



# Trends and challenges for CDN

James Kretchmar  
CTO EMEA

# The Akamai Distributed Platform

## Highly distributed deployment

- 170,000+ servers
- 2,700+ locations
- 1,300+ networks
- 700+ cities
- 100+ countries
- Within 1 hop of 90% of Internet end-users

Commodity hardware

On the public Internet

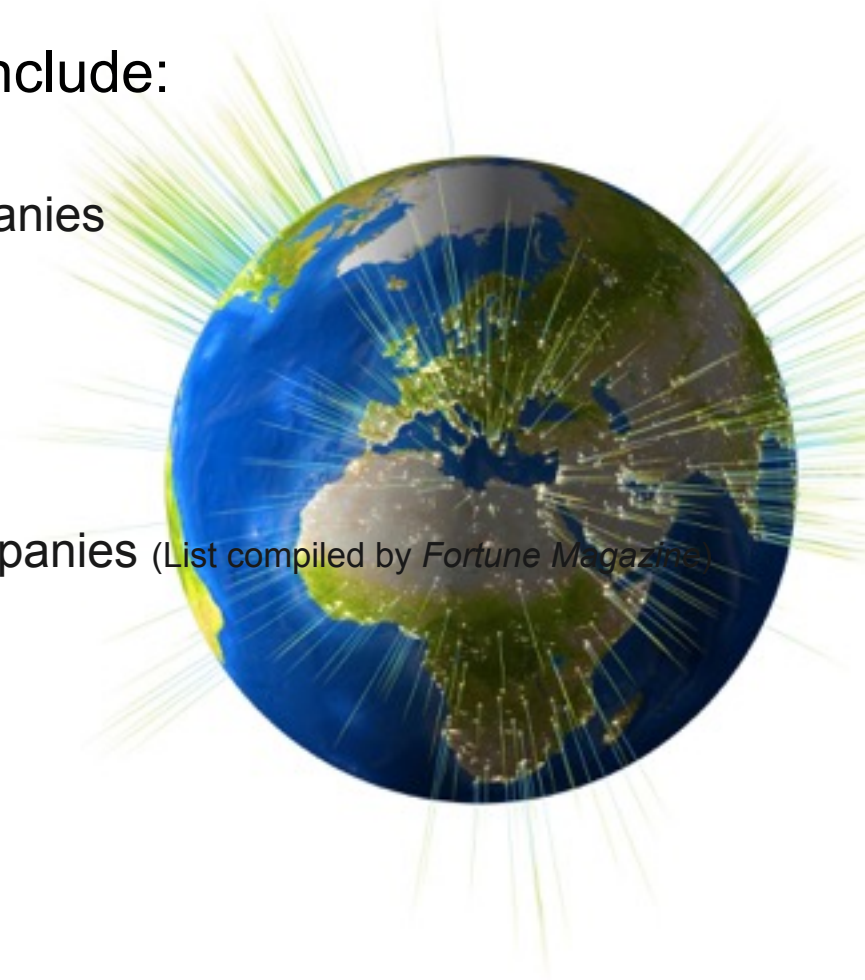
Delivering up to 30% of all Web traffic



# Akamai Customers

## Customers on the Akamai platform include:

- All 20 top global eCommerce sites
- The top 30 media & entertainment companies
- 7 of the top 10 banks (Source: *The Banker*)
- 9 of the top 10 largest newspapers
- 9 out of 10 top social media sites
- All of the top anti-virus companies
- One out of every three Global 500® companies (List compiled by *Fortune Magazine*)



# Early Internet Problems

By 1998 “www” had become the “World Wide Wait”

- Peering links congested at peak times
- Network bottlenecks around single deployments for popular content
- Server and network capacity don't scale to flash crowds

Goals for Akamai's 1<sup>st</sup> generation CDN technology:

- Eliminate slowness due to congestion, network distance, bottlenecks
- Offload origin web servers
- Provide high availability

Three-fold solution:

- Serve content from close to end-users
- Dynamically assign users to best server using real-time network conditions
- Use smart algorithms instead of big hardware for 100% availability

## End-user networks

Serving a user on-network avoids congested peering points

Need to cover as many end-users as possible

How many networks does it take to reach 90% of eyeballs?

# of Top Networks	Internet End-Users Covered
10	30%
100	70%
1,000	93%

# Dynamic Mapping

- A highly distributed deployment is only as good as the ability to use it
- Anycast?
  - Far too coarse to find all deployments
  - Little to no load control, need to over capacitate
  - No load feedback
  - No accounting for real-time network conditions
- Solution:
  - Use DNS to assign users to servers
  - Make the assignments based on real-time network conditions
  - And real-time load feedback from servers

# Cluster Scalability & Reliability

Within a deployment of servers:

- Servers coordinate use of disk space
- Dynamically use more servers if object load is high
- If a server fails:
  - Reassign in a minimally impacting way (consistent hashing)
  - DNS reflects the new assignments
  - A “buddy” machine grabs the IP



## Deployment

- Acts like one big cache
- Multiple machines can fail with no bad effect

## 2<sup>nd</sup> Generation Technology

- Good at removing delays due to network for static HTTP(S) content
- By 2005 looking to accelerate of dynamic content and generalized IP traffic
- SureRoute:
  - Akamai deployments are so widespread, we can find faster paths than BGP
  - By bouncing traffic from Akamai server to Akamai server
  - Also BGP doesn't account for congestion or loss
- Methodology
  - Measure the network performance between Akamai deployments
  - Compute where better than Internet to hop from one deployment to another
  - Send the traffic along that path (multiple or single)
- Applications
  - Faster delivery of cold content
  - Fast delivery of dynamic (uncachable) content
  - IP applications sensitive to packet loss or latency



# GRAND CHALLENGES



Instant web performance for any device, anywhere



The web must be secure



The ultimate in video quality over IP at scale



Fast, affordable IP access in the workplace and branch office

# **Trend 1: Performance**

# Web Pages Are Slowing Down!

	2012	2013	2014
Typical Page Size	788 Kb	1,081 Kb	1,622 Kb
Typical # of Objects	88	101	112
Mobile Penetration	9%	19%	30%
<b>Avg. Page Load Time (Sec.)</b>	<b>6.6</b>	<b>7.2</b>	<b>10.7</b>

## And It's Getting Even Slower!



- # of 3rd party domains up **56%** in 2 years
- JavaScript transfer size up **40%** in 2 years
- More round trips = **more wait** time
- **Demand** is far outstripping growth in capacity



Heavier  
Content



3<sup>rd</sup> Party  
Content



Complex  
Experiences



Chatty  
Protocols



Cellular  
Congestion


# Front End Optimization (FEO) - Focus On What's Important

## Back-end

## Front-end

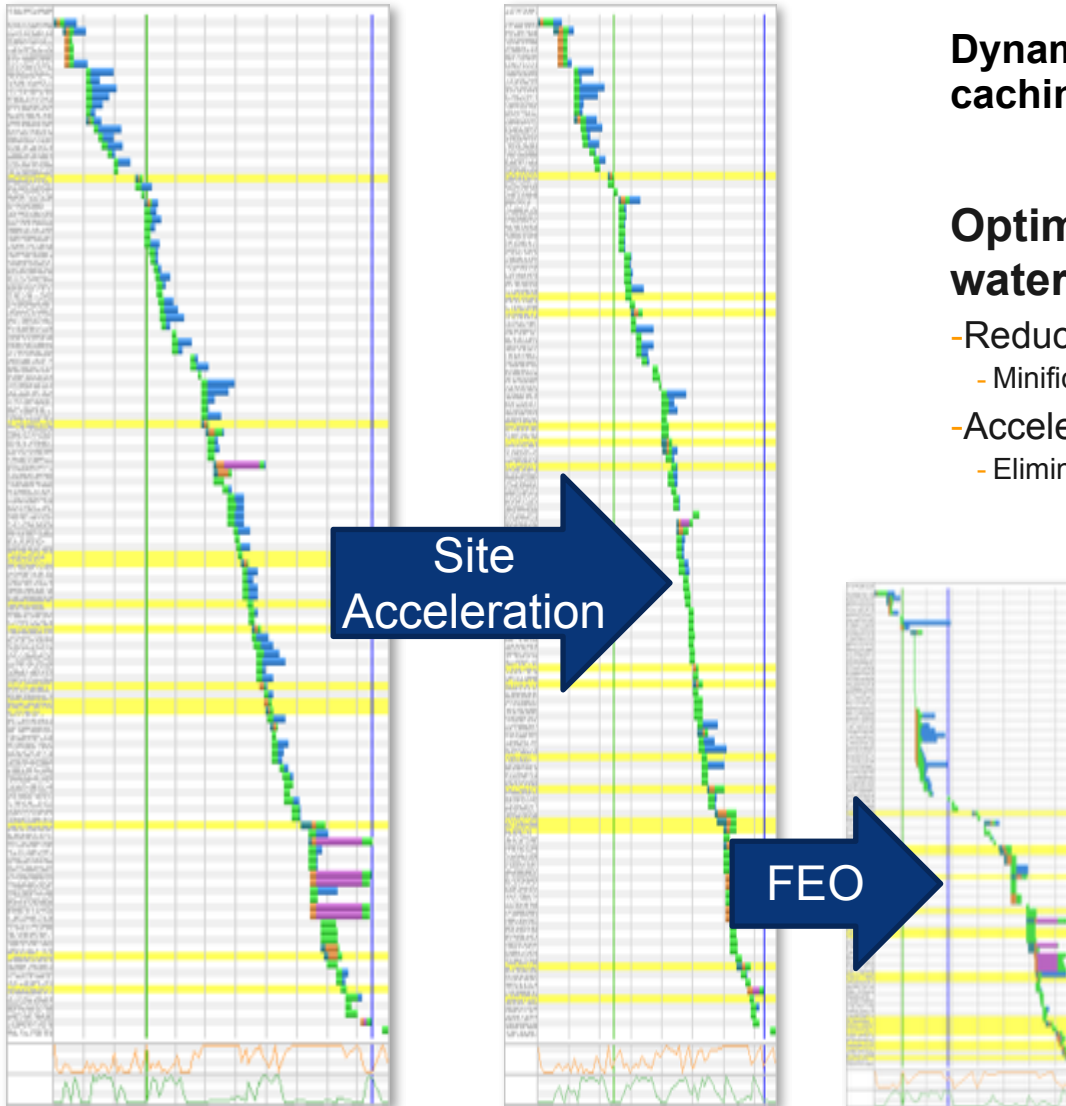
Backend – the time it takes the server to return the first byte back to the client

Frontend – Everything else



**Souders' Golden Rule**  
80-90% of the end-user response time is spent on the frontend.

# Front End Optimization



**Dynamic site acceleration & caching “shrink” waterfall**

**Optimization shrinks & overhauls waterfall**

- Reduces bytes & requests
  - Minification, Consolidation, better client side caching, etc.
- Accelerates rendering
  - Eliminates blocking, defers invisible content, etc.

# Many Optimizations at Play

Optimization	Reduce Requests	Reduce Bytes	Accelerate Rendering
HTML5 adaptive consolidation	✓		
Just in time (JIT) image loading	✓		
Embed small images and JS	✓		
Version files for long term cache ability	✓		
HTML5 persistent cache	✓		
Resize images to display size		✓	
Match image size to HTML dimensions		✓	
Lossless image compression		✓	
Optimize image resolution		✓	
Serve static files from cookie-less domain		✓	
JS/CSS minification		✓	
Streaming consolidation			✓
Asynchronous JavaScript Execution			✓
Optimize JS, CSS load order			✓
Domain sharding			✓
Pre-fetch/Pre-render next page			✓


# Adapt To The End-Users' Situation

Original Image



Quality: 100%  
Size: 60KB

Network Conditions



Adapted Image

Throughput: High  
Size: 54KB (90%)



Throughput: Med  
Size: 40KB (67%)




Throughput: Low  
Size: 24KB (40%)





# **Trend 2: Security**

## Major Attacks on Akamai Customers




Typical Attack Size  
**10s of Gbps**

Large Attack Size  
**100s of Gbps**

Attacks are originating from all geographies and are moving between geographies during the attack

# Attack Size

- Attack size (Gigabits per second)
- Attack size (Million packets per second)



# 2014: Year of Major Security Vulnerabilities

## Heartbleed

- Vulnerability in OpenSSL can divulge private keys

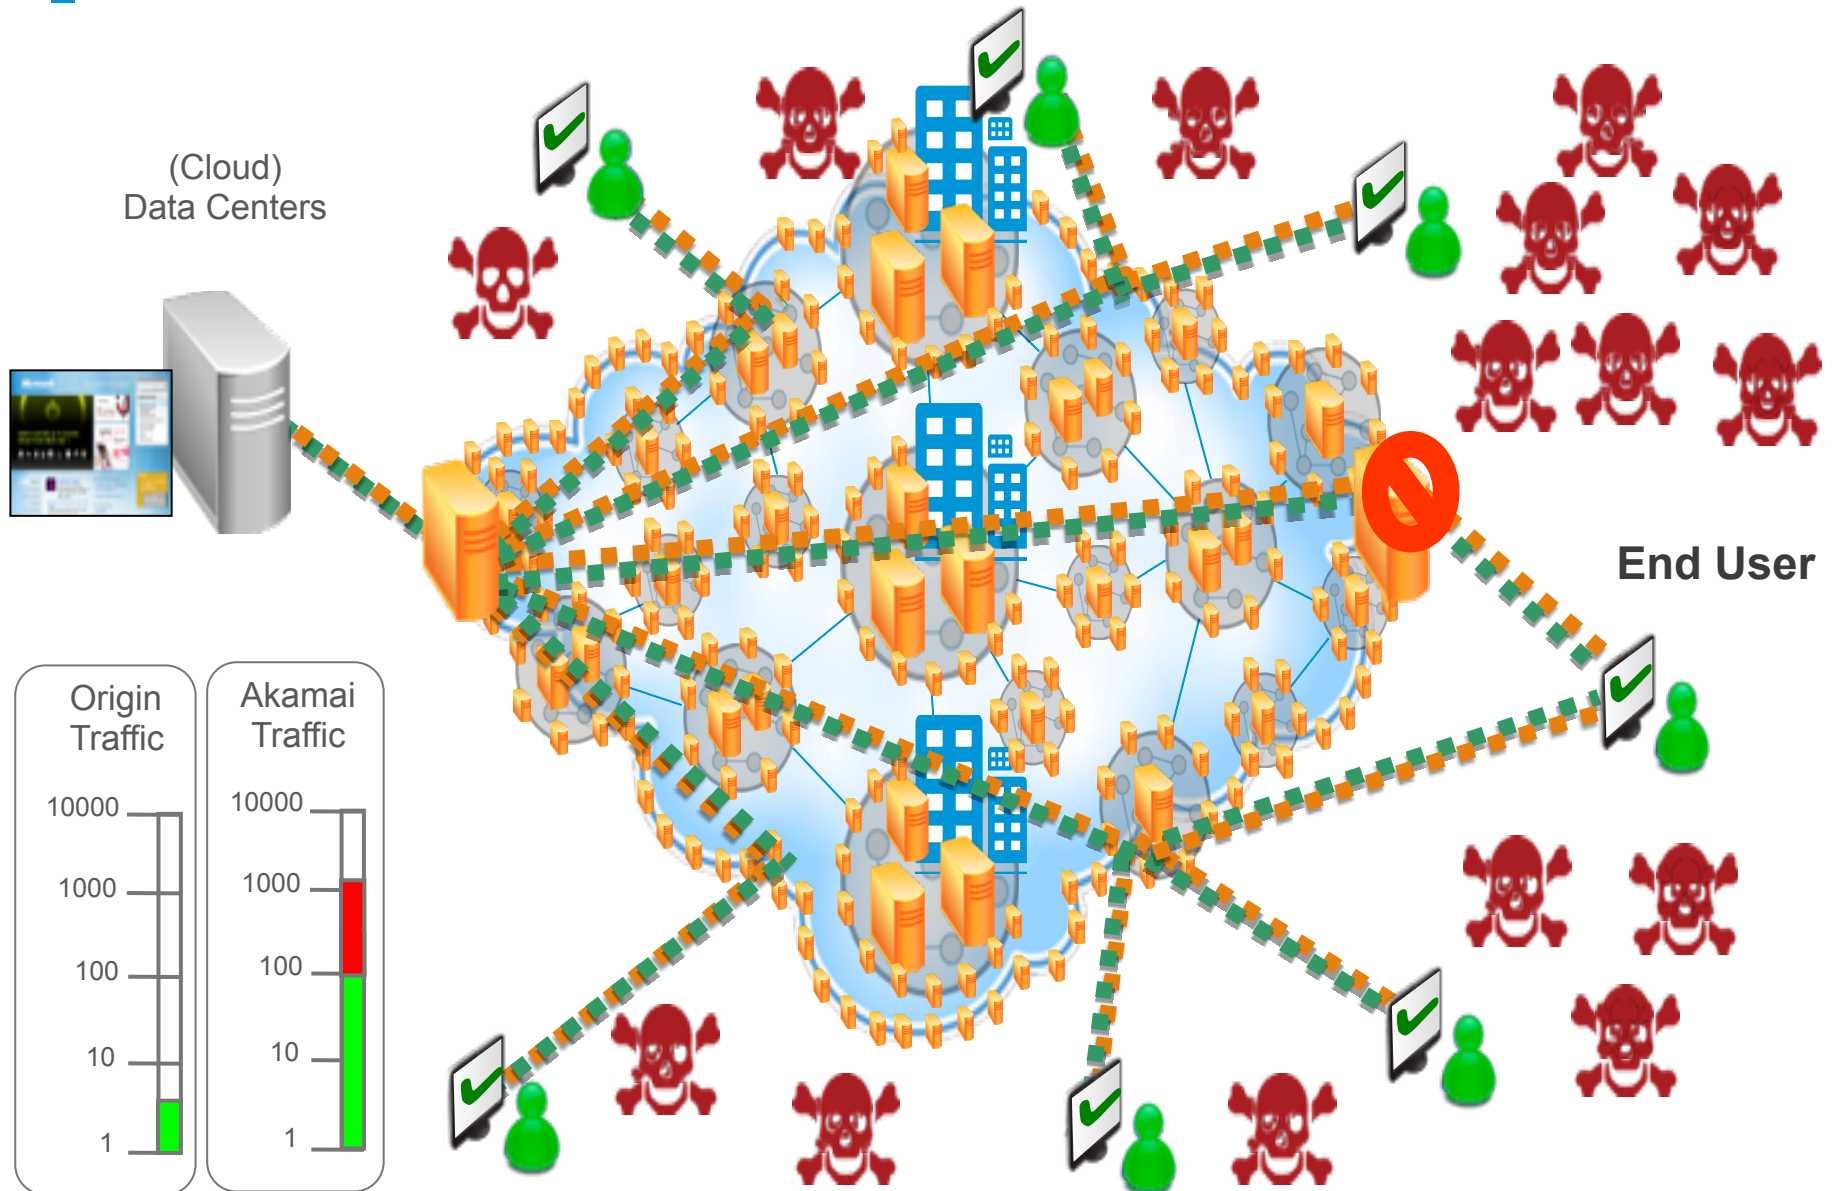
## Shellshock

- Vulnerability in bash shell grants remote execution to web servers

## Poodle

- Fallback to SSLv3 plus padding oracle attack can decrypt connections

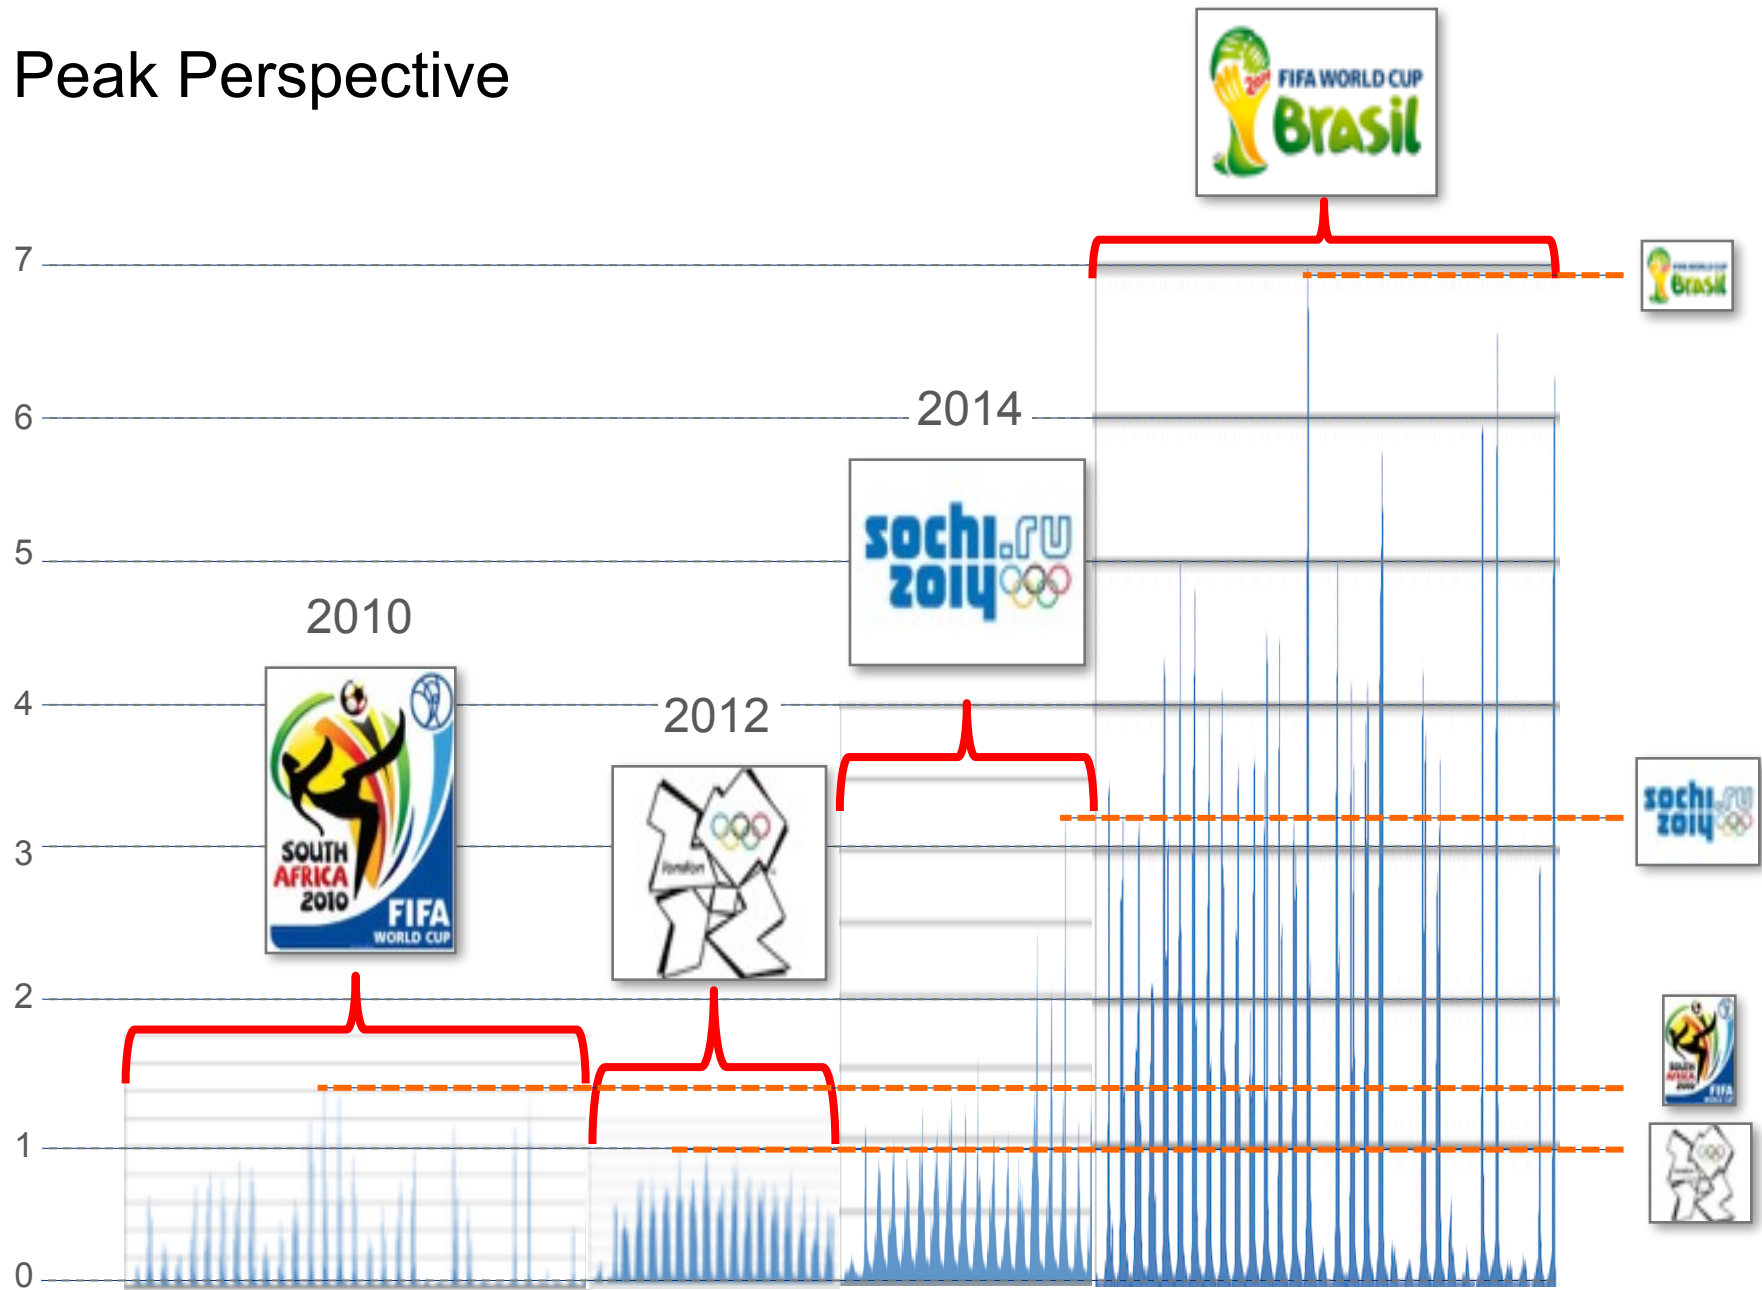
# Cloud Security




# **Trend 3: Media Growth**

# Peak Perspective

Terabits per second

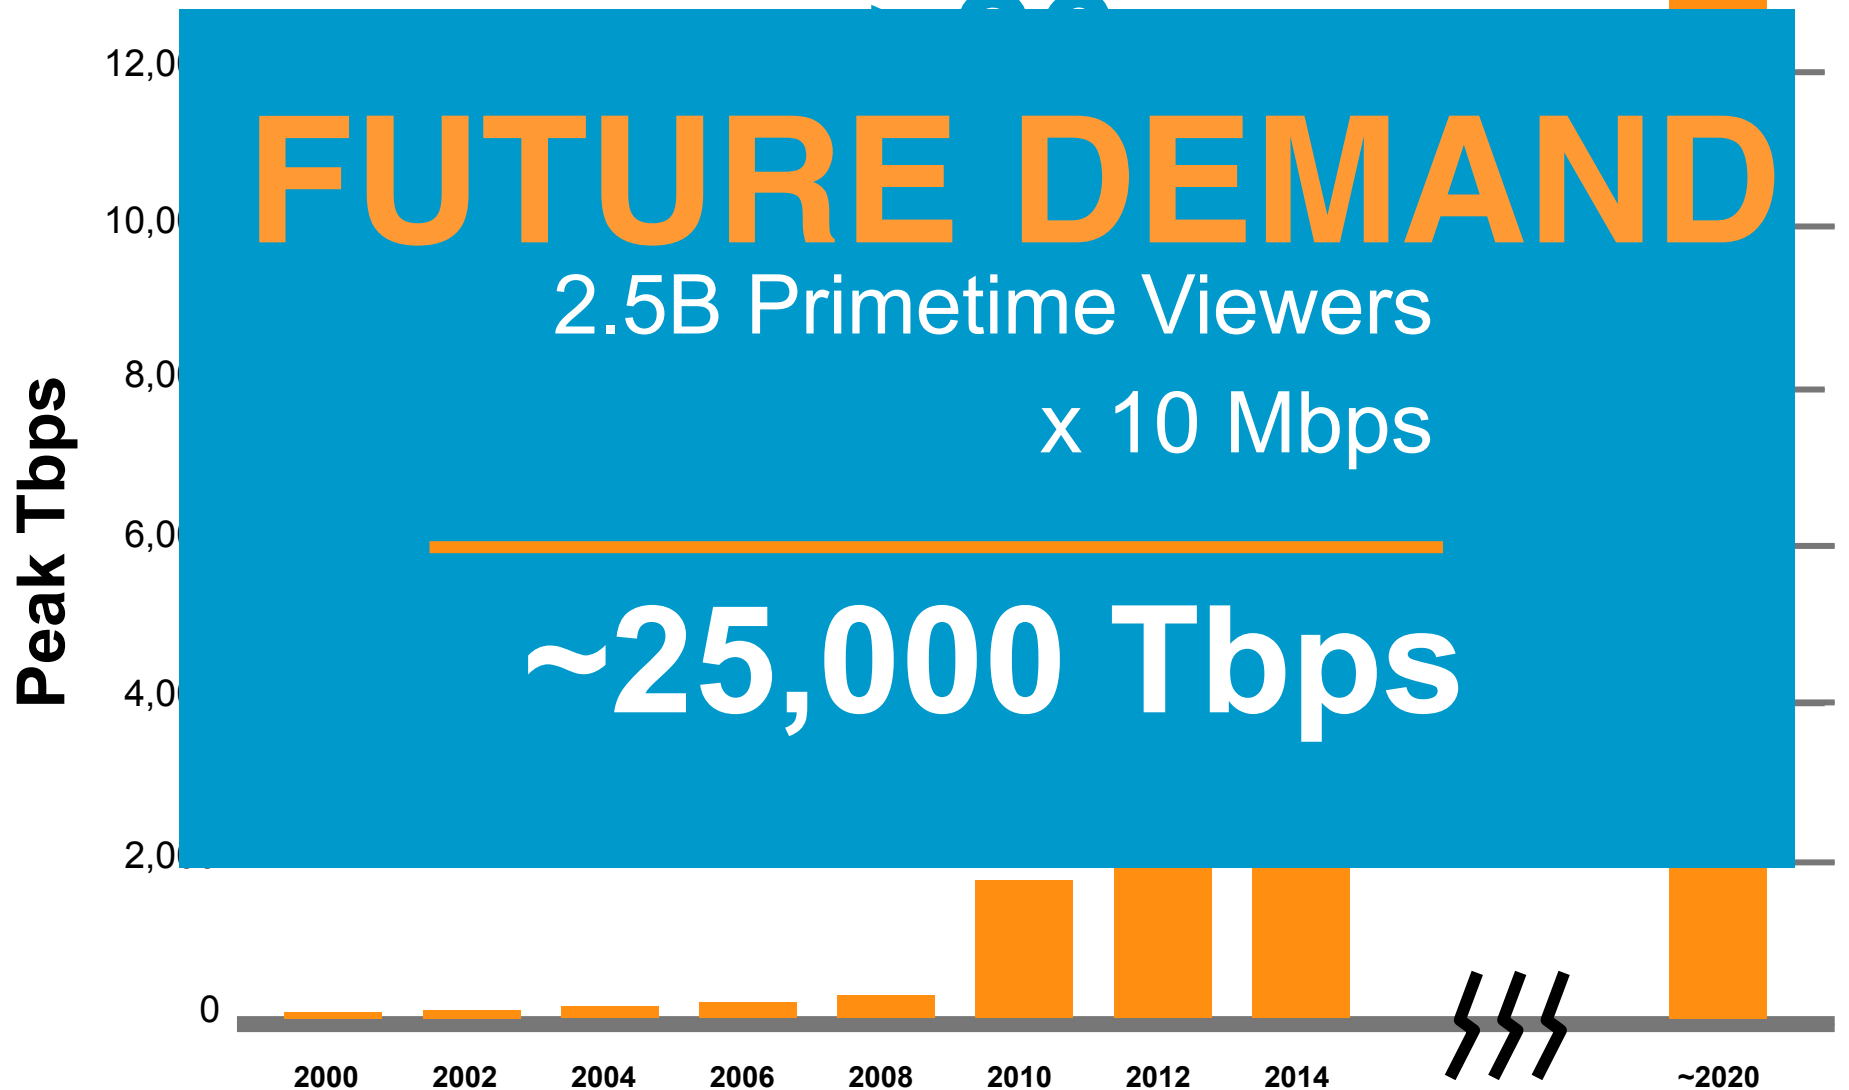


# Relative Event Size







# Traffic Growth on Akamai's Platform




# Edge Strategy Today: Inside Over 1,300 Networks



# Extending the Edge: Moving Into Devices

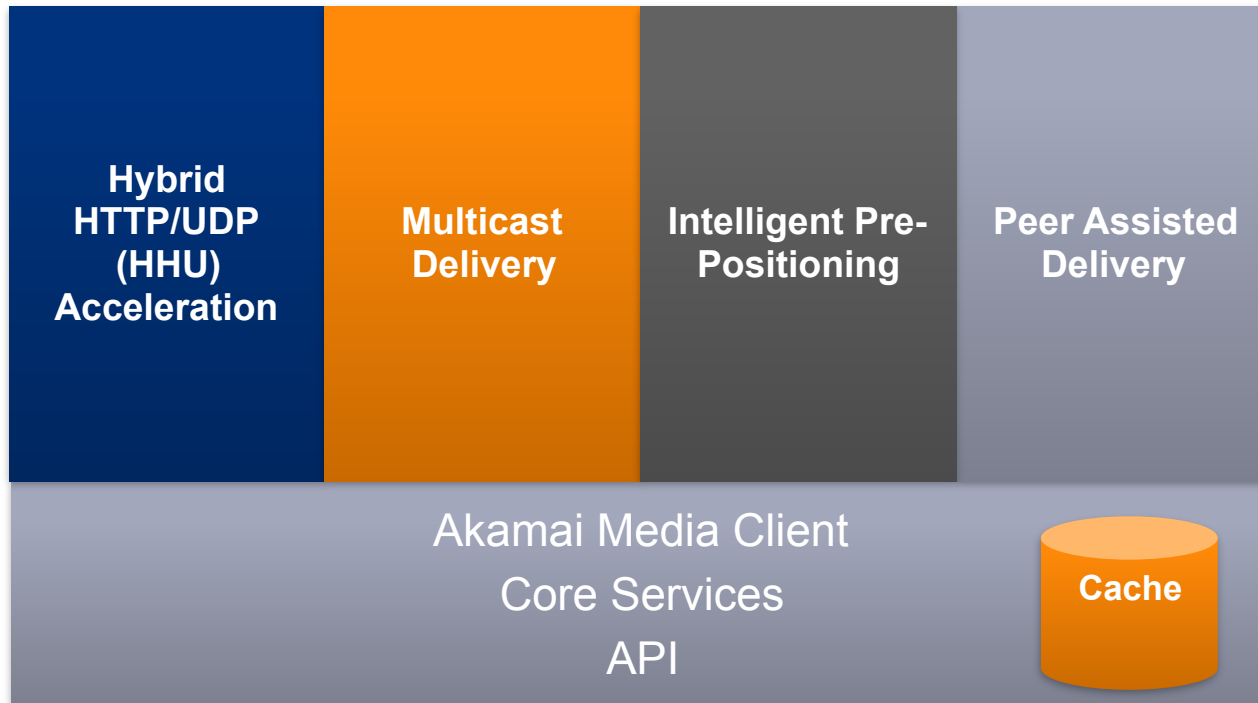


# Extending the Edge: Moving Into Homes



# Delivery Innovation: Akamai Media Client Technology

A **toolkit** for the next-generation network

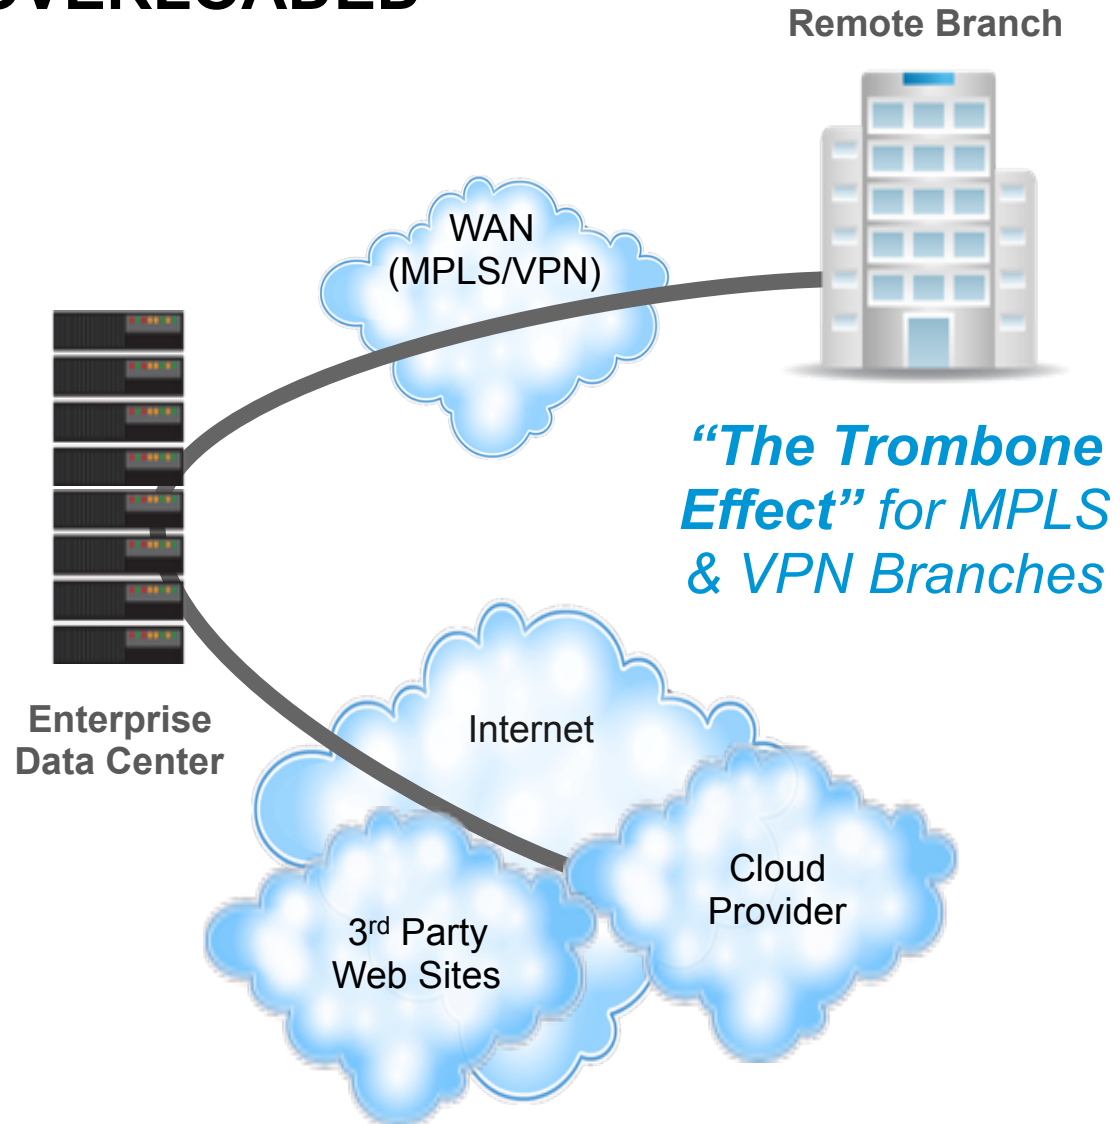


# **Trend 4: Hybrid WAN**

# Branch Networks are **OVERLOADED**

**10-100X**  
increase  
in capacity  
needed

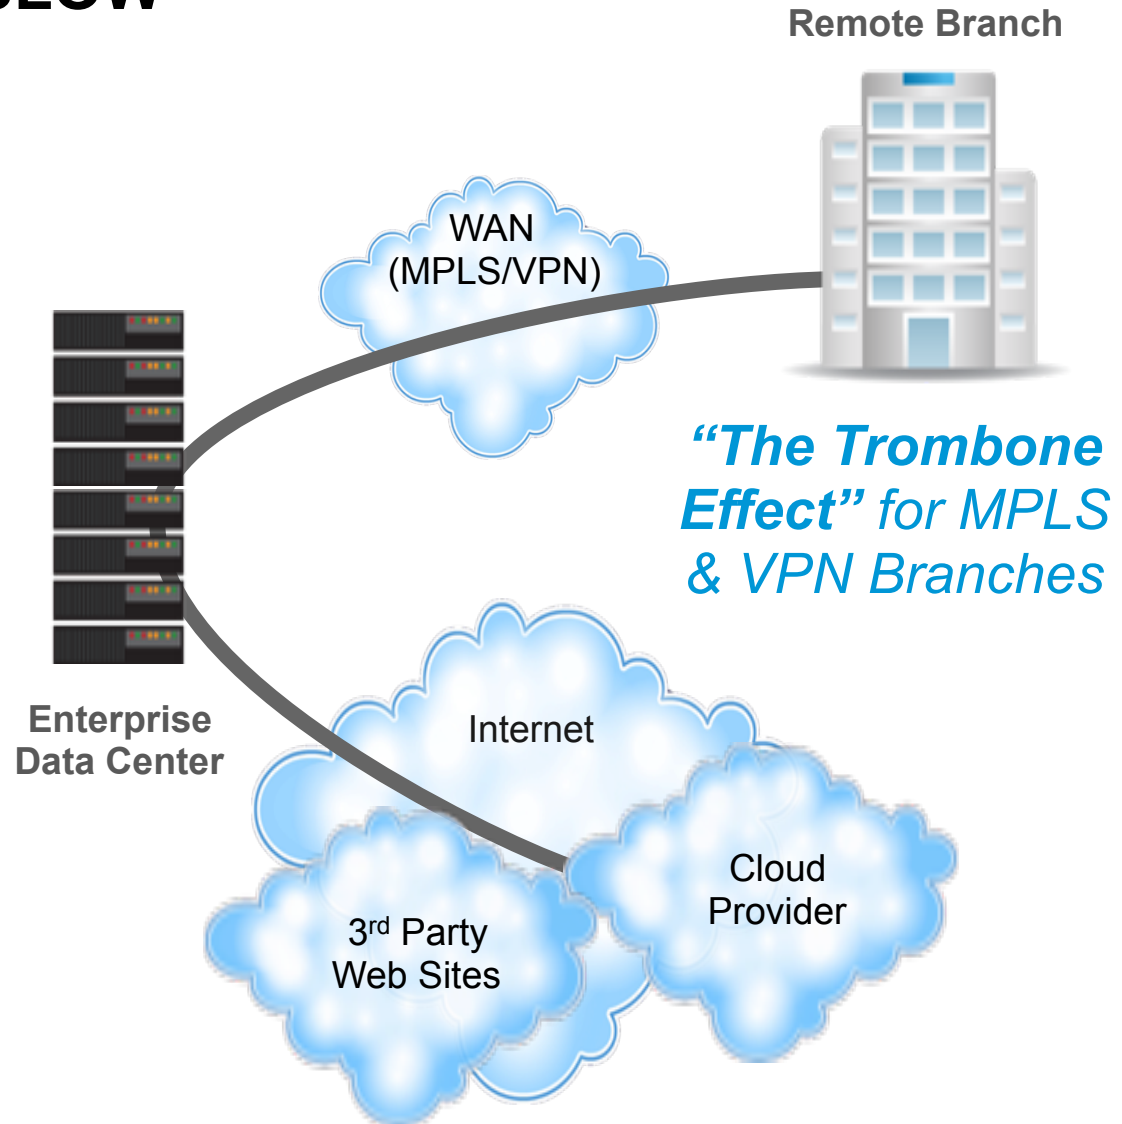
video, virtual desktops,  
collaboration tools,  
digital displays, backup/  
archiving, catalogue  
and software downloads,  
web apps, PowerPoint



Branch Networks are **SLOW**

**LONG  
DISTANCE**

Between  
branch and  
web applications

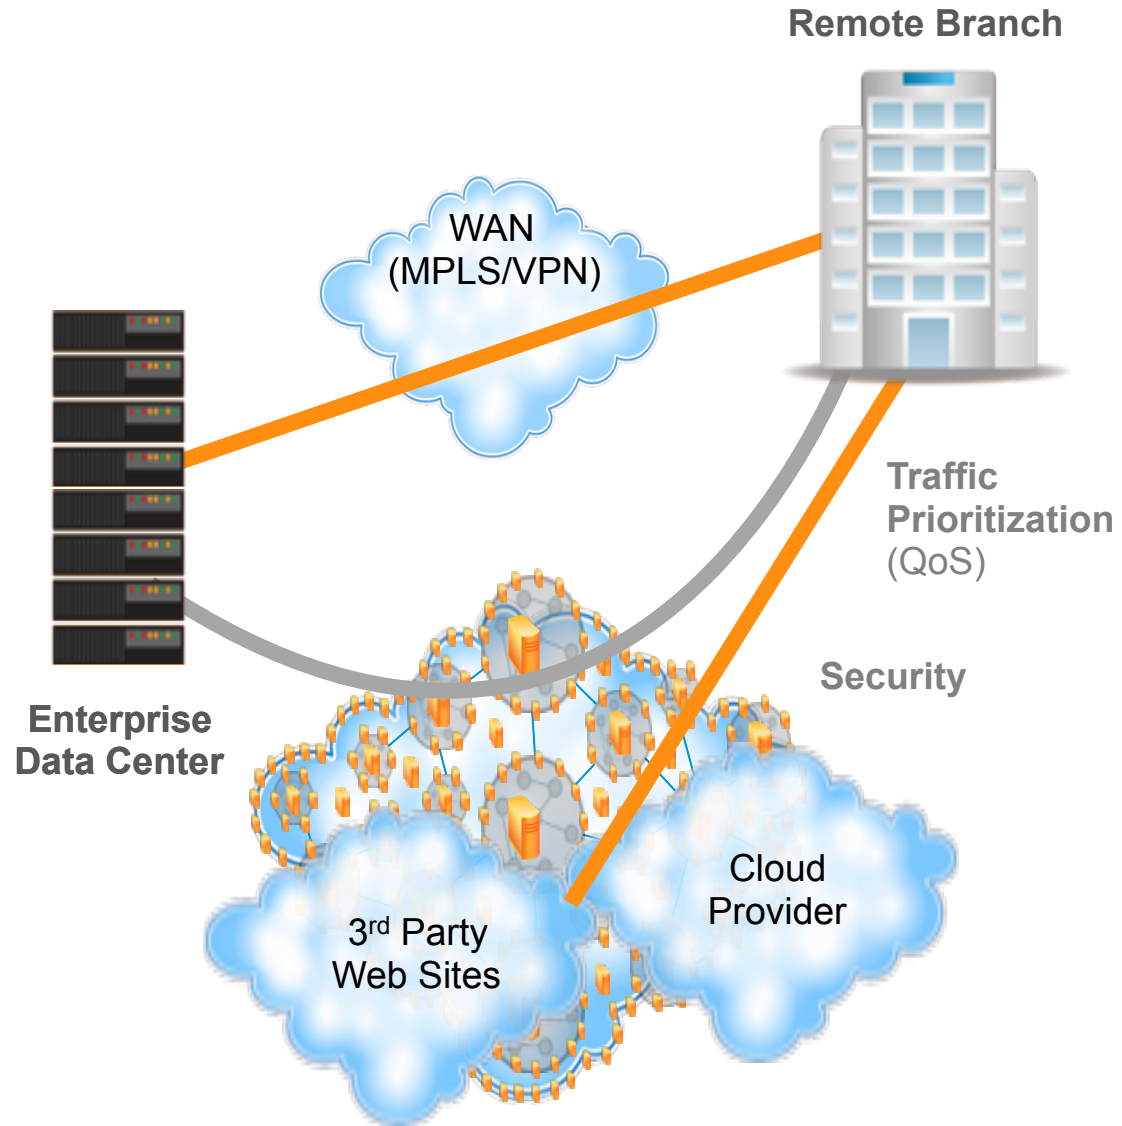




Branch of the **Future**

# Hybrid WAN

Leverage public and private networks together.



# GRAND CHALLENGES



Instant web performance for any device, anywhere



The web must be secure



The ultimate in video quality over IP at scale



Fast, affordable IP access in the workplace and branch office



# Trends and challenges for CDN

James Kretchmar  
CTO EMEA