Traditional Applications

- World Wide Web
  - The World Wide Web has been so successful and has made the Internet accessible to so many people that sometimes it seems to be synonymous with the Internet.
  - In fact, the design of the system that became the Web started around 1989, long after the Internet had become a widely deployed system.
  - The original goal of the Web was to find a way to organize and retrieve information, drawing on ideas about hypertext—interlinked documents—that had been around since at least the 1960s.

Traditional Applications

- World Wide Web
  - The core idea of hypertext is that one document can link to another document, and the protocol (HTTP) and document language (HTML) were designed to meet that goal.
  - One helpful way to think of the Web is as a set of cooperating clients and servers, all of whom speak the same language: HTTP.
  - Most people are exposed to the Web through a graphical client program, or Web browser, like Safari, Chrome, Firefox or Internet Explorer.

Traditional Applications

- World Wide Web
  - When you ask your browser to view a page, your browser (the client) fetches the page from the server using HTTP running over TCP.
  - Like SMTP, HTTP is a text oriented protocol.
  - At its core, HTTP is a request/response protocol, where every message has the general form

    START_LINE <CRLF>

    MESSAGE_HEADER <CRLF>

    <CRLF>

    MESSAGE_BODY <CRLF>
  - where as before,<CRLF>stands for carriage-return-line-feed.The first line (START LINE)
  - indicates whether this is a request message or a response message.

Traditional Applications

- World Wide Web
  - Request Messages
    - The first line of an HTTP request message specifies three things: the operation to be performed, the Web page the operation should be performed on, and the version of HTTP being used.
    - Although HTTP defines a wide assortment of possible request operations—including “write” operations that allow a Web page to be posted on a server—the two most common operations are GET (fetch the specified Web page) and HEAD (fetch status information about the specified Web page).

Traditional Applications

- World Wide Web
  - Request Messages

Operation	Description
OPTIONS	Request information about available options
GET	Retrieve document identified in URL
HEAD	Retrieve metainformation about document identified in URL
POST	Give information (e.g., annotation) to server
PUT	Store document under specified URL
DELETE	Delete specified URL
TRACE	Loopback request message
CONNECT	For use by proxies

HTTP request operations

Traditional Applications

- World Wide Web
  - Response Messages
    - Like request messages, response messages begin with a single START LINE.
    - In this case, the line specifies the version of HTTP being used, a three-digit code indicating whether or not the request was successful, and a text string giving the reason for the response.

Traditional Applications

- World Wide Web
  - Response Messages

Code	Type	Example Reasons
1xx	Informational	request received, continuing process
2xx	Success	action successfully received, understood, and accepted
3xx	Redirection	further action must be taken to complete the request
4xx	Client Error	request contains bad syntax or cannot be fulfilled
5xx	Server Error	server failed to fulfill an apparently valid request

Five types of HTTP result codes

Traditional Applications

- World Wide Web
  - TCP Connections
    - The original version of HTTP (1.0) established a separate TCP connection for each data item retrieved from the server.
    - It’s not too hard to see how this was a very inefficient mechanism: connection setup and teardown messages had to be exchanged between the client and server even if all the client wanted to do was verify that it had the most recent copy of a page.
    - Thus, retrieving a page that included some text and a dozen icons or other small graphics would result in 13 separate TCP connections being established and closed.

Traditional Applications

- World Wide Web
  - TCP Connections
    - To overcome this situation, HTTP version 1.1 introduced *persistent connections*— the client and server can exchange multiple request/response messages over the same TCP connection.
    - Persistent connections have many advantages.
      - First, they obviously eliminate the connection setup overhead, thereby reducing the load on the server, the load on the network caused by the additional TCP packets, and the delay perceived by the user.
      - Second, because a client can send multiple request messages down a single TCP connection, TCP's congestion window mechanism is able to operate more efficiently.
        - This is because it's not necessary to go through the slow start phase for each page.

Traditional Applications

- World Wide Web
  - TCP Connections

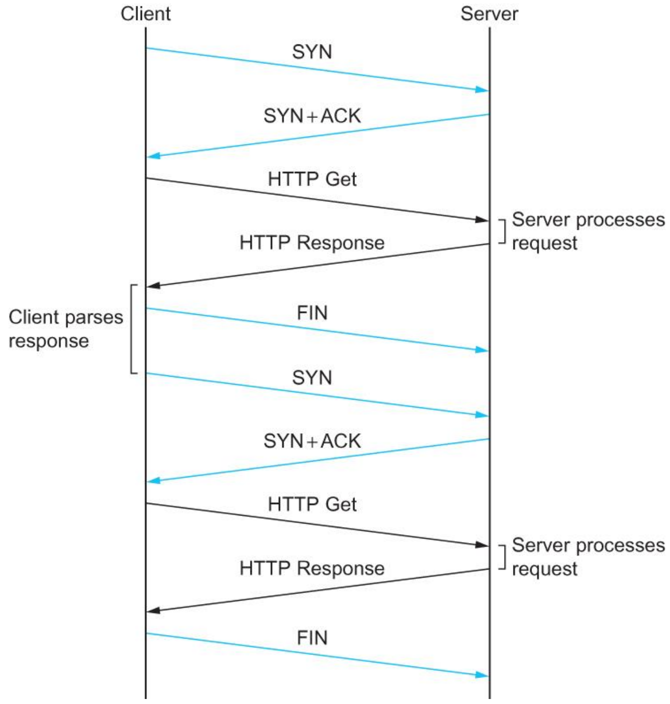

HTTP 1.0 behavior

Traditional Applications

- World Wide Web
  - TCP Connections

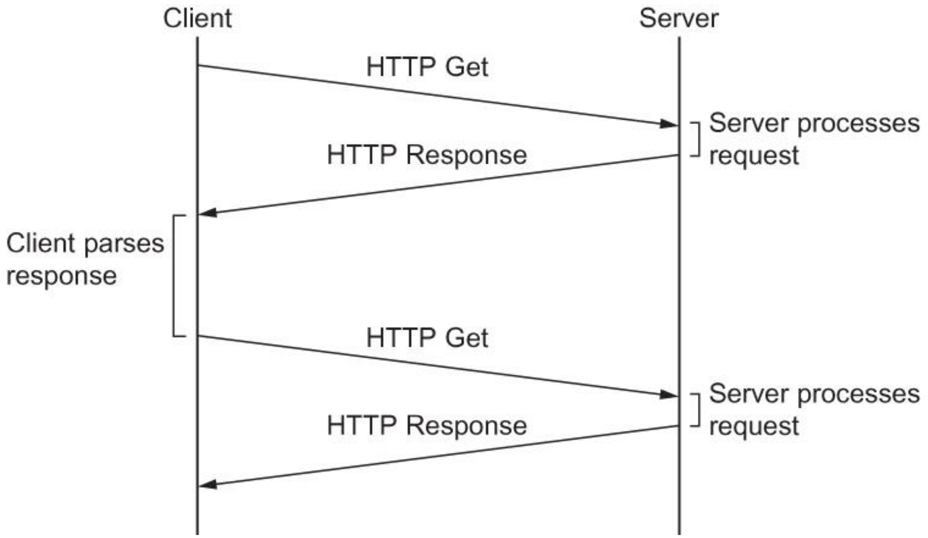

HTTP 1.1 behavior with persistent connections

Traditional Applications

- World Wide Web
  - Caching
    - One of the most active areas of research (and entrepreneurship) in the Internet today is how to effectively cache Web pages.
    - Caching has many benefits. From the client's perspective, a page that can be retrieved from a nearby cache can be displayed much more quickly than if it has to be fetched from across the world.
    - From the server's perspective, having a cache intercept and satisfy a request reduces the load on the server.

Traditional Applications

- World Wide Web
  - Caching
    - Caching can be implemented in many different places. For example, a user's browser can cache recently accessed pages, and simply display the cached copy if the user visits the same page again.
    - As another example, a site can support a single site-wide cache.
    - This allows users to take advantage of pages previously downloaded by other users.
    - Closer to the middle of the Internet, ISPs can cache pages.
    - Note that in the second case, the users within the site most likely know what machine is caching pages on behalf of the site, and they configure their browsers to connect directly to the caching host. This node is sometimes called a *proxy*

Multimedia Applications

- SIP

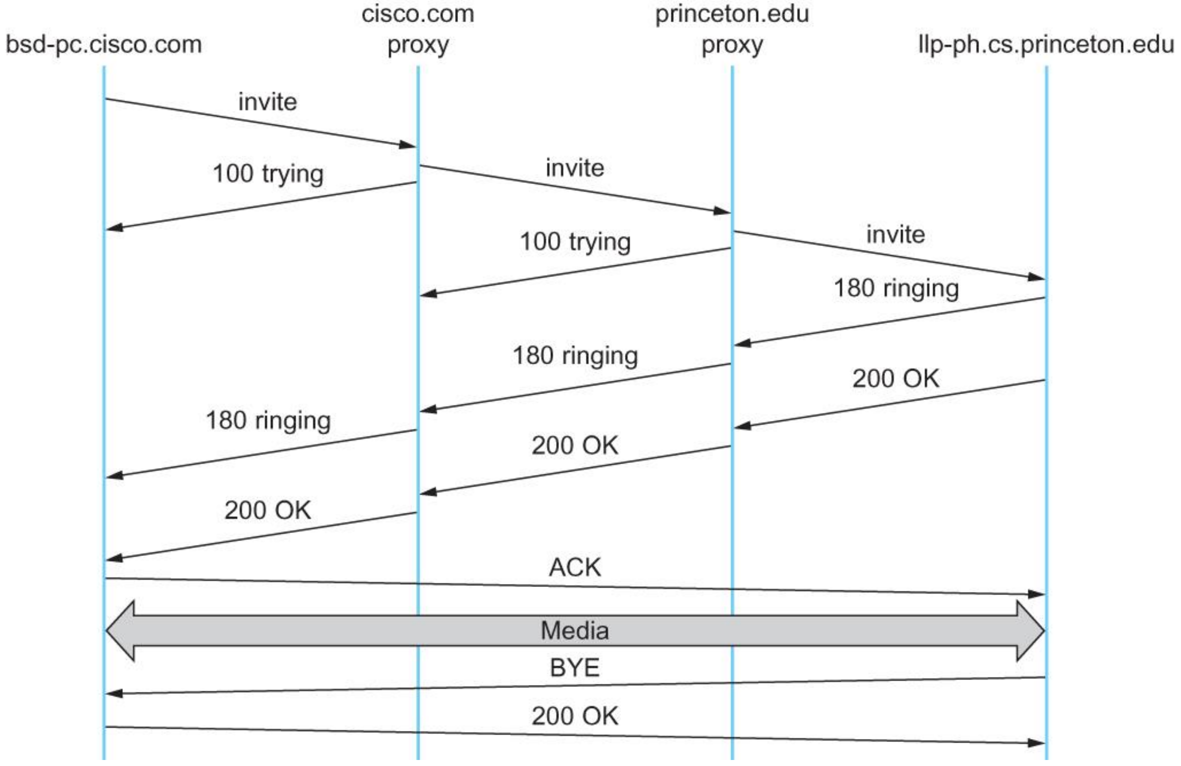

Message flow for a basic SIP session

Multimedia Applications

- H.323
  - The ITU has also been very active in the call control area, which is not surprising given its relevance to telephony, the traditional realm of that body.
  - Fortunately, there has been considerable coordination between the IETF and the ITU in this instance, so that the various protocols are somewhat interoperable.
  - The major ITU recommendation for multimedia communication over packet networks is known as H.323, which ties together many other recommendations, including H.225 for call control.
  - The full set of recommendations covered by H.323 runs to many hundreds of pages, and the protocol is known for its complexity

Multimedia Applications

- H.323

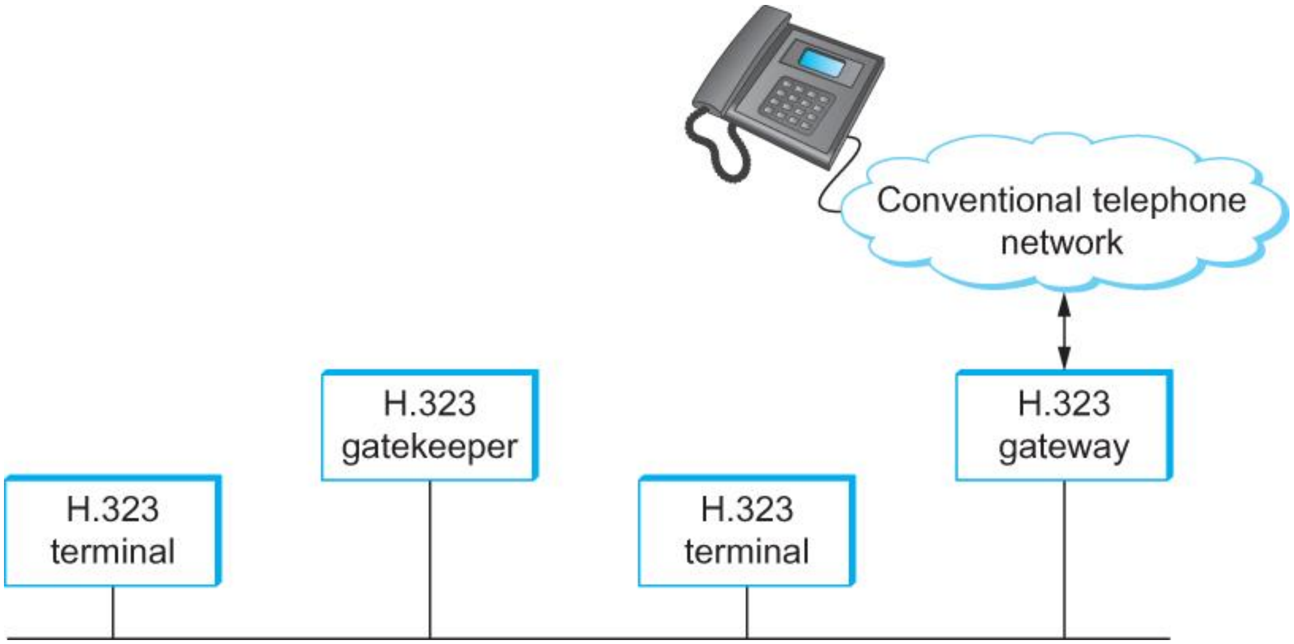

Devices in an H.323 network.

Infrastructure Services

- Name Service (DNS)
  - In most of this book, we have been using addresses to identify hosts.
  - While perfectly suited for processing by routers, addresses are not exactly user-friendly.
  - It is for this reason that a unique *name is also typically assigned to each host in a network.*
  - Host names differ from host addresses in two important ways.
    - First, they are usually of variable length and mnemonic, thereby making them easier for humans to remember.
    - Second, names typically contain no information that helps the network locate (route packets toward) the host.

Infrastructure Services

- Name Service (DNS)
  - We first introduce some basic terminology.
    - First, a *name space defines the set of possible names.*
      - A name space can be either *flat (names are not divisible into components), or it can* be *hierarchical (Unix file names are an obvious example).*
    - *Second, the naming system* maintains a collection of *bindings of names to values. The value can be anything we* want the naming system to return when presented with a name; in many cases it is an address.
    - Finally, a *resolution mechanism is a procedure that, when invoked with a* name, returns the corresponding value. A *name server is a specific implementation of a* resolution mechanism that is available on a network and that can be queried by sending it a message.

Infrastructure Services

- Name Service (DNS)

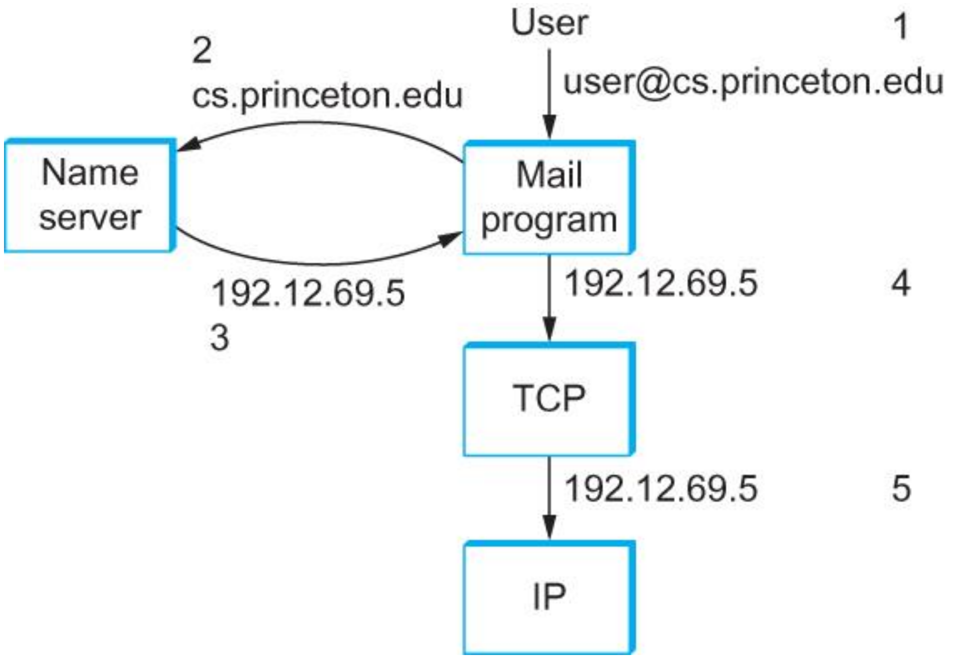

Names translated into addresses, where the numbers 1–5 show the sequence of steps in the process

Infrastructure Services

- Domain Hierarchy
  - DNS implements a hierarchical name space for Internet objects. Unlike Unix file names, which are processed from left to right with the naming components separated with slashes, DNS names are processed from right to left and use periods as the separator.
  - Like the Unix file hierarchy, the DNS hierarchy can be visualized as a tree, where each node in the tree corresponds to a domain, and the leaves in the tree correspond to the hosts being named.

Infrastructure Services

- Domain Hierarchy

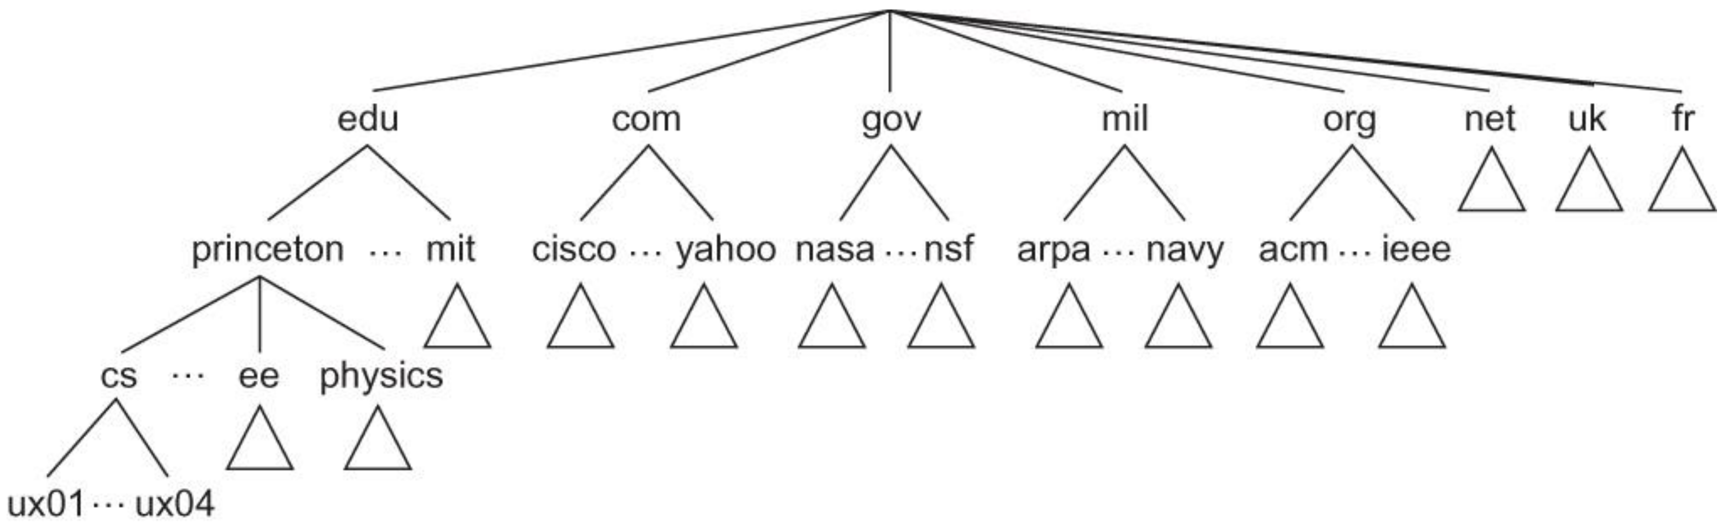

Example of a domain hierarchy

Infrastructure Services

- Name Servers
  - The complete domain name hierarchy exists only in the abstract.
  - We now turn our attention to the question of how this hierarchy is actually implemented.
  - The first step is to partition the hierarchy into subtrees called *zones.*
  - Each zone can be thought of as corresponding to some administrative authority that is responsible for that portion of the hierarchy.
  - For example, the top level of the hierarchy forms a zone that is managed by the Internet Corporation for Assigned Names and Numbers (ICANN).

Infrastructure Services

- Name Servers
  - Each name server implements the zone information as a collection of *resource records.*
  - In essence, a resource record is a name-to-value binding, or more specifically a 5-tuple that contains the following fields:
  - <Name, Value, Type, Class, TTL >

Infrastructure Services

- Name Servers
  - The Name and Value fields are exactly what you would expect, while the Type field specifies how the Value should be interpreted.
    - For example, Type = A indicates that the Value is an IP address. Thus, A records implement the name-to-address mapping we have been assuming. Other record types include
    - NS: The Value field gives the domain name for a host that is running a name server that knows how to resolve names within the specified domain.
    - CNAME: The Value field gives the canonical name for a particular host; it is used to define aliases.
    - MX: The Value field gives the domain name for a host that is running a mail server that accepts messages for the specified domain.

Infrastructure Services

- Name Resolution

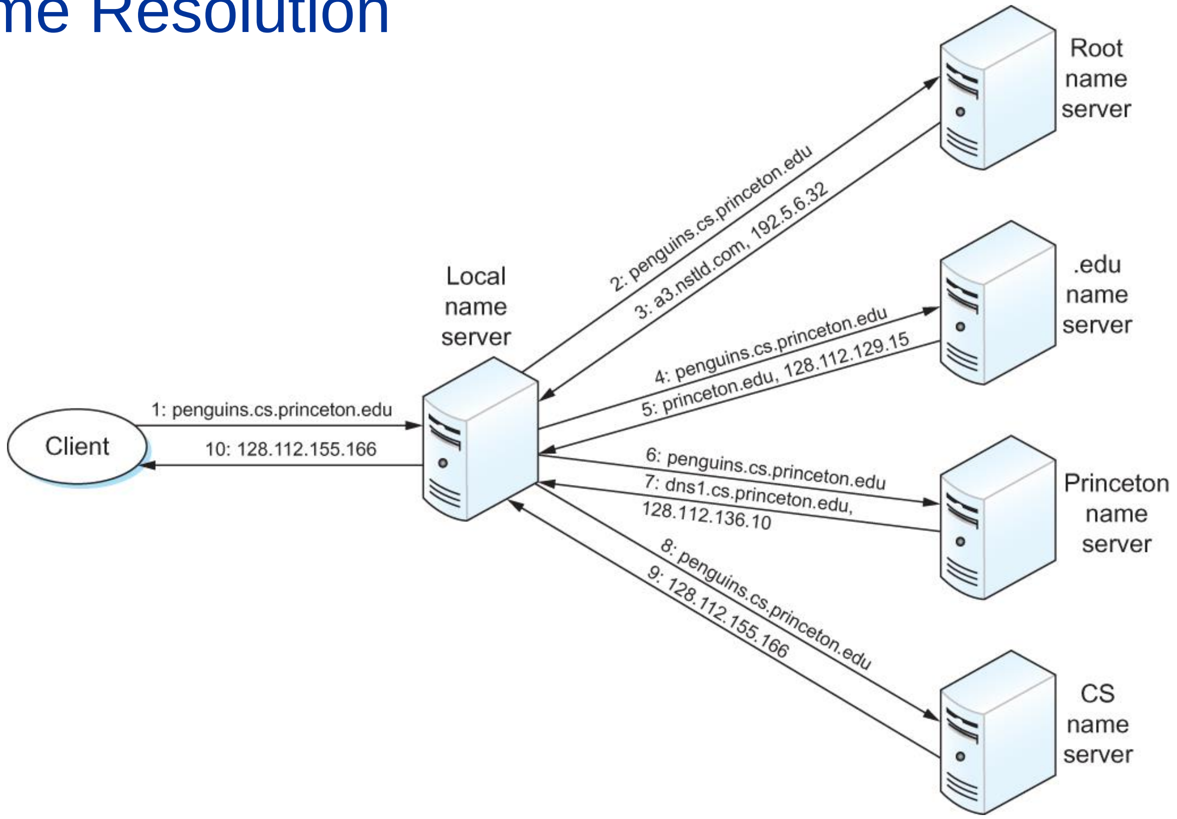

Name resolution in practice, where the numbers 1–10 show the sequence of steps in the process.

Infrastructure Services

- Network Management
  - The most widely used protocol for this purpose is the Simple Network Management Protocol (SNMP).
  - SNMP is essentially a specialized request/reply protocol that supports two kinds of request messages: GET and SET.
  - The former is used to retrieve a piece of state from some node, and the latter is used to store a new piece of state in some node.
  - SNMP is used in the obvious way.
    - A system administrator interacts with a client program that displays information about the network.
    - This client program usually has a graphical interface. Whenever the administrator selects a certain piece of information that he or she wants to see, the client program uses SNMP to request that information from the node in question. (SNMP runs on top of UDP.)
    - An SNMP server running on that node receives the request, locates the appropriate piece of information, and returns it to the client program, which then displays it to the user.

Overlay Network

- Content Distribution Network (CDN)
  - The idea of a CDN is to geographically distribute a collection of *server surrogates* that cache pages normally maintained in some set of *backend servers*
    - Akamai operates what is probably the best-known CDN.
  - Thus, rather than have millions of users wait forever to contact www.cnn.com when a big news story breaks—such a situation is known as a *flash crowd*—*it is possible to spread this load* across many servers.
  - Moreover, rather than having to traverse multiple ISPs to reach www.cnn.com, if these surrogate servers happen to be spread across all the backbone ISPs, then it should be possible to reach one without having to cross a peering point.

Overlay Network

- Content Distribution Network (CDN)

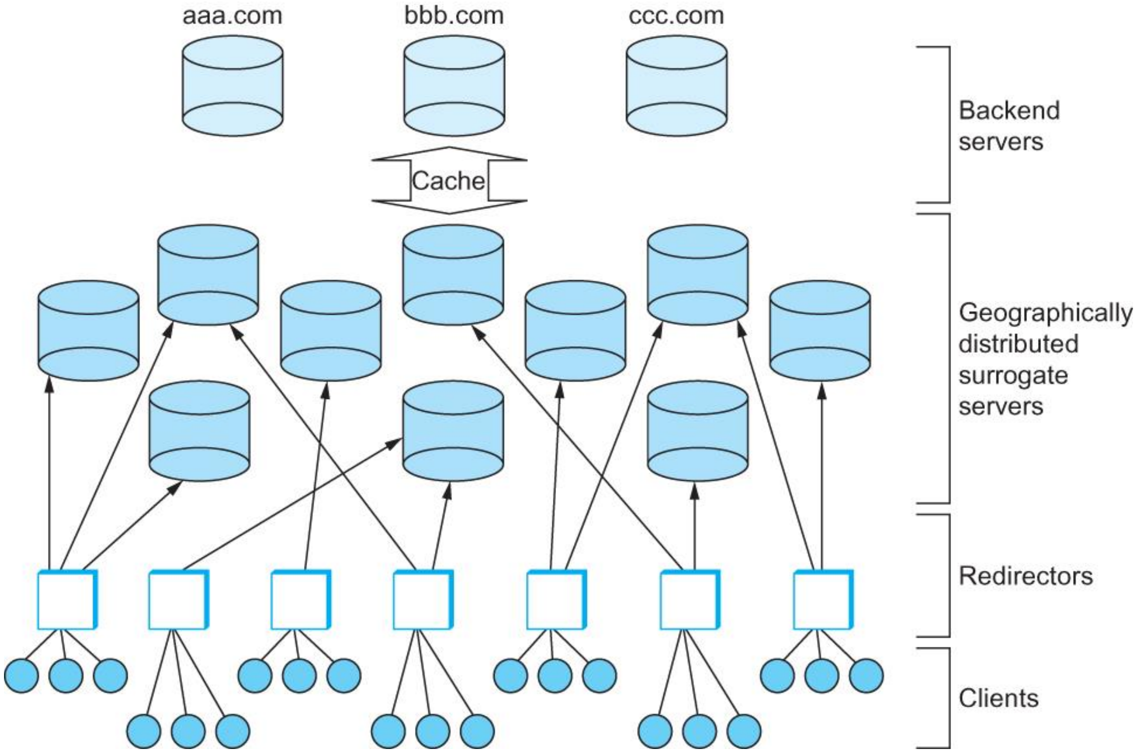

Components in a Content Distribution Network (CDN).