


Content Distribution

- How to distribute content without requiring centralized, heavy-duty servers?
- Examples:
 - Bittorrent
 - Peer-to-peer content distribution
 - Akamai
 - Content distribution service

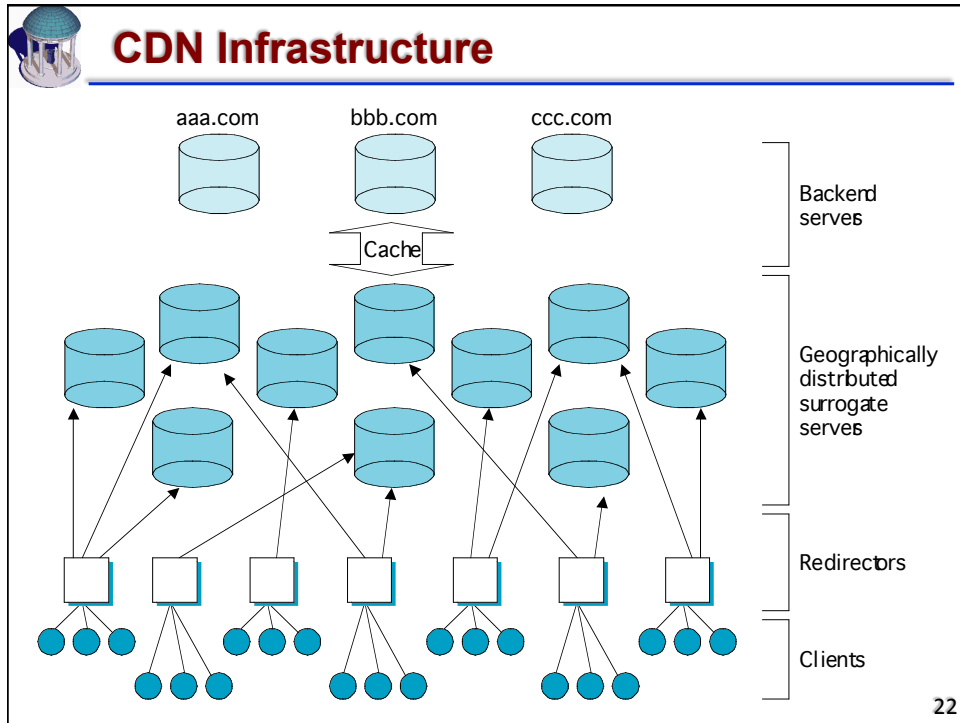
20




Bottlenecks for Web Downloads

- 3 potential bottleneck locations:
 - The first mile
 - User's access connection may be limited
 - The last mile
 - Link connecting server to Internet can get overloaded by too many requests
 - Peering points
 - ISPs have little motivation to provide high-capacity connectivity to their peers
- Idea: replication can address the 2nd and 3rd issues
 - Only end-user can solve the 1st one
- Content Distribution Networks
 - Systems that provide such replication

21



-
- Content Distribution vs Caching**
- ❑ **Explicitly manage cache content**
 - Sell as service to web site owners for performance
 - ❑ **“Push” content to caches at major Internet providers**
 - Make content appear “nearby” (low latency) no matter where the requester is located
 - ❑ **Change cache content when necessary**
 - Operates as a proxy cache to refresh content
- 23



Akamai Content Distribution Network

http://www.akamai.com/html/about/facts_figures.html

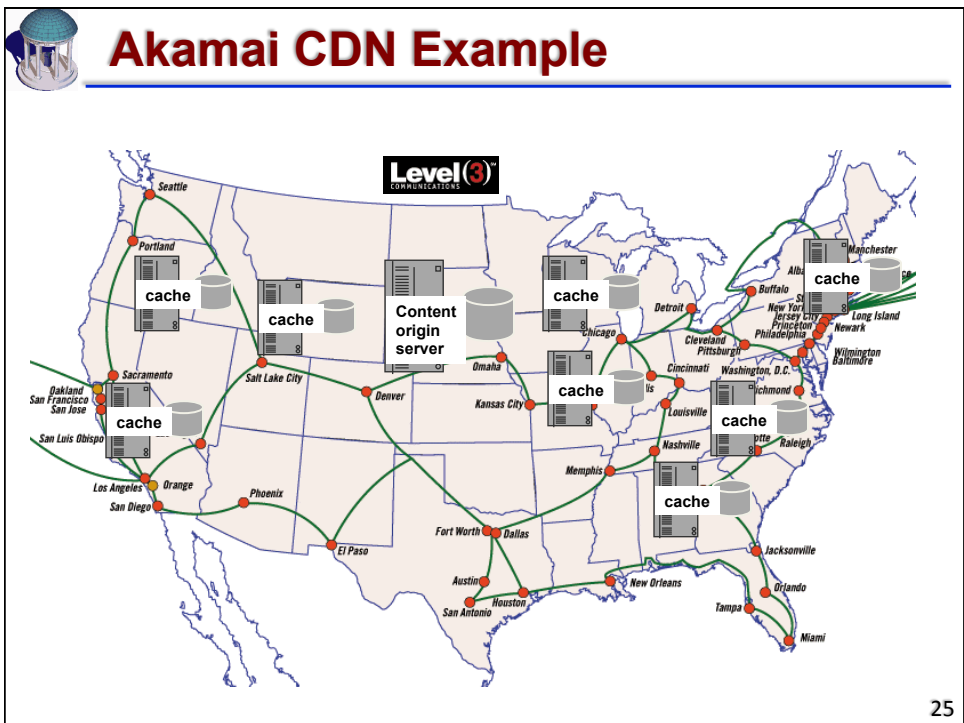
- ❑
56,000 servers in 70 countries within nearly 1000 networks
 - 85% of the world's Internet users have a single "network hop" to an Akamai server

- ❑
Delivers 15 - 20% of all Web traffic
 - Web traffic at times reaching more than 2 Terabits/sec
 - Hundreds of billions of daily Internet interactions

Representative Customers

1-800-Flowers	L'Oréal
Adobe	MRO Software
All Nippon Airways (ANA)	MTV Networks
Analog Devices Inc.	MySpace
Apple Computer	Nintendo
AUDI AG	NPR
Best Buy	Pacific Sunwear
BMW Japan	The Princeton Review, Inc.
Cabela's	Qantas Airways
Canadian Broadcasting	QVC
Cathay Pacific Airline	Red Hat Inc.
Caterpillar	Rediff.com
Clear Channel Online Music	Reebok International Ltd.
Cognos	Reuters
Corel Corp.	Ross Simons
Department of Defense	Saatchi & Saatchi
The Dow Chemical Co.	The Seattle Times
eBags	Sony Ericsson Mobile Communications
Europages Group	Starz Entertainment
Federal Emergency Mngmt	Terra Lycos
FedEx Corp.	Thomson Financial
Food & Drug Administration	Ticketmaster Online-CitySearch
Fujitsu	Toyota Motor Sales
FUJI TV	Travelocity.com
General Motors	TrendMicro
Harley-Davidson, Inc.	Turner Sports
Hitachi	Urban Outfitters
IBM Corp. U.S. Geological Survey	Victoria's Secret
IBS	XM Satellite Radio
International Hotels Group	J.C. Penney Company, Inc.
J.C. Penney Company, Inc.	Yahoo!
Jet Airways	The Yomiuri Shimbun
Lands' End Inc.	
Logitech	

24





Akamai Content Types

- **Static (HTML, images, PDF, etc.)**
 - Expiration time (0 – infinity) assigned by customer
 - On-demand cache invalidation available to customer
 - Special features (authenticated access, transfer encodings, etc.)
- **Dynamic**
 - Assembles cacheable and non-cacheable elements of page at cache (contacts origin only for non-cacheable)
- **Streaming**
 - Uses redundant streams and jitter-control to ensure quality playback

26



Web Site Redirection to Akamai

- **DNS CNAME aliases**
 - e.g., images.pcworld.com
CNAME=images.pcworld.com.edgesuite.net
- **Modified URLs (“Akamized”)**
 - Prefix with domain name in Akamai
 - e.g., http://a1694.g.akamai.net/8675/images.pcworld.com..

27

“Akamai-izing” Web Pages

An “Akamaized” URL

Staging Server → *RENAME Script Inserted* → **Live Servers**

YourCo Home Page → **YourCo Home Page**

<http://www.yourco.com/images/weather.gif>
<http://serial#.akamai.com/fingerprint/www.yourco.com/images/weather.gif>

serial#: maps content to a large virtual space – allows for optimal and load balanced mapping to physical servers
fingerprint: cryptographic hash of data used to ensure freshness of content

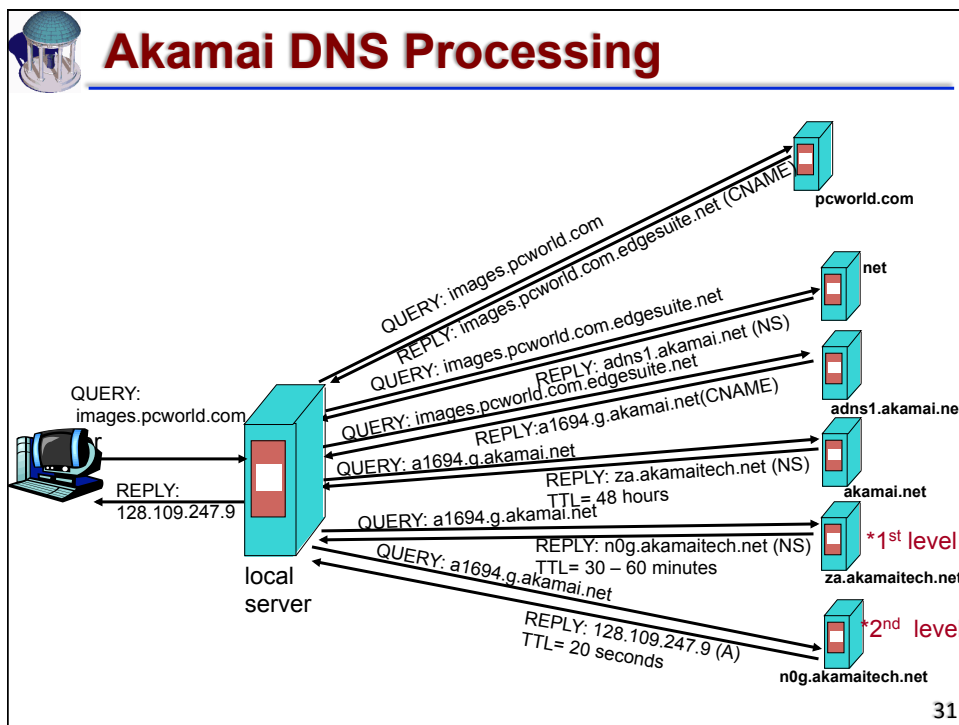
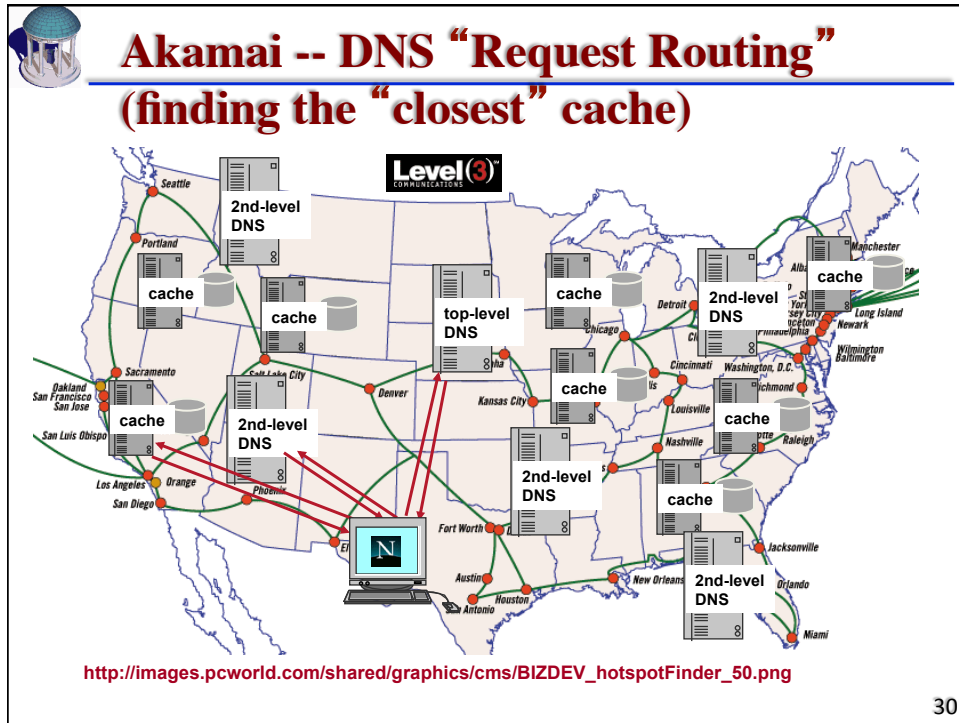
4


28

Akamai URL Example (WRAL-TV)


Image: http://a1844.g.akamai.net/7/1844/95/22/html.wral.com/images/structures/menus/nav_curve.gif
 Image: <http://a1844.g.akamai.net/7/1844/95/22/html.wral.com/images/structures/spacer.gif>
 Image: <http://a1844.g.akamai.net/7/1844/95/22/html.wral.com/images/structures/spacer.gif>
 Image: http://a1844.g.akamai.net/7/1844/95/22/html.wral.com/images/structures/misc/personalize_left.gif
 Image: <http://a1844.g.akamai.net/7/1844/95/22/html.wral.com/images/structures/spacer.gif>
 Image: http://a1844.g.akamai.net/7/1844/95/22/html.wral.com/images/structures/misc/personalize_right.gif
 Image: <http://a1844.g.akamai.net/7/1844/95/22/html.wral.com/images/structures/spacer.gif>
 Image: <http://a1844.g.akamai.net/7/1844/95/22/html.wral.com/images/structures/spacer.gif>
 Image: <http://scrooge.ibsys.com/s/ral.gif?rnd=999208094182&version=2>
 Image: <http://a1844.g.akamai.net/7/1844/95/22/html.wral.com/images/structures/spacer.gif>
 Image: <http://a1844.g.akamai.net/7/1844/95/22/html.wral.com/images/structures/spacer.gif>
 Image: <http://images.ibsys.com/2001/0126/435616.gif>
 Image: <http://images.ibsys.com/2001/0126/435612.gif>
 Image: <http://images.ibsys.com/2001/0126/435612.gif>
 Image: <http://images.ibsys.com/2001/0126/435612.gif>
 Image: <http://a1844.g.akamai.net/7/1844/95/22/html.wral.com/images/structures/spacer.gif>
 Image: <http://a1844.g.akamai.net/7/1844/95/22/html.wral.com/images/structures/misc/morestories.gif>
 Image: <http://a1844.g.akamai.net/7/1844/95/22/html.wral.com/images/structures/spacer.gif>
 Image: <http://a1844.g.akamai.net/f/1844/95/10m/images.ibsys.com/2001/0830/940459.jpg>
 Image: http://a1844.g.akamai.net/f/1844/95/10m/images.ibsys.com/2001/0830/940452_120X90.jpg
 Image: <http://a1844.g.akamai.net/7/1844/95/22/html.wral.com/images/structures/spacer.gif>

29






Akamai Network Operations Center



- **Monitors all proxy servers and end-to-end conditions on paths from 1000s of network locations to servers**
- **DNS servers updated with new loads and maps every few seconds**
- ◆ **Monitored conditions**
 - Server load and operational status
 - Routing topology
 - Latency and packet losses per route
 - Available bandwidth per route

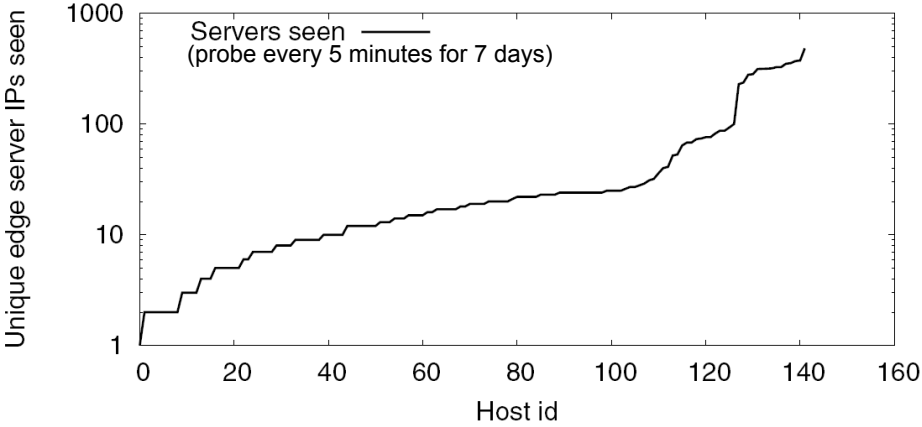
32



How Much Server Diversity Exists?

A. Su, et. al., "Drafting Behind Akamai (Travelocity-Based Detouring)", SIGCOMM 2006

Yahoo!



Unique edge server IPs seen

Host id

Servers seen (probe every 5 minutes for 7 days)

33

