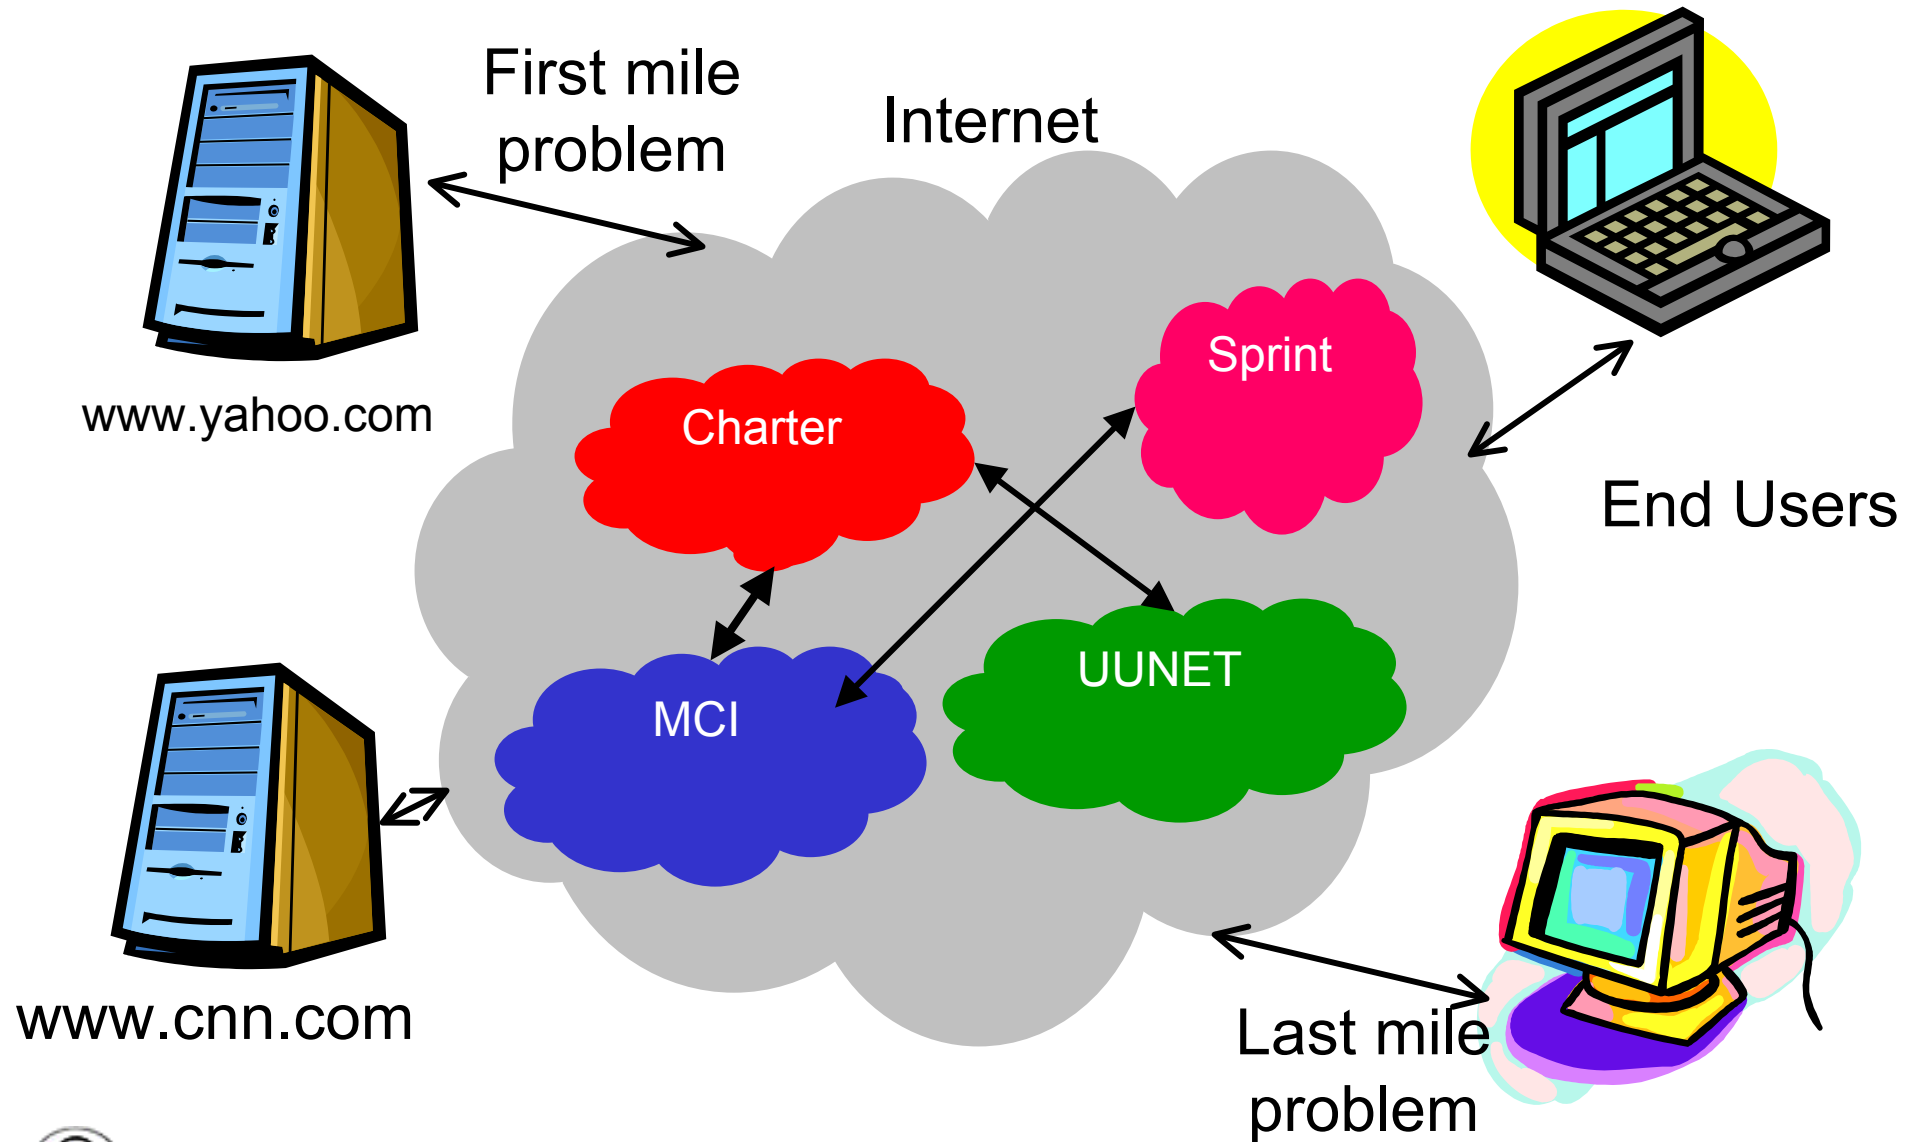


Outline

- Naming
- Applications
 - Central Server
 - Hierarchical
 - Peer-to-peer (Cont)

Internet




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 propagate 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>

(c)2000 Clip2.com, Inc.



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



Internet RFC

- Request for comment (RFC) (started in 1969)
 - In informal way to publish new ideas/protocols
 - First publish Internet Drafts before publishing as RFC
 - RFC can be obsoleted by newer RFCs, not modified
 - Some of them become standards
- Check out www.rfc-editor.org
 - RFC [0133](#) **File Transfer and Recovery** R.L. Sundberg [Apr 1971]
 - RFC [1889](#) **RTP: A Transport Protocol for Real-Time Applications** Audio-Video Transport Working Group, H. Schulzrinne, S. Casner, R. Frederick, V. Jacobson [January 1996]
 - RFC [1945](#) **Hypertext Transfer Protocol -- HTTP/1.0** T. Berners-Lee, R. Fielding, H. Frystyk [May 1996]



Multimedia streams

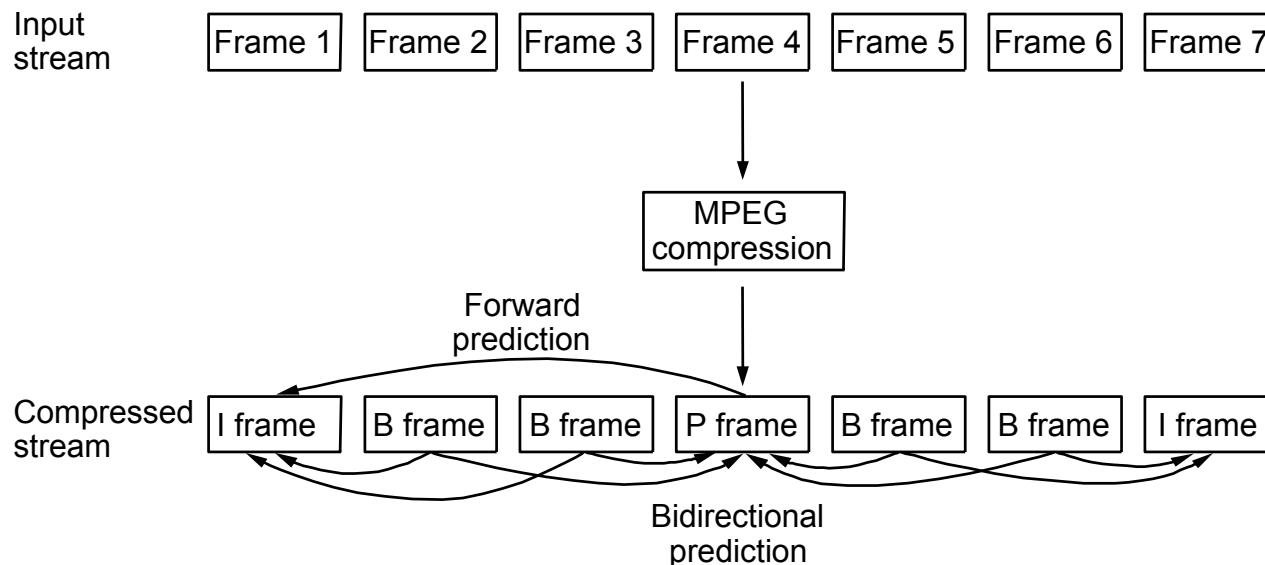
- Variable bit rate requirements
 - The data rate requirements changes with time
- Real time requirements to avoid jitter (packet not arriving on time)
- Need to efficiently transmit lots of data to many clients (multicasting)

MPEG

- Motion Picture Expert Group
- Lossy compression of video
- First approximation: JPEG on each frame
- Also remove inter-frame redundancy

MPEG (cont)

- Frame types
 - I frames: intrapicture
 - P frames: predicted picture
 - B frames: bidirectional predicted picture



- Example sequence transmitted as I P B B I B B

MPEG (cont)

- B and P frames
 - coordinate for the macroblock in the frame
 - motion vector relative to previous reference frame (B, P)
 - motion vector relative to subsequent reference frame (B)
 - delta for each pixel in the macro block
- Effectiveness
 - typically 90-to-1
 - as high as 150-to-1
 - 30-to-1 for I frames
 - P and B frames get another 3 to 5x

RTP

- Application-Level Framing
- Data Packets
 - sequence number
 - timestamp (app defines “tick”)
- Control Packets (send periodically)
 - loss rate (fraction of packets received since last report)
 - measured jitter

Real-Time Scheduling

- Priority
- Earliest Deadline First (EDF)
- Rate Monotonic (RM)
- Proportional Share
 - with feedback
 - with adjustments for deadlines