The Economics of Content Delivery

Bruce Maggs

Duke University
Emerald Innovations
(and formerly Akamai Technologies)

In the beginning…

- Akamai deployed servers around the world and offered to deliver images that appeared on web pages.
- The goal was to speed up page load times
- … and reduce infrastructure investment for content providers.

Impact of PLT on conversion for e-commerce web sites

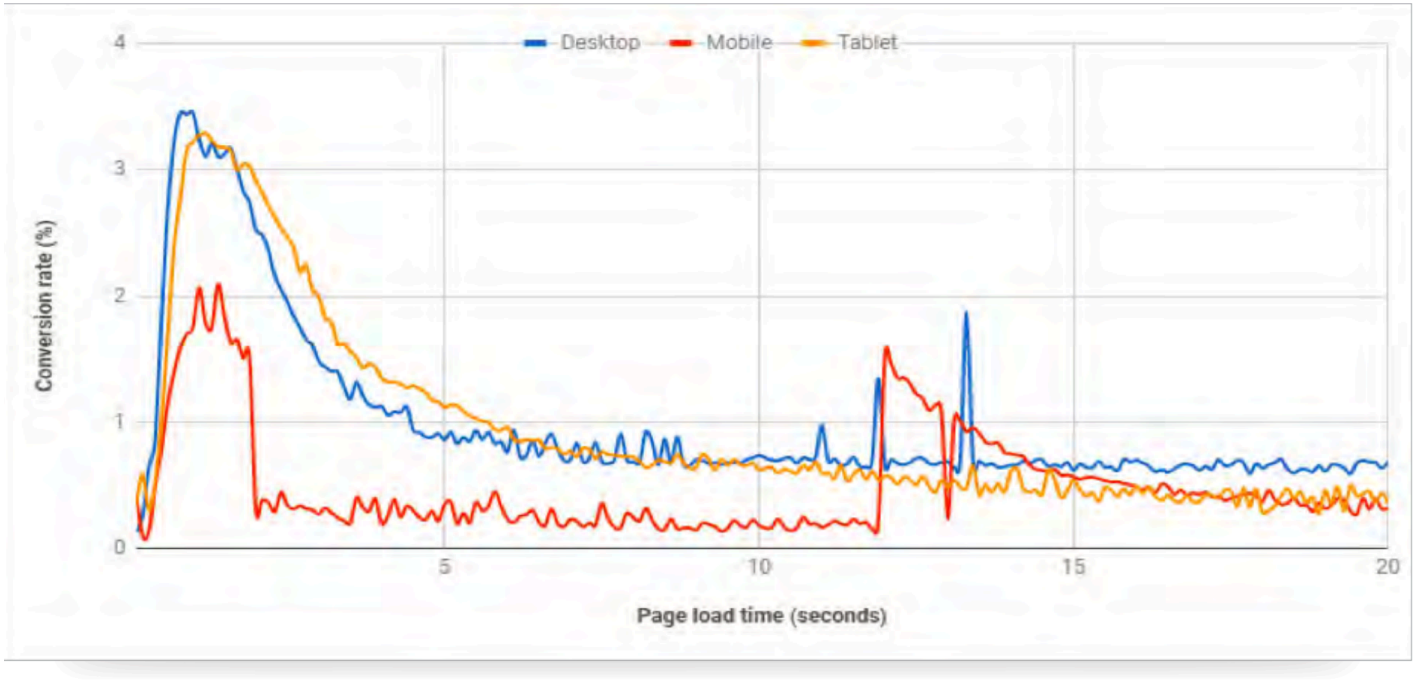

Impact of page load times on conversion rates for desktop, mobile, and tablet.

Source: Akamai, State of Online Retail Performance, 2017 Holiday Perspective

TECH • FORTNITE

New Fortnite Release Downloads Shatter Record for Internet Traffic

By CHRIS MORRIS July 13, 2018

Obama's inauguration? The 2016 presidential election? They've got nothing on *Fortnite*.

The release of the game's fifth season, a notable update that players have been anxiously awaiting, broke Internet game traffic records, according to Akamai. Actually, no ... it shattered those records.

37Tbps

To put things in perspective, the 2016 U.S. Presidential election saw peak traffic on Akamai's platform of 7.5 Tbps. Yesterday's *Fortnite* update peaked at 37 Tbps. That's nearly five times the bandwidth.

Embedded Image Delivery

Embedded URLs are Converted to ARLs

```
<html>
<head>
<title>Welcome to xyz.com!</title>
</head>
<body>
<img src="http://www.xyz.com/logos/logo.gif">
<img src="http://www.xyz.com/jpgs/background.jpg">
<h1>Welcome to our Web site!</h1>
<a href="page2.html">Click here to enter</a>
</body>
</html>
```

ak

Move to Full-Site Delivery

88 of Alexa top 500 home pages delivered by Akamai

Motivated in part by split-TCP/TLS optimization

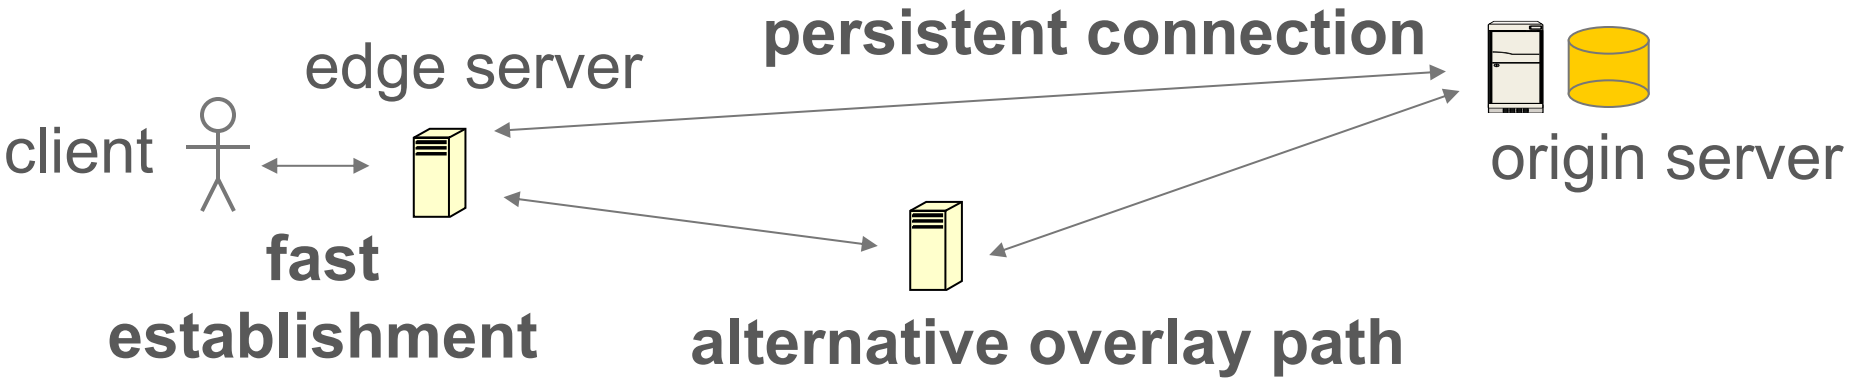

Motivated in part to protect web site

Akamai operates DNS for main domain, holds private keys!

The Akamai Platform and Services

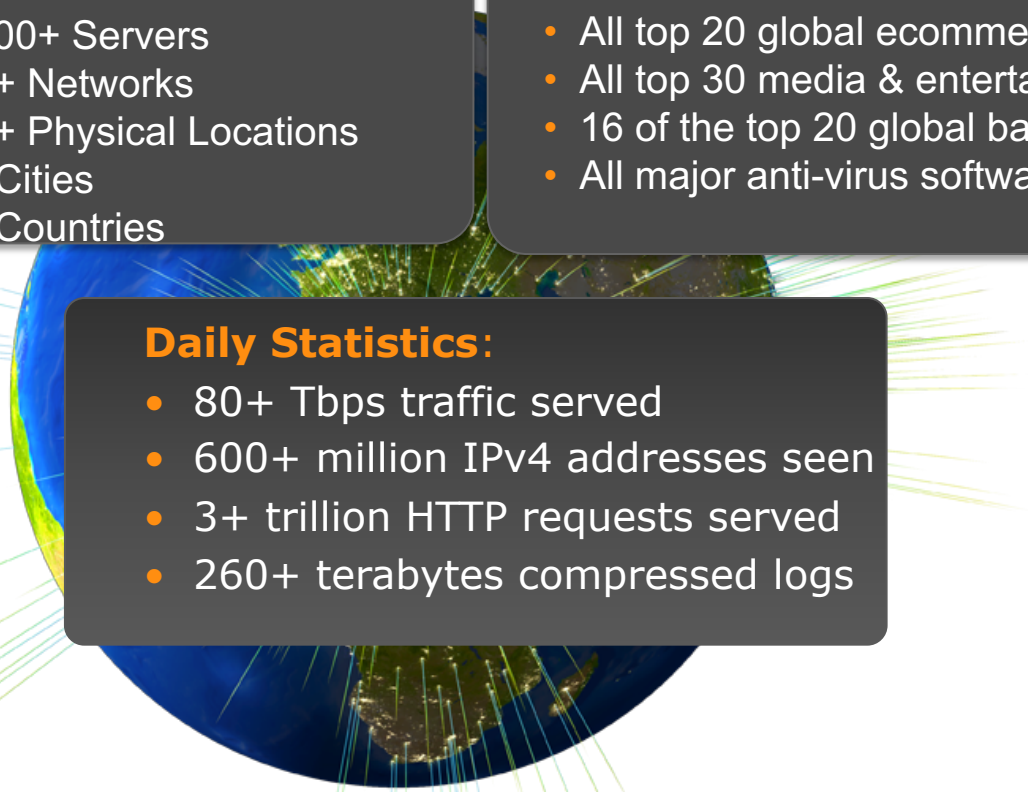

A Global Platform:

- 240,000+ Servers
- 1,700+ Networks
- 3,300+ Physical Locations
- 750+ Cities
- 130+ Countries

Delivering Content for 130,000+ Domains

- All top 20 global ecommerce sites
- All top 30 media & entertainment companies
- 16 of the top 20 global banks
- All major anti-virus software vendors

Daily Statistics:

- 80+ Tbps traffic served
- 600+ million IPv4 addresses seen
- 3+ trillion HTTP requests served
- 260+ terabytes compressed logs


Akamai

Flow of Payments

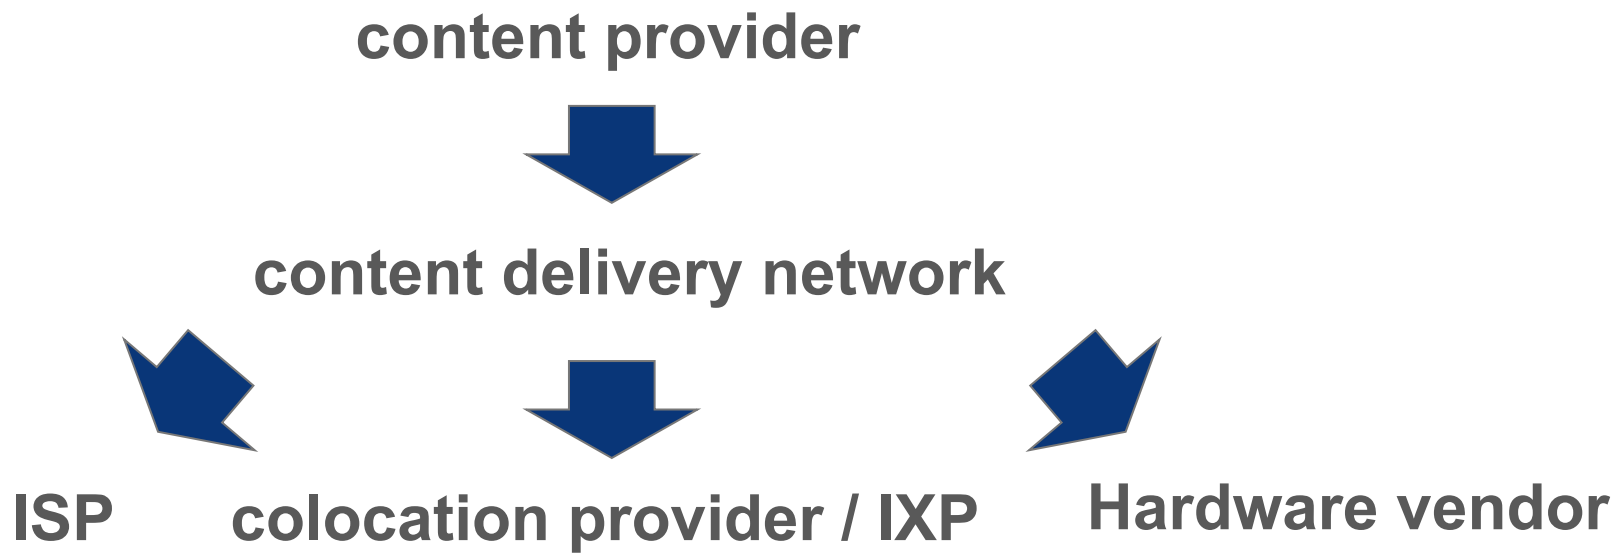

Cost of Goods Sold (CoGS)

25% bandwidth (95/5 rule)
25% colocation
25% server depreciation (now over 4 years)
25% other

colocation breakdown:
40% rack space
40% energy (now metered)
20% other (e.g., interconnect within facility)

Mapping Clients to Akamai Clusters

- Many “no brainers” – servers co-located with clients, e.g., on university campuses
- Otherwise, make decisions based on active measurements

Wide Area Network Measurement

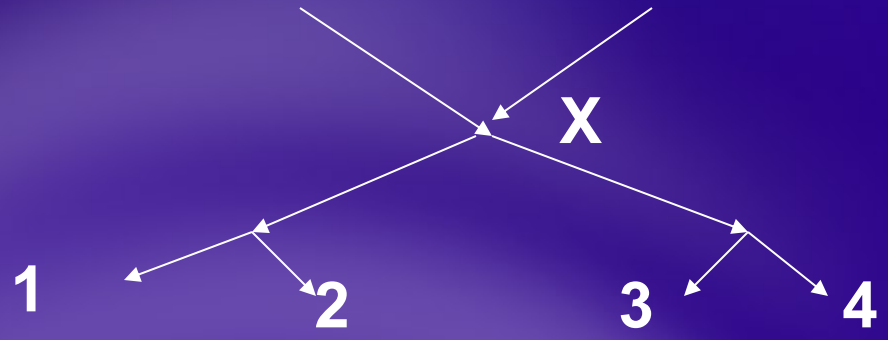

- Traceroute to all important name servers
- Core point X is the first router at which all paths to name servers 1, 2, 3, and 4 intersect.
- X can be viewed as the straddling the core and the edge of the network.

Identifying Core Points

500,000 nameservers

Reduced to

30,000 core points

7,000 account for 95% end-user load

ping these continuously

Bandwidth Breakdown

40% free from edge servers to clients

40% paid from edge servers to clients

20% from edge servers to origin servers

Next Service: Streaming Media

- Proprietary servers and players
  Real, Windows Media, Quicktime, Flash
- UDP-based protocols
- Network for live streaming built for Steve Jobs.

Live Streaming Architecture

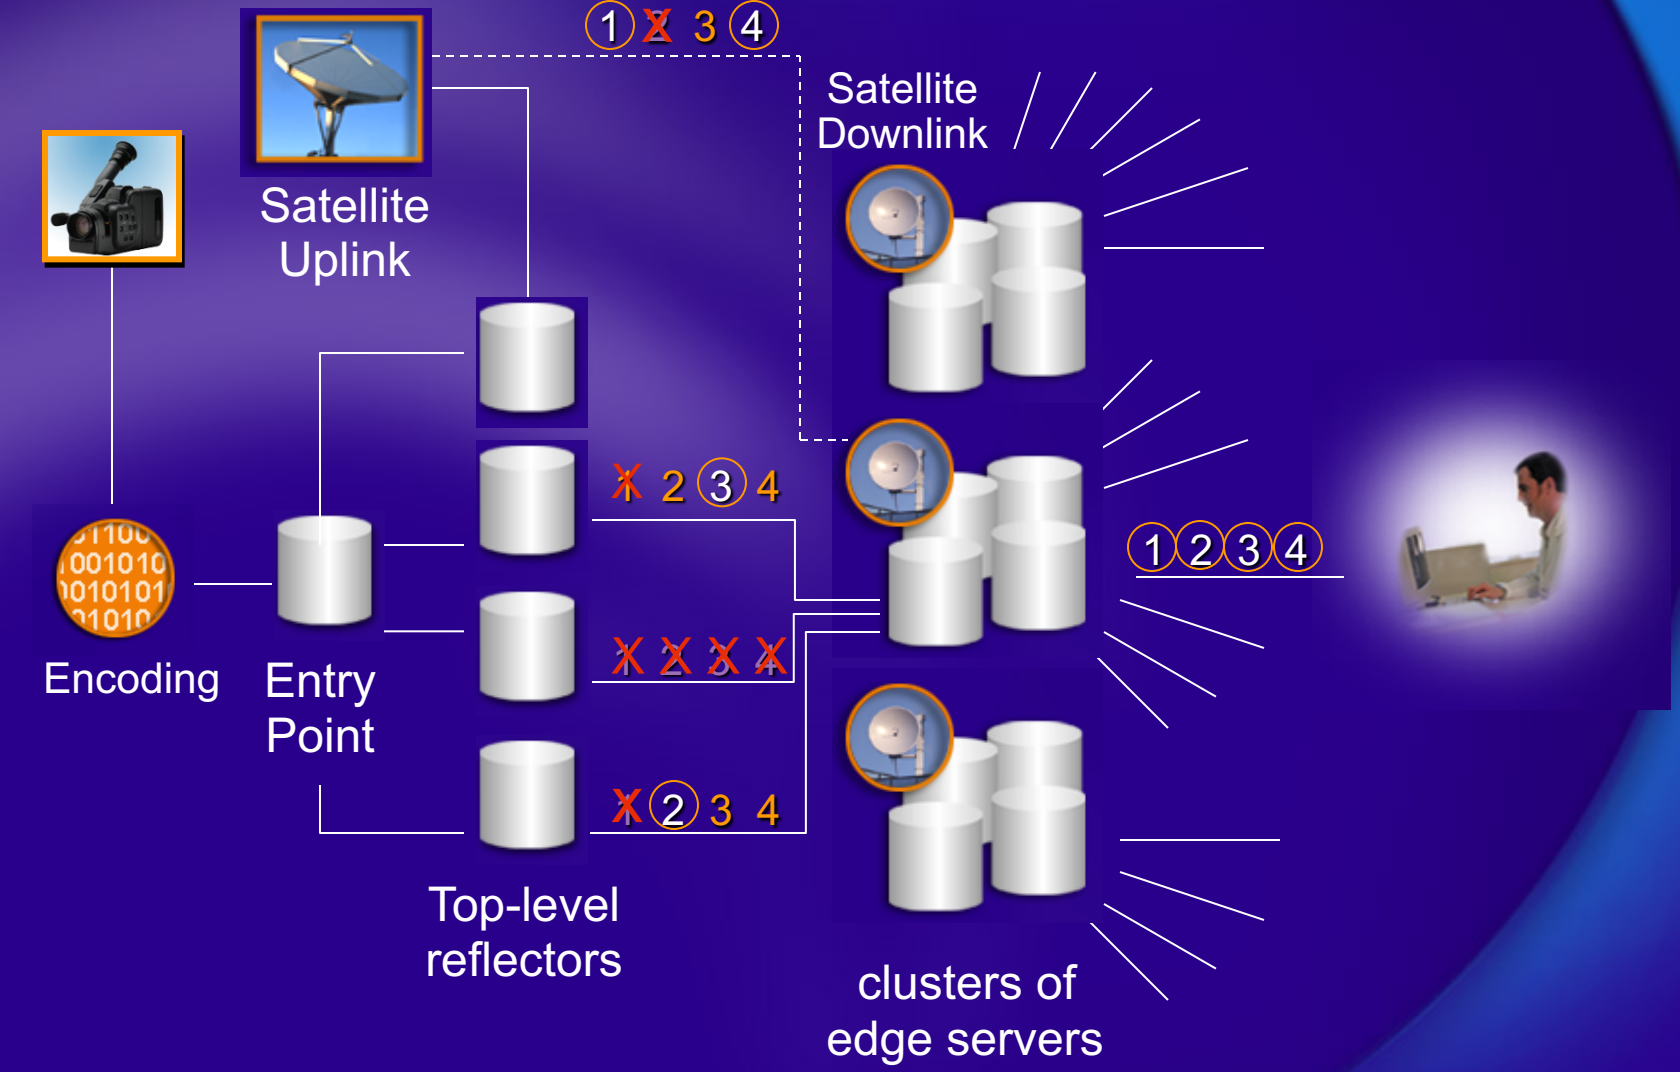

Streaming Today

- Everything delivered over TCP/HTTP(S)
- Client (typically JavaScript) makes all decisions
- “Brokers” may choose from multiple CDNs
- Vast majority of bytes served by CDN
- Highly commoditized, low margin business

High Margin Services

Low latency content delivery

Bells & whistles (e.g., automatic A/B testing)

Protection from denial of service attacks

Defending Web-Based Services from Attacks

Two Architectures:

Traditional CDN protects web sites

Packet filtering data centers provide generic protection for network-based services

Distributed Denial of Service (DDOS) Attacks

The attacker hopes to overwhelm the content provider's resources with requests for service.

Sometimes the attacker issues requests through a "bot army" of compromised or rented machines.

The attacker looks for "amplification" where an easy-to-generate request requires a large or difficult-to-generate response.

Largest DDOS Attacks by Year

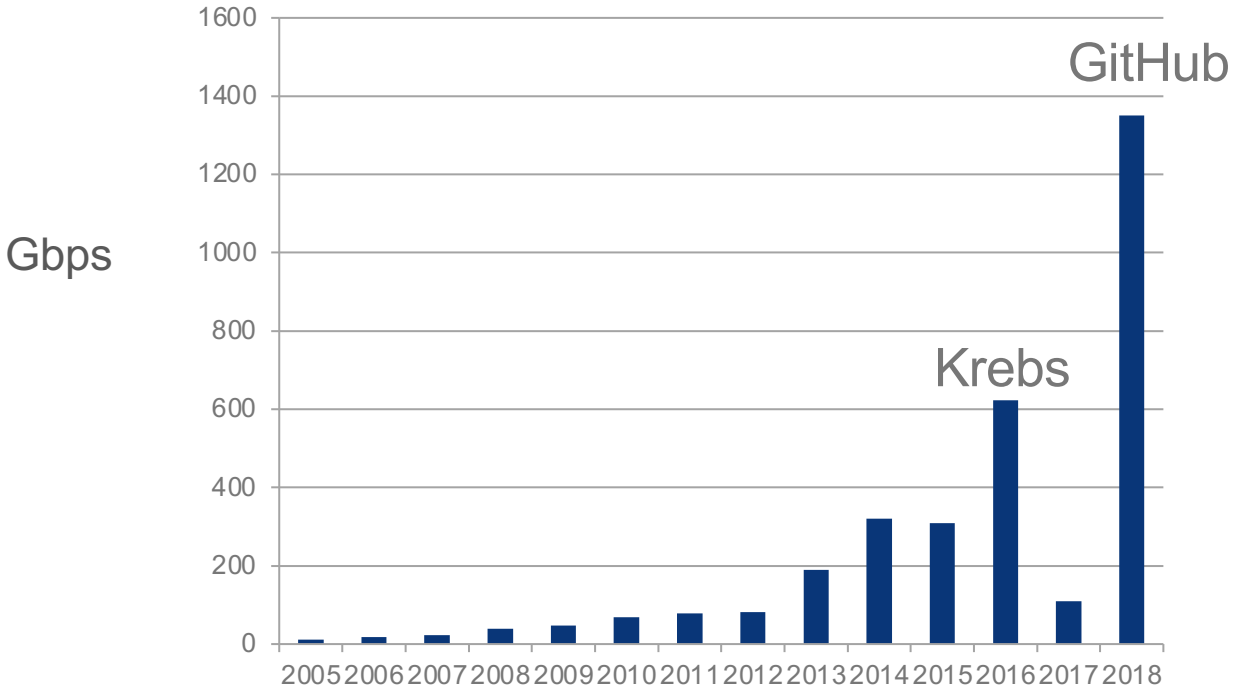

Akamai Kona Site Defender

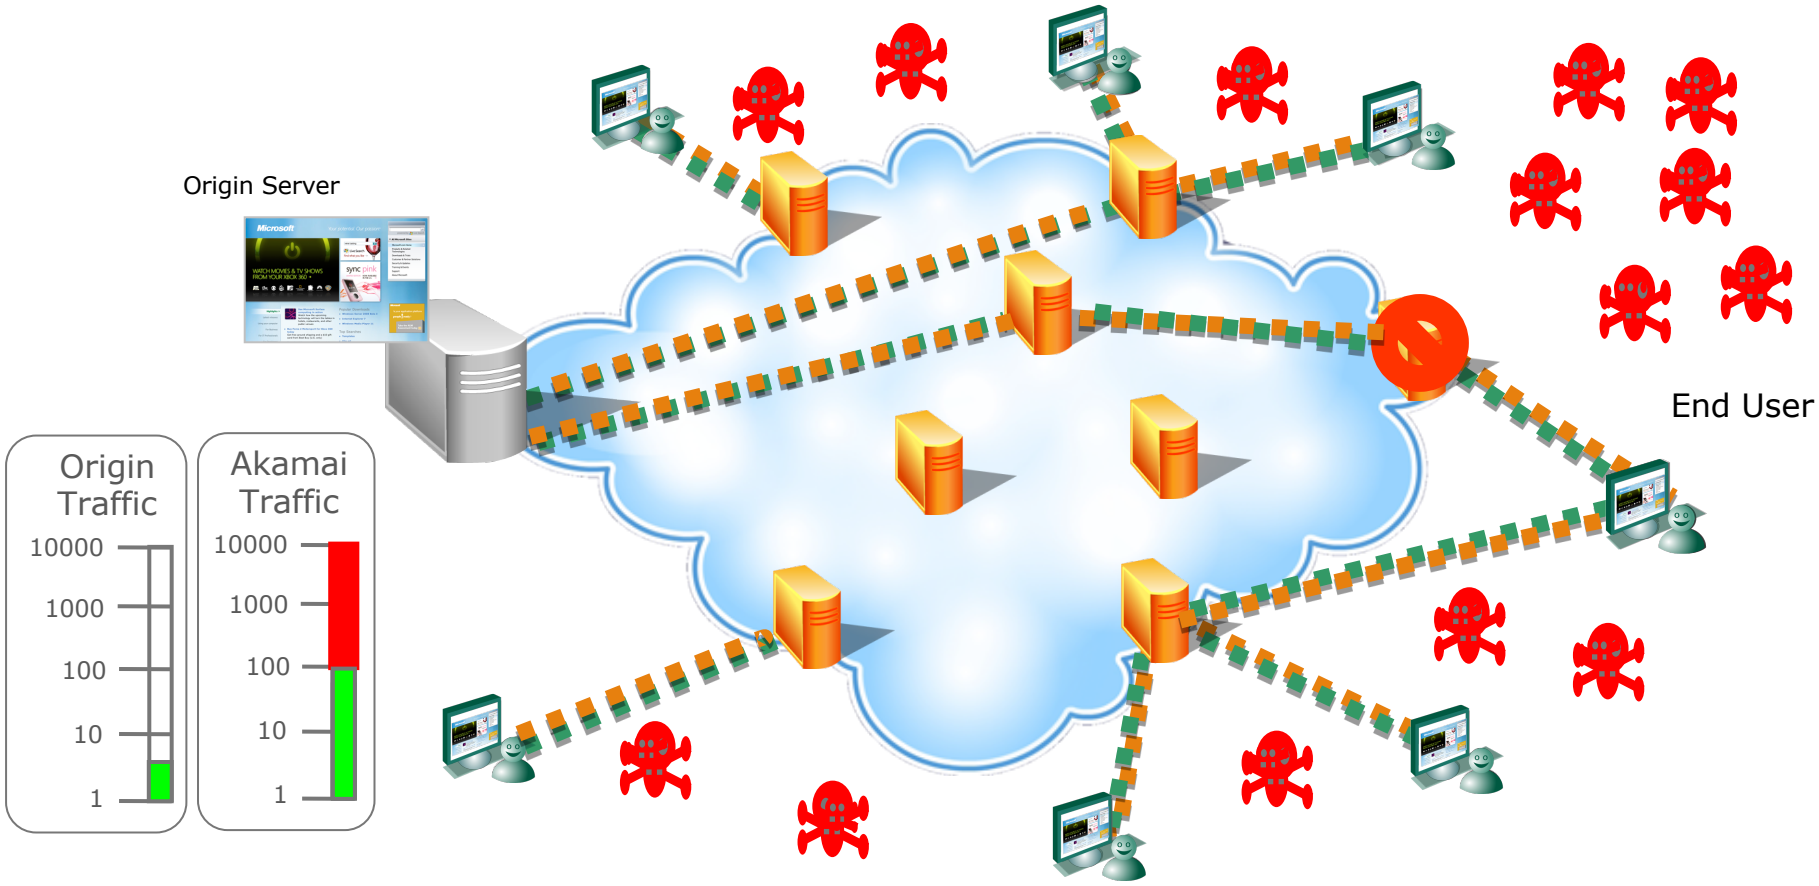

Malicious login attempts

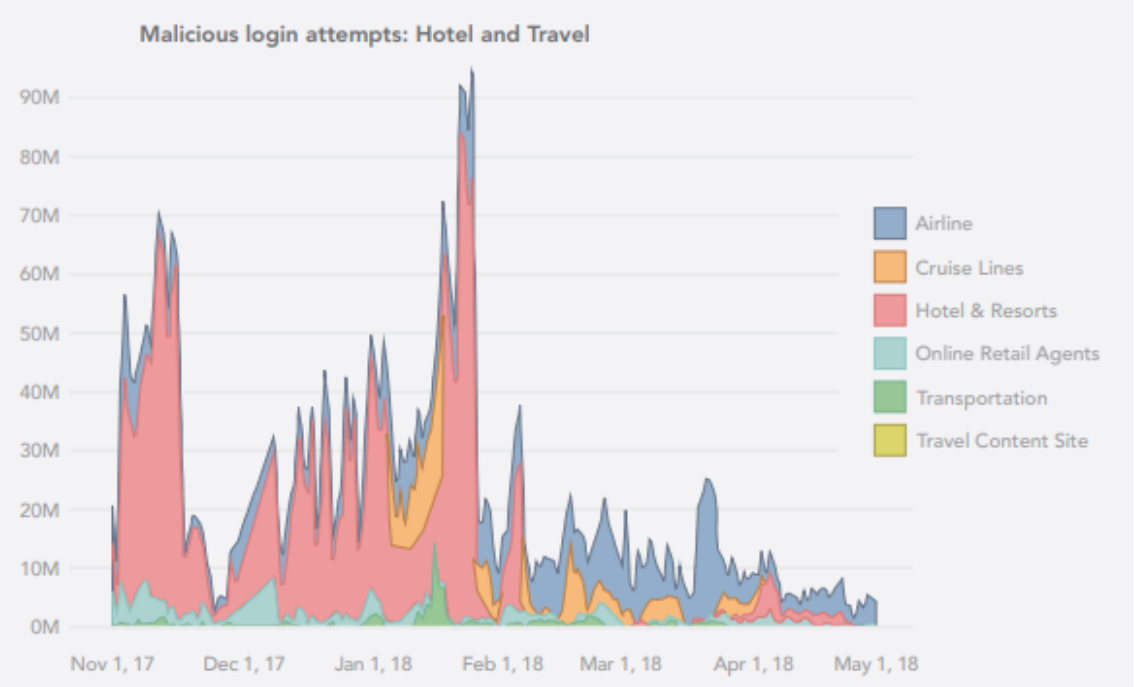

Prolexic IP Anycast Scrubbing Centers

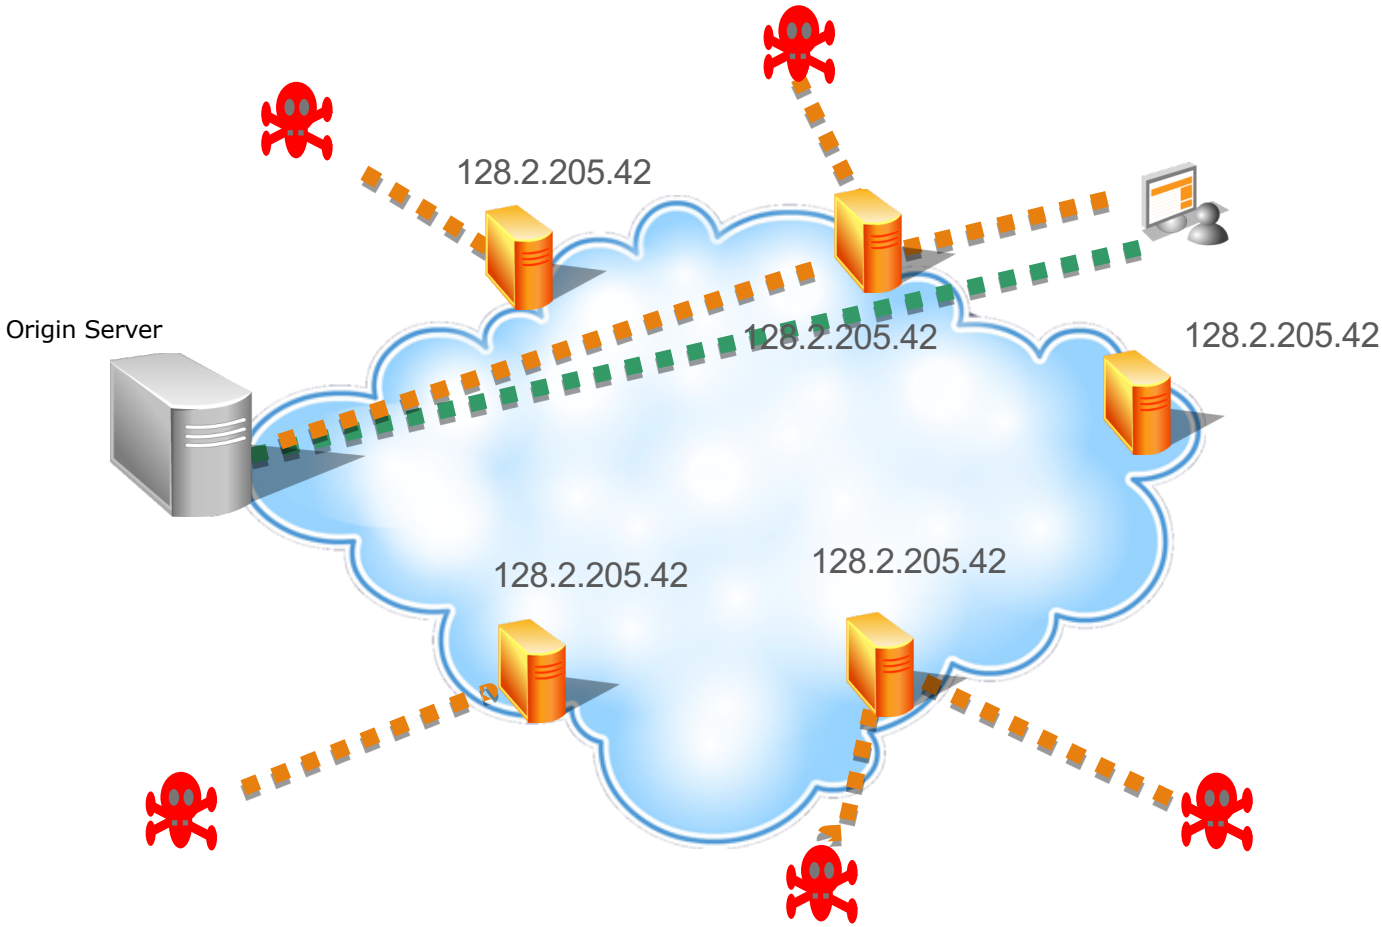

Moving Web Applications to the Edge

Standard three-tiered architecture

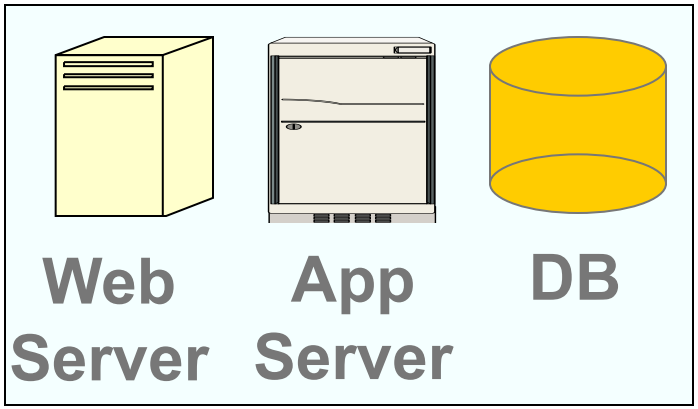

Original CDN Application

Clients

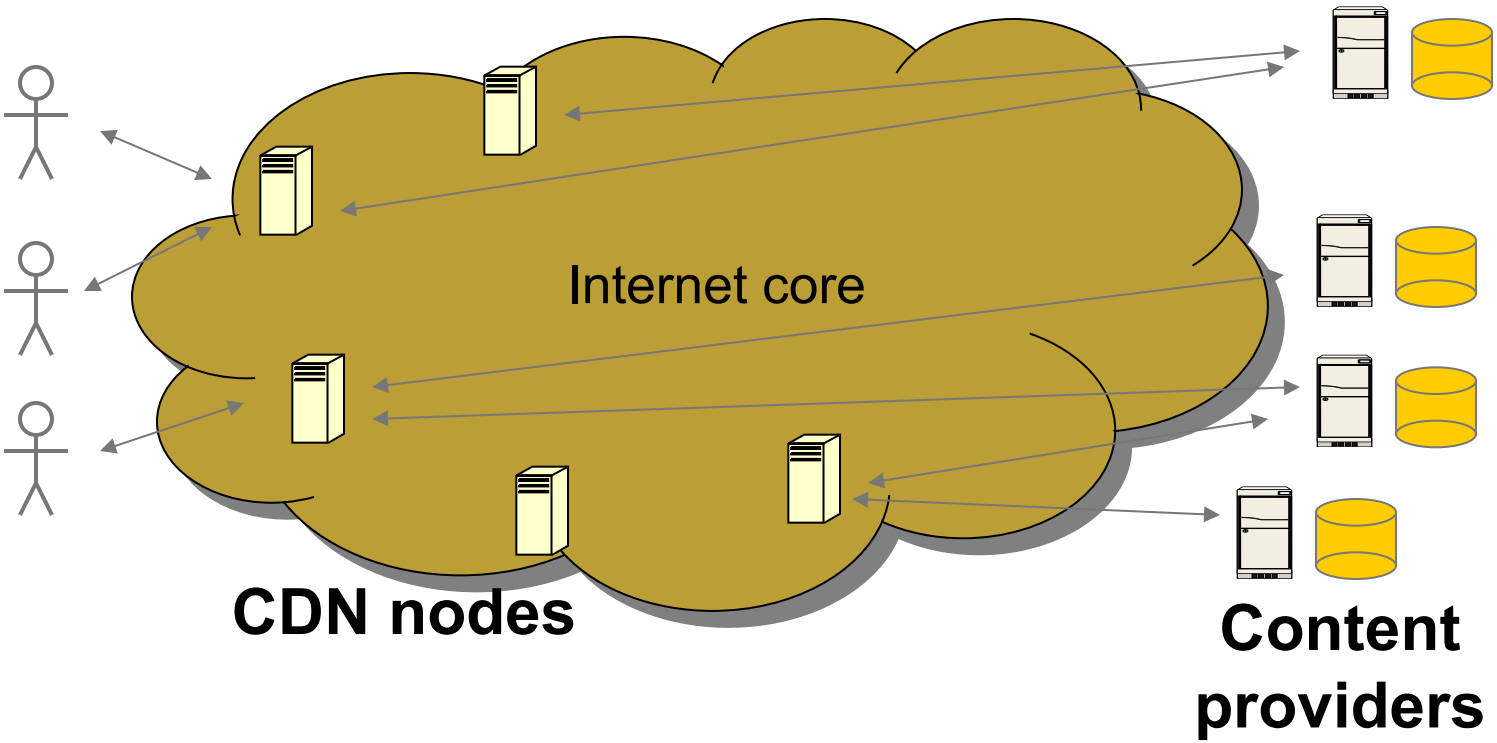

CDN Application Services

Huge investment to support IBM WebSphere at edge

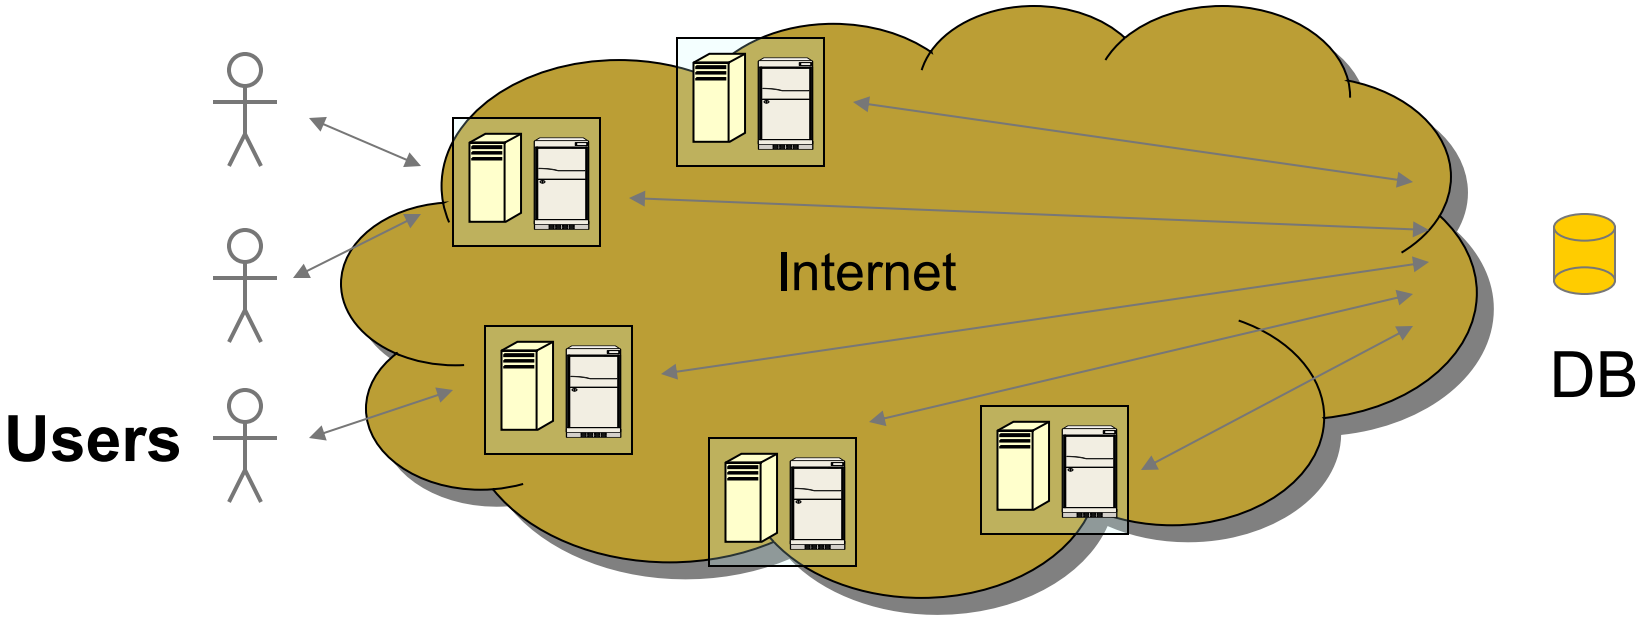

But… high latency to database server, and for data-intensive applications database server becomes the bottleneck.

Peer-to-Peer Assisted Content Delivery

Hybrid between a fixed-infrastructure CDN and a pure peer-to-peer delivery system

Based on technology developed by

- Red Swoosh (acquired by Akamai in 2007)

Goal is to deliver large files at lower cost

Architecture

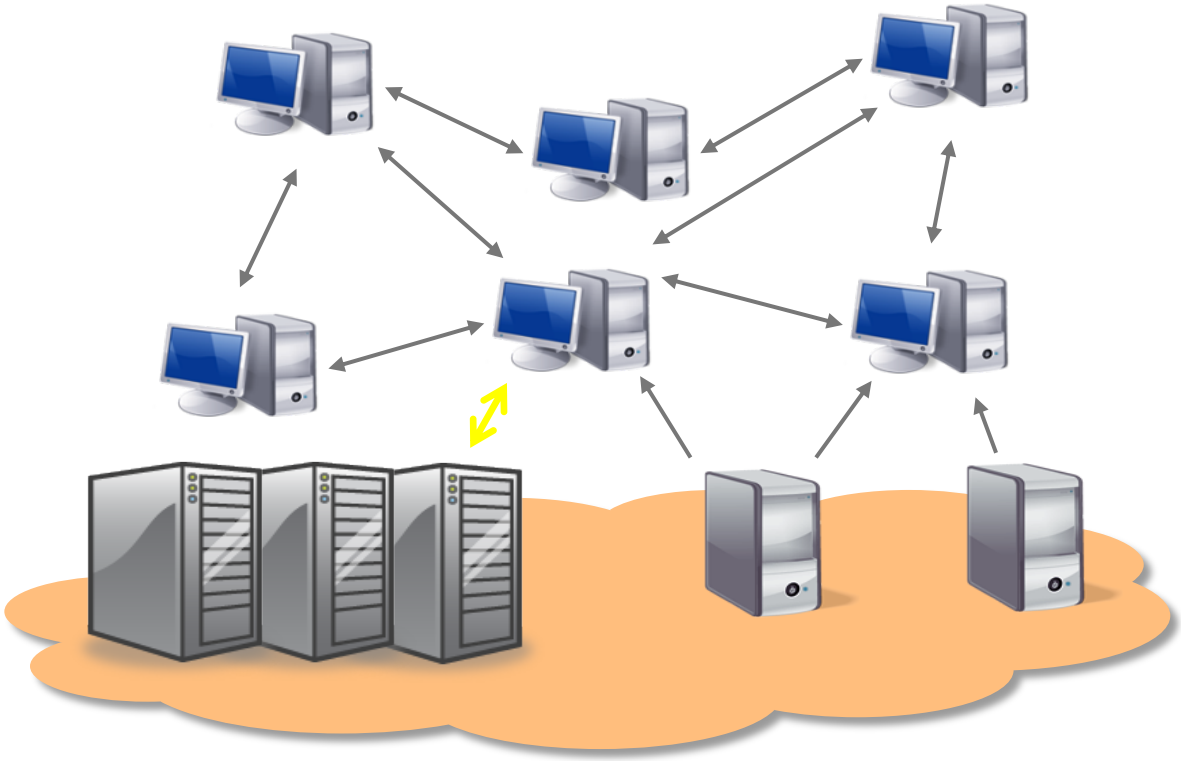

BitTorrent-like protocol with control nodes serving as “trackers” and ***assigning*** peers

CDN acts as a backstop

Growth in Number of Installations

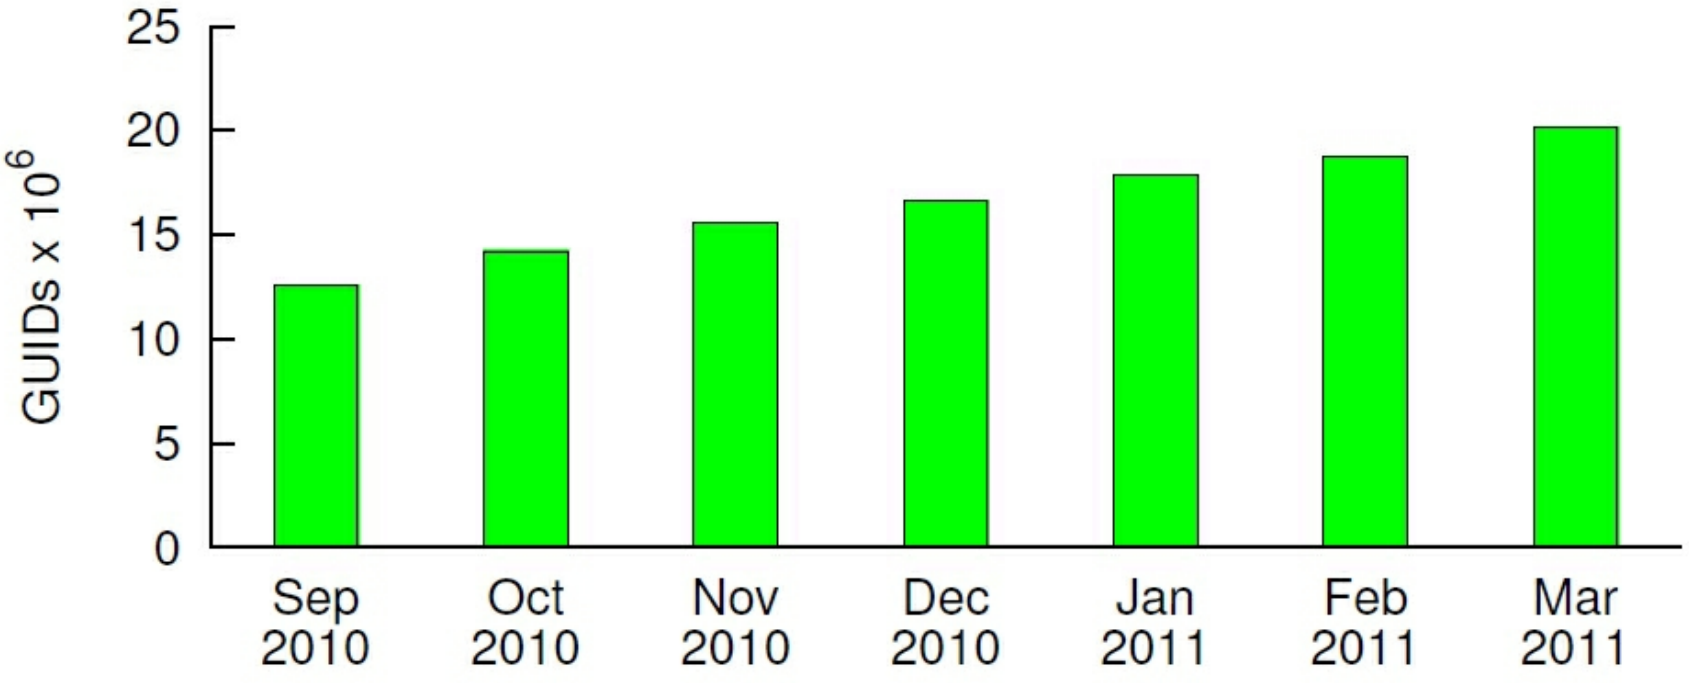

Locations of Clients per EdgeScape

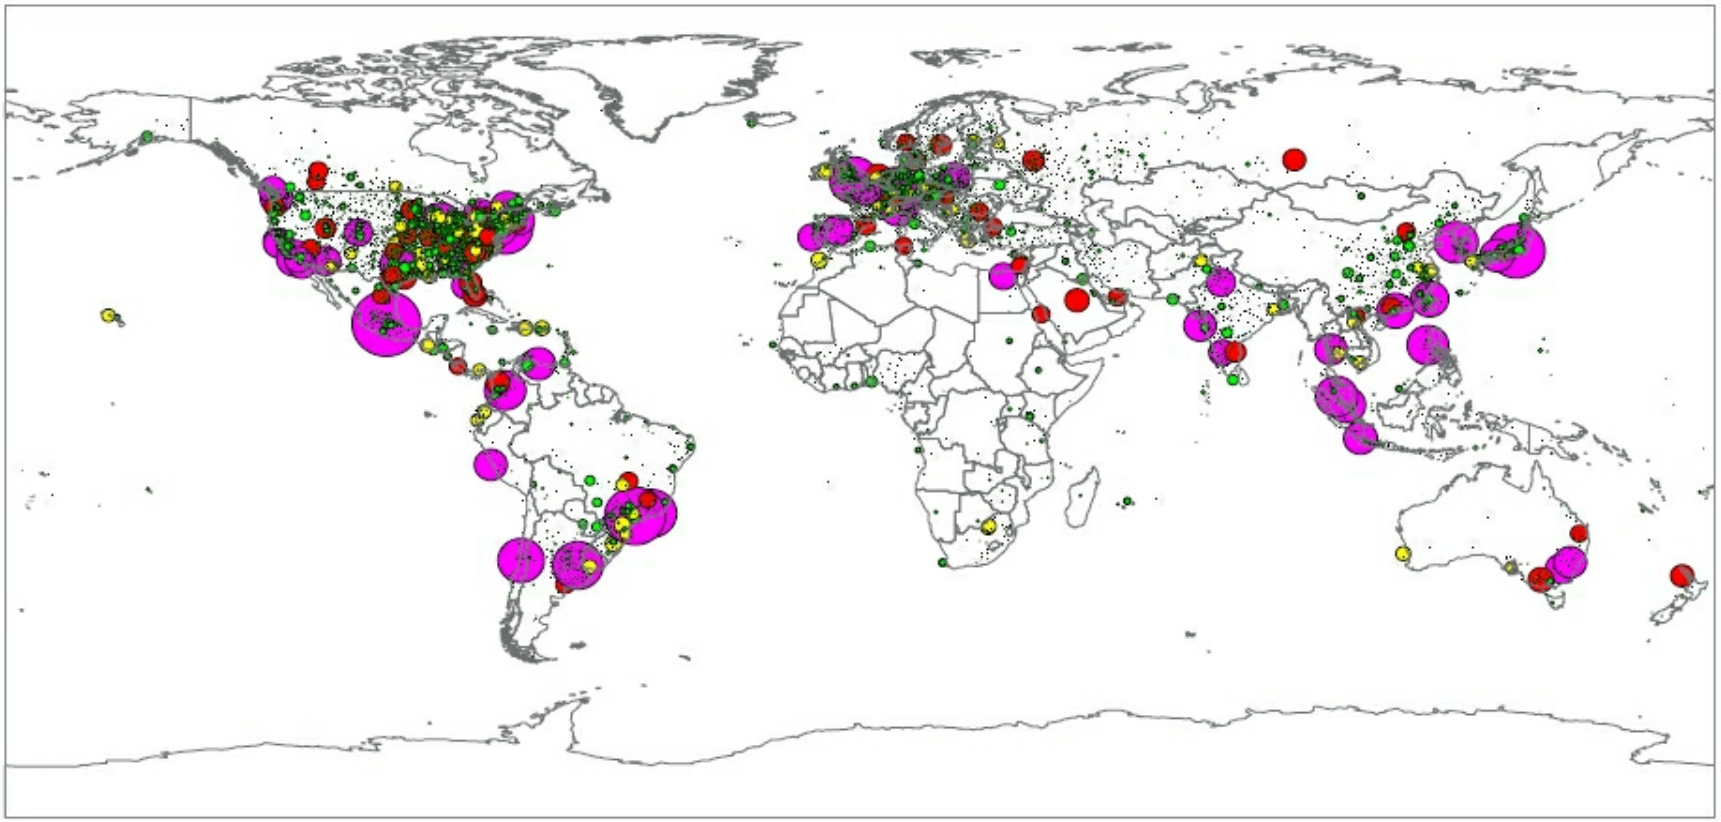

P2P Efficiency for Largest Enabled Customers

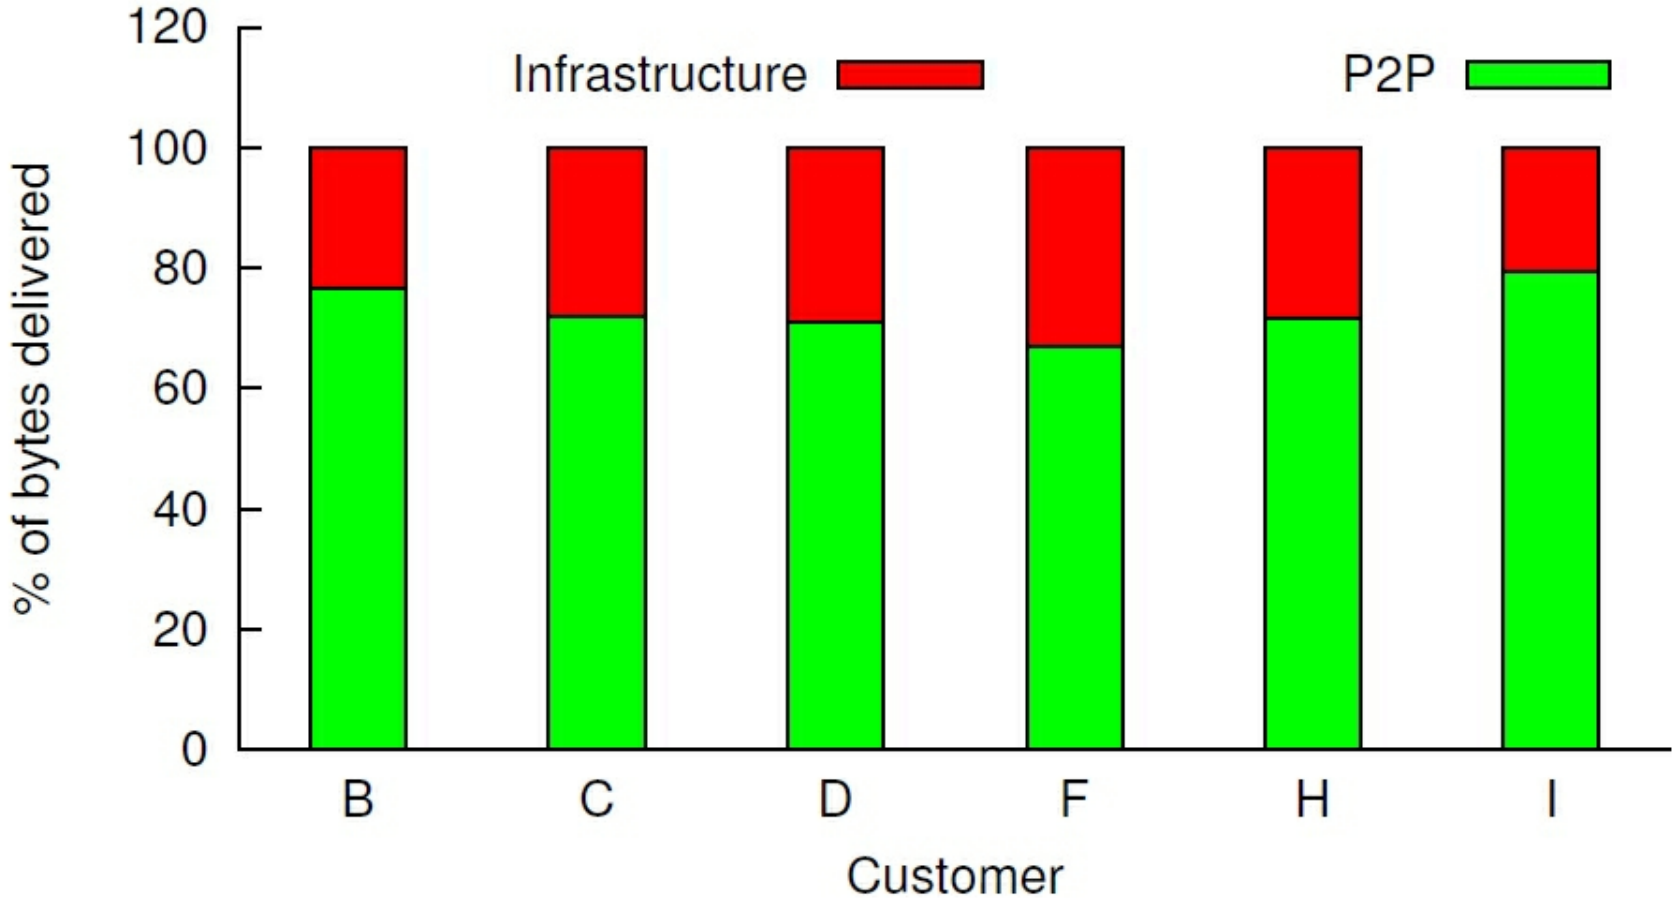

Performance

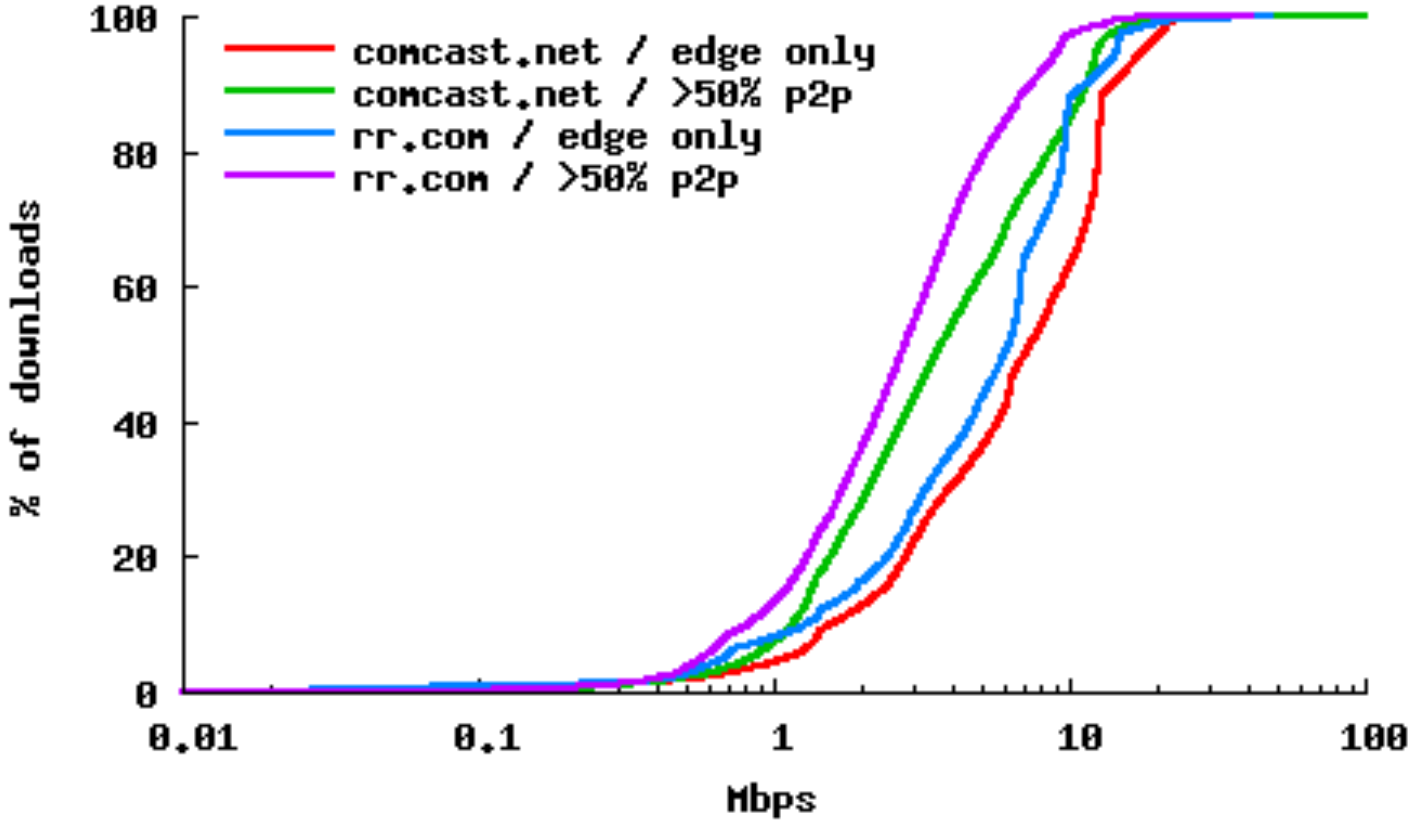

Implications for ICN

Incentives and costs must be aligned.

Performance benefits from caching:

- Reducing latency is valuable to content providers.
- Reducing network load is valuable to ISPs.

Caching in the core in addition to caching at the edge – how much additional benefit, and who pays?

Delivering everything over HTTPS/TLS makes traffic opaque, but recent history demonstrates that new protocols (e.g., QUIC) can be introduced.