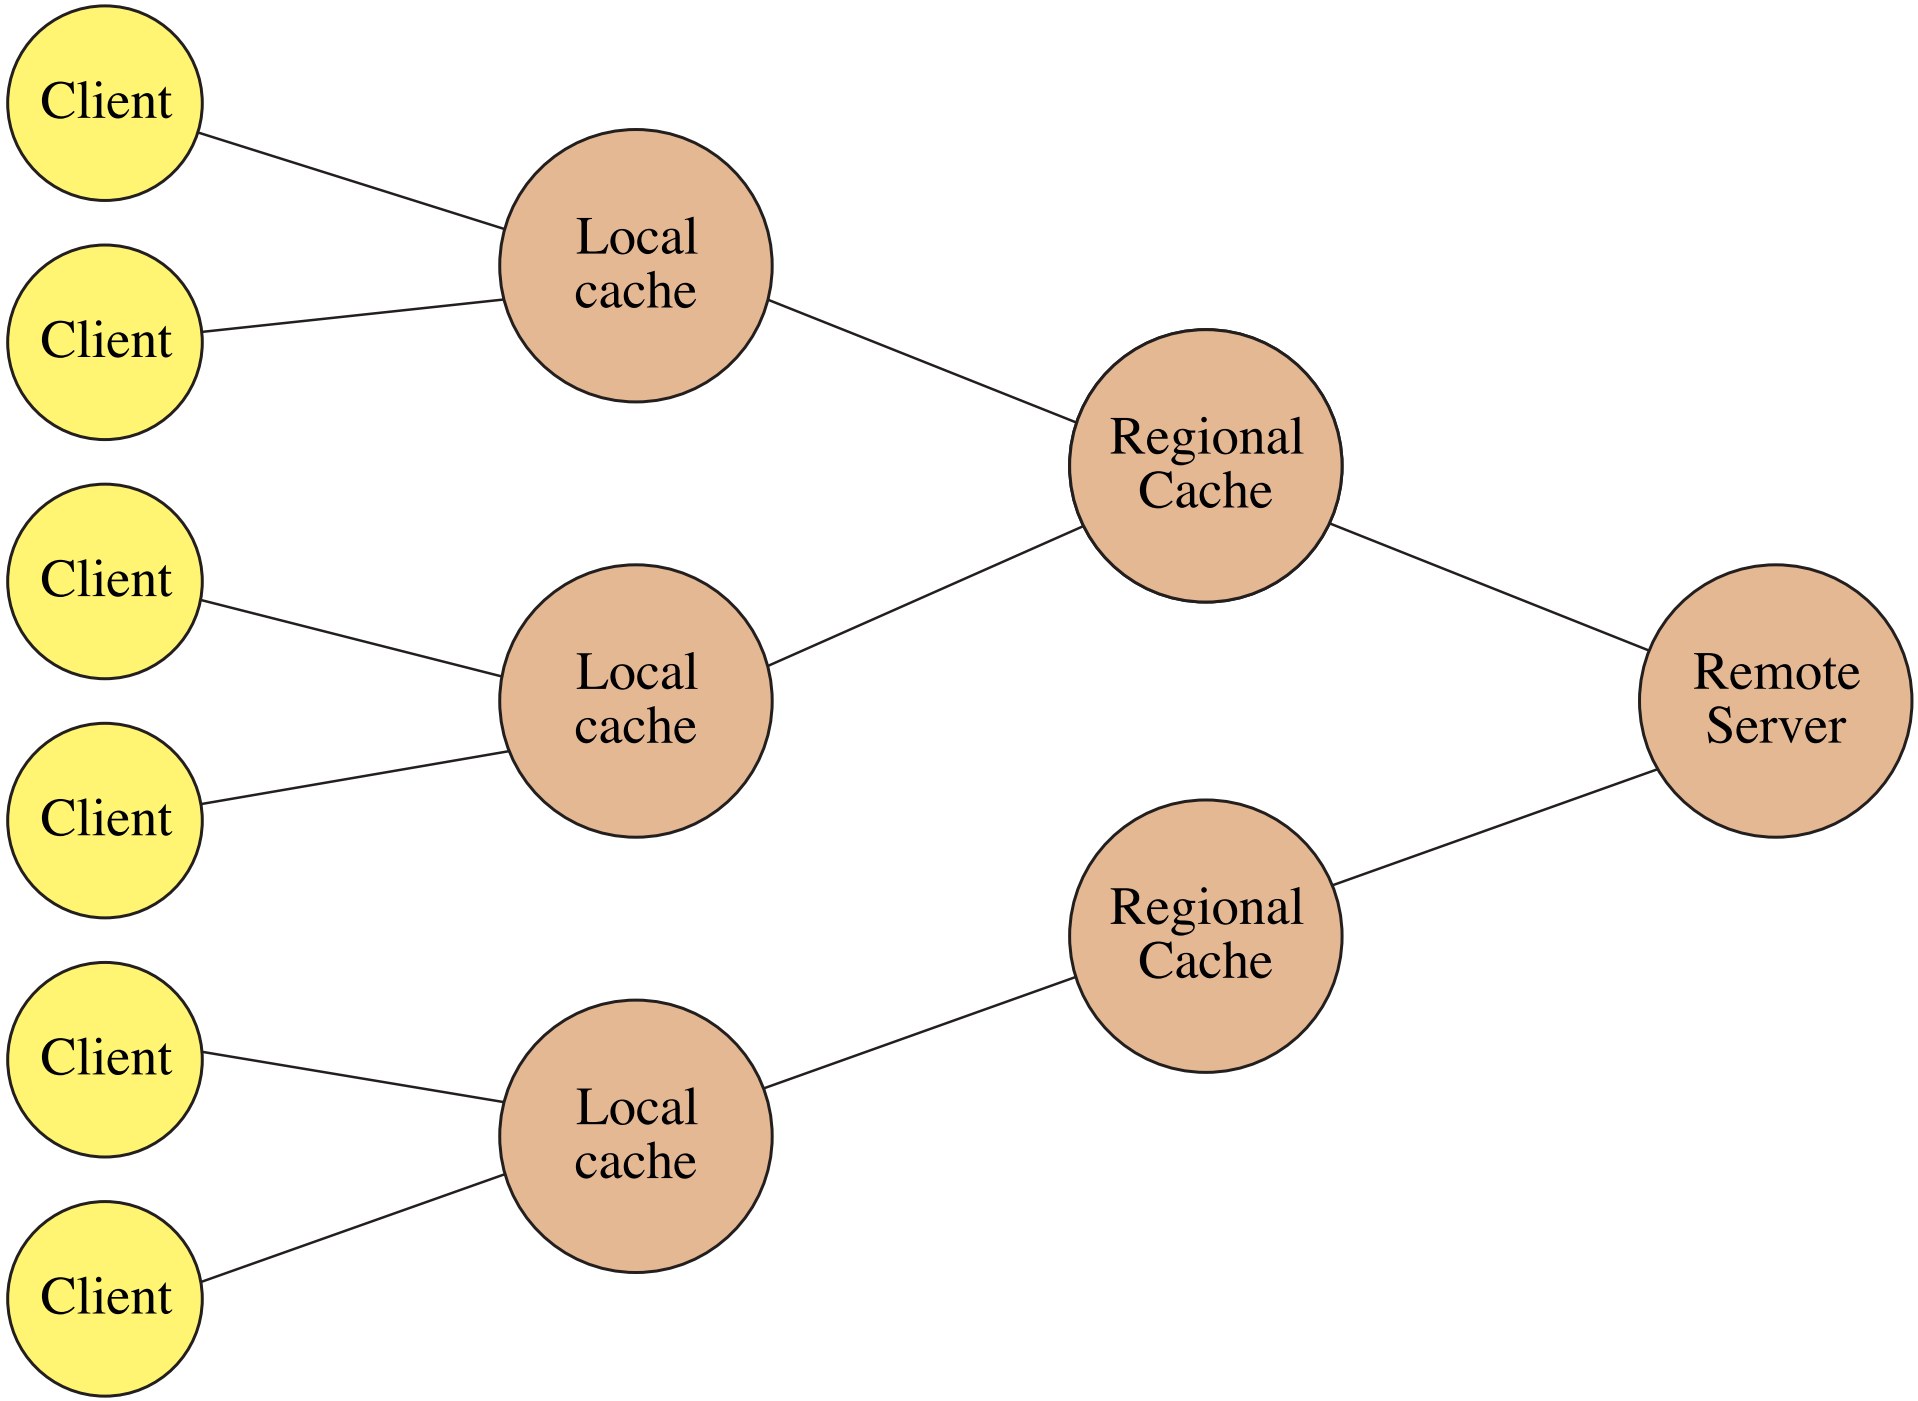
Client
Client
Client
Client
Client
Client
Local cache
Local cache
Local cache
Regional Cache
Regional Cache
Remote Server

Negative caching

Caching of the fact that something does not exist, to avoid trying to get it several times from a remote source

Problems with caching

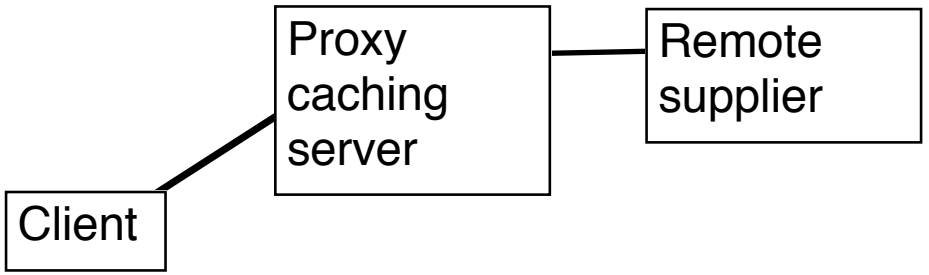

User gets out-of-date result

- Example: Cartoon changed daily, proxy caching server updated graphics only every 14 days!

Copyright violation?

Solutions:

- Best solution: Controlled mirroring
- Cacher checks for changes
- Provider supplies refresh time
- Guess at refresh based on last update
- User pushes refresh button
  - Important: User refresh request must go all the way

Access statistics not correct

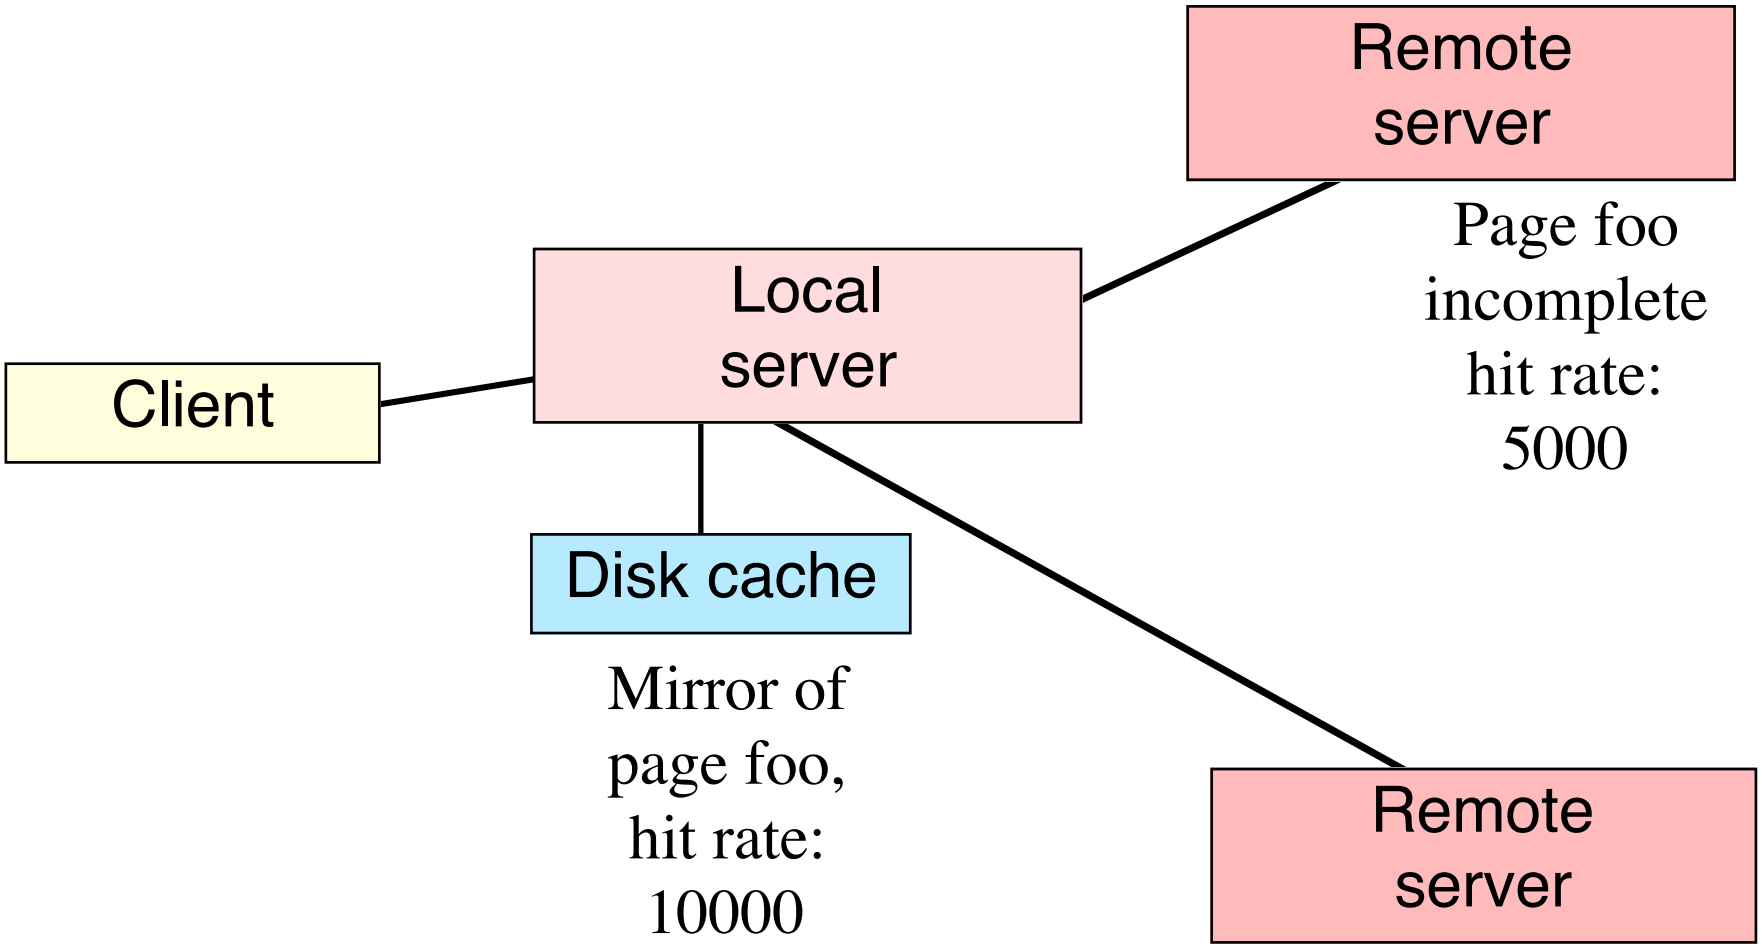

Security and caching

IP authentication defeated - maybe an advantage?

Crypthographic encryption and authentication will work

Cache manager of course a security risk

Cache must be programmed not to support private pages to non-authorised raders

Designing web page to utilize caching

Many small graphics, reuse the same graphic several times

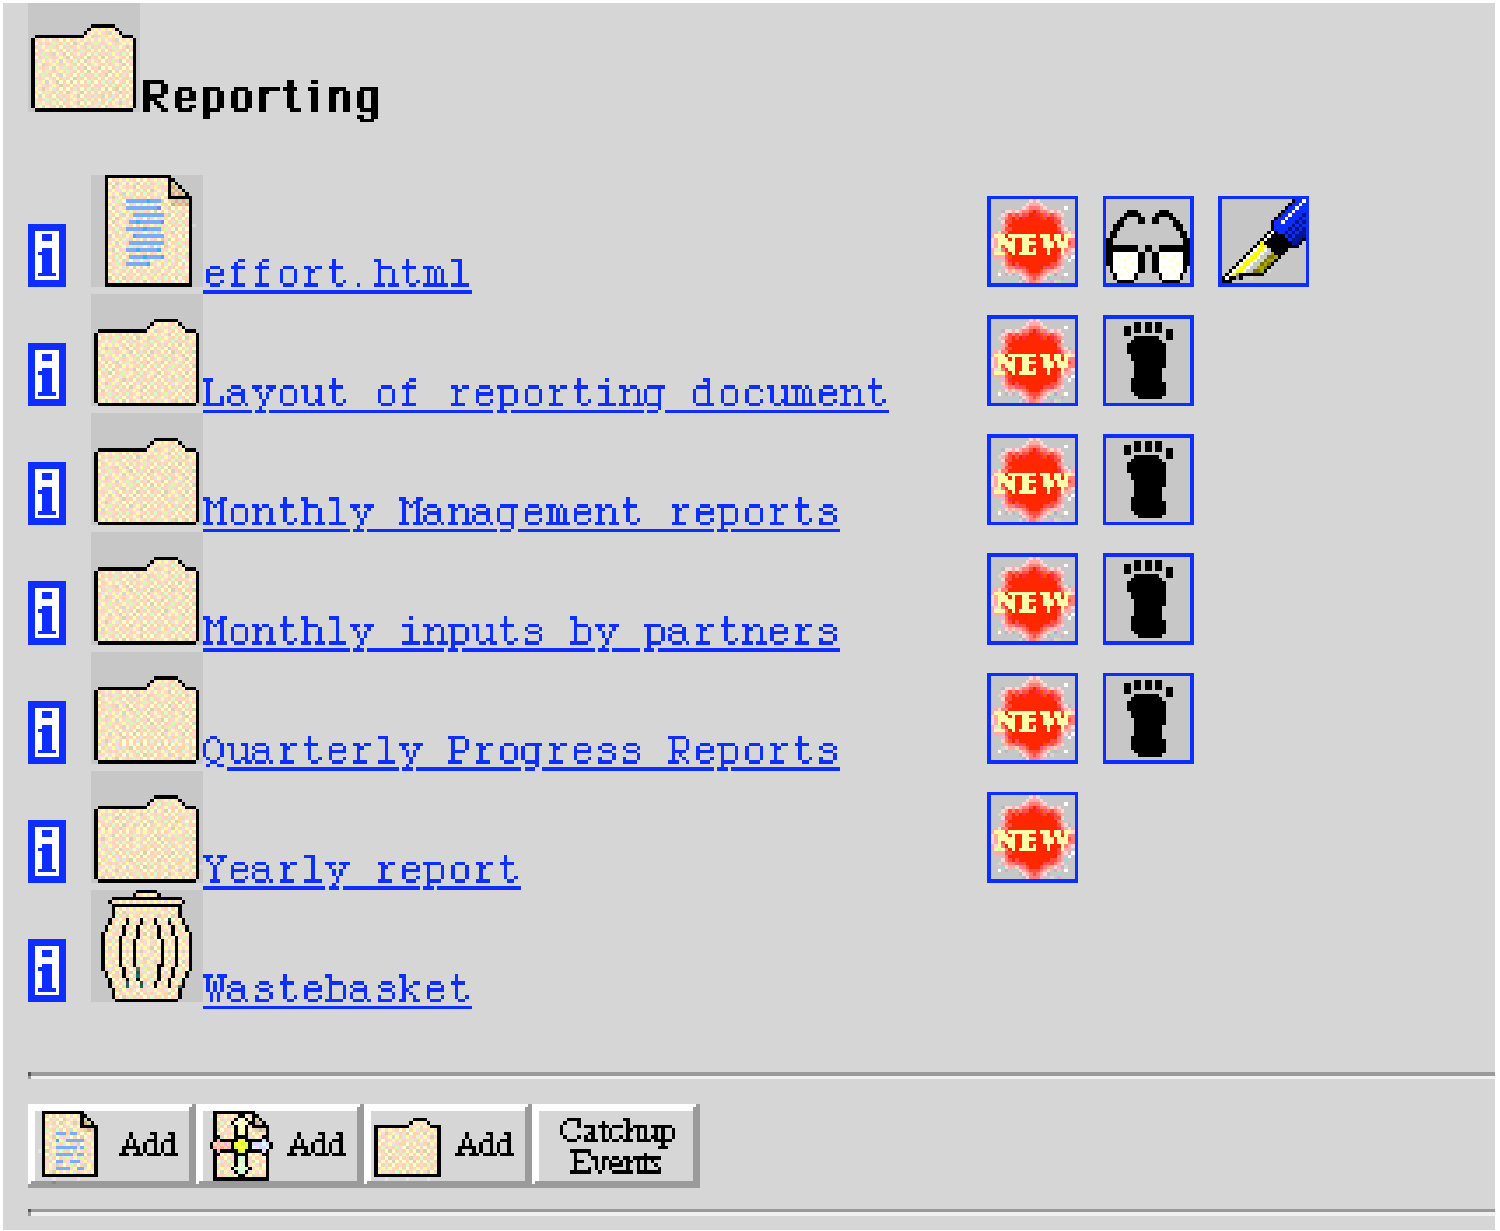

Locating nearest copy

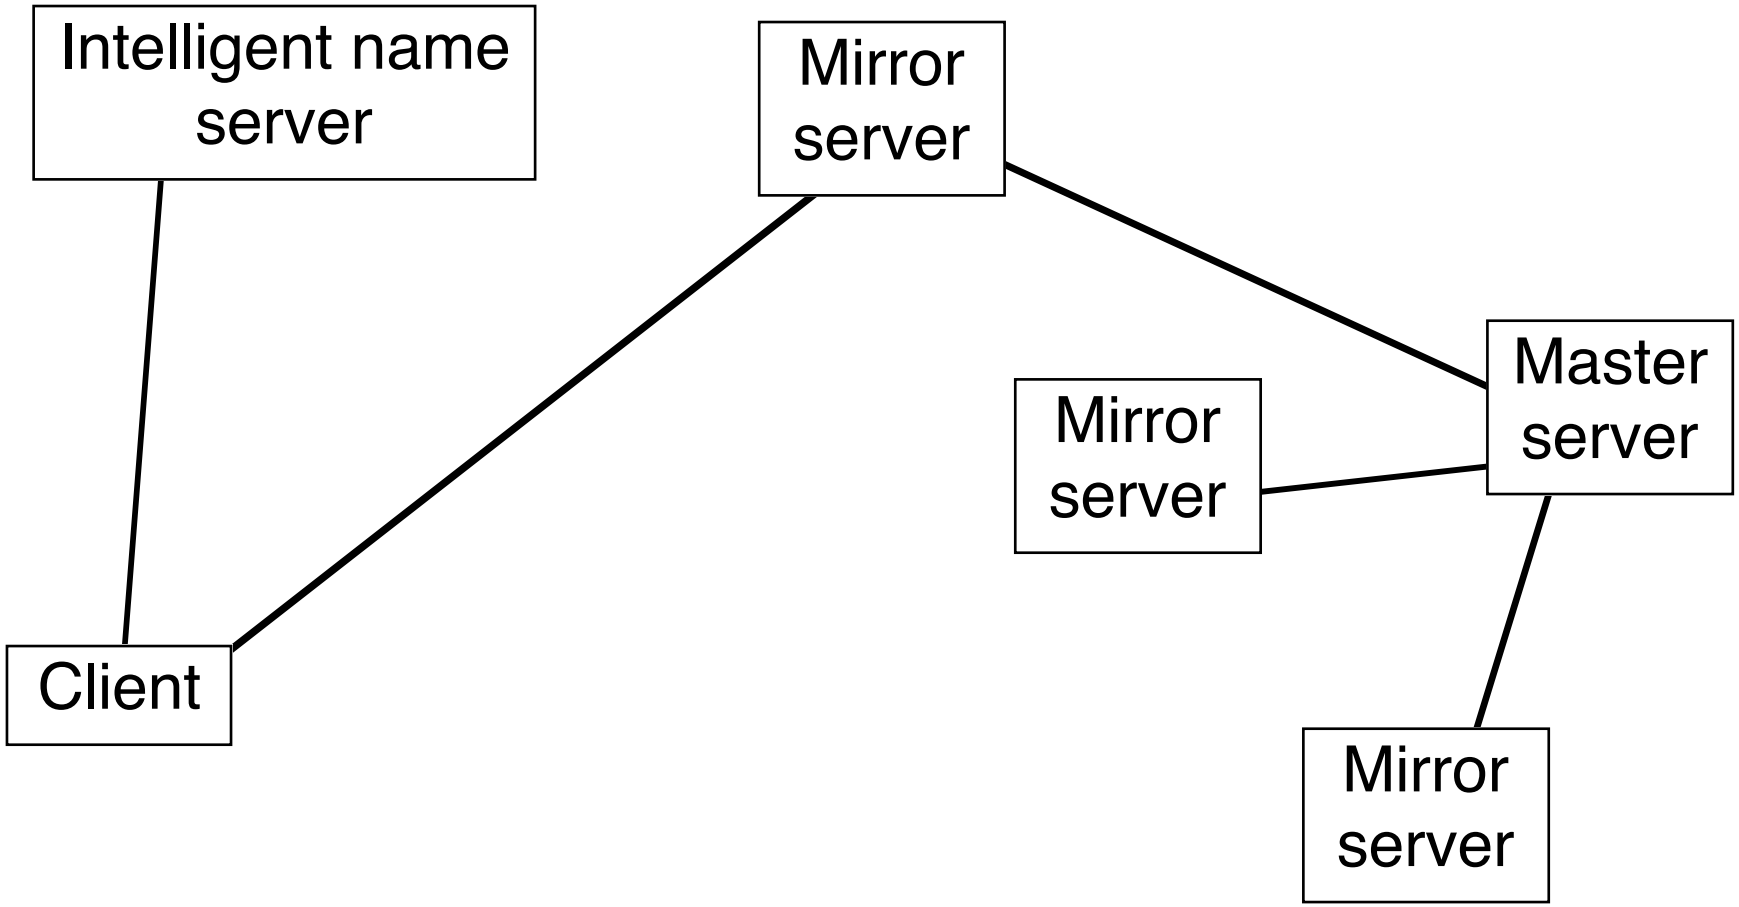

IETF Standards terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119.

Classes of standards

- Experimental standard
- Proposed standard
- Draft standard
- Standard
- Historical
- Informational
- BCP

The First Golden Rule:

Be liberal in what you accept,
be conservative in what you produce.

Does this mean a different protocol and syntax for what you produce and what you accept?

How do you know what (in excess to the standard) you should accept, and what (included in the standard you should not produce)?

Example; e-mail: Do not use blanks in e-mail names

Example; e-mail: Accept

John T. Smith <jsmith@foo.bar.net>

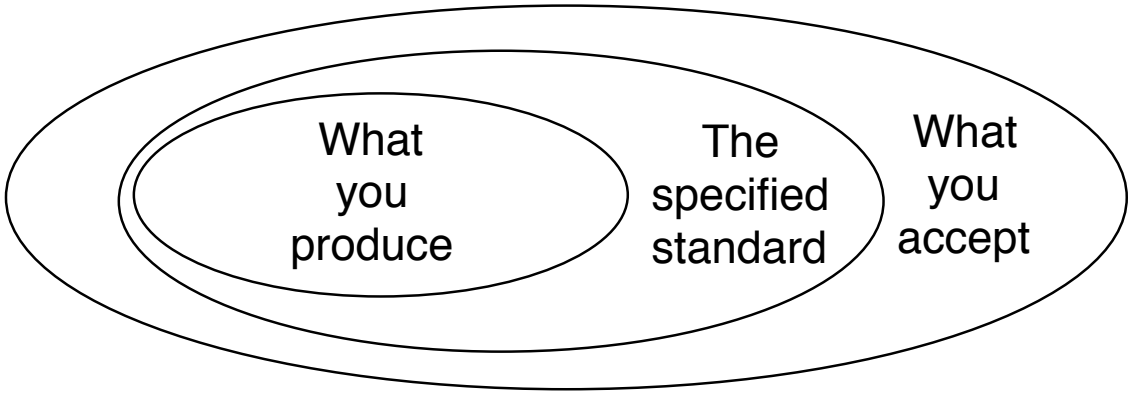

Golden Rules

(1) Be liberal in what you accept, be conservative in what you produce

Use a narrow produce syntax and a wide accept syntax

(2) Do no harm

What may be good in your special case, may in other cases cause harm

(3) Do not munge

Munge = Modify what other network modules has produced

Names in the Internet

Physical net addresses, example: 130.237.161.10

Domain names, example: ester.dsv.su.se, eies2.njit.edu

E-mail-addresses: example:
president@whitehouse.gov

DNS = Domain Naming Service translates domain names to physical net addresses. Can be accesses through the client "nslookup" (RFC 1034, RFC 1035)

People seldom see the physical net addresses, since translation from domain names to physical net addresses is done by the application programs used.

Domain naming tree structure

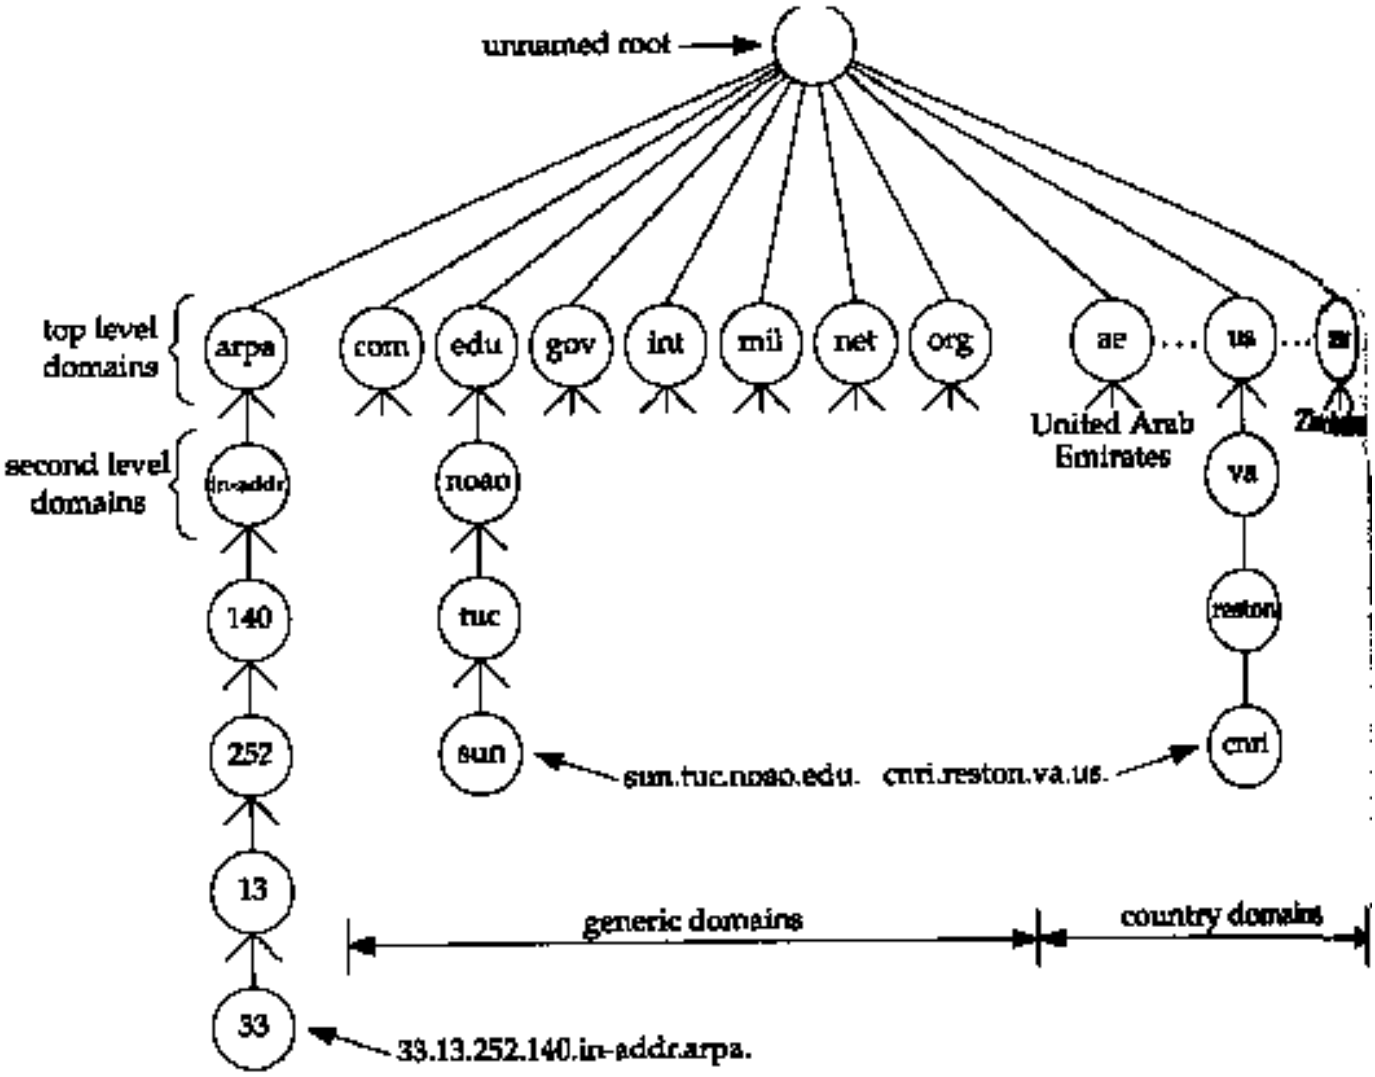

Figure 14.1 Hierarchical organization of the DNS.

Picture from Stevens: TCP/IP Illustrated, Volume 1 page 188

Top level domains

Organizational groups (U.S. or international)

COM	for commercial companies
EDU	for schools and universities
GOV	for other government agencies
INT	for international and multinational organizations
MIL	for military organizations
NET	network providers and gateways
ORG	for organizations

New TLDs?

FIRM	Businesses, firms
STORE	Offering goods to purchase
WEB	Activities related to the WWW
ARTS	Cultural and entertainment
REC	Recreational/entertainment
INFO	Information services
NOM	Personal names, etc.

Countries

US	U.S.A.
UK	for United Kingdom
SE	for Sweden
FR	for France
etc.	

Conceptual arrangement of name server hiearchy

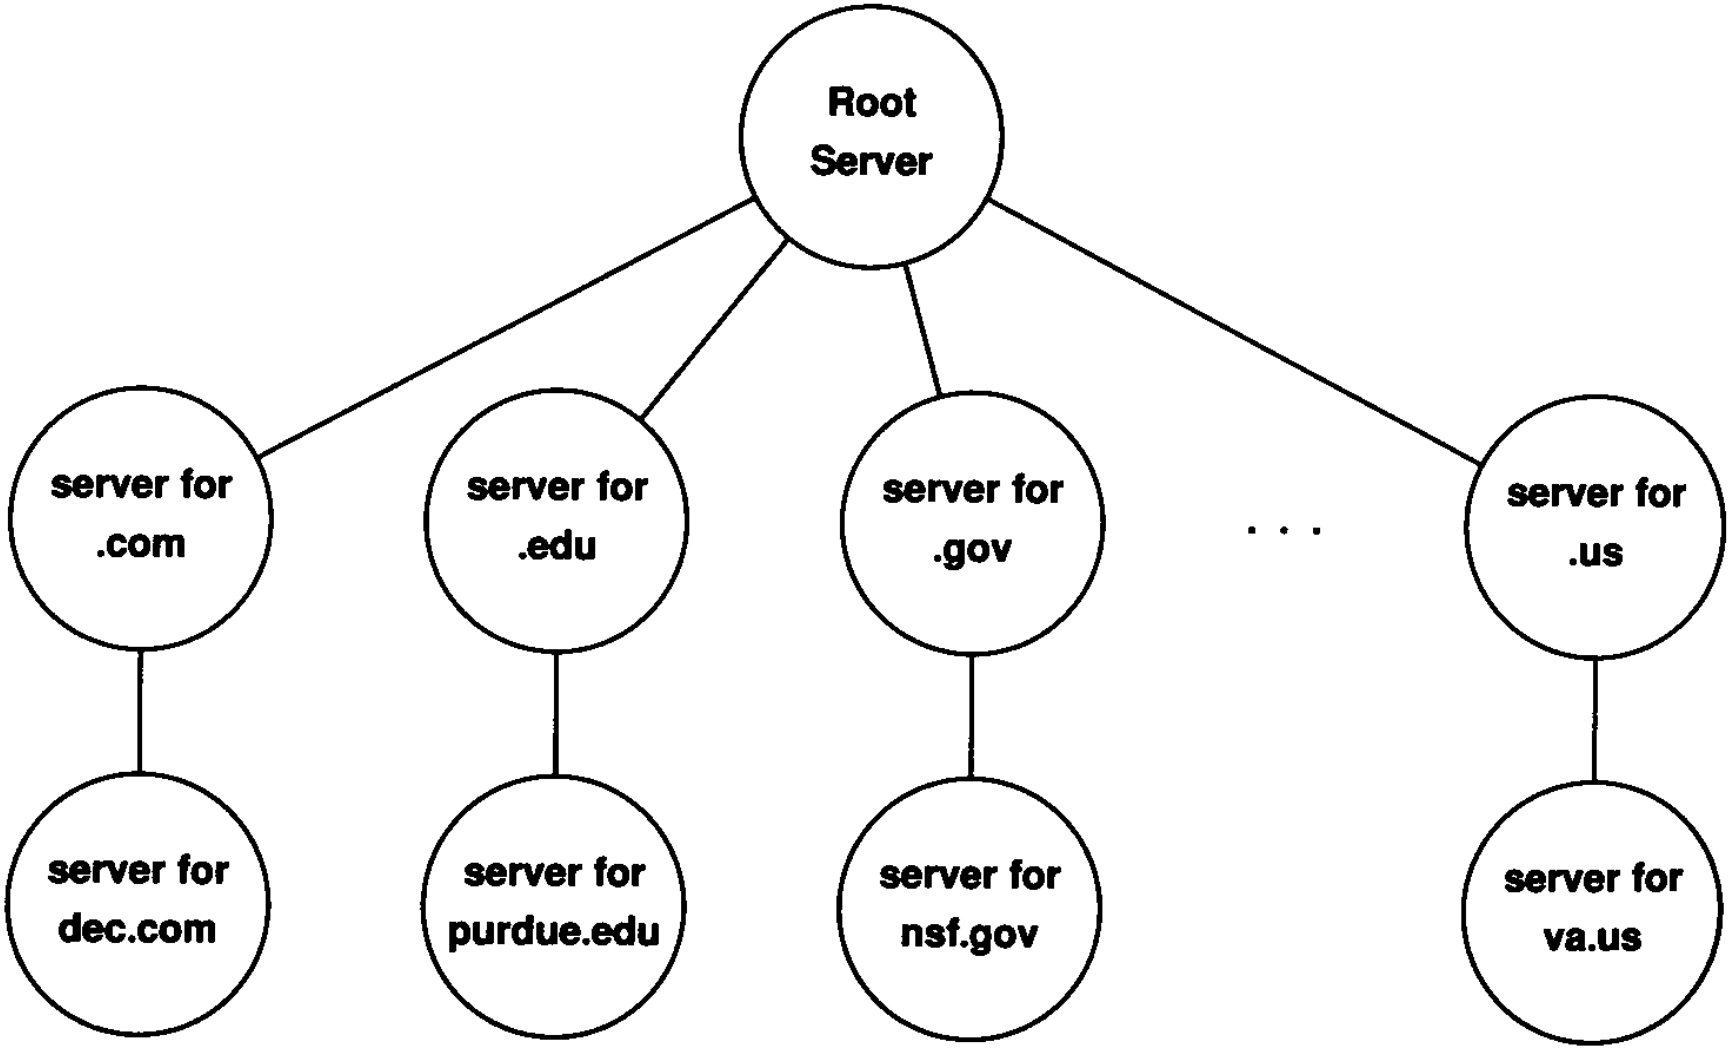

Picture from Comer: Internetworking with CTP/IP, Volume 1, page 392

Architecture: Simplest variant

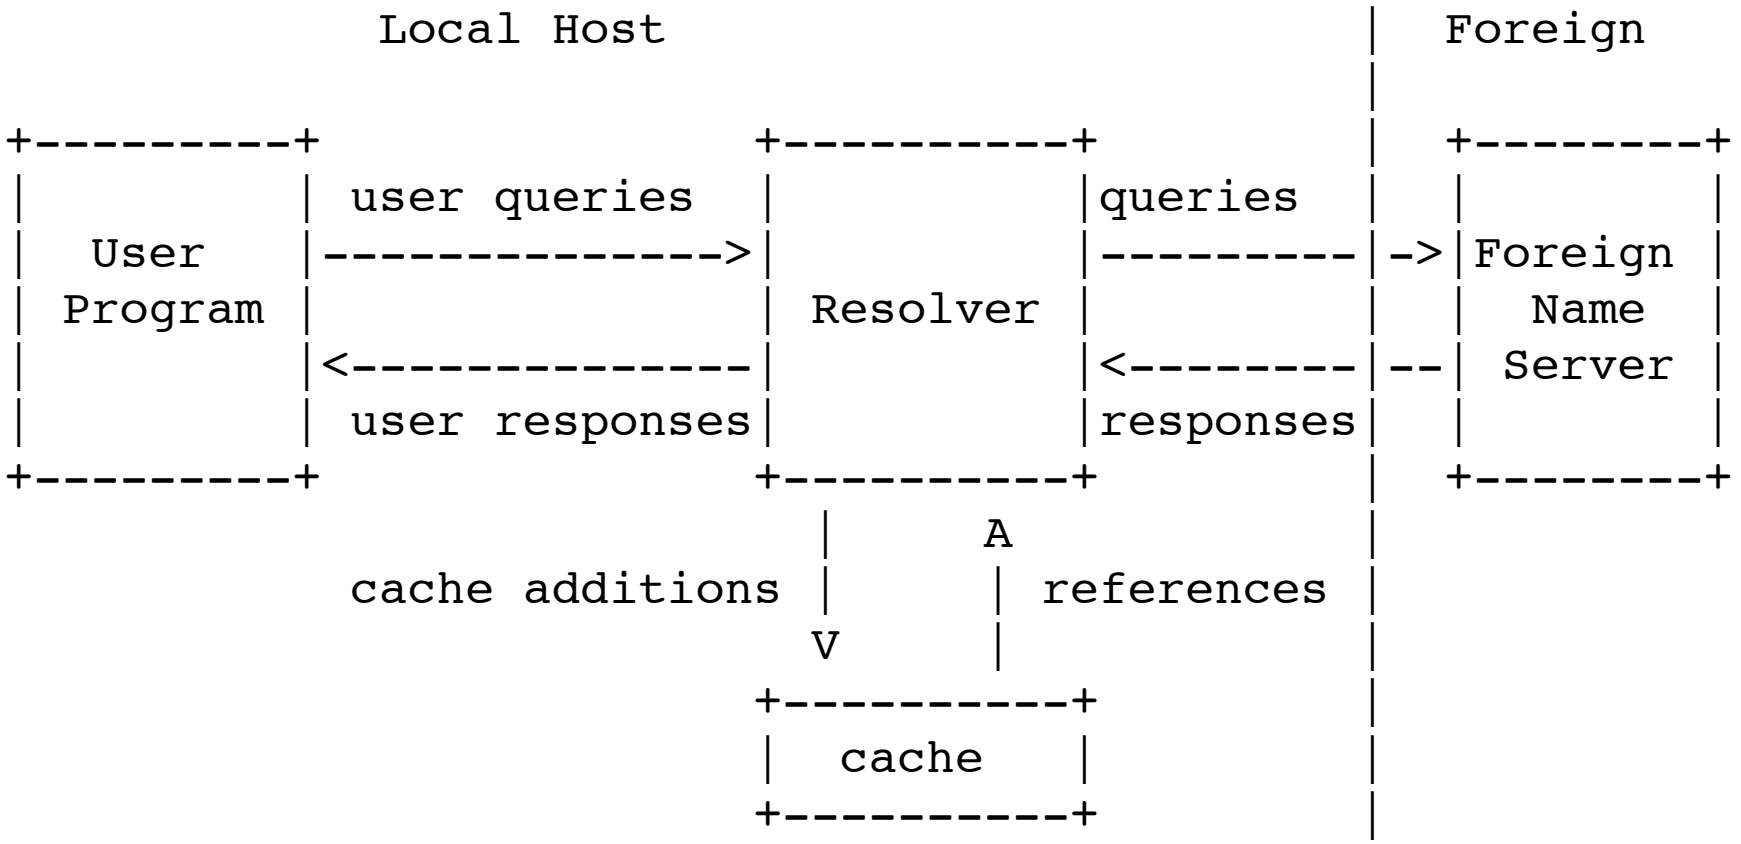

Picture from RFC 1035

Architecture: Maintenance

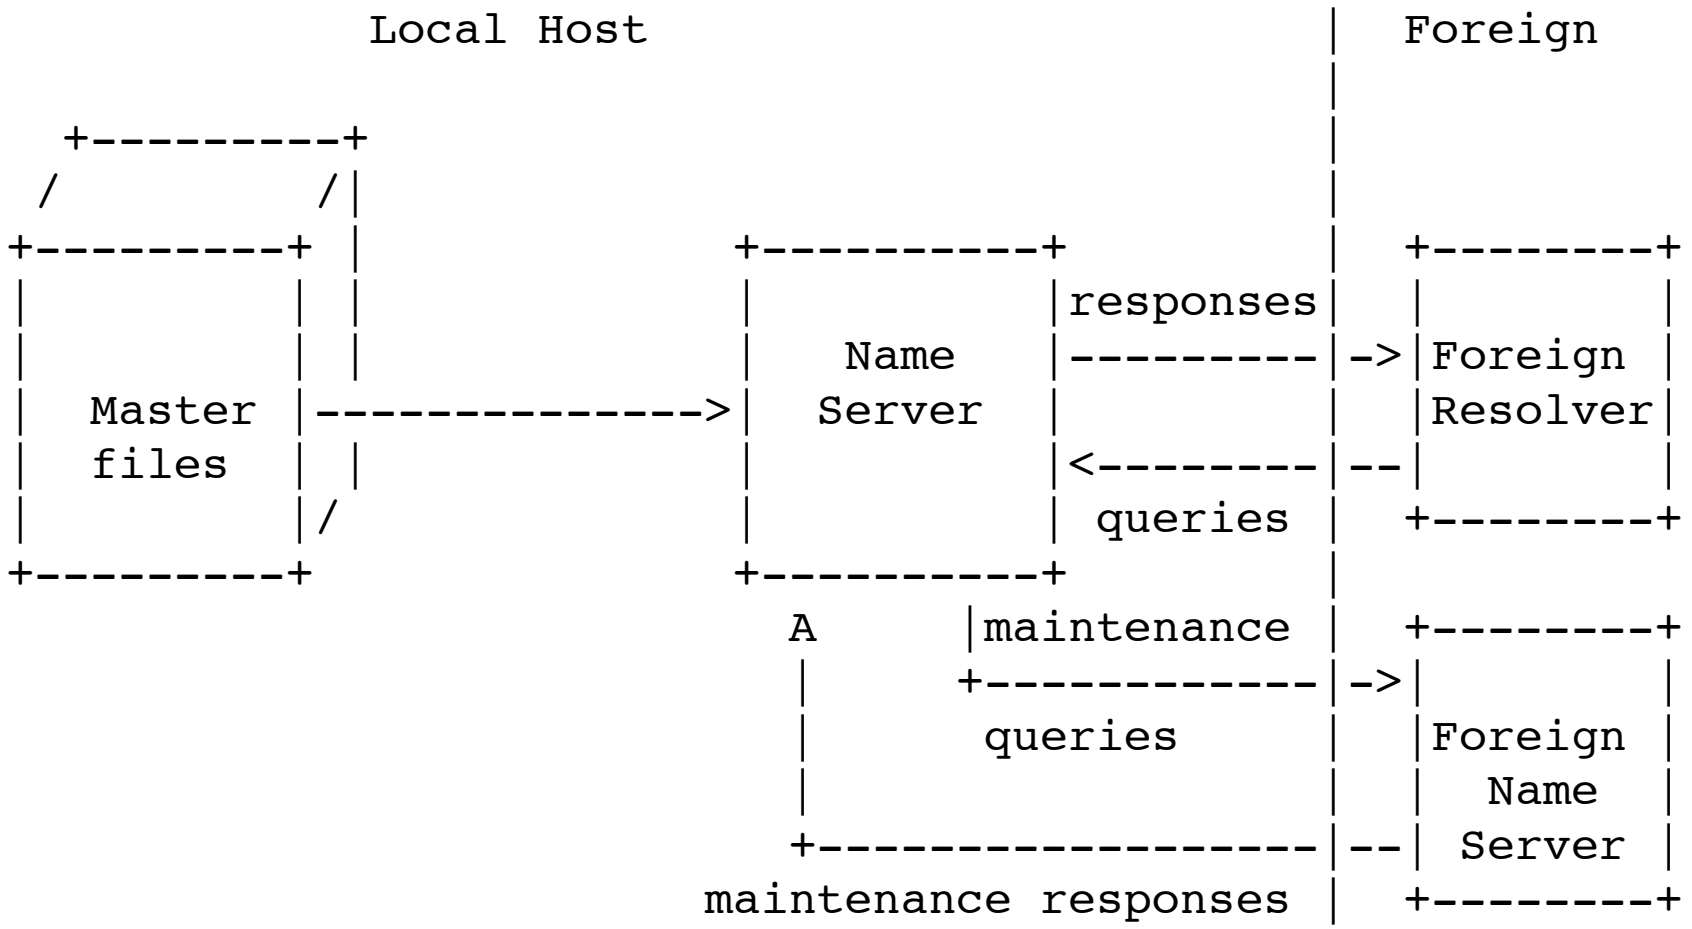

Picture from RFC 1035

Use of multiple name servers

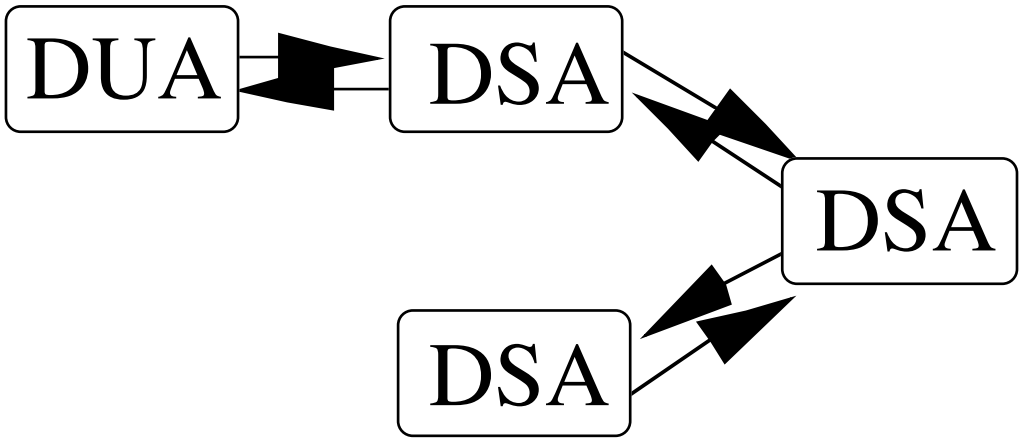

Chaining
(recursive
look up)

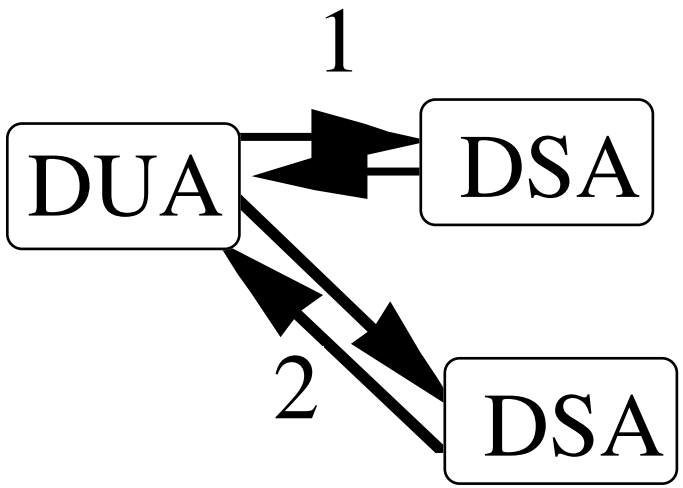

Referral
(iterative
look up)

Resource Records

A	IP address
PTR queries	Pointer record used for pointer (address to domain)
CNAME	Canonical name (for an alias)
HINFO	CPU and operating system of host
MX	Mail exchange record

MX records preference value

```
host -tv mx dsv.su.se
mars.dsv.su.se    86400IN      MX     0
jupiter.dsv.su.se 86400IN      MX     10
sunic.sunet.se    86400IN      MX     20
```

DNS message format

Picture from Comer: Internetworking with TCP/IP page 397

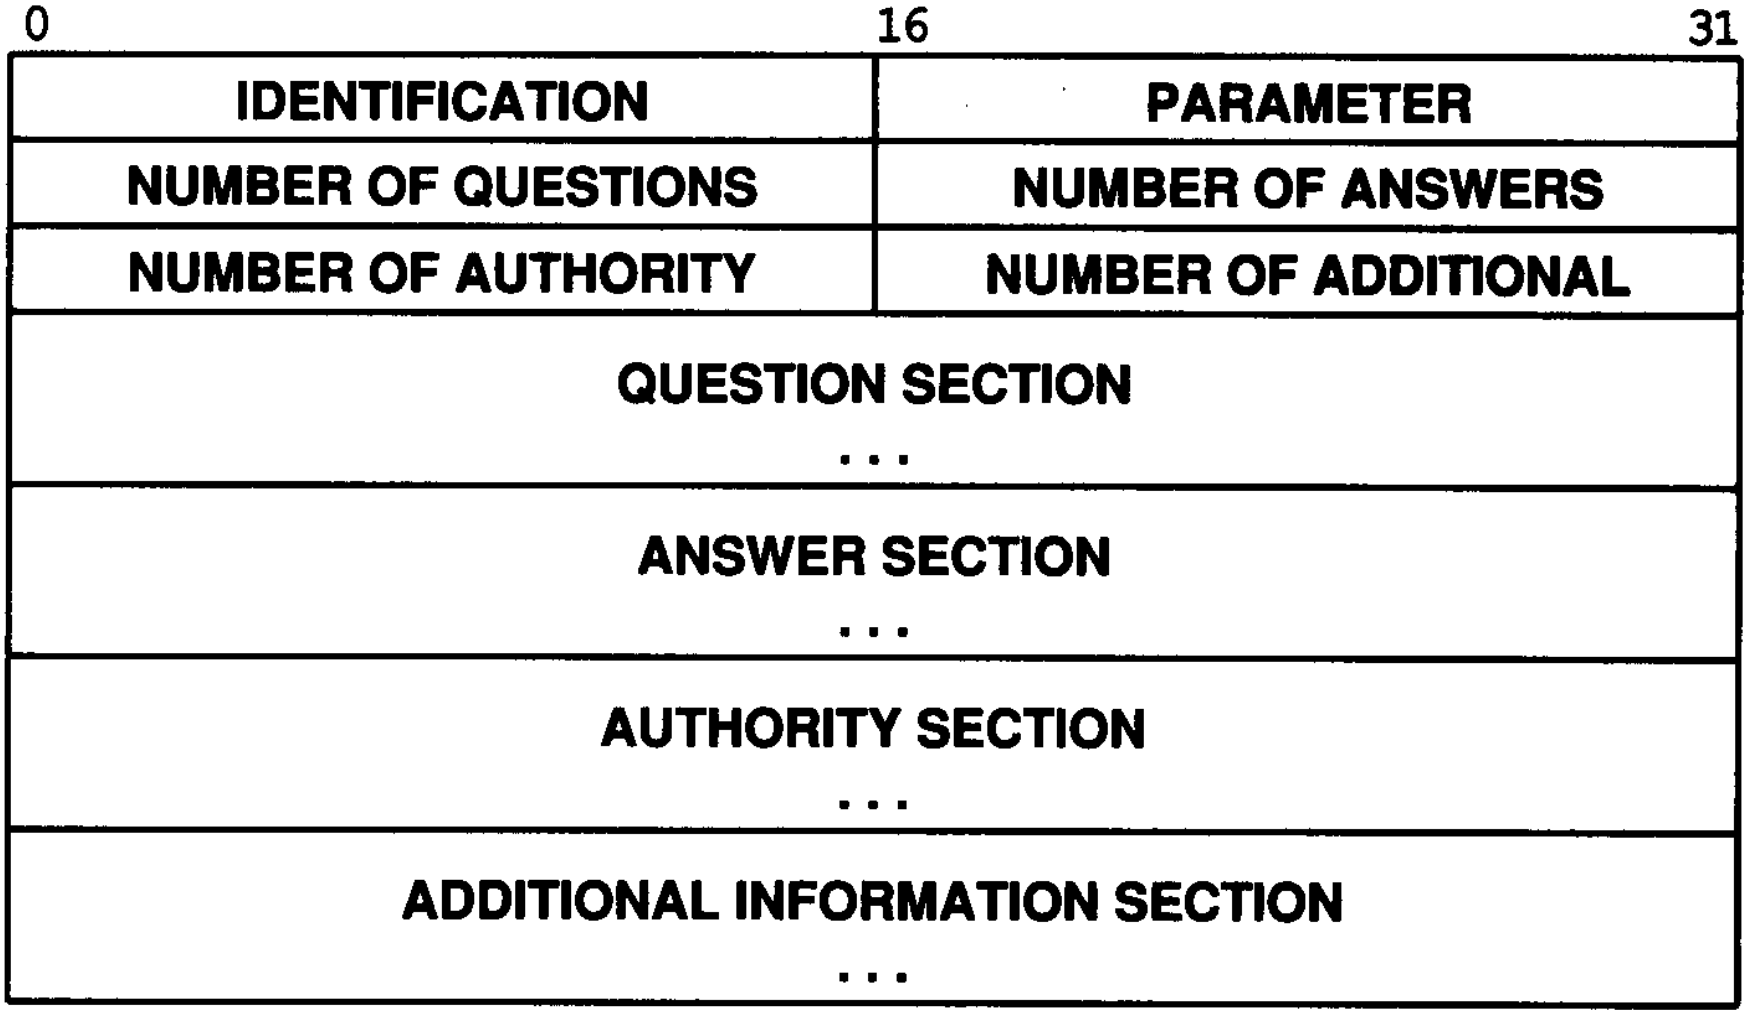

Figure 22.5 Domain name server message format. The question, answer, authority, and additional information sections are variable length.

DNS message format fields

identification	used to match queries to responses
flags	Query or response
	Authorative answer (from authorative server)
	Truncated
	Recursion desired
	Recursion available
	return code (no error or error code)
no. of questions	(must be > 0 for a query)
no. of answers	(must be > 0 for an answer)
question	name being looked up, query type
answer	domain name
	type
	time-to-live
	resource-data-length
	resource data

Representation of a domain name in the DNS protocol

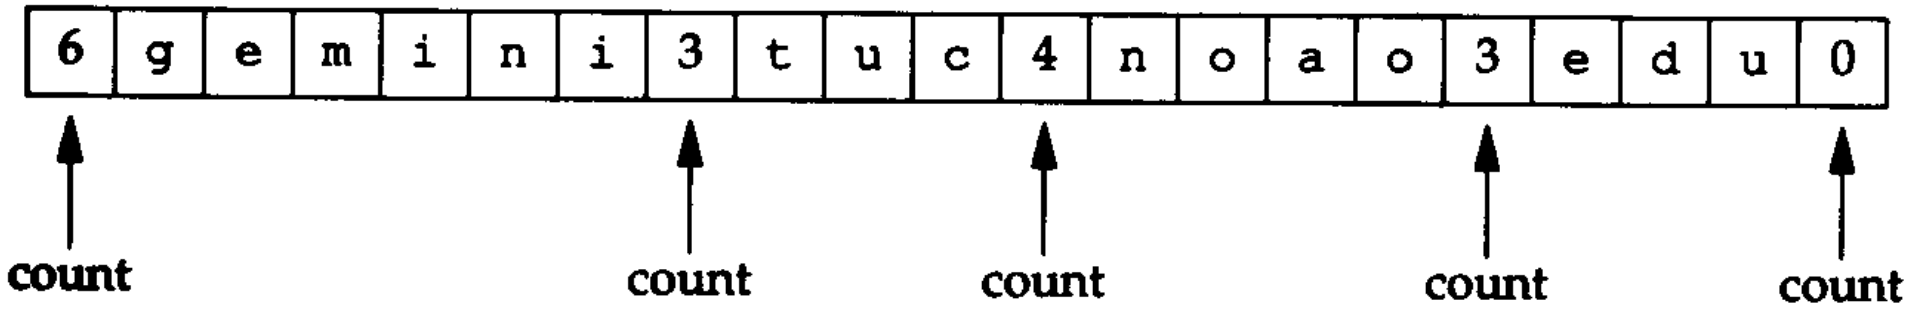

Figure 14.6 Representation of the domain name `gemini.tuc.noao.edu`.

Picture from Stevens: TCP/IP Illustrated, Volume 1, page 193

Weak verification of client's domain

(will only work for clients which have a registered domain address)

If the client has indicated a domain address

Look up this domain address in the DNS

Check that the IP address which the access came from is registered with this domain

If the client has indicated an IP address

Check if the IP address indicated by the client aggrees with the IP address the access came from

If no match, sometimes you might want to make a reverse DNS lookup to find a domain address for that IP address

Make a non-reverse DNS on the found domain; can return several IP addresses

Check that the IP address which the access came from is one of the IP addresses registered with this domain

Basic security services:

Identification (authentication), Authorization, Seals, Signatures, Envelopes (Encryption)

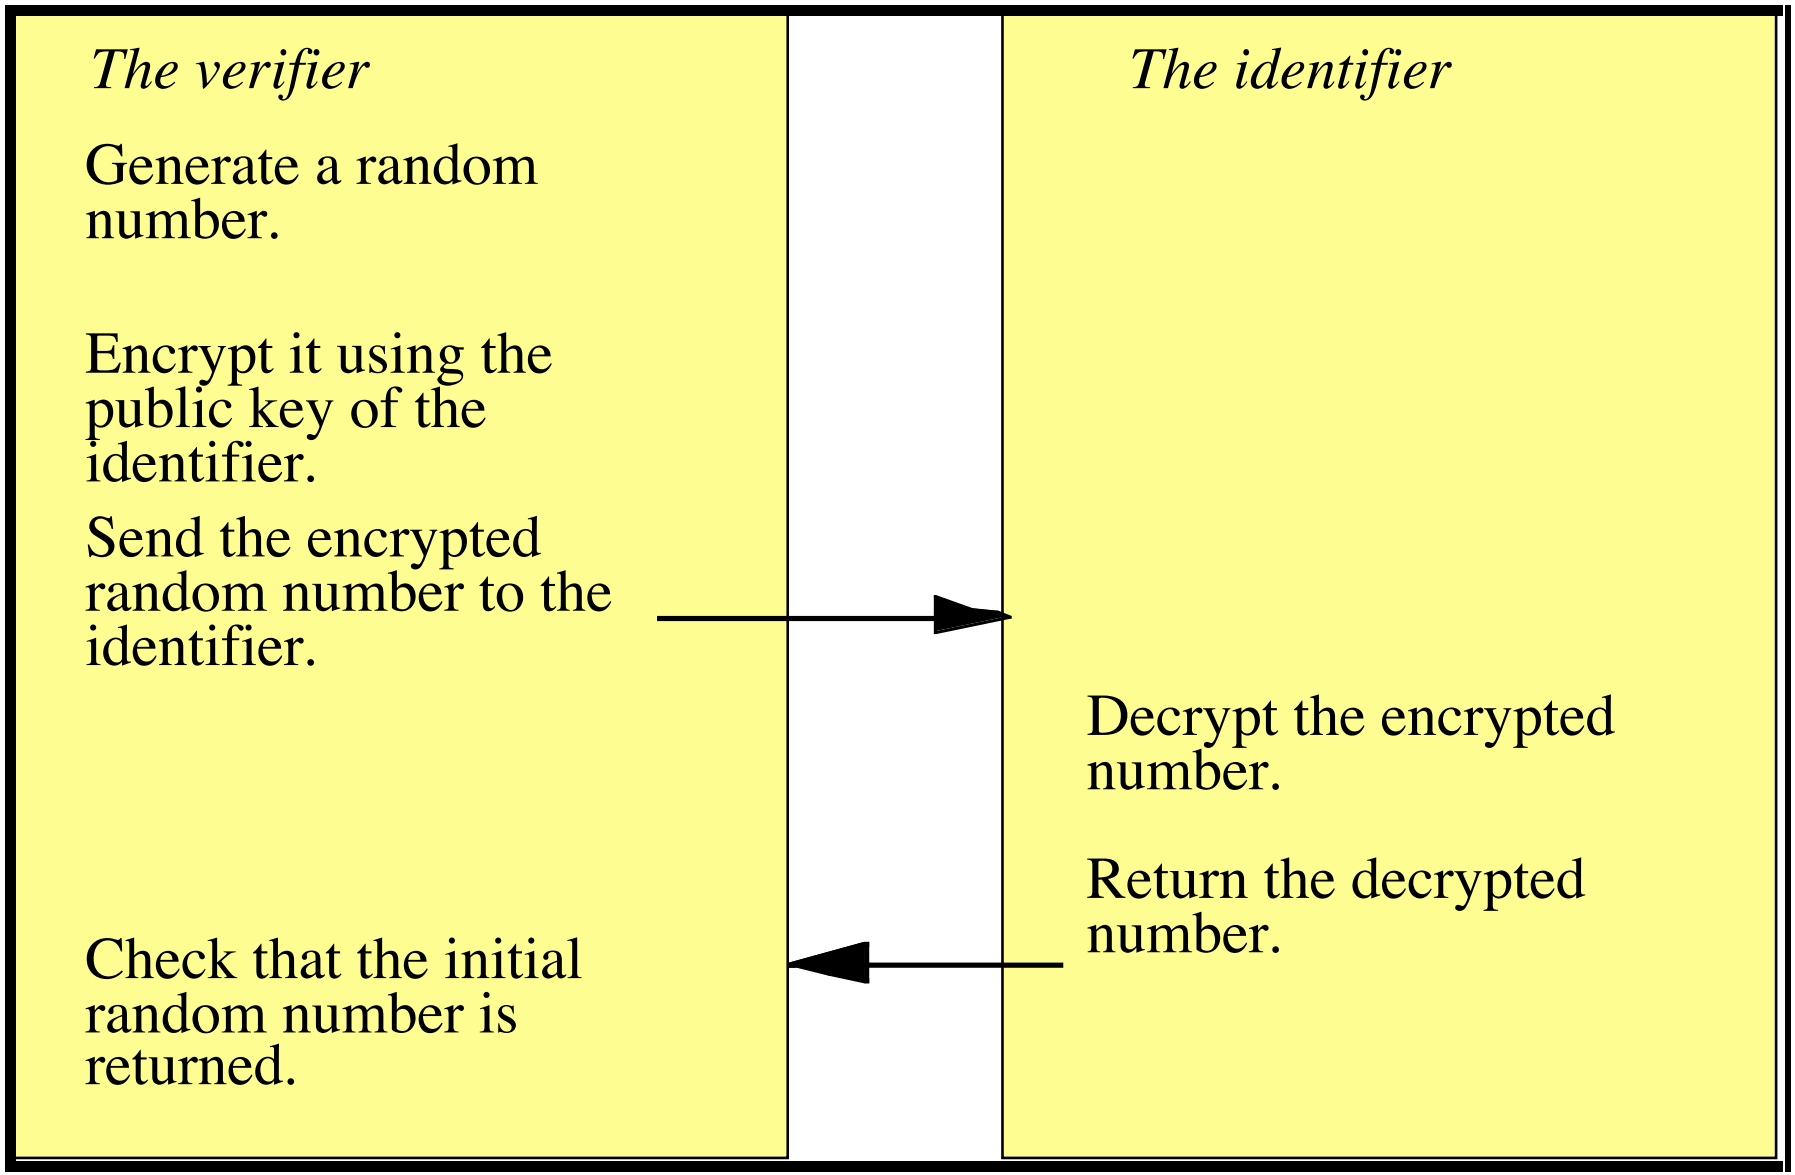

Use of e-mail for authentication

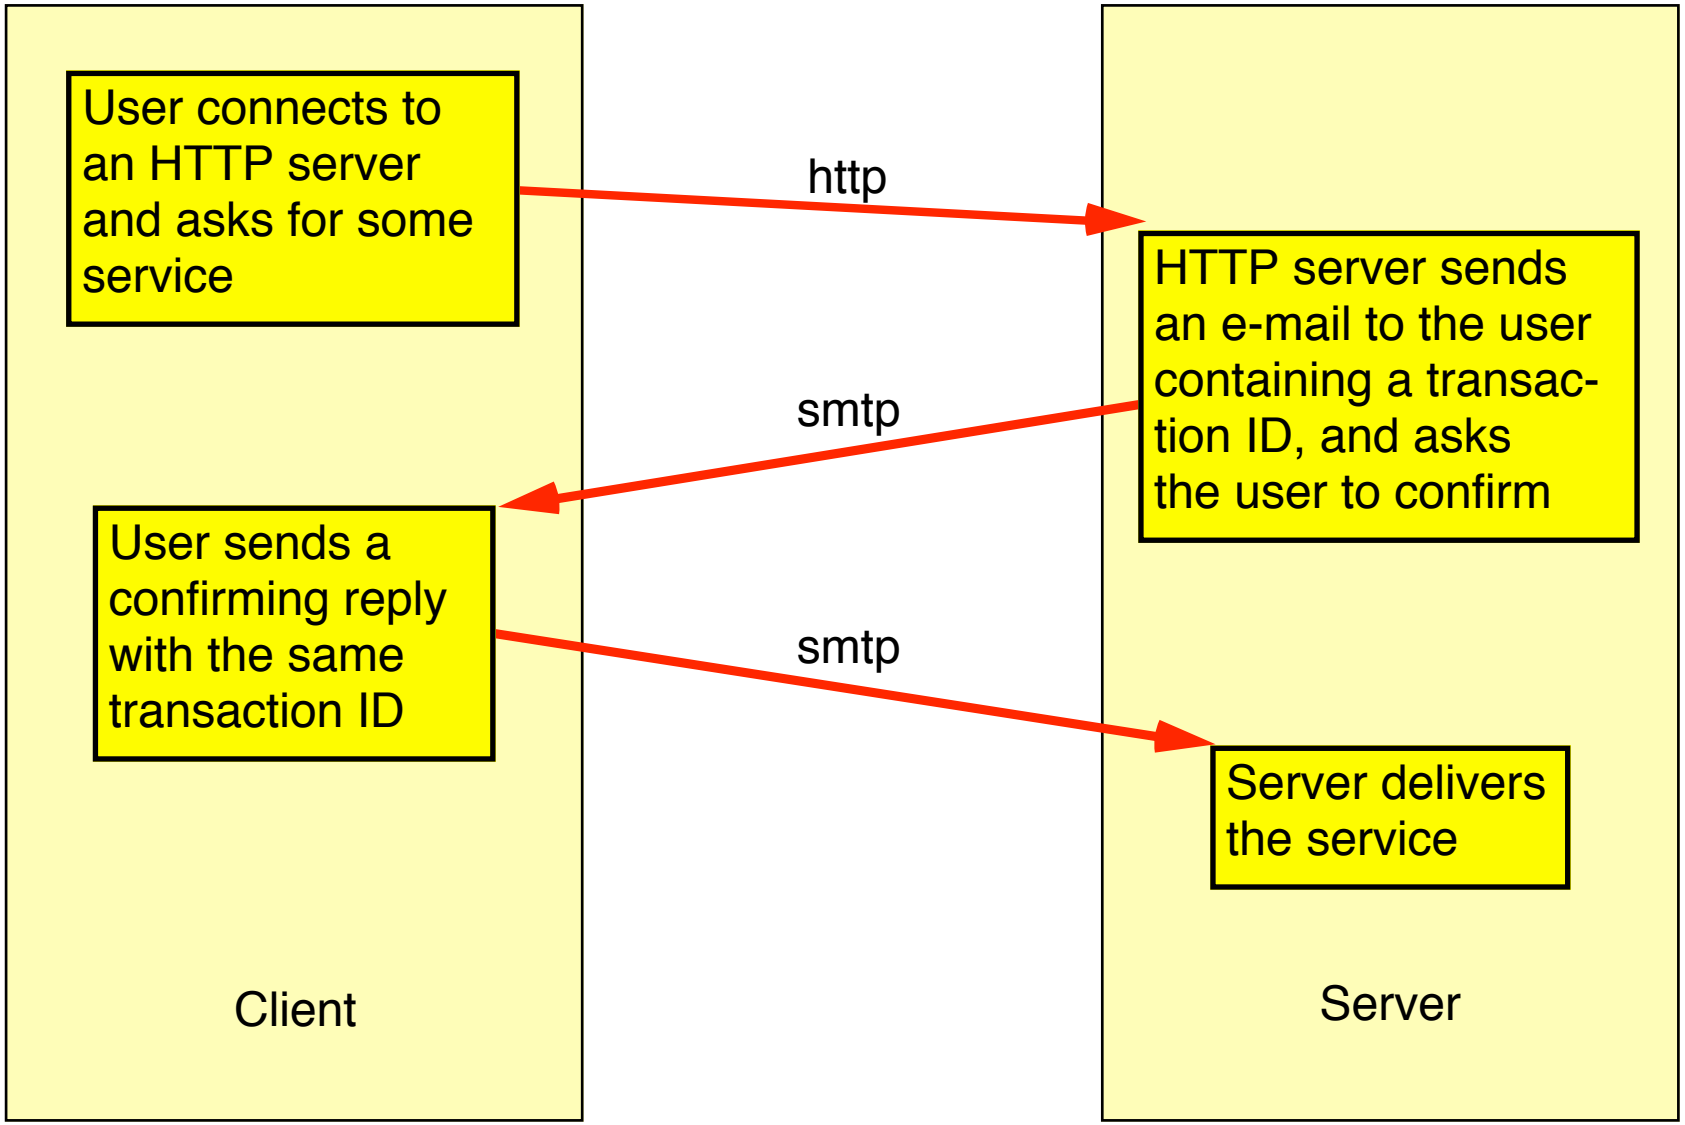

IETF Standards terminology

Classes of standards

- Experimental standard
- Proposed standard
- Draft standard
- Standard
- Historical
- Informational

BCP (Best Current Practive, cannot go through compatibility testing)

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 *(see next page)*.

1. MUST This word, or the terms "REQUIRED" or "SHALL", mean that the definition is an absolute requirement of the specification.

2. MUST NOT This phrase, or the phrase "SHALL NOT", mean that the definition is an absolute prohibition of the specification.

3. SHOULD This word, or the adjective "RECOMMENDED", mean that there may exist valid reasons in particular circumstances to ignore a particular item, but the full implications must be understood and carefully weighed before choosing a different course.

4. SHOULD NOT This phrase, or the phrase "NOT RECOMMENDED" mean that there may exist valid reasons in particular circumstances when the particular behavior is acceptable or even useful, but the full implications should be understood and the case carefully weighed before implementing any behavior described with this label.

5. MAY This word, or the adjective "OPTIONAL", mean that an item is truly optional. One vendor may choose to include the item because a particular marketplace requires it or because the vendor feels that it enhances the product while another vendor may omit the same item. An implementation which does not include a particular option MUST be prepared to interoperate with another implementation which does include the option, though perhaps with reduced functionality. In the same vein an implementation which does include a particular option MUST be prepared to interoperate with another implementation which does not include the option (except, of course, for the feature the option provides.)