Outline

- Applications
  - Central Server
  - Hierarchical
  - Peer-to-peer

Networked distributed system architectures

Central Server based

- Web servers

Hierarchical Services

- Domain Name System – DNS

Peer-to-Peer Systems

- Napster, gnutella

Central Server based

- A central server provides service
  - Reliability and fault tolerance
    - If server shuts down, then no service
  - Scalability
    - Performance bottle neck
    - E.g. if everyone accesses Microsoft.com from the east coast (new release of web browser), accesses to Yahoo.com in California might be slow because we share the same link from east coast till Utah (say)
  - Easy to deploy, administer

Domain Name Service (DNS)

- Provides Internet domain name to IP address translation
  - Domain name translation (uga.edu)
  - Hostname translation (greenhouse.cs.uga.edu)
  - Service location (MX records, mail service for UGA)

```
$ nslookup –query=mx home.com
home.com        preference = 100, mail exchanger = mx-d-rwc.mail.home.com
home.com        preference = 150, mail exchanger = mx-a-rwc.mail.home.com
home.com        preference = 100, mail exchanger = mx-c-tx.mail.home.com
home.com        preference = 150, mail exchanger = mx-a-tx.mail.home.com
home.com        preference = 175, mail exchanger = mx-a-va.mail.home.com
home.com        preference = 50, mail exchanger = mx-rr.home.com
```

- Hierarchical
  - Decentralized administration of name space
  - Hierarchy of authority and trust

Domain Naming System Hierarchy

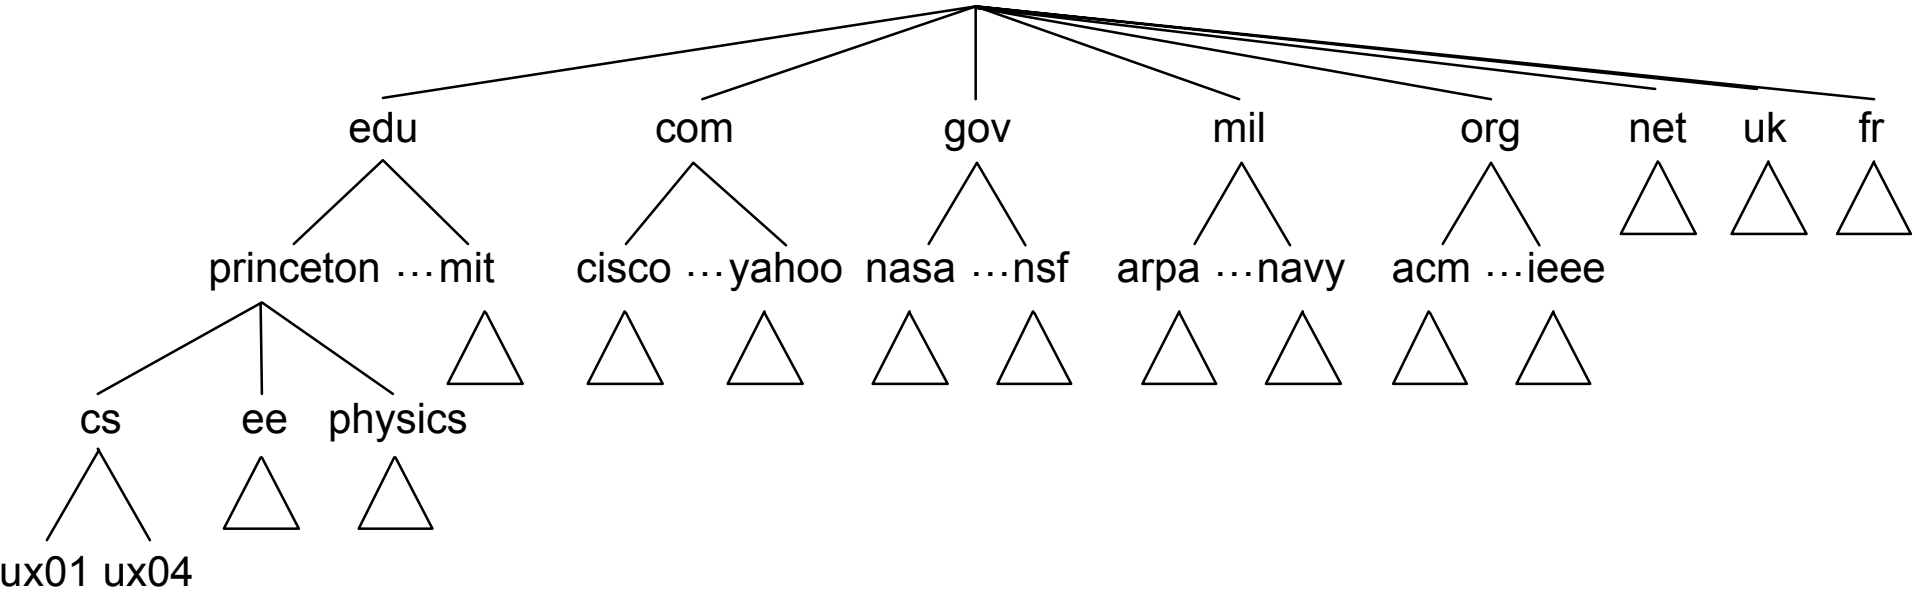

DNS hierarchy

DNS name space is *hierarchical*:

- fully qualified names are “little endian”
  - scalability
  - decentralized administration
  - domains are naming *contexts*

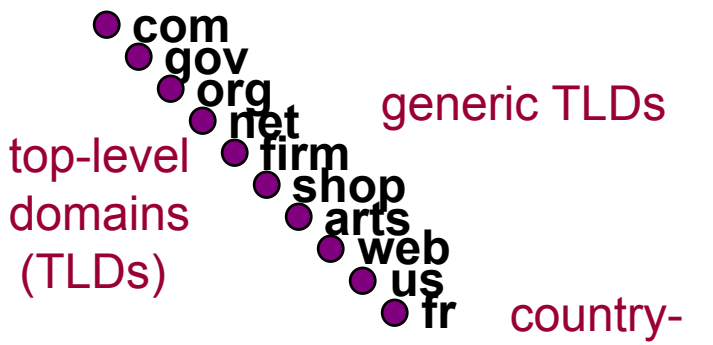

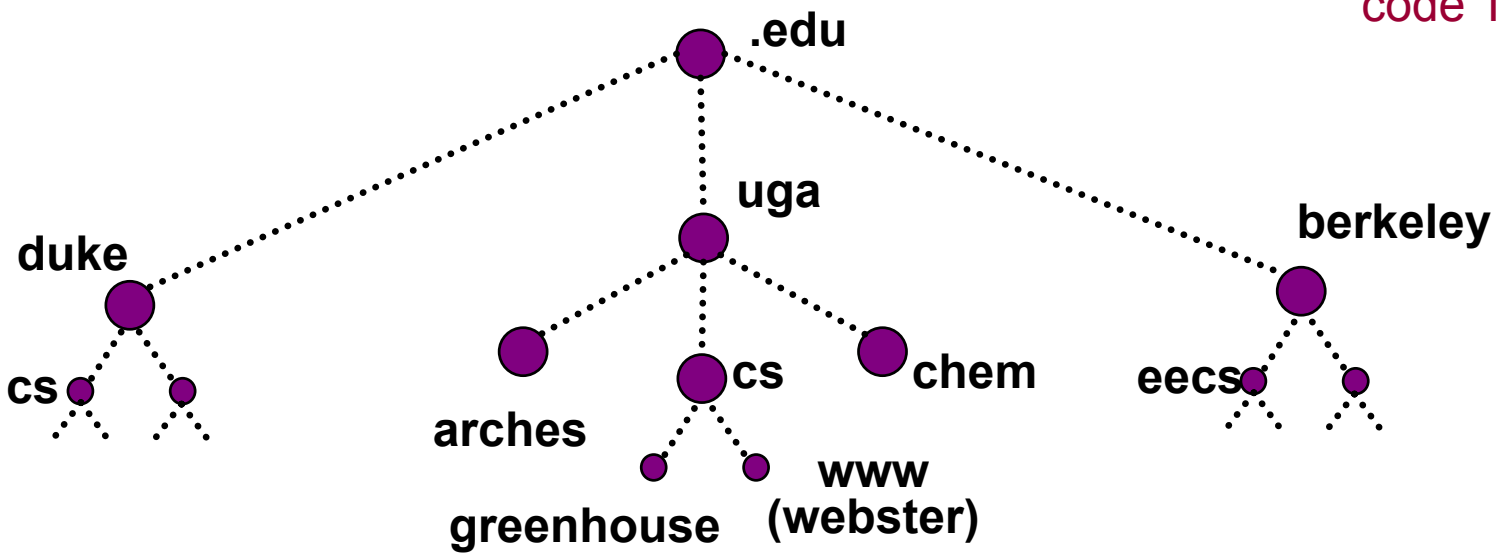

Source: Jeff Chase

DNS Protocol

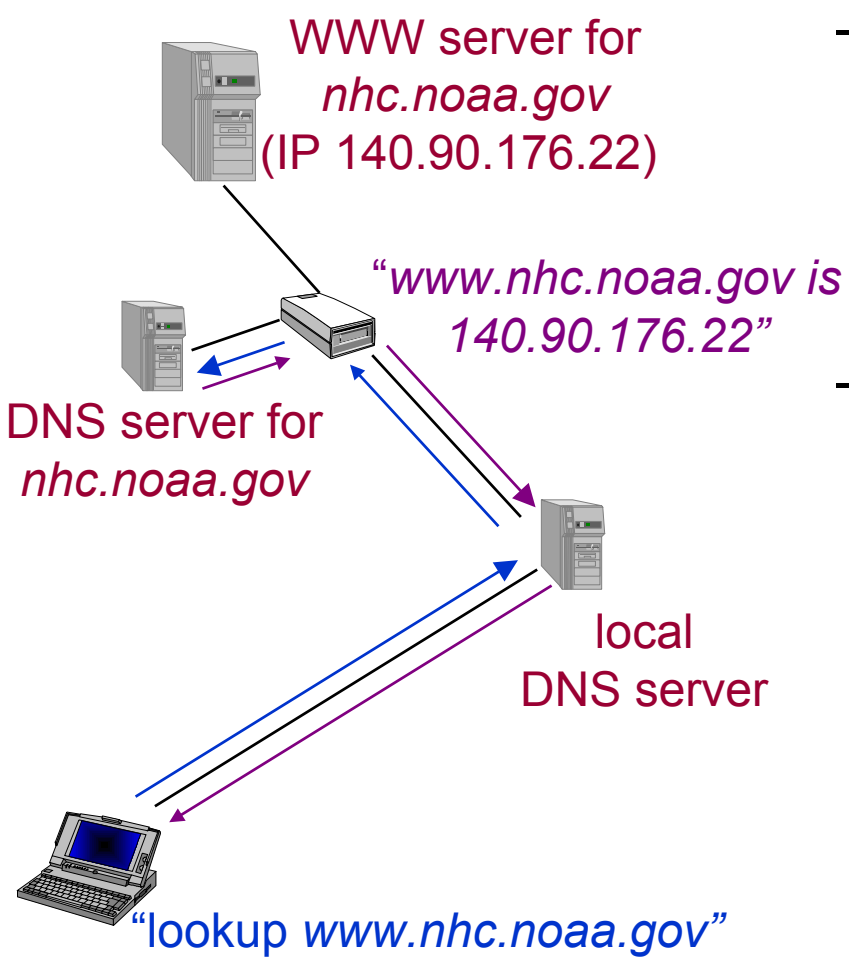

- UDP-based client/server
  - client-side *resolvers*
    - typically in a library
    - *gethostbyname*, *gethostbyaddr*
  - cooperating servers
    - query-answer-referral model
    - forward queries among servers
    - server-to-server may use TCP ("zone transfers")

Source: Jeff Chase

DNS Name Server Hierarchy

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

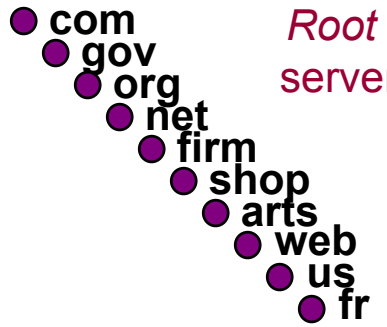

Root servers list servers for every TLD.

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

Subdomains correspond to organizational (*admininstrative*) boundaries, which are not necessarily geographical.

Servers may delegate management of *subdomains* to child name servers.

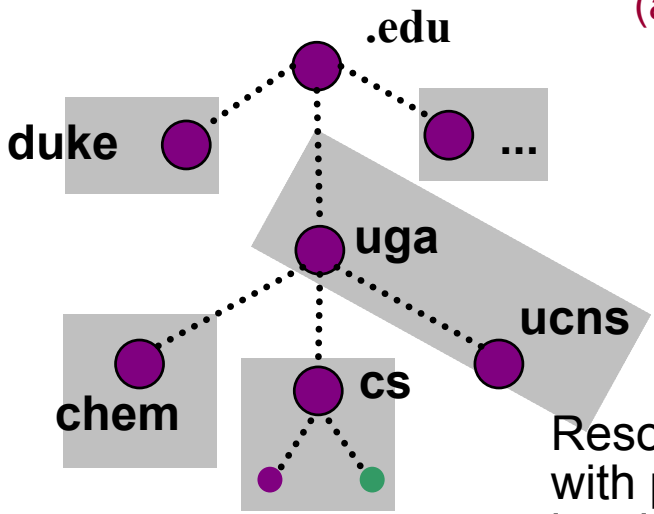

Servers are bootstrapped with pointers to selected peer and parent servers.

Parents refer subdomain queries to their children.

Resolvers are bootstrapped with pointers to one or more local servers; they issue *recursive* queries.

Source: Jeff Chase

Server selection problem

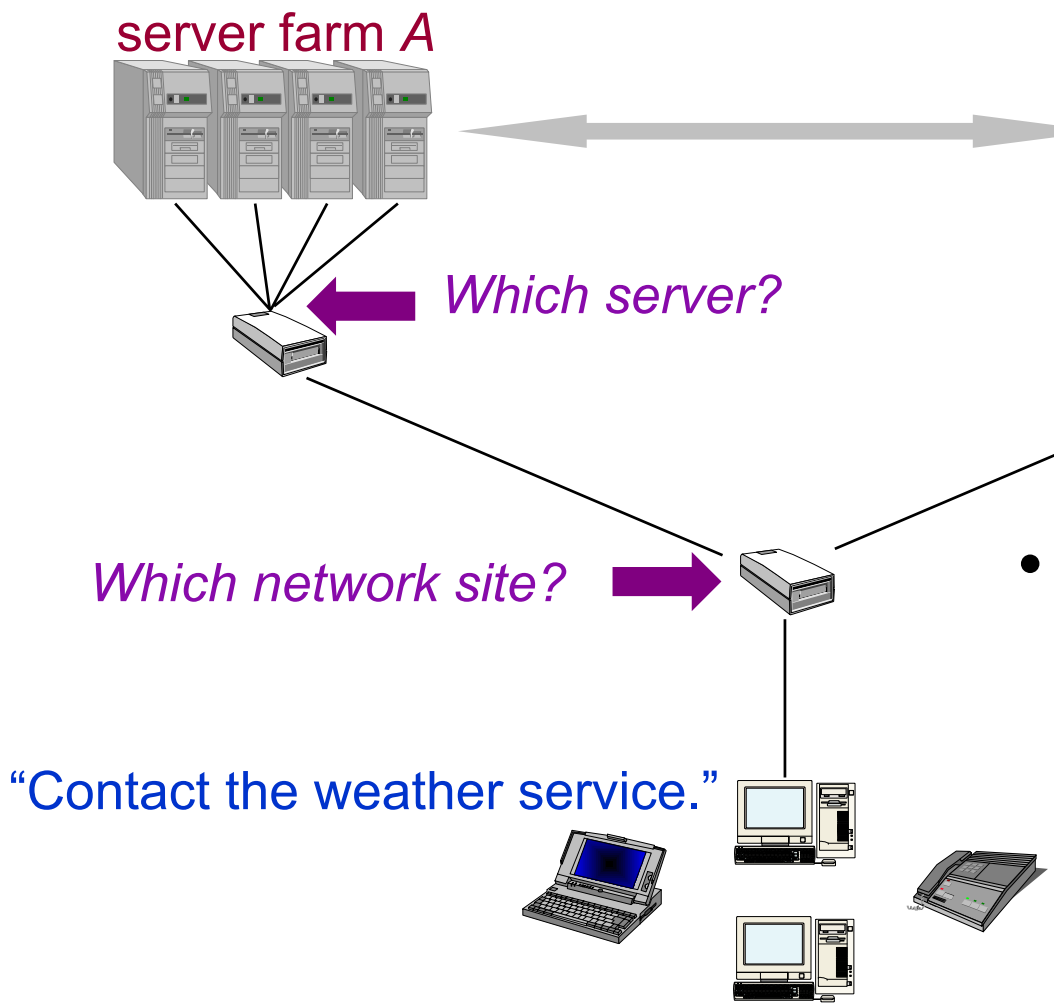

- Avoid the scaleability problems of central servers by "distributing" load

Source: Jeff Chase

DNS round robin

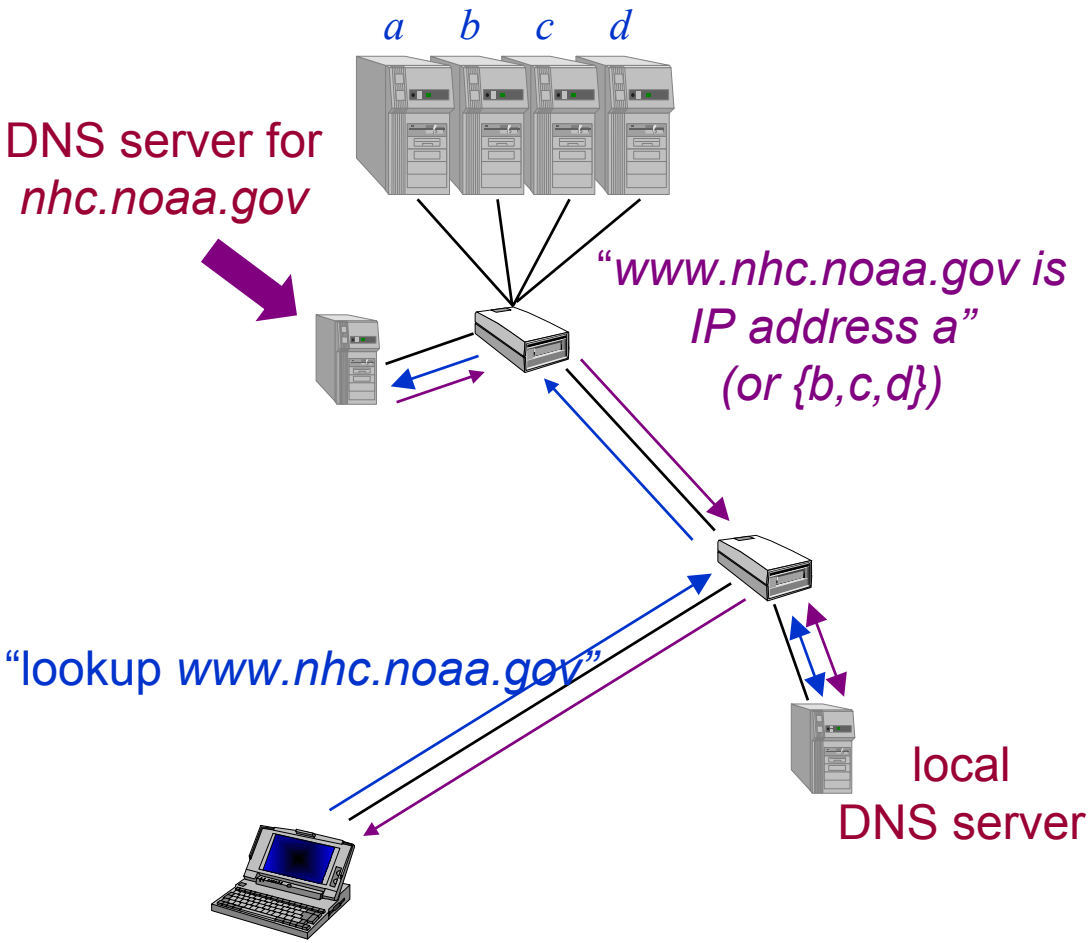

DNS server returns one of multiple addresses based on load
e.g. www1.aol.com
www2.aol.com

Source: Jeff Chase

DNS record for www.yahoo.com

```
www.yahoo.com.          1002    IN      CNAME   www.yahoo.akadns.net.
www.yahoo.akadns.net.   292     IN      A       64.58.76.223
www.yahoo.akadns.net.   292     IN      A       64.58.76.224
www.yahoo.akadns.net.   292     IN      A       64.58.76.225
www.yahoo.akadns.net.   292     IN      A       64.58.76.227
www.yahoo.akadns.net.   292     IN      A       64.58.76.228
www.yahoo.akadns.net.   292     IN      A       64.58.76.229
www.yahoo.akadns.net.   292     IN      A       64.58.76.176
www.yahoo.akadns.net.   292     IN      A       64.58.76.177
www.yahoo.akadns.net.   292     IN      A       64.58.76.178
www.yahoo.akadns.net.   292     IN      A       64.58.76.179
www.yahoo.akadns.net.   292     IN      A       64.58.76.222

;; AUTHORITY SECTION:
akadns.net.             984     IN      NS      ZF.akadns.net.
akadns.net.             984     IN      NS      ZG.akadns.net.
akadns.net.             984     IN      NS      ZH.akadns.net.
akadns.net.             984     IN      NS      ZA.akadns.net.
akadns.net.             984     IN      NS      ZB.akadns.net.
akadns.net.             984     IN      NS      ZC.akadns.net.
akadns.net.             984     IN      NS      ZD.akadns.net.
akadns.net.             984     IN      NS      ZE.akadns.net.
```

Internet

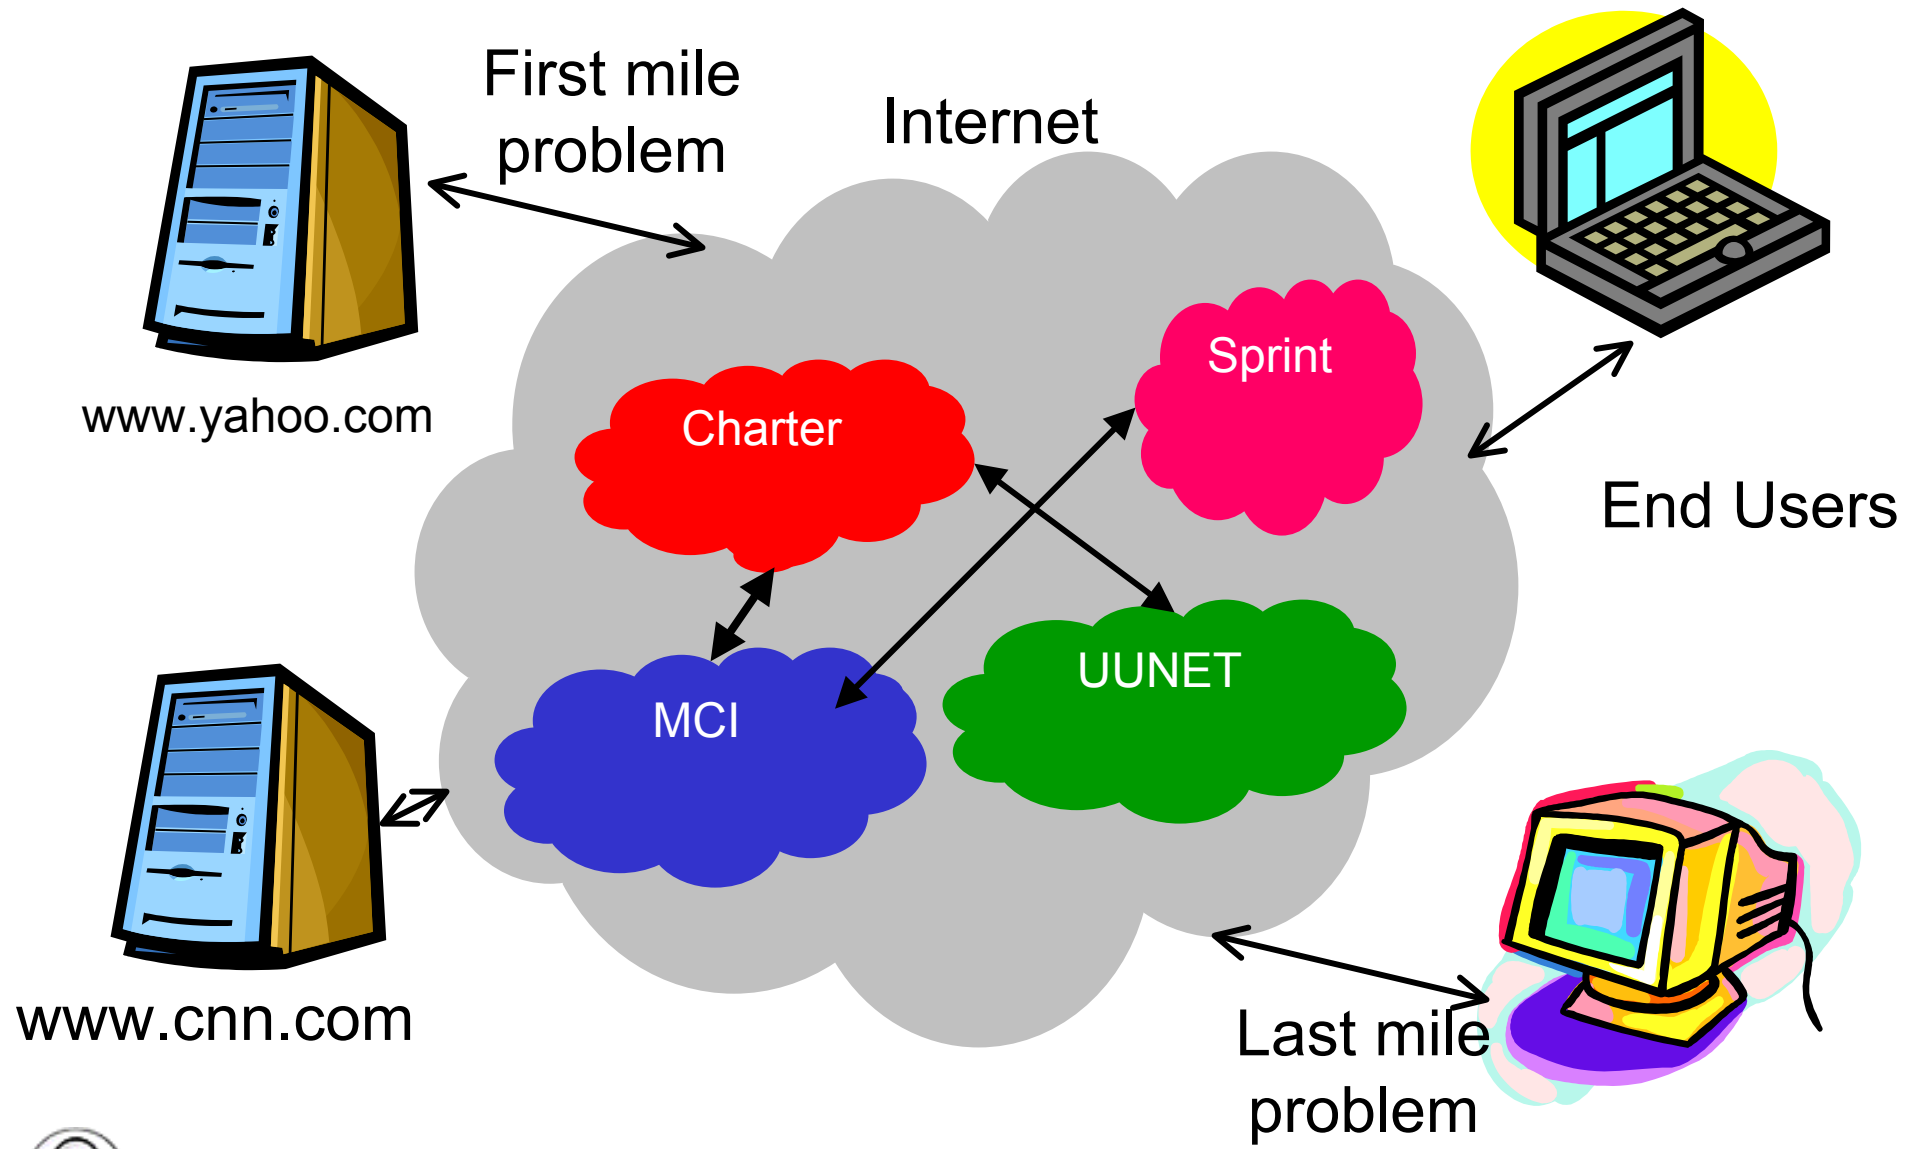
First mile problem
Internet
www.yahoo.com
Charter
Sprint
End Users
MCI
UUNET
www.cnn.com
Last mile problem

Performance bottlenecks

- First mile problem:
  - Server to the Internet
  - Everyone wants to access one popular service
- Last mile problem:
  - End user to the Internet
  - Broadband (cable, DSL), T1, T3, dialup, 2G cellular (slow)
- Peering problem:
  - Data goes through multiple networks and service providers at peering points
- Backbone problem:
  - The information highway for data traffic

Content delivery network

- Move server replicas to the edge

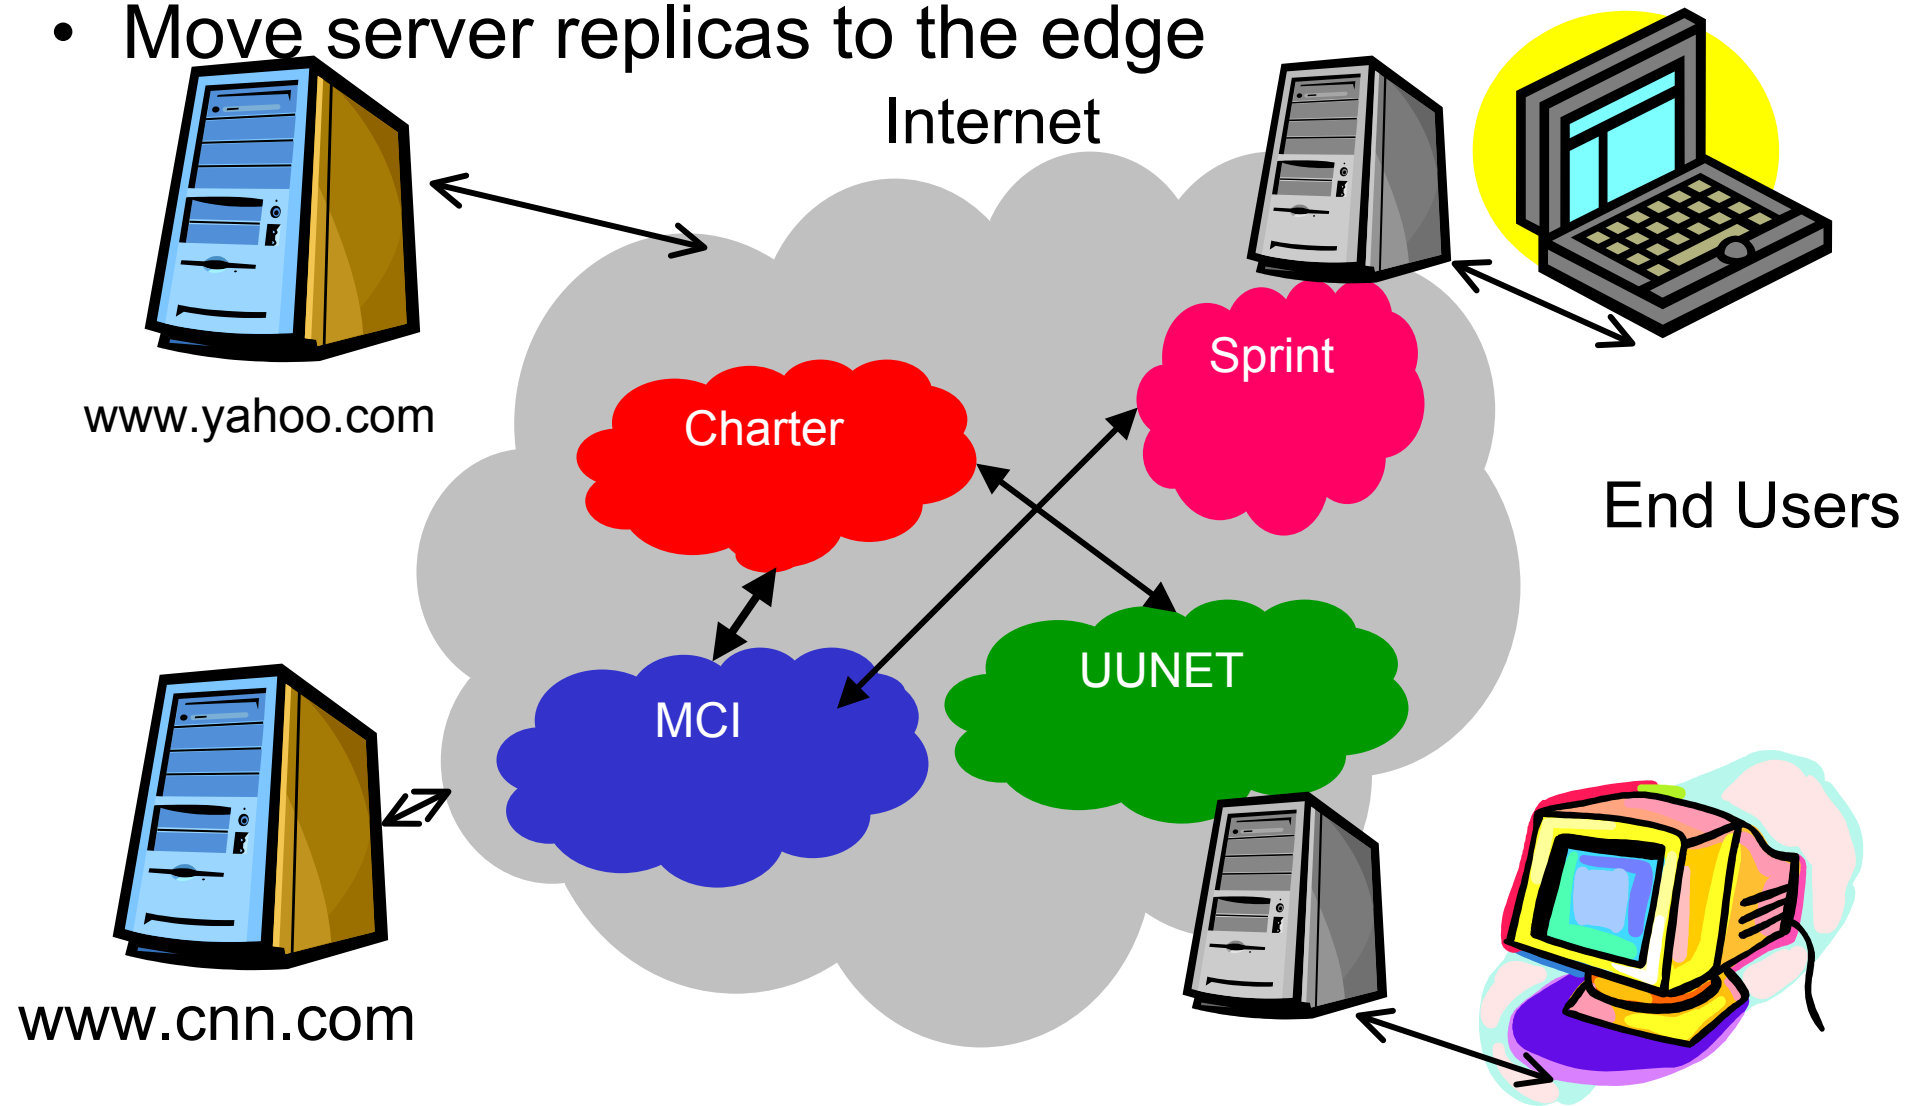

Content Delivery Network

- CDN (e.g., Akamai) creates new domain names for each client content provider.
  - e.g., *a128.g.akamai.net*
- The CDN's DNS servers are authoritative for the new domains.
- The client content provider modifies its content so that embedded URLs reference the new domains.
  - "Akamaize" content
  - e.g.: http://www.cnn.com/image-of-the-day.gif becomes

    http://a128.g.akamai.net/image-of-the-day.gif
- Using multiple domain names for each client allows the CDN to further subdivide the content into groups.
  - DNS sees only the requested domain name, but it can route requests for different domains independently.

Source: Jeff Chase

Akamai with DNS hooks

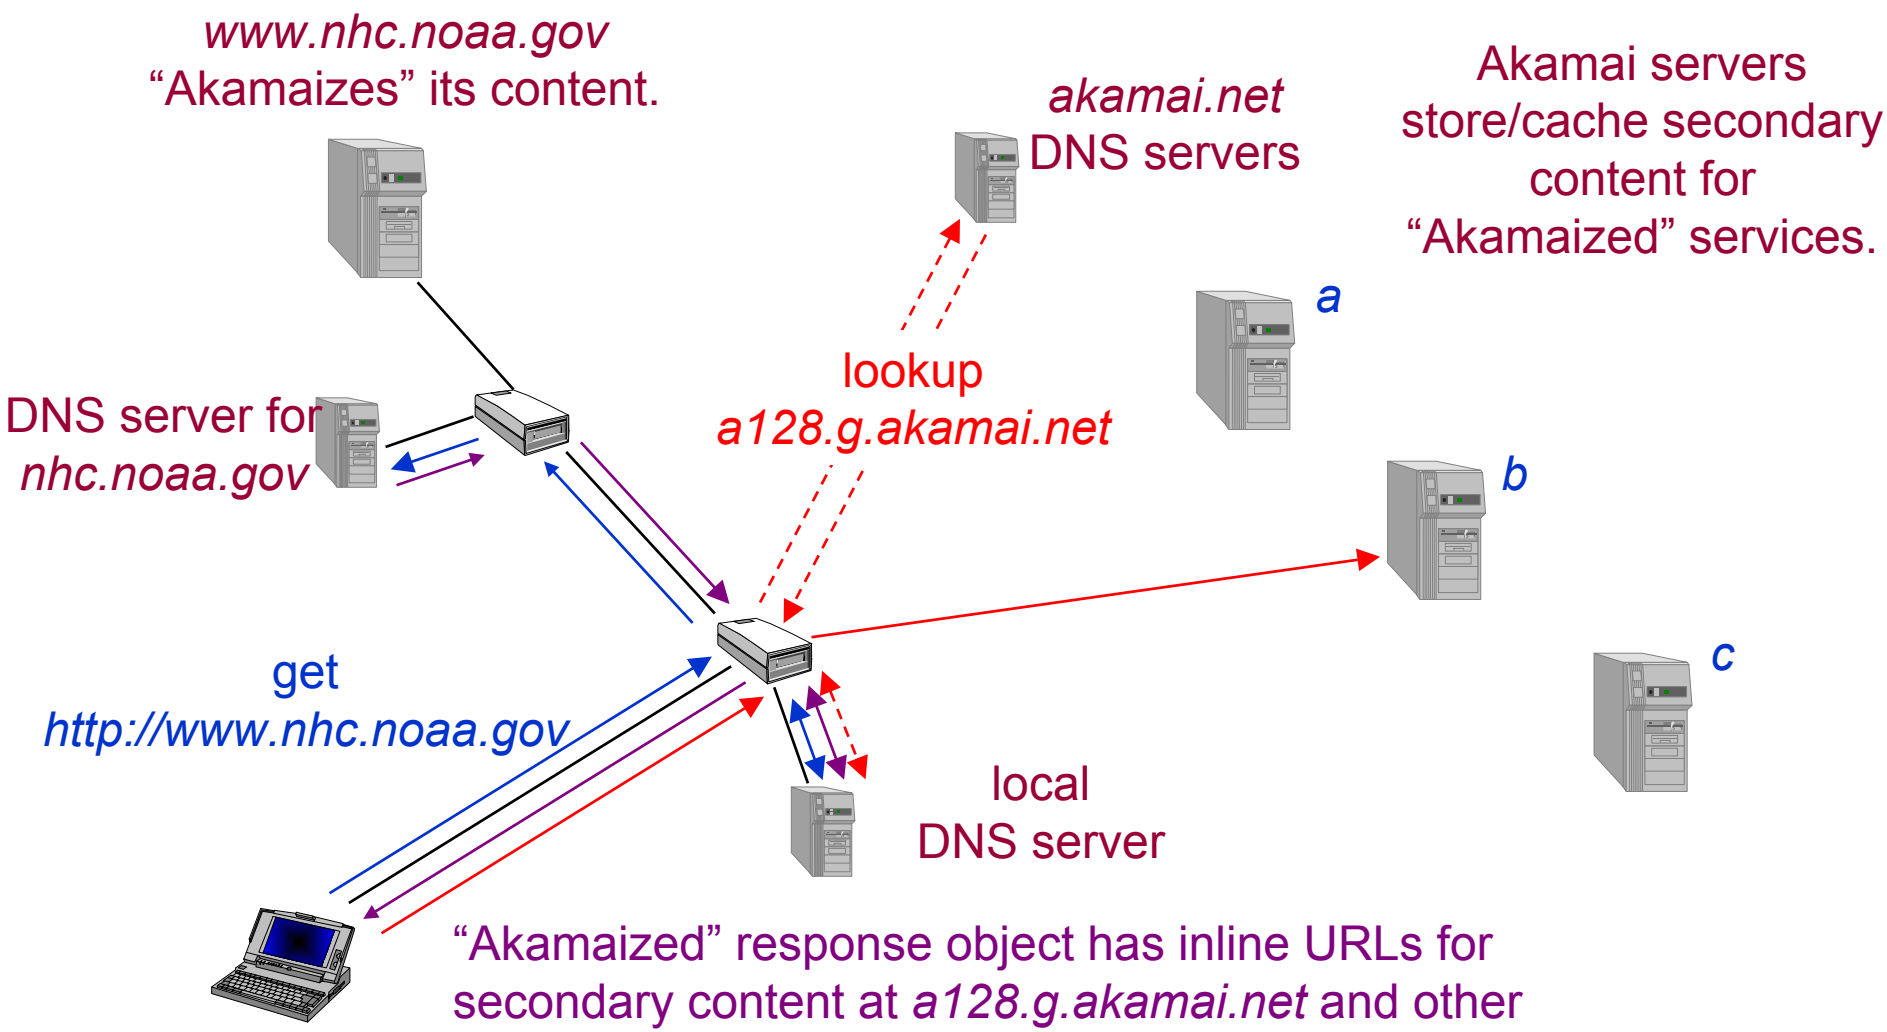

Source: Jeff Chase

Peer-to-peer systems

- Decentralized, no "server"
- Robust – no single point of failure
- "Will perform work for others since they will work for us" computing
- Can scale up

- Locating resources harder
- E.g. napster (has a central directory server)
  gnutella

Gnutella

- Queries issued by a servant at a given node propagate out to neighbor nodes
- The neighbors propage the query to their neighbors, and so on, for a given number of hops.
- Depending on where a user's query is first issued, it may or may not reach a node that has the file sought by the user.

Partial Map of Gnutella Network - 7/27/00

Clip2 Distributed Search Services

http://dss.clip2.com

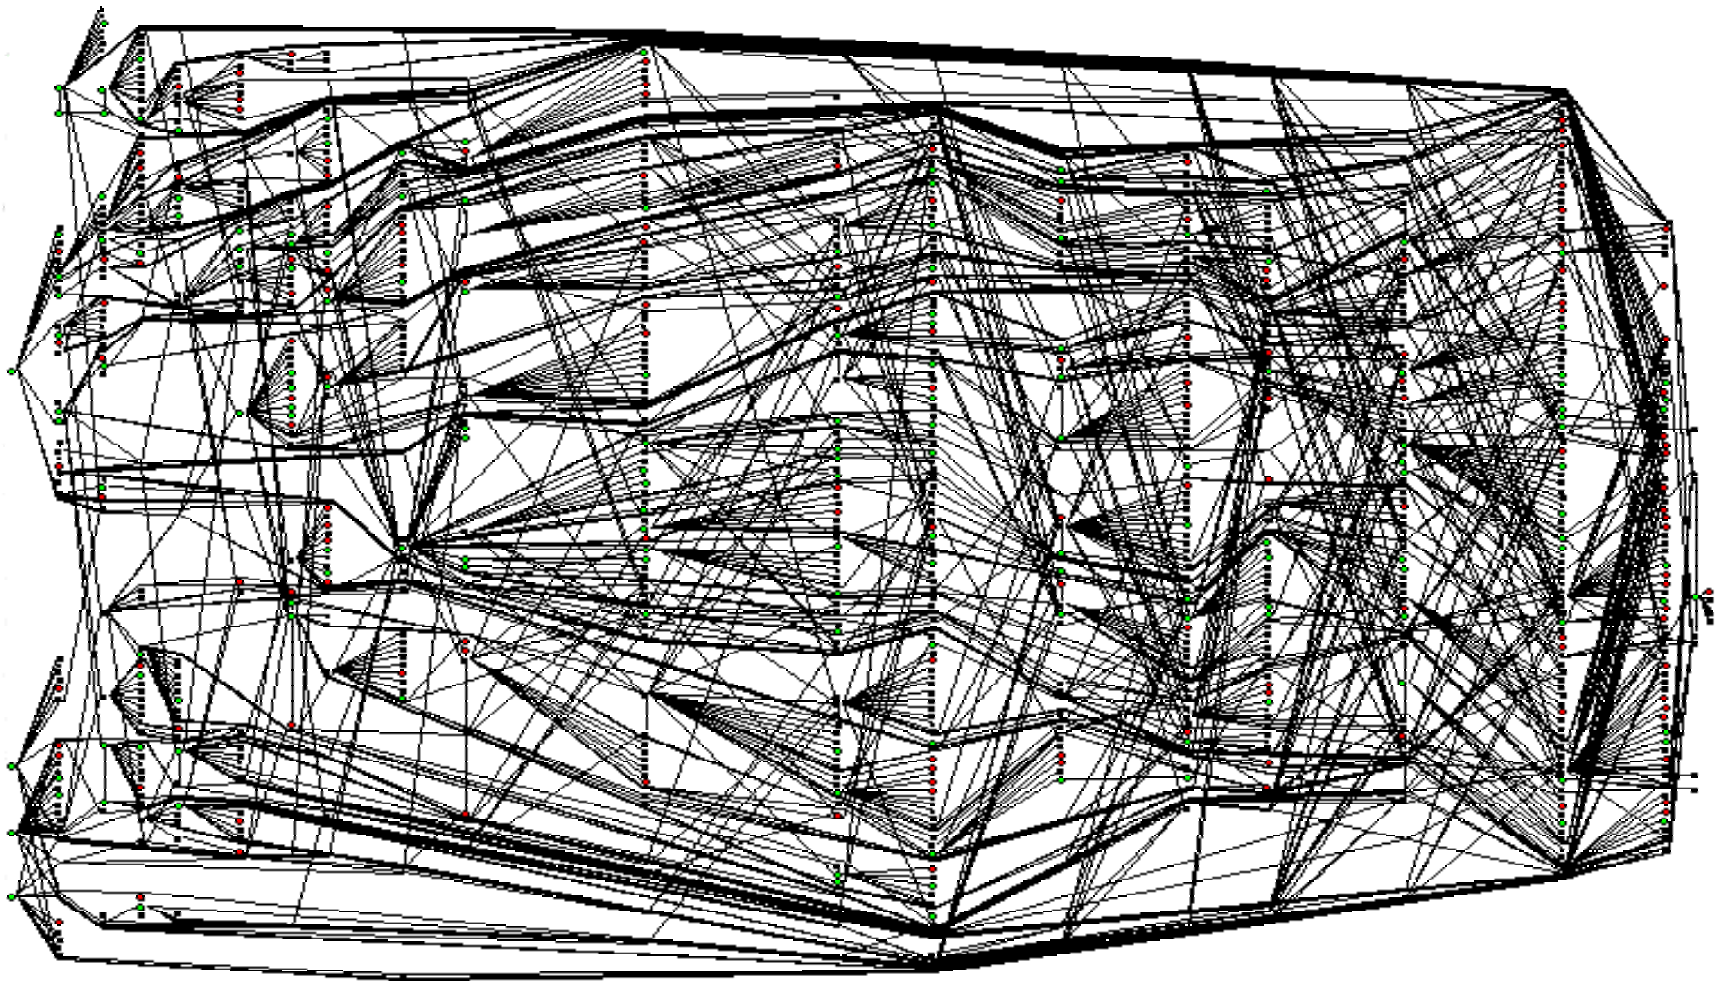

Scalability

- The scalability of a Gnutella network to accommodate more users performing more searches is limited by the lowest bandwidth links prevalent within the network
- For dial-up users it is 10 requests per second and has been reached