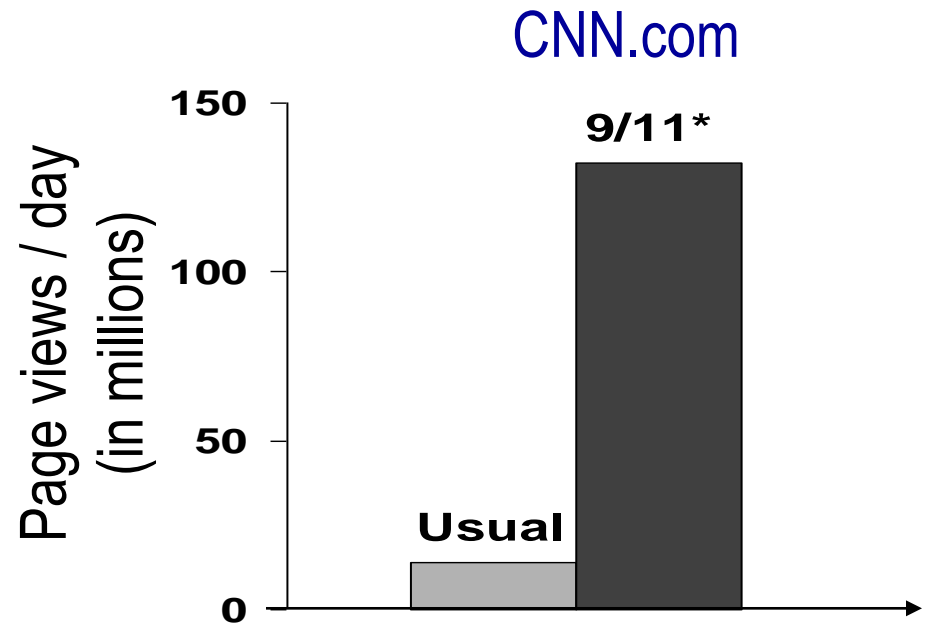


Content Delivery Networks

Credits:
Lucy Cherkasova,
HP Research Labs Palo Alto

Website Requests Unpredictable

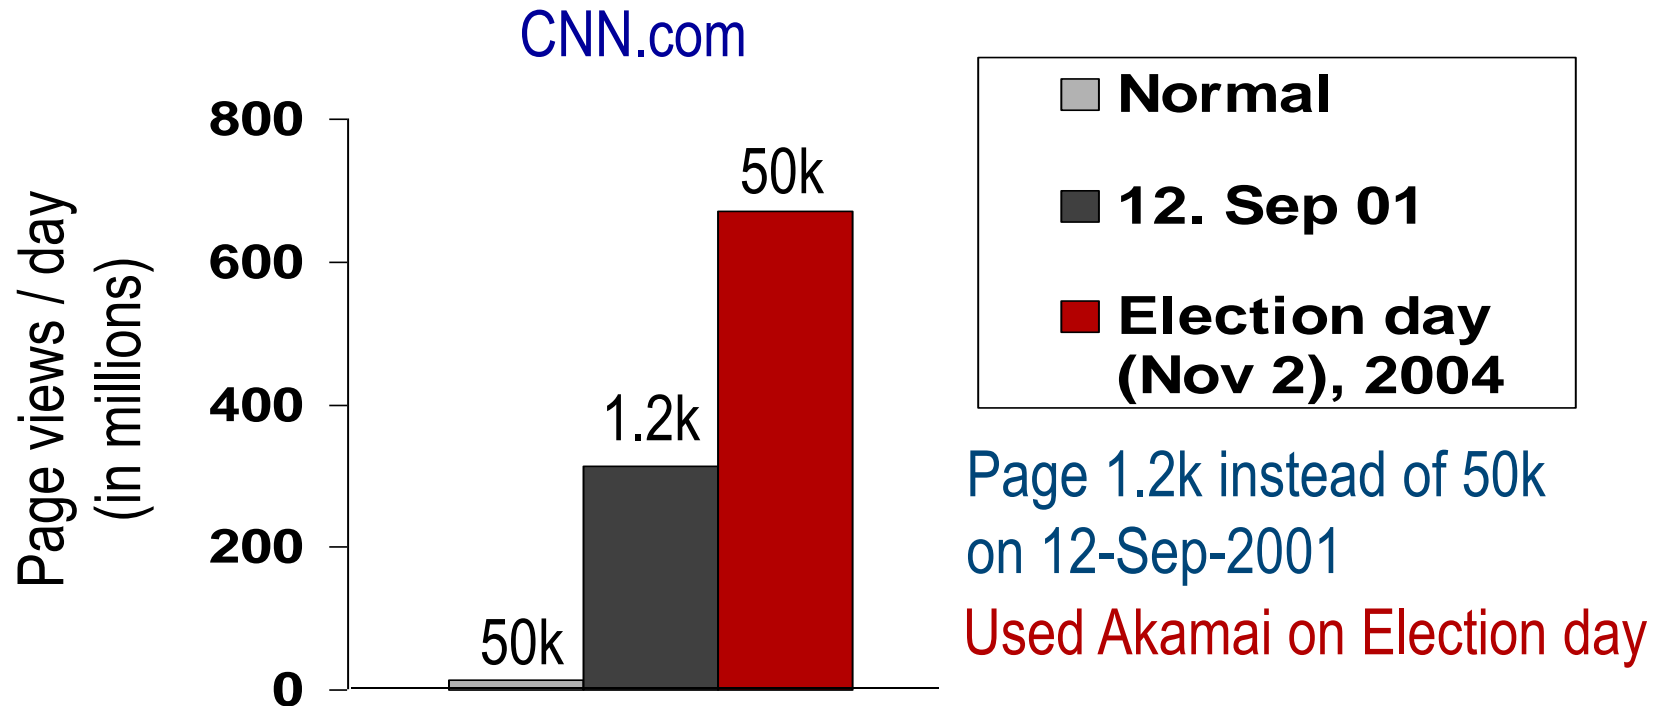
CNN, NY Times, ABC News
unavailable from 9-10 AM
(Eastern Time)



Content providers' dilemma: how many resources to provision?

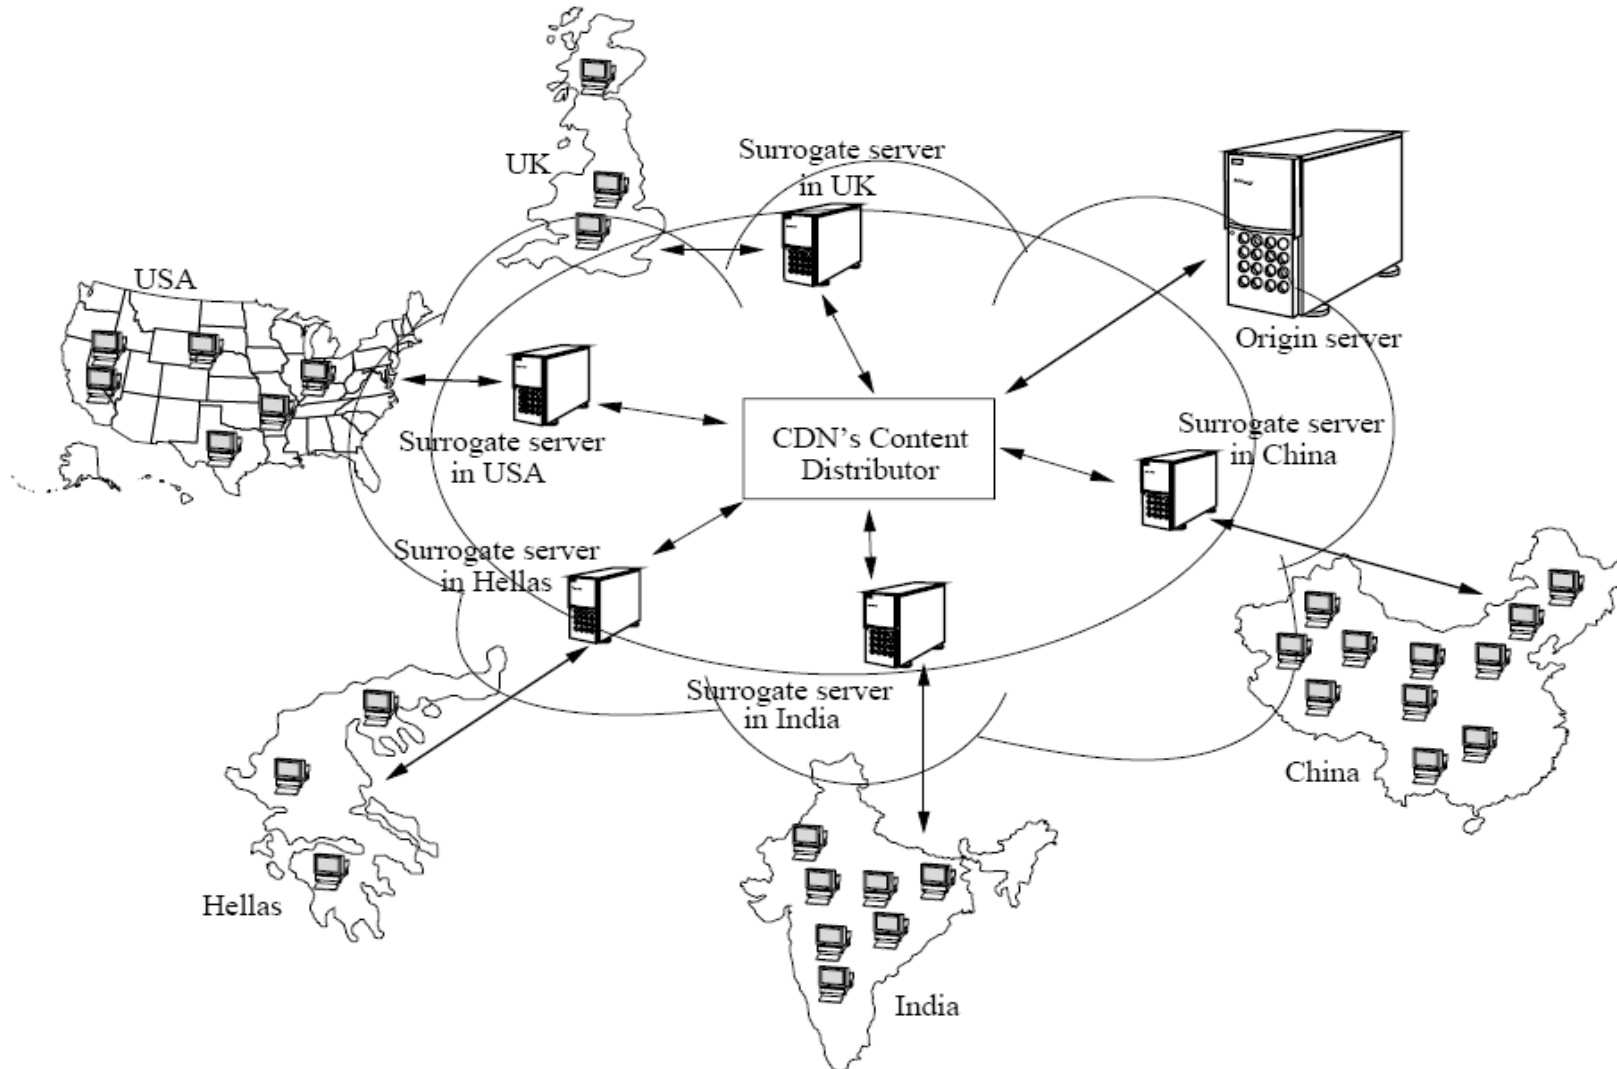
Need on-demand scalability

Content Delivery Networks (CDN)



Source: <http://www.tcsa.org/lisa2001/cnn.txt>
<http://www.akamai.com/en/html/about/press/press479.html>

CDN Architecture



CDN, Explained

- Goal: serve content to end-users with high **availability**, high **performance**
- Synonyms:
content delivery network = content distribution network (CDN)
- distributed **system of servers** deployed in multiple data centers
- CDNs serve large fraction of Internet today
 - web objects (text, graphics and scripts)
 - downloadable objects (media files, software, documents)
 - applications (e-commerce, portals)
 - live streaming / on-demand streaming media
 - social networks, ...

Also: minimize hops for minimizing „man in the middle“ sniffing, attacks

Mechanisms

- URL rewriting
 - ``
 - ``
- DNS redirection
 - Transparent, no content modification
 - Typically: two-level DNS lookup - choose most appropriate edge server


name -> list of edge servers

selected list item -> IP address

Transformations in CDNs

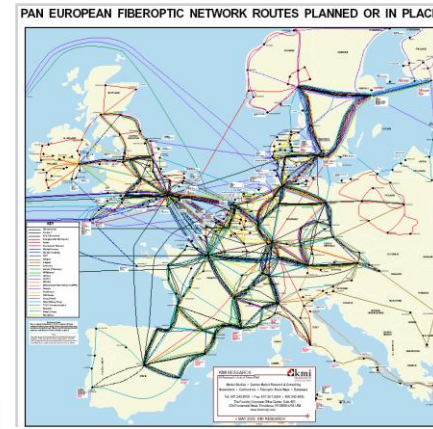
- Delivered contents are usually **modified or transformed** by proxies
 - Modify sizes and resolutions of multimedia files
 - Customize dynamic web pages based on client preferences
- Data transformations may involve multiple proxies
- Security issue: who allowed to do what?

Ex: 2-Step Data Transformations




Edge Devices


- = **entry point (ie: router)** into enterprise or service provider **core networks**
- Translating between heterogeneous network types
 - Ethernet, Token Ring, ATM, ISDN, ...
- Normally authenticated
- CDNs use edges as Point of Presence (PoP)
 - Often 10s of thousands



[www.lboro.ac.uk/gawc/rb/rb136.html]



[img: wikipedia]



Strategy Parameters

- How to determine optimal number of edge servers & placement?
- Two different approaches:
 - **Co-location**: placing servers closer to the edge (Akamai)
 - **Network core**: server clusters in large data centers near main network backbones (Limelight, AT&T)
- Content **placement**
- Needs large-scale system **monitoring & management**
 - gather evidence as a basis for design decisions

Business Model

- CDN pays ISP, carriers, network operators
- Advantage:
 - Less transmission costs: data closer to user
 - Some protection against DoS attacks
- Examples:
 - Akamai; as of 2009: 56,000 servers in 950 networks in 70 countries; deliver 20% of all Web traffic - ex: CNN
 - Microsoft Azure CDN; Amazon CloudFront; Amazon S3 – online storage (DropBox!)

Challenges

- Efficient large-scale content distribution
 - large files, video on demand, streaming media
 - low latency, real-time requirement
 - *FastReplica for CDNs*
 - *BitTorrent (general purpose)*
 - *SplitStream (multicast, video streaming)*
- Update propagation
 - Privacy: delete propagation

Fog Computing

- Fog Computing = Cloud Computing + Edge Computing:
 - dynamic localization of services on user demand
 - across Internet
 - cf CDNs: data + services close to user
- Manifold applications:
 - user devices & routers; Smart Grid; Smart Traffic Lights / connected vehicles; Wireless Sensor & Actuator Networks; Decentralized Smart Building Control; ...
 - Swarms!
 - cf. [ORBiDANSe project](#): Array Database on board an EO satellite

Discussion

- “Flash Crowd” problem
 - L. Niven: Flash crowd. In: The Flight of the Horse. Ballantine Books, 1971
- Goal: High availability + responsiveness key factors for business Web sites
 - overcome server overload for popular sites
 - minimize network impact in delivery path
- CDN: large-scale distributed network of servers
 - Surrogate servers (proxy caches) located closer to edges of Internet
 - edge servers