



Akamai's V6 Rollout Plan and Experience from a CDN Point of View

Christian Kaufmann
Director Network Architecture
Akamai Technologies, Inc.

Agenda



- About Akamai
- General IPv6 transition technologies
- Challenges of reaching IPv6 users via transition technologies
- Challenges of reaching IPv6 users natively
- Akamai's IPv6 development
 - Engineering: Mapping, EdgePlatform, Back End, etc.
 - Network
- Akamai's IPv6 Transition Phases
 - Early Adoption Phase
 - End User Transition Phase
 - Origin Transition Phase
- What we did so far...
- IPv6 World Day
- Experience with IPv6
 - DNS / Mapping
 - Server OS
 - Network
 - IP allocation

The Akamai System

The world's largest on-demand, distributed computing platform delivers all forms of Web content and applications for over 3,400 customers and 20,000 domains

The Akamai EdgePlatform:

95,000+
Servers

1,795
POPs

1000+
Networks

660+
Cities

79
Countries

Resulting in traffic of:

6.1 Tbps peak traffic

49 petabytes / day

1100+ billion hits / day

471+ million unique clients IPs / day



General IPv6 transition technologies



- Tunneling
 - encapsulates IPv6 packets inside of IPv4 packets
 - Tunneling is used to connect v6 islands with no connectivity together
 - Can cause performance and reliability issues
- CGN/NAT64
 - NAT64 provides connectivity between IPv6 hosts and IPv4 hosts by translating between IPv6 and IPv4 packet streams
 - Might cause performance issues or be a single point of failure
- Dual Stack
 - running both the IPv4 and IPv6 protocol stacks at the same machine
 - Must be supported by hard- and software
 - Requires enough IPv4 address space for all dual-stacked devices
- Each transition technology has its limitations and drawbacks, but they are necessary until the long IPv6 migration is complete.

Challenges of reaching IPv6 users via transition technologies

- Broken location-aware services
 - Caused by CGN's and Tunneling
 - Licensing content under geographic restrictions
 - *abuse mitigation*—IP blacklisting or whitelisting
 - Web site analytics relying on geo-location
- Broken applications
 - VoIP
 - Peer to Peer traffic
 - Complex troubleshooting
- Poor performance
 - Because of limited CGN deployments (latency, congestions)
 - Because of tunneling (MTU, latency)
- Worst case scenario, multiple NAT and CGN in a row

Challenges of reaching IPv6 users natively

- Delivering IPv6 natively during the transition phase will face several key challenges:
 - Often lower performance, because of less direct routes and tunnels
 - Reliability issues, because of fewer network interconnections and less monitoring
 - Denial-of-service because of browser and OS bugs, broken home gateways
V6 end-user asks for a AAAA but has no working v6 connectivity
 - Whitelists create management and scalability challenges
Restricting AAAA responses can cause issues, as there is no direct correlation between DNS IPv6 support and end user IPv6 connectivity.
 - Lack of IPv6 reporting and analytics
 - IPv6 "support" often does not mean performance parity with IPv4

Challenges of reaching IPv6 users natively – Performance Examples

- A test with more than 9k name servers worldwide showed:
 - For North America tests showed a median IPv4 latency of 48 ms compared to a median IPv6 latency of 87 ms
 - For tunneled IPv6 traffic, the difference in latencies was even more significant
- Download tests from www.akamai.com showed similar results:
 - The median speed via IPv4, was almost 2x faster than via IPv6.
 - Of the 4000+ clients who were able to reach us via IPv6 in this test:
 - only 10.6% of them had a native IPv6 addresses
 - the rest came over tunnels.

Akamai IPv6 Engineering: Mapping, EdgePlatform, Back End, etc.

- Native IPv6 Support in the EdgePlatform Servers and throughout the system
 - Including for example log files for our customer
- Mapping the IPv6 Internet
 - Significant effort in making the mapping system v6 capable
 - The evolving topology of the IPv6 Internet required iterative surveying and mapping efforts
- Understanding the performance differences between IPv4 and IPv6
 - Real-time mapping topology for IPv6 is independent of IPv4
 - End-user will get the content via v4 or v6 depending on which protocol has a better performance
- Making the IPv6 Internet Reliable

Akamai's IPv6 development: Network

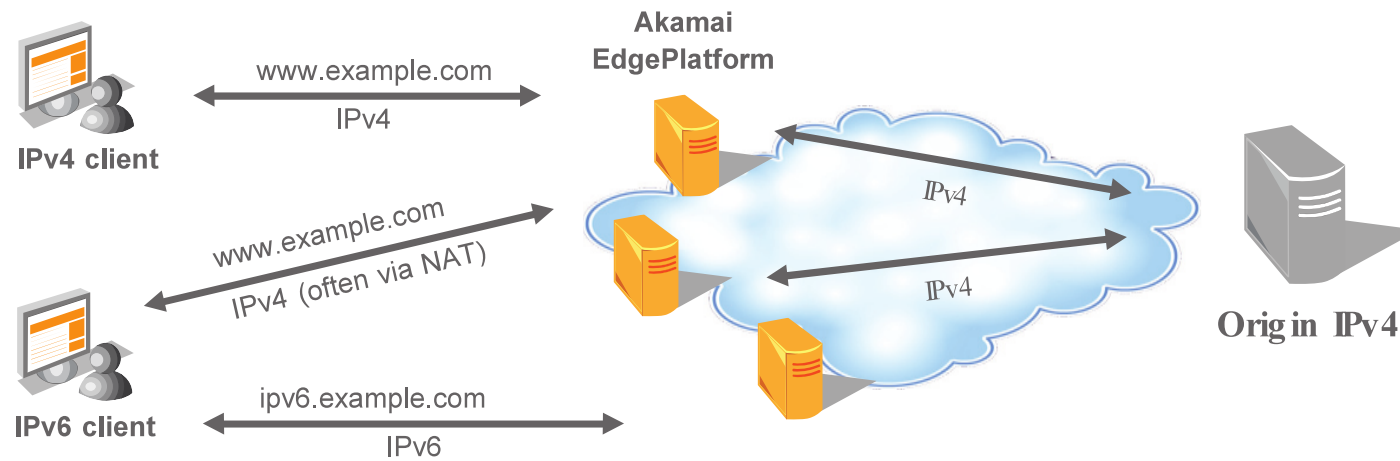
- Dualstack for all IX router
 - 42 out of 52 IX routers globally are dual stack.
- Requesting IPv6 blocks from all network partners
 - Already started

Akamai's IPv6 Early Adoption Phase



- Early Adoption Phase – Q1 2011
 - Selected set of customer
 - IPv6 only site and/or beacon
 - Customers who are interested in joining should contact us.
- Akamai will enable early adopters to offer IPv6 content through a separate hostname, without changing their IPv4 networking infrastructure and without affecting existing users.
 - For example: <http://ipv6.akamai.com>
 - This allows customers to ensure that client IPv6 addresses work end-to-end: in their reporting, IDS, log processing, etc.
 - prior to enabling it for their full production site, but without needing IPv6 connectivity at the network level.
- Customers can offer their services on IPv6 immediately
 - Without spending time or money on training, software, hardware, etc.
 - For what is currently a small percentage of the user base
 - Gives customer time to build up expertise on IPv6 in a controlled manner

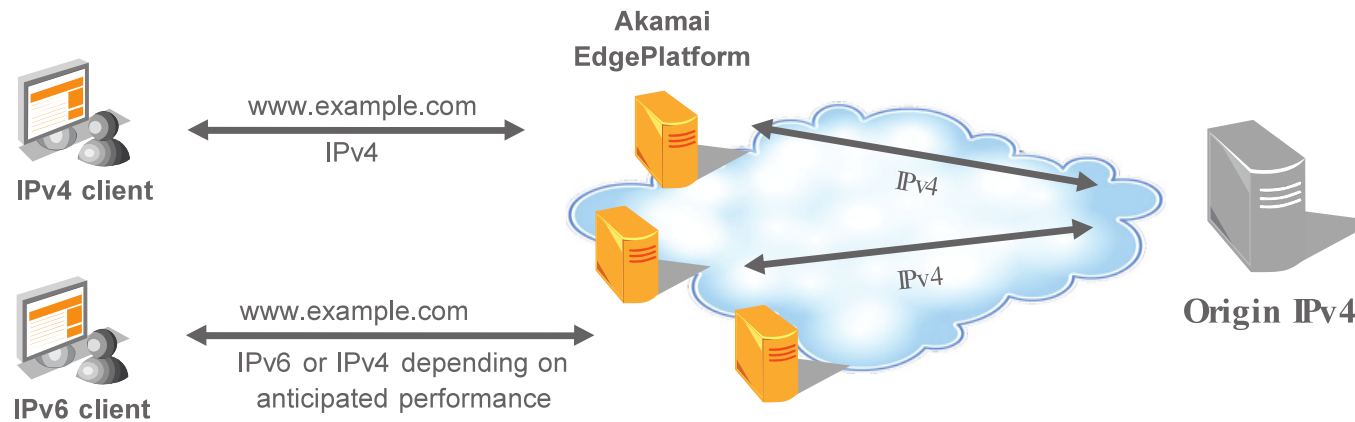
Akamai's IPv6 Early Adoption Phase



Akamai's IPv6 End User Transition Phase

- We are having Beta service available in 2H 2011, with Limited Availability in 1H 2012.
- Customers will be able to deliver optimized experiences to all of their users across the hybrid Internet on a single hostname, without changing their IPv4 network infrastructure.
- Avoid CGN, tunnels, indirect routes, and other performance and reliability bottlenecks
- Example:
 - Dual stacked user makes a request to a customer site
 - Mapping system will determine if the user can be served via IPv6 or via IPv4
 - IPv6: User will receive an IPv6 or an IPv4 address of an optimal server
 - IPv4: Only an IPv4 addresses will be returned
 - customer can choose how aggressively they prefer v6 over v4

Akamai's IPv6 End User Transition Phase

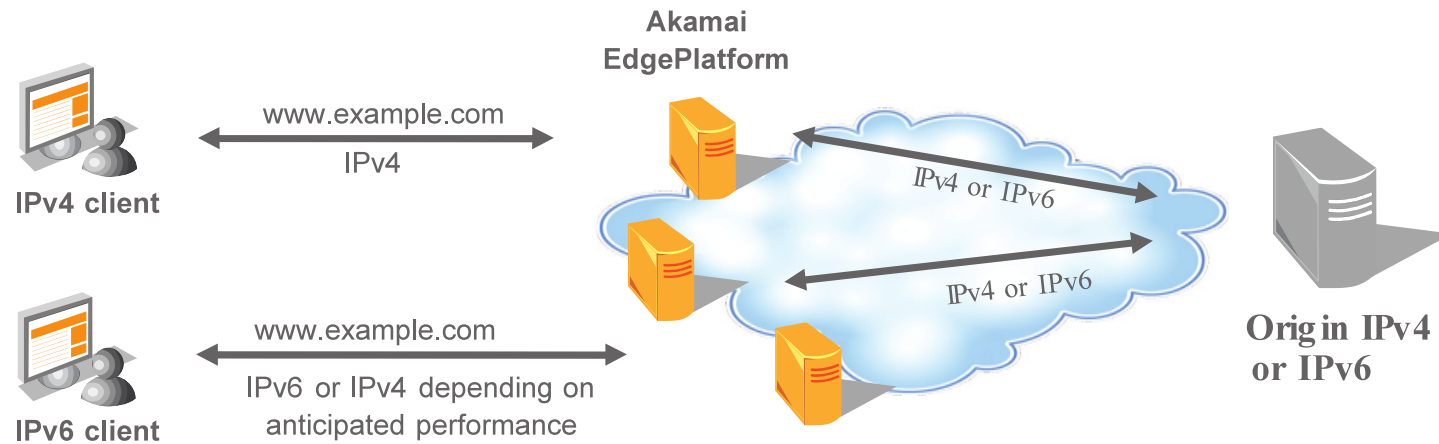


Akamai's IPv6 Origin Transition Phase



- Akamai will support dual stack or IPv6 only origin infrastructures so that customers may transition their origin infrastructures when they are ready to.
- Provide seamless origin server transition to IPv6—no impact to end users

Akamai's IPv6 Origin Transition Phase



What we did so far...



- Akamai is already delivering IPv6 traffic out of some clusters for testing purposes
 - deployment is limited but growing
 - Some advanced features are not ready yet
- IPv6 beacon on www.akamai.com to check for IPv6 capable end-users
- ISOC's IPv6 World Day at the 8th June for interested customer

IPv6 World Day



- Akamai had over 20 customers participating with over 50 properties
- Served over 8 million requests over IPv6 during the course of the day to 280,000 unique IPv6 addresses
 - 50% of the end-uses came via 6rd
 - 40% via native IPv6 connections
 - 10% split across 6to4 (~9%) and Teredo (~1%).
- IPv6-preferring end-users is still very small ranging from 0.2% to 0.5% of users
- End-user networks and customers didn't have noticeable increases in support calls

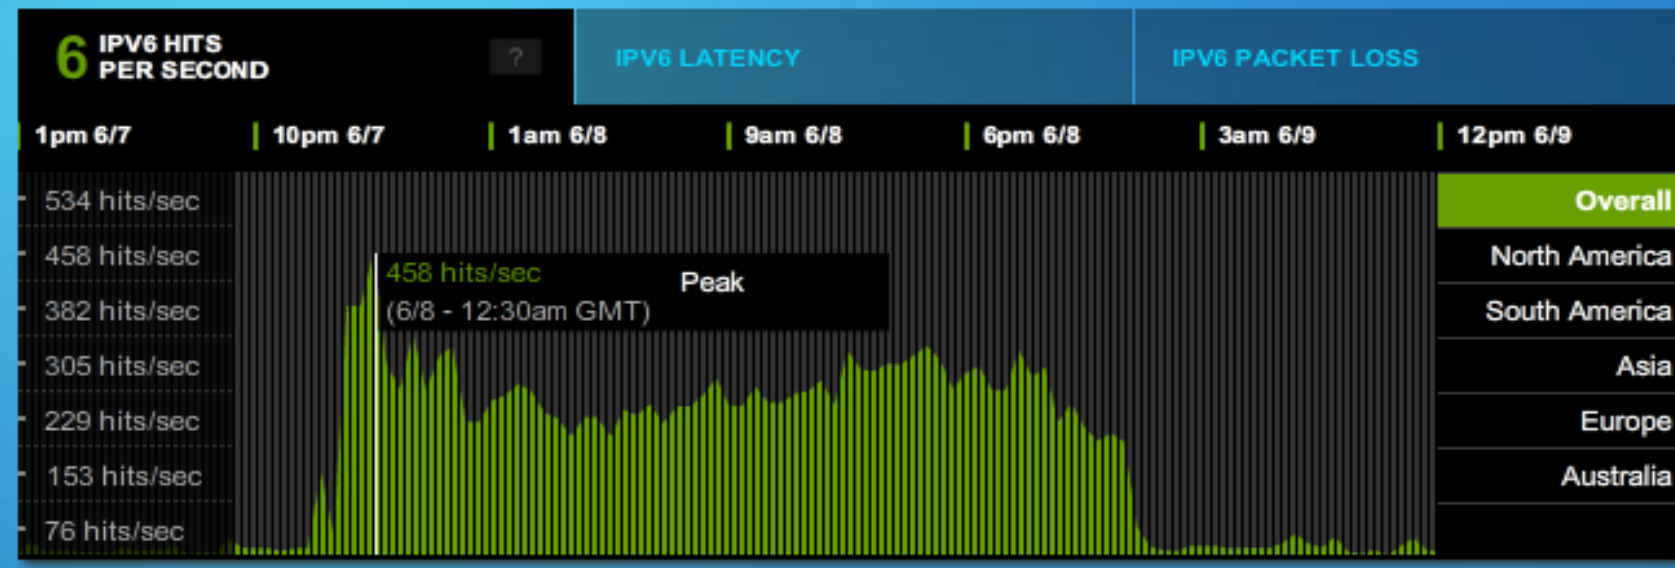
IPv6 World Day



IPv6



Akamai IPv6 Statistics



Experience with IPv6 – DNS / Mapping



- Handling an "IP" in our software as an abstraction
 - rather than have separate data types or separate fields for IPv4 vs IPv6).
 - This is more work up-front, but it often ends up being cleaner in the end.
- Still finding v6 bugs in Client-OS but not yet on our underlying Linux Server OS.

Experience with IPv6 – Network



- Router OS
 - No big problem with current OS version
- Configurations
 - Standard configurations
 - A lot of examples and best practices are online available
- Network Interconnects
 - Still not every transit provider can offer v6 in every pop.
 - Present at 46 IX in the world with thousands of sessions, a lot of admin / manual work to get the same connectivity like in v4.

Experience with IPv6 – IP allocation



- ARIN space:
 - 2600:1400::/27
 - Multiple discrete network policy
 - Announcing a /32 per peering or transit cluster
- RIPE space
 - 2a02:26f0::/32
 - Announcing a /48 per peering or transit cluster
- This will have an impact on the v6 routing table size
- Using a /64 - /48 per on-net deployment from the partner network

- Credits: David Belson, Erik Nygren, Matthew Levine
- URL: <http://www.akamai.com/ipv6>
- Questions: ck@akamai.com