

The Challenges of Delivering Content on the Internet

Tom Leighton

**Chief Scientist
Akamai Technologies**



Outline

How the Web Works

Akamai's Services

Technology Overview

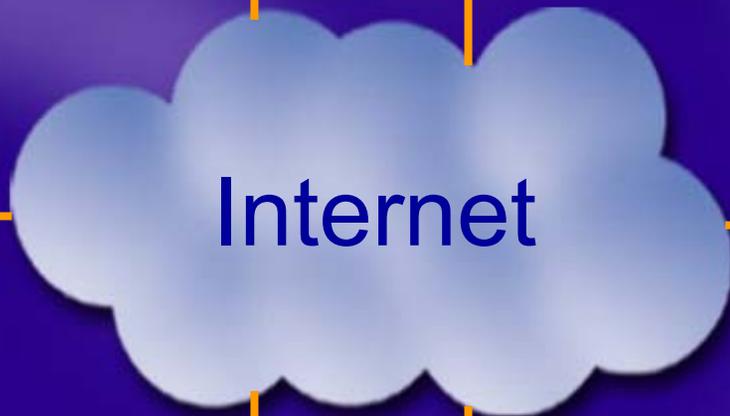
Technological Challenges

The Future



The Web: Simple on the Outside...

Content
Providers



End
Users



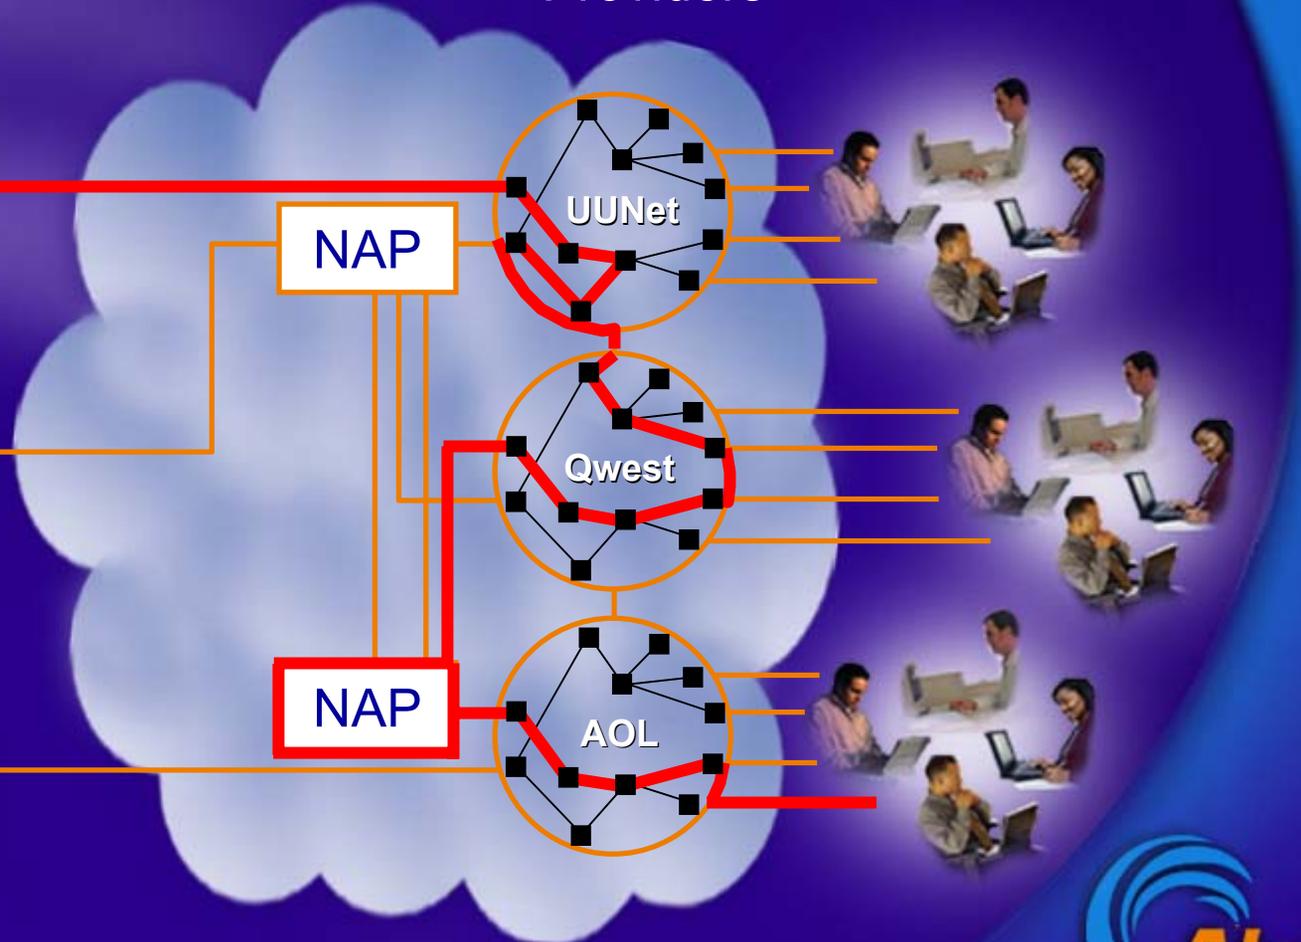
...But Problematic on the Inside

Content Providers

Peering Points

Network Providers

End Users



The First Mile Bottleneck

- Content Provider now has worldwide audience; centralization doesn't work



400 million potential users
= huge infrastructure problem

The Problems with Peering

- The Internet consists of over 7000 networks
- No single network controls a large % of access traffic—biggest has 6% share



The Problems with Peering (cont.)

- Economic considerations limit peering capacity—results in loss of routes and congestion
- Routing algorithms (BGP) ignore congestion!
- Routing algorithms are also subject to accidental loss of routes (or introduction of incorrect routes)

Bottleneck Implications

- Slow downloads
 - Content must traverse multiple backbones and long distances
- Unreliable performance
 - Content may be blocked by congestion or backbone peering problems
- Not scalable
 - Usage limited by bandwidth available at master site
- Inferior streaming quality
 - Packet loss, congestion, and narrow pipes degrade stream quality
- Broadband doesn't help
 - As broadband becomes ubiquitous, the disaster of centralized solutions becomes more obvious— not better



Outline

How the Web Works

Akamai's Services

Technology Overview

Technological Challenges

The Future

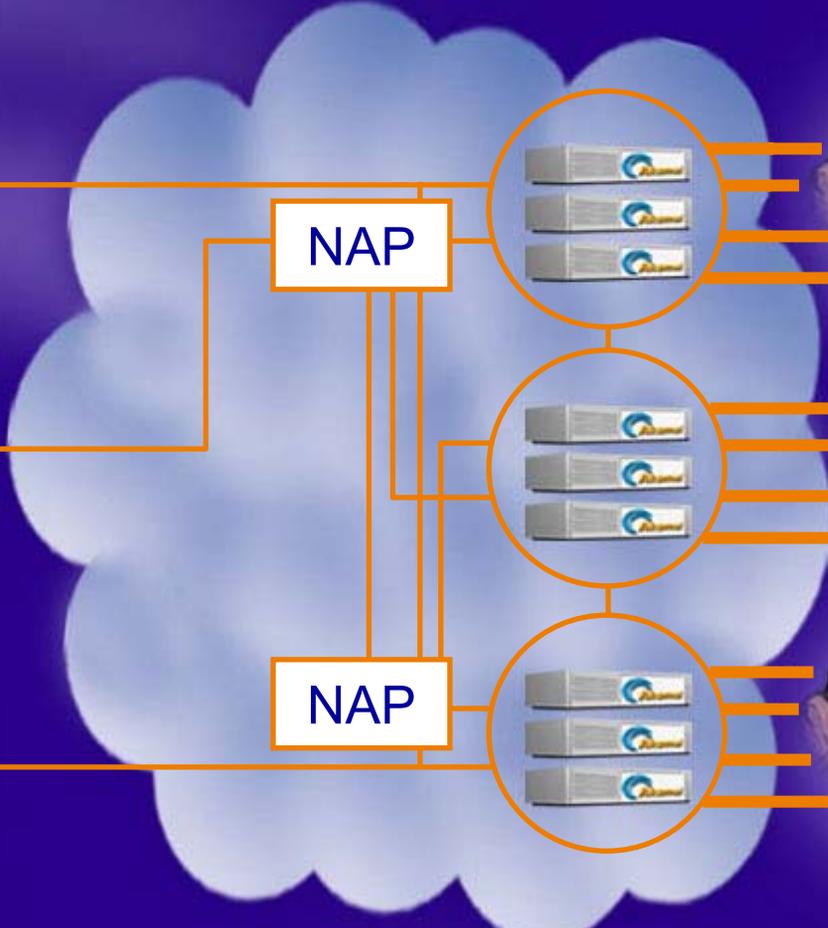


The Akamai Solution

Content Providers



Akamai Servers at Network Edge



End Users



Akamai's Network Deployment



13,500+
Servers

1000+
Leading
Carriers

Hosting, Access,
Backbones,
Satellite & Broadband

60+
Countries



Thousands of Websites Are Akamaized



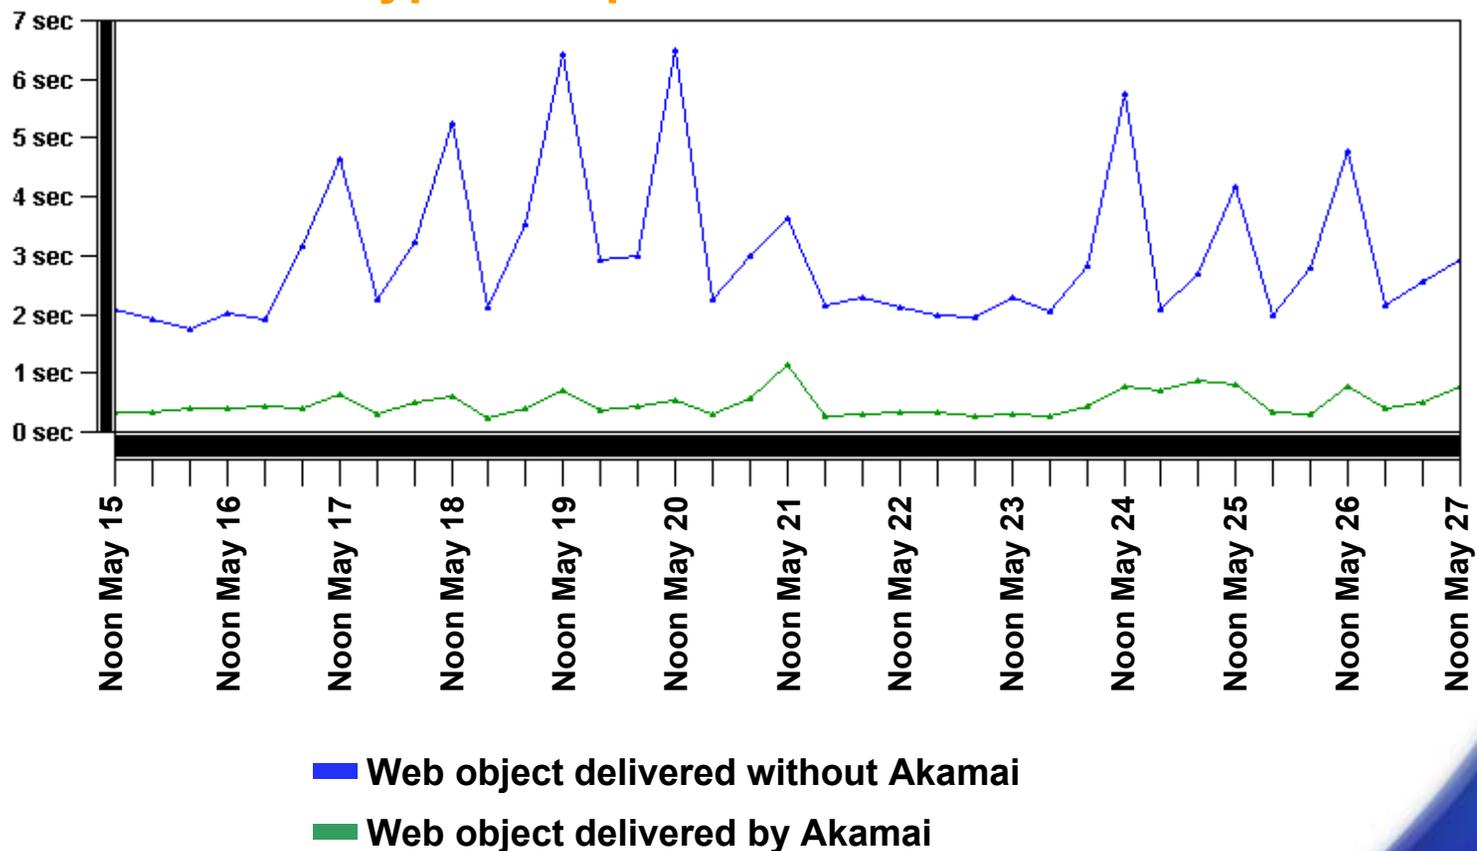
Advantages

- Fast
 - Content is served from locations near to end users
- Reliable
 - No single point of failure
 - Automatic failover
- Scalable
 - Master site no longer requires massive available bandwidth



Keynote™ Results

Web Site Performance Typical Improvement with Akamai



Akamai's Service Offerings

- **FreeFlow:** delivers objects (gifs, jpgs, etc.) and rich graphics to end users from the edge of the Internet
- **FreeFlow Streaming:** delivers streaming content to viewers worldwide with dramatic improvements in quality and reliability
- **Akamai Conference:** a reseller offering that uses streaming media to extend the reach and functionality of ordinary conference calls
- **Akamai Forum:** enables businesses to produce live, interactive Webcasts

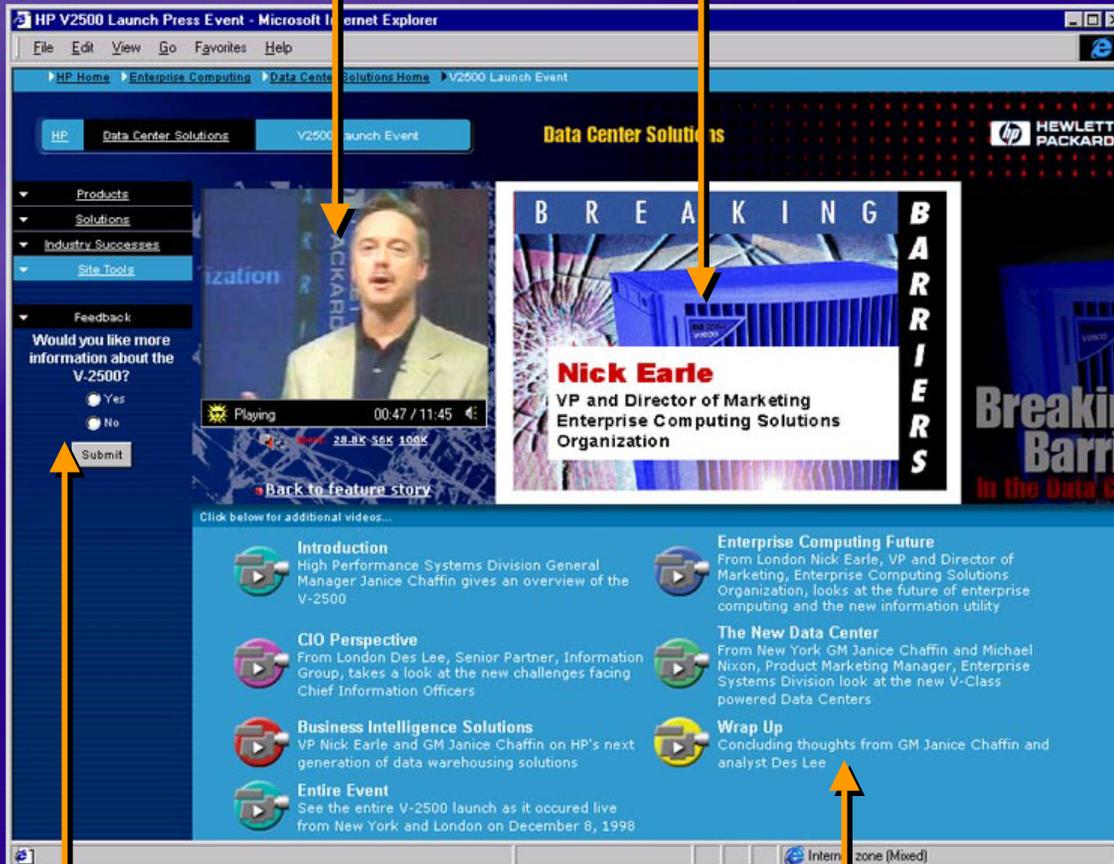


Akamai Forum

No special client software

Live or On-Demand Streaming Video

Speaker Support e.g. PowerPoint



Dynamic Surveys & Profiling

Indexed Program Schedule

Other Features:

- Ask a Question
- Live Audience Phone-in
- Viewer Registration
- E-mail promotion
- Download Slides
- Searchable Content

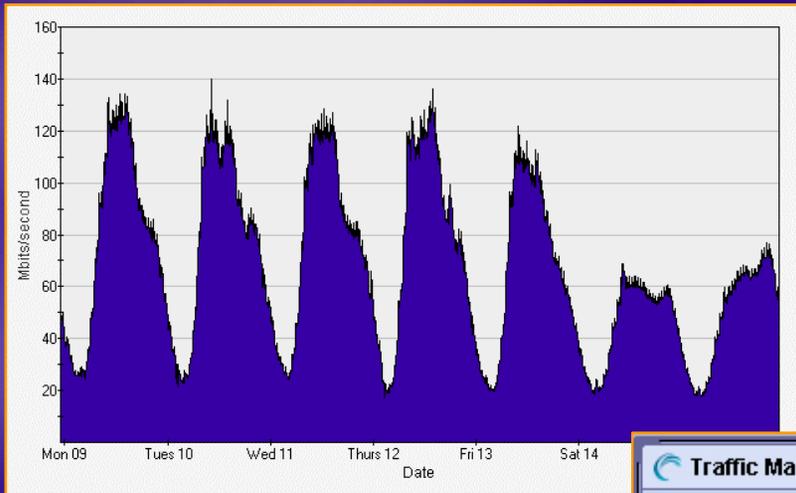


Akamai's Service Offerings

- **FirstPoint:** a global traffic management service for content providers with geographically distributed Web servers
- **EdgeScape:** allows customization of content based on user's geographic location and connection bandwidth
- **Reporter and Traffic Analyzer:** provide historical and real-time Web site usage data



Reporter and Traffic Analyzer

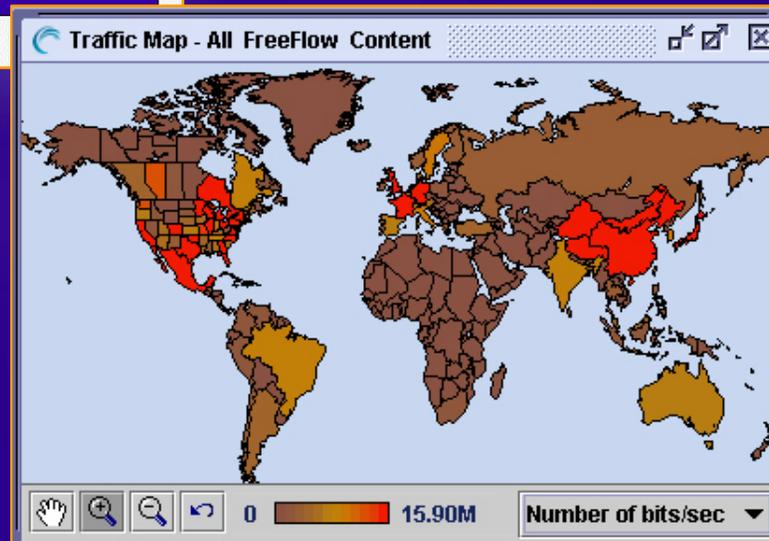


Reporter:

- For viewing of historical logs
- Customized data-mining of customer traffic

Traffic Analyzer:

- Real-time viewing of customer traffic
- Reports geographic distribution of traffic



Akamai's Service Offerings

- **ACS:** storage management service that persistently stores content delivered to end users via Akamai's network
- **Digital Parcel Service:** a comprehensive digital distribution and rights management solution
- **EdgeSuite:** enables dynamic assembly of personalized content at Akamai's edge servers



Outline

How the Web Works

Akamai' Services

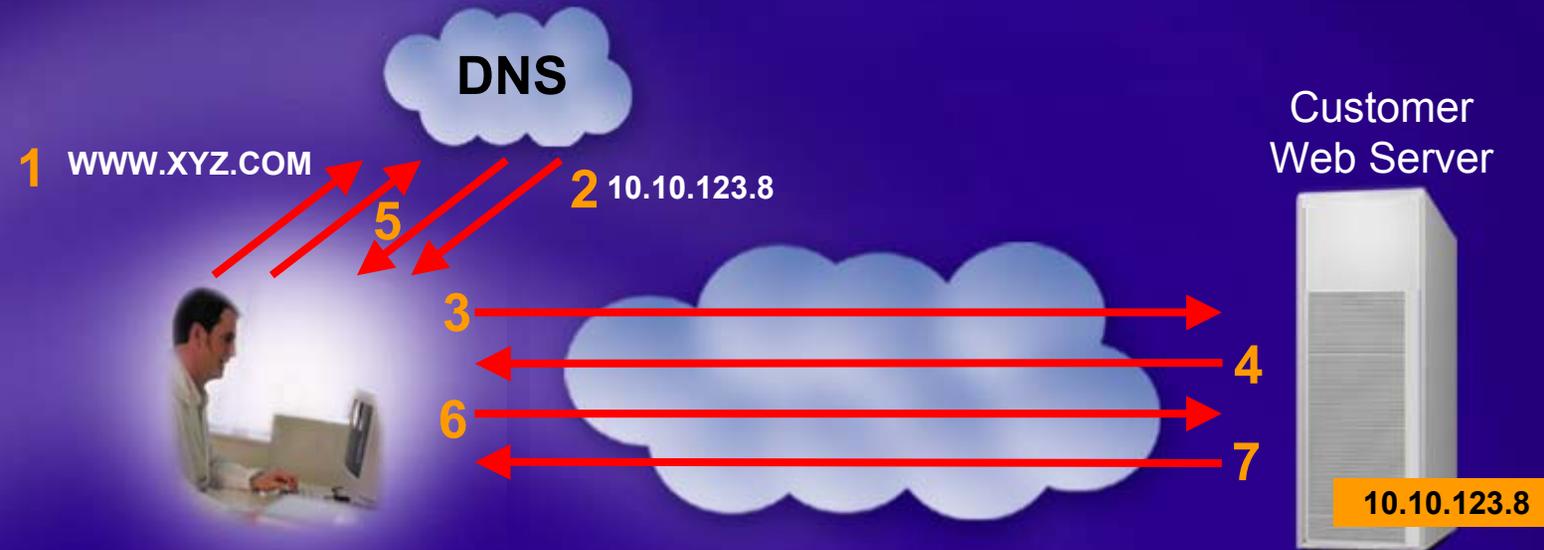
Technology Overview

Technological Challenges

The Future

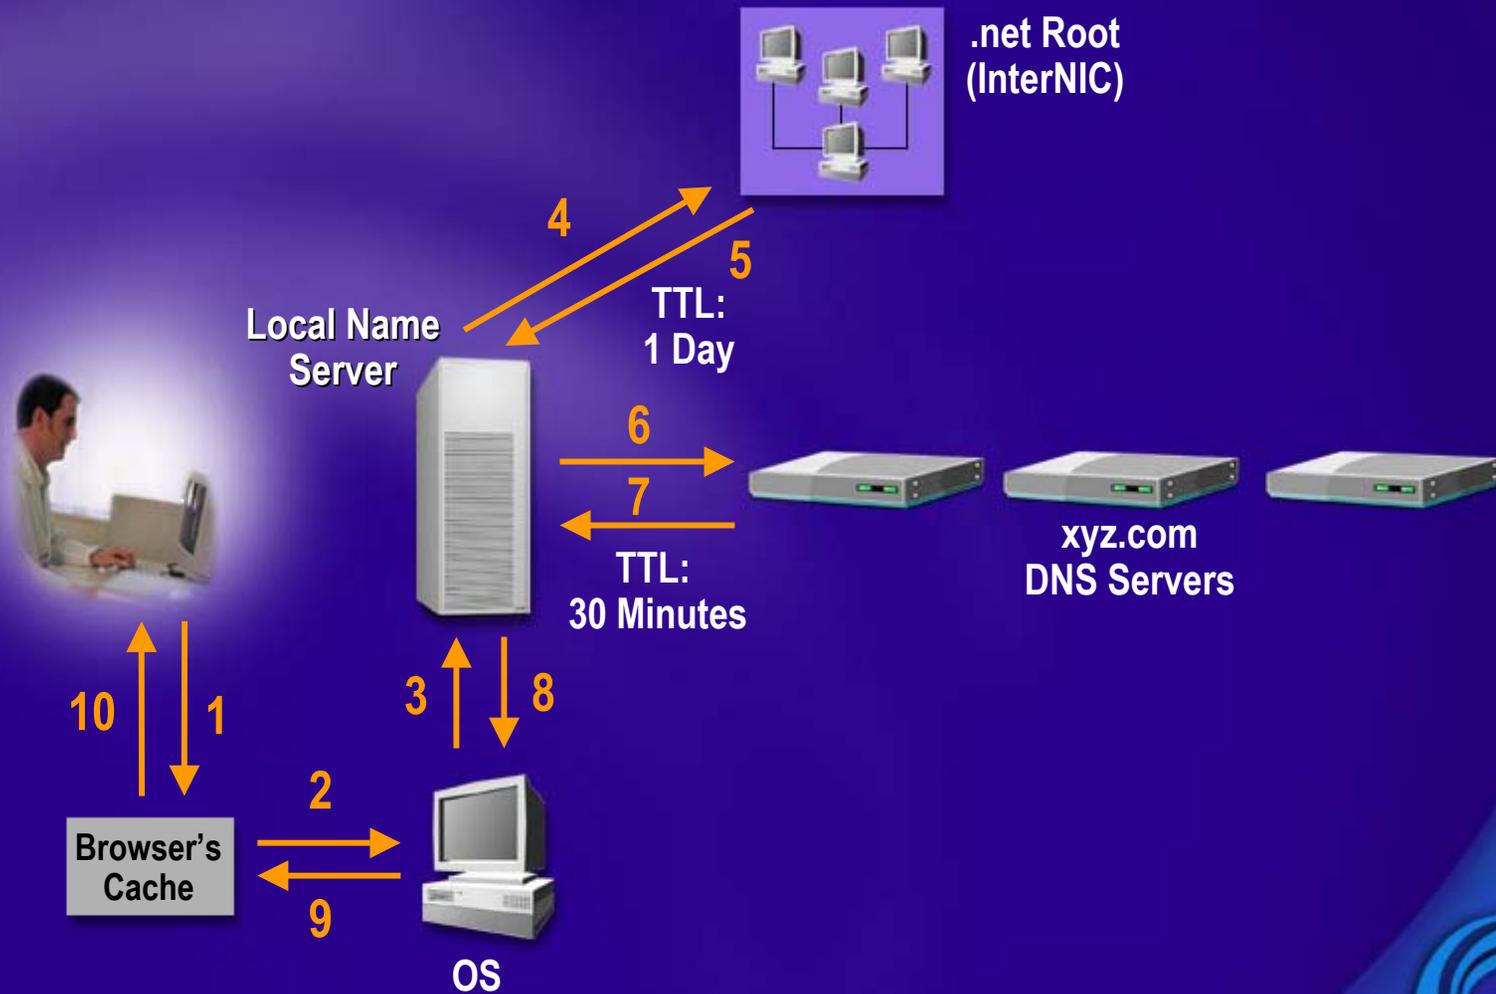


Downloading www.xyz.com — The Old Way

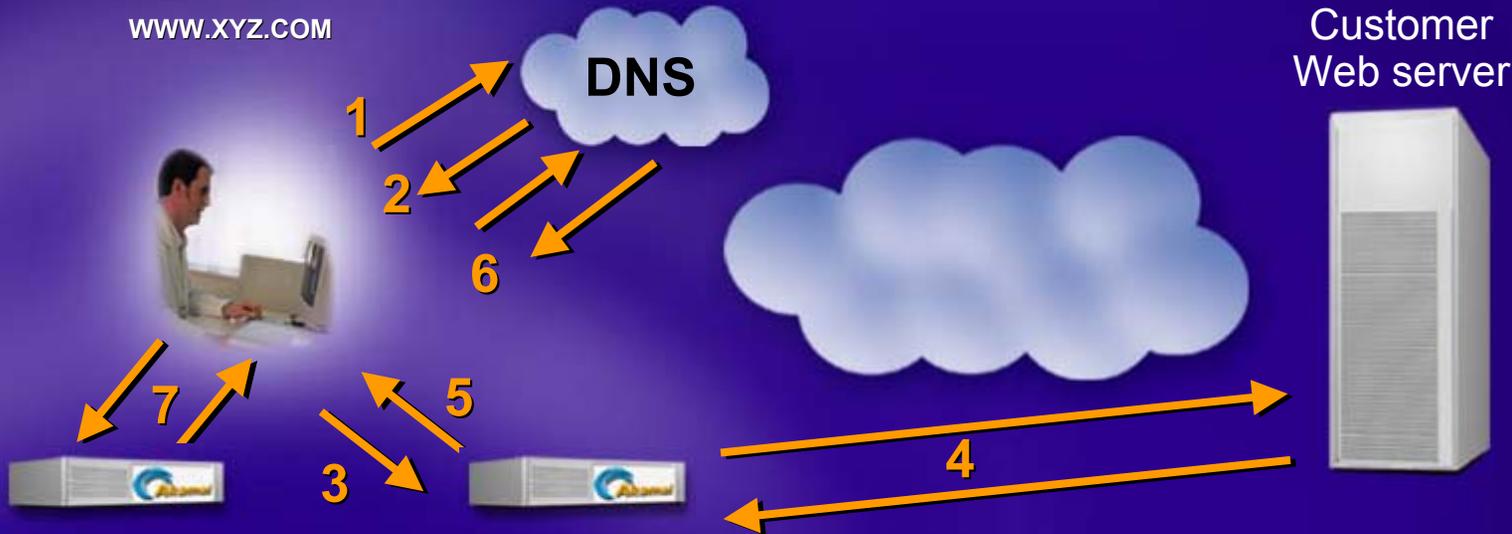


- User enters `www.xyz.com`
- Browser requests IP address for `www.xyz.com`
- DNS returns IP address
- Browser requests HTML
- Customer Web server returns HTML
- Browser obtains IP address for domain of embedded objects for page
- Browser requests embedded objects
- Customer's Web server returns embedded objects

Finding the IP Address for `www.xyz.com` — The Old Way



Downloading www.xyz.com with Akamai's EdgeSuite



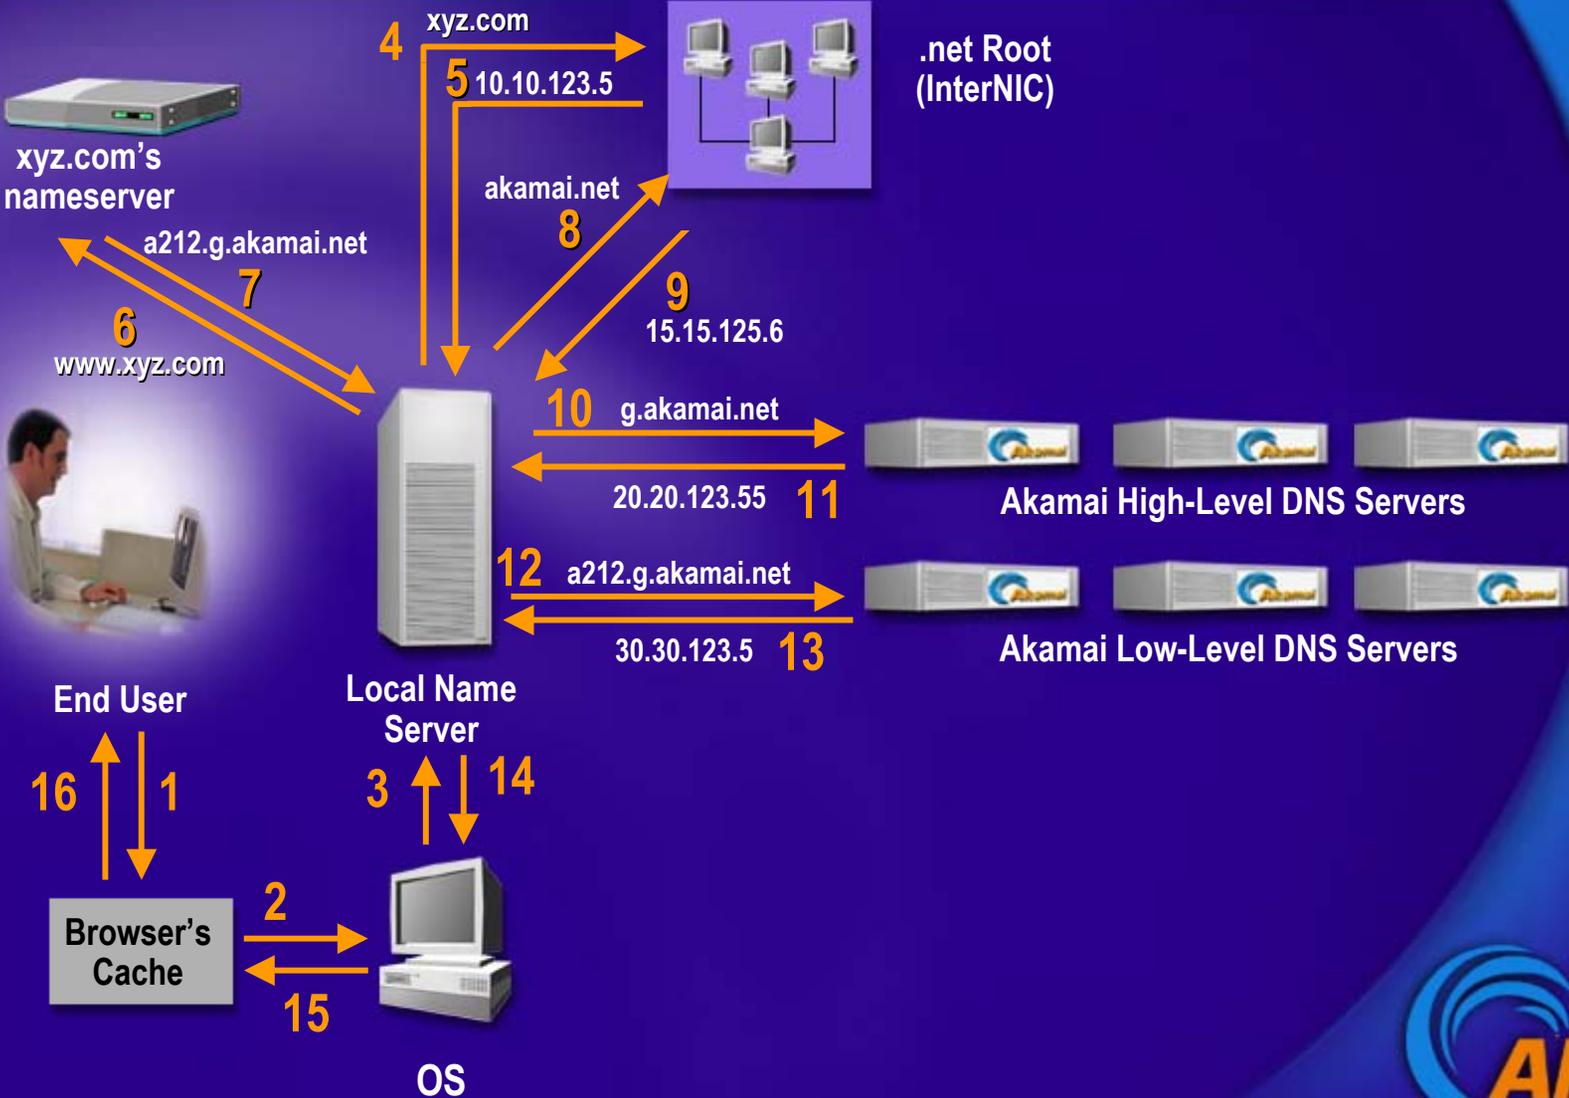
- User enters www.xyz.com
- Browser requests IP address for www.xyz.com
- DNS returns IP address of **optimal Akamai** server
- Browser requests HTML
- Akamai server assembles page, contacting customer Web server if necessary

- **Optimal Akamai** server returns Akamaized HTML
- Browser obtains IP address of **optimal Akamai** servers for embedded objects
- Browser obtains objects from **optimal Akamai** servers

Key Components

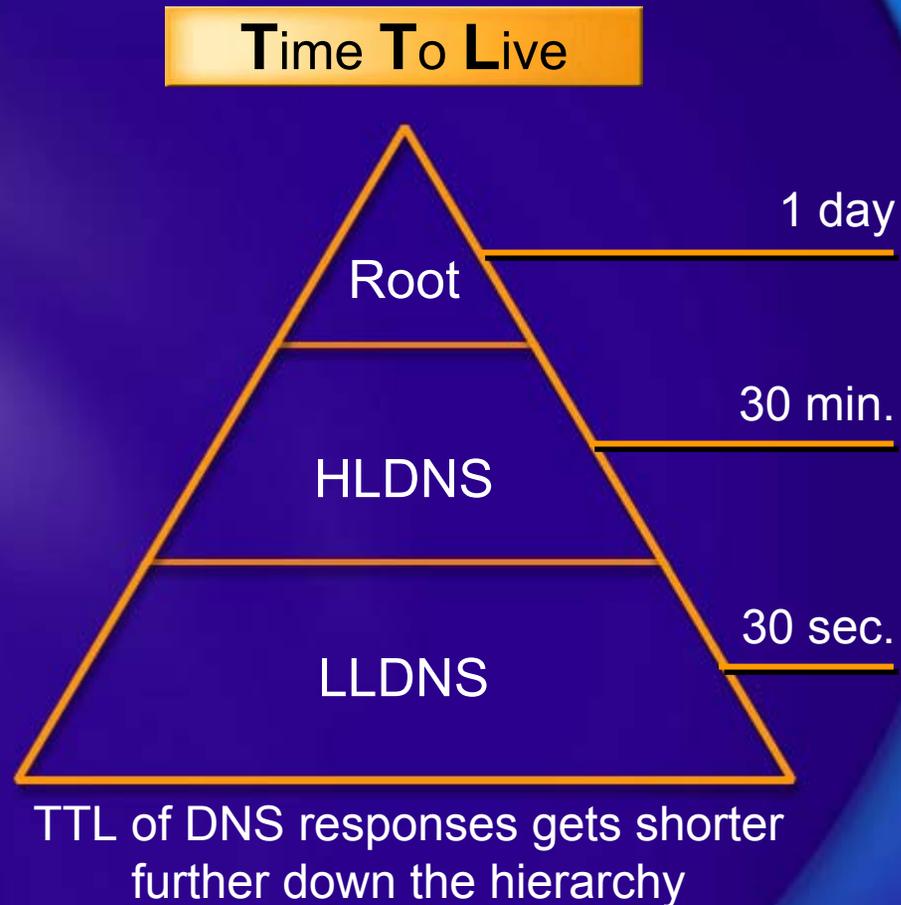
- DNS Resolution: Finding the IP address for `www.xyz.com`
- Page Assembly
- Connecting from the edge to the source

Finding the IP Address: The Akamai Way



DNS Maps & Time-To-Live

- Maps created using info on:
 - Internet congestion
 - System loads
 - User demands
 - Server status
- Maps are constantly recalculated:
 - Every few minutes for HLDNS
 - Every few seconds for LLDNS



Page Assembly



Container Page
[TTL=5d]

Site owners create container pages that can be populated with varying content

Page Assembly



If gender = male and geography = New York then show article on NY Giants Sportswear

EdgeSuite enables Web sites to build and deliver customized content at the edge

Syntax Example

- Similar to SSI so it's easily understood— but also breaks performance bottlenecks when distributed across 12,000 servers
- Used as an API to third-party applications on Akamai's network

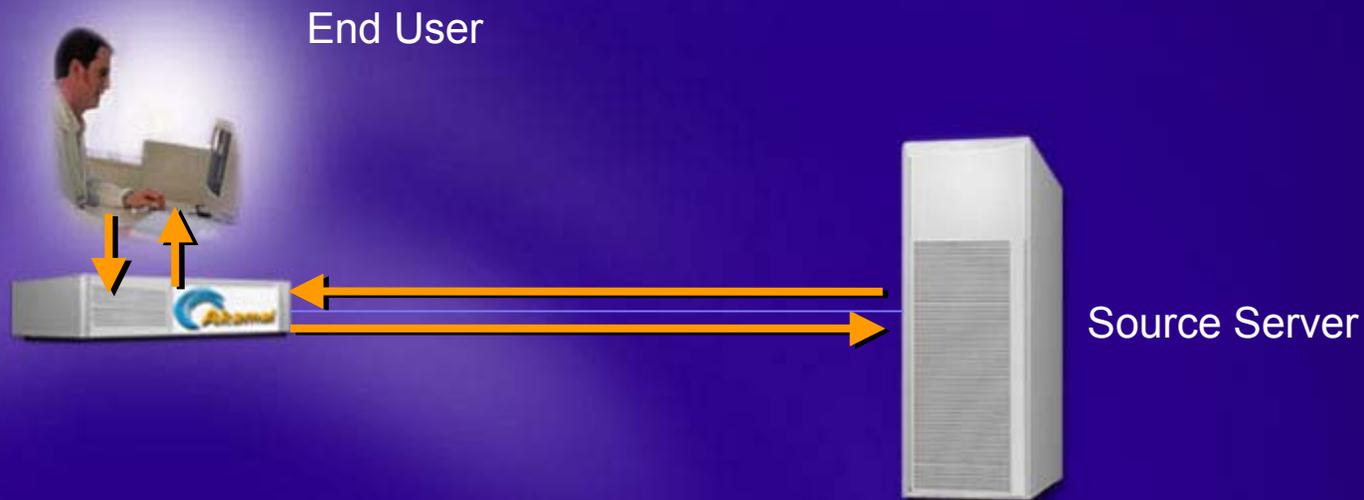
```
<html>
<asi version = "1.0">
<head><title>Akamai Weather Forecast
</title></head>

<asi if= "$geo == dc" include="dcwthr.htm"/>
<asi else-if= "$geo == boston" include="bostonwthr.htm"/>
<asi else if = "$geo == nyc " include= "nycwthr.htm"/>
<asi else == " " include= "uswthr.htm"/>
</asi>

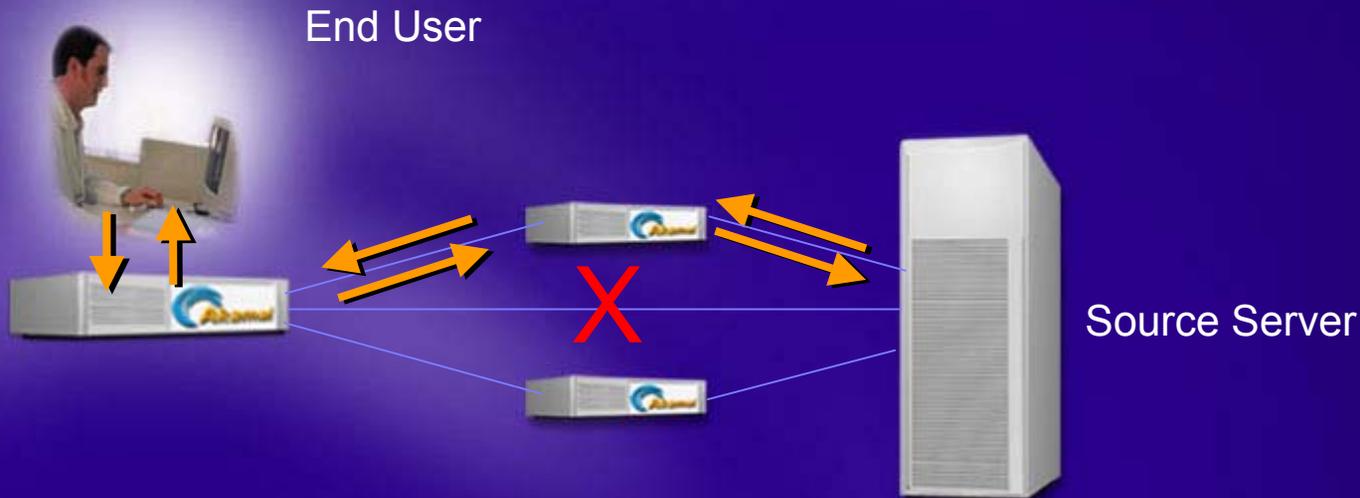
</html>
```



Connecting from the Edge to the Source

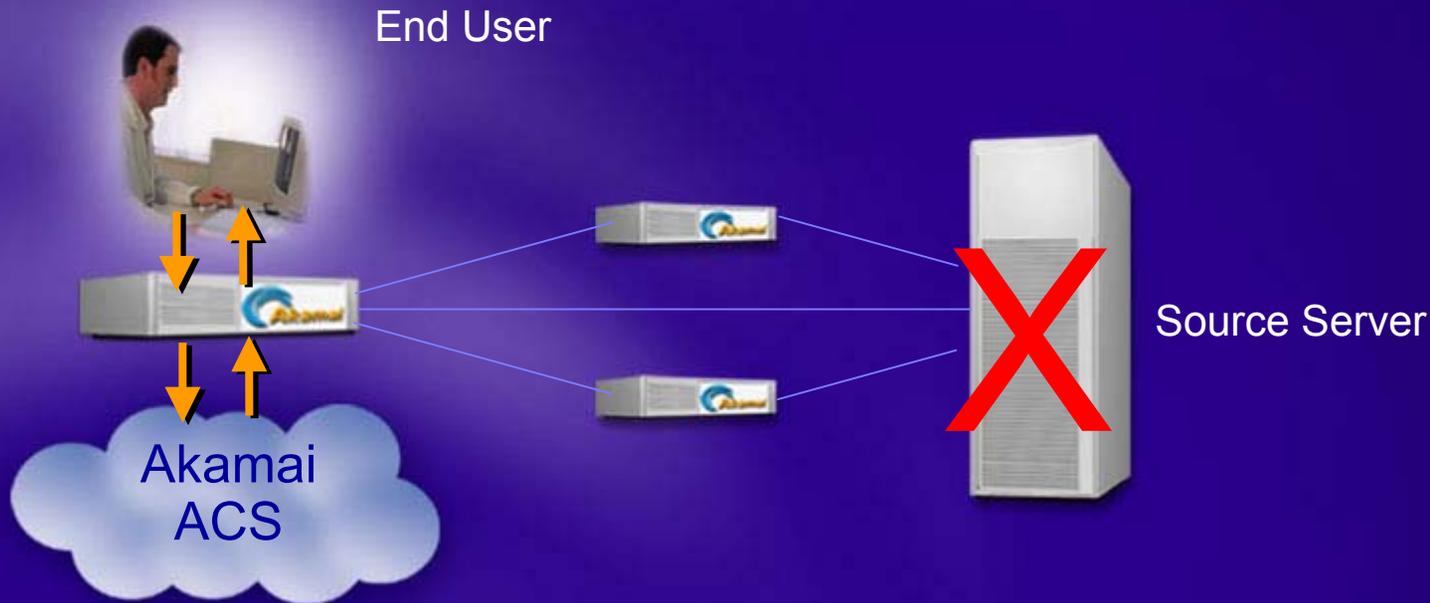


Connecting from the Edge to the Source



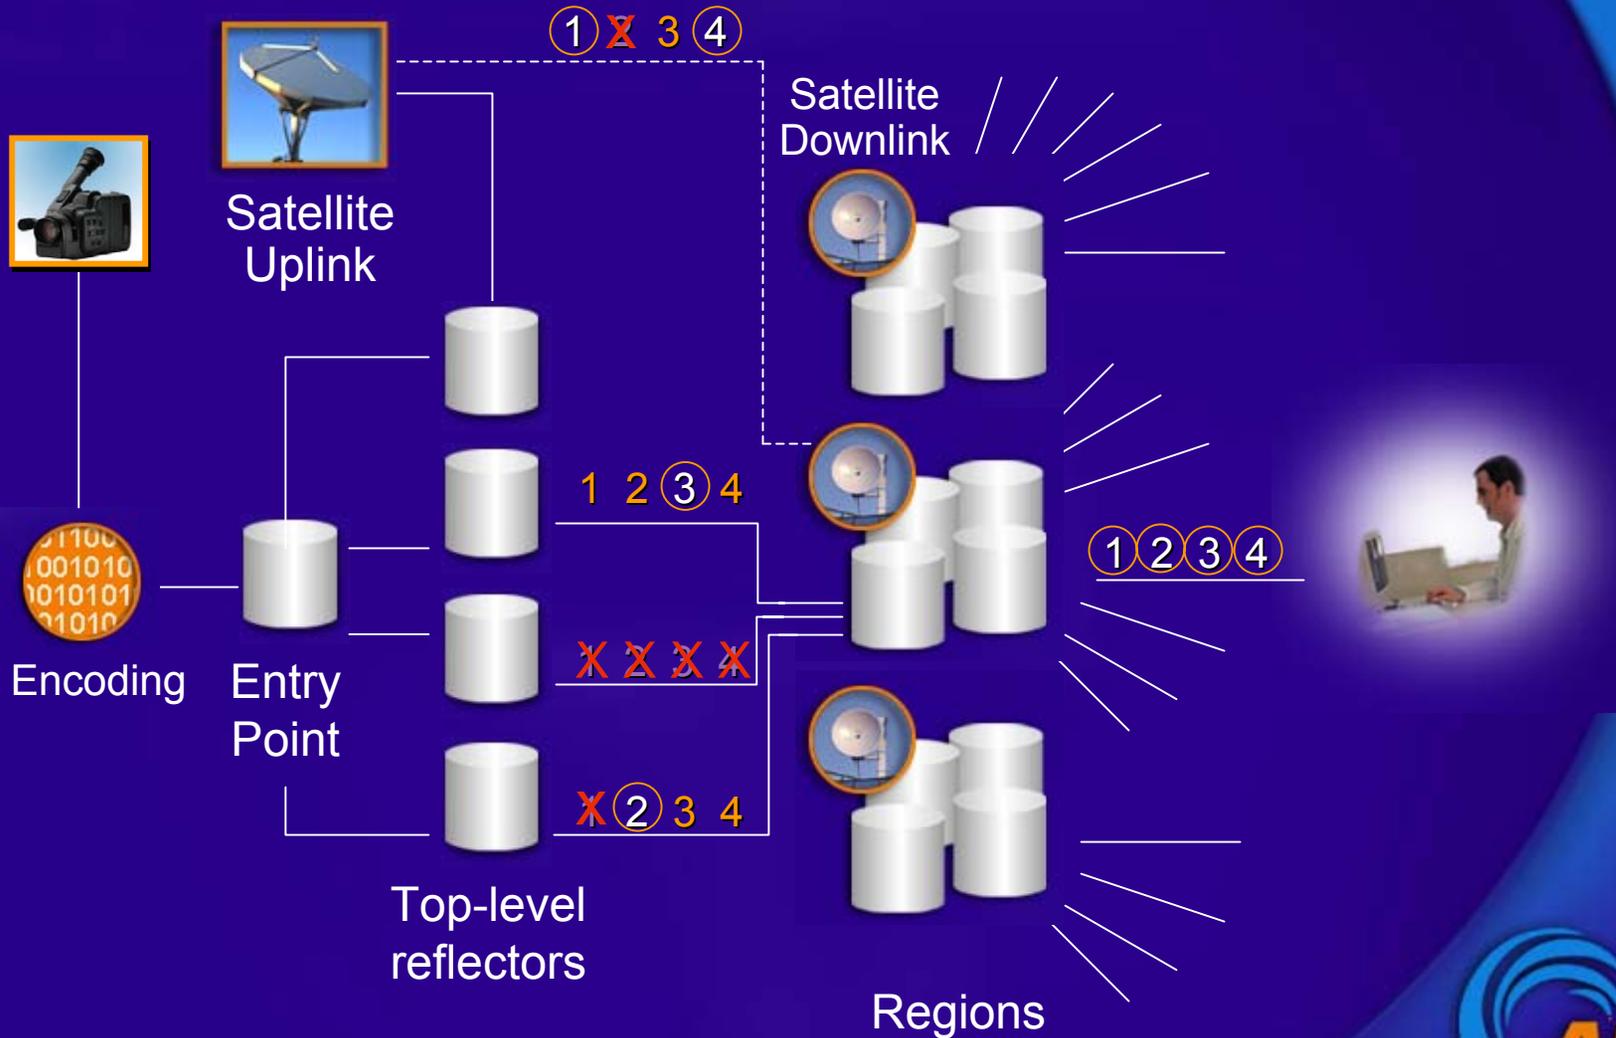
- Maintain path performance data so that the optimal path can be used to reach optimal customer location (Akarouting)

Connecting from the Edge to the Source



- Maintain path performance data so that the optimal path can be used to reach optimal customer location
- If site is unreachable, then retrieve authoritative default information from Akamai's ACS service

Differences for Live Streaming



Outline

How the Web Works

Akamai's Services

Technology Overview

Technological Challenges

The Future



Technological Challenges

1) Mapping and Server Assignment

- Huge scale
 - Hundreds of millions of users
 - Tens of thousands of servers
 - Thousands of locations
 - Thousands of customers
- Must monitor Internet conditions & instantly respond to changes
 - Internet congestion and failures are widespread and unpredictable
- Must load balance widely varying kinds of traffic, optimize multiple kinds of resources, and minimize various costs
- Must tolerate large numbers of component failures without ever disrupting service
- Control algorithms must be distributed and work with imperfect information
- DNS responses must be given in milliseconds



Technological Challenges

2) Logging, Reporting, and Billing

- Several billion hits per day and growing rapidly
- Real-time reporting of data to customers with user profiling and customized data mining
- Real-time monitoring of system for NOCC with meaningful alerts and performance metrics
- Support for real-time SQL queries to the system

Technological Challenges

3) Operations

- Huge distributed network that cannot go down, despite frequent software changes and constant growth
- System needs to be secure against attacks as well as buggy third-party software

Technological Challenges

4) Content Freshness and Accuracy

- Stale content can never be served
- Fast correction procedures in case of customer errors
- Allow flexibility and ease of customer control over content

Technological Challenges

5) Management of Live Streaming and Webcasting

- Information dispersal to handle packet loss
- Automatic swapping on connections to improve performance
- Support for interactive and personalized messaging; e.g., Q&A
- Real-time data aggregation for polling, etc.
- Synchronized delivery of audio, video, and slides

Outline

How the Web Works

Akamai's Services

Technology Overview

Technological Challenges

The Future



The Premier African-American Portal

“On Monday our traffic doubled, so we added two new servers with no effect.

Tuesday afternoon we called Akamai. Tuesday night we were Akamaized and instantly 6-10 times faster.”

Craig Maccubbin
CTO of BET.com

BET.com Akamaized 90% of Each Web Page with FreeFlow:

- Improved site performance (6-10 times)
- Quadrupled page view capacity
- Postponed 2nd data center build out
- Preserved graphic-rich page design
- Improved user experience
- Net first year benefits of \$1.3 million

Questions?

