



World IPv6 \$NEXT: Impact on Japan

Lorenzo Colitti
Erik Kline

All data and analysis is to the best of our ability as this time.

Overview



- Google's IPv6 measurements
- State of IPv6 in Japan and World IPv6 Day
- World IPv6 \$NEXT and impact on Japan
- Possible solutions

Measuring IPv6

Google

Methodology



- We conduct ongoing measurements of the IPv6 connectivity of Google users
- Use Javascript inserted into Google properties
 - YouTube, Google ad network
- Two HTTP requests in random order
 - One to an IPv4-only host, one to a dual-stack host
- "Sentinel" request after 30s to report results
- Measurements over 15s are considered failures
 - 15s represents unacceptable performance

- Dual-stack impact
 - Dual-stack failure rate - IPv4 failure rate
- Latency impact
 - Time taken by dual-stack request minus time taken by IPv4 request
- Native IPv6 adoption
 - Measurements where dual-stack probe is over IPv6
- Aggregated by DNS server, AS, country, ...
- Impact can be negative if IPv6 is better

IPv6 in Japan

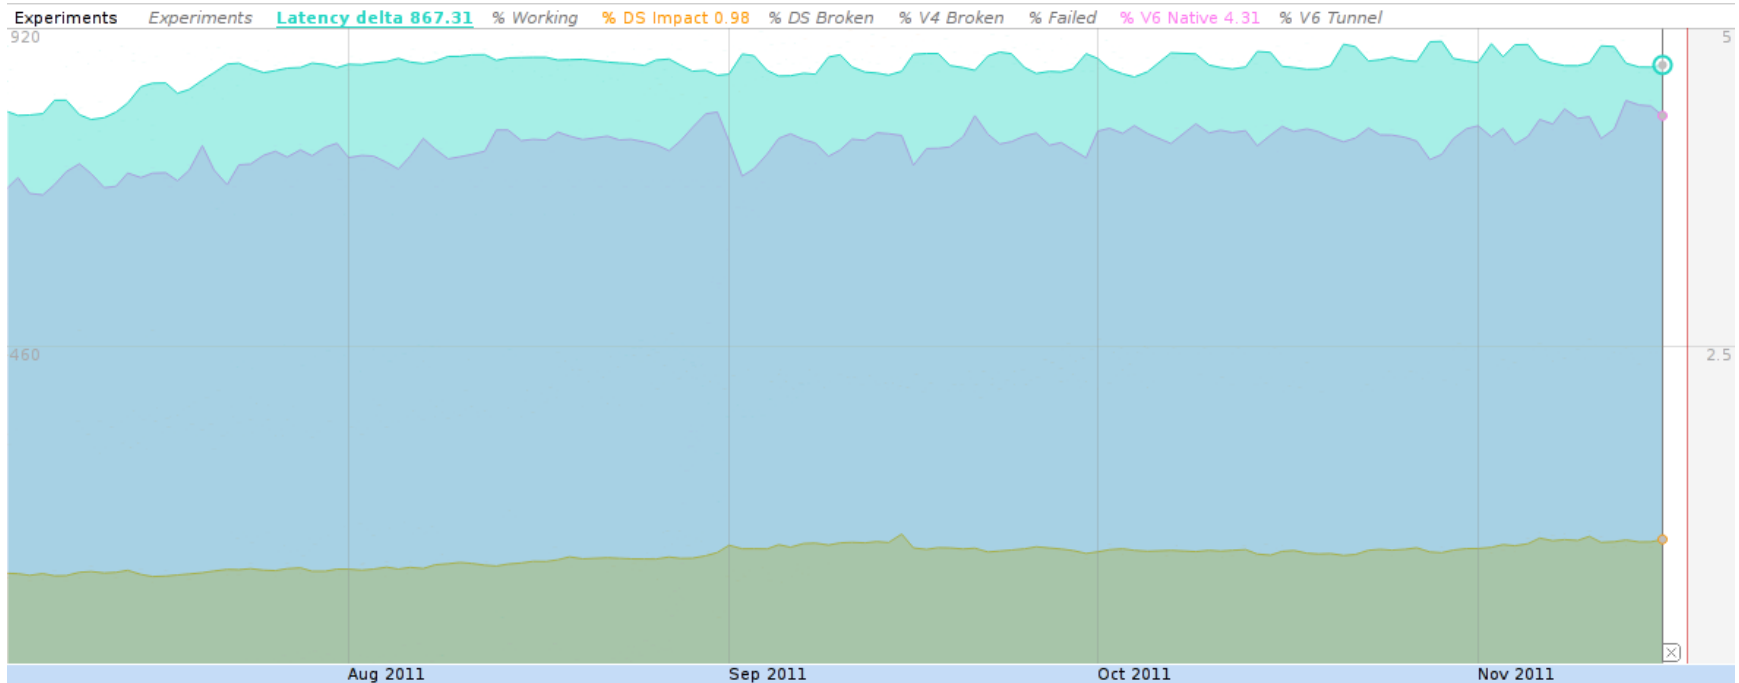
Google

A leader, in its own direction




- Wide-scale IPv6 deployment
 - KDDI one of the largest IPv6 deployments in world
 - NTT B-FLETS / NGN brings IPv6 to all users
- Closed networks
 - NTT B-FLETS / NGN not connected to the Internet
 - Options for IPv6 Internet access (Type 2/4) complex
- Result:
 - 1.4% of users have IPv6 Internet (#2, after France)
 - O(30%) users have impaired IPv6 (highest in world)
 - Enormous untapped IPv6 deployment potential

Current state



- Graph of IPv6-connected hosts in Japan
 - Represents ~33% of traffic (Rest Of World: ~8%)
 - Failure rate: ~1% (ROW: 0.6%)
 - Latency impact: ~870ms (ROW: ~0ms)

Closed network, TCP resets




Source: APNIC 32 meeting, 2011-08 [[slides](#)]

IPv6 in B-FLETS/NGN




- Host has default route, but only partial reachability
 - Doesn't know it can't reach Internet, so tries to
- Network resets IPv6 TCP connections
 - Relies on IPv6-to-IPv4 fallback for things to work
 - Not mandated by IETF
 - Not implemented by all applications
 - Causes failures on all platforms
 - e.g., if RST is lost?
 - On Windows also causes a 1+ second delay
 - Security feature (protects against RST spoofing)
- UDP doesn't work

World IPv6 Day in Japan



- Clear dent in the graph
- Some ISPs used AAAA filtering to mitigate the problem
- Different ISPs reacted differently
- Let's take a look in more detail


Example 1: OCN



No AAAA filtering


(Graph excludes Plala / Hikari TV, which are also in AS4713)

Example 2: Biglobe




AAAA filtering on World IPv6 Day only

Example 3: Nifty



AAAA filtering a week before and a month after



- Rolled out IPv6 to 15% of AS2516 users
- Made possible by their own fiber, not B-FLETS / NGN
- No latency issues, except for a minority of NGN users

World IPv6 \$NEXT

Google

World IPv6 \$NEXT



- Follow up to World IPv6 Day
 - Currently under discussion
 - Tentative date: June 2012
- Three categories:
 - Content: enable IPv6 for websites **permanently**
 - Access: bring IPv6 to **x% of active users** ($x \geq 1$)
 - Home router: **enable IPv6 by default** on product line
- Tentative participants:
 - Google, Facebook, Yahoo! (intl.), Bing, Akamai, ...
 - Free, KDDI, AT&T, Comcast, Time Warner Cable, ...
 - Linksys, D-Link, ...

Impact on websites



- Simulated impact of RSTs on various websites
- Windows client PC connected through Linux router
- Linux router simulated TCP resetter using iptables
 - `iptables -A FORWARD -p tcp -j REJECT --reject-with tcp-reset`
- Browser restarted and cache cleared before each run
- Measurements taken using stopwatch, in seconds
 - Dual-stack website with working IPv6
 - Dual-stack website with IPv6 TCP resets

Firefox 7 on Windows 7



	IPv6 (seconds)	IPv6 to IPv4 fallback (seconds)	Slowdown
Google homepage	1.5	3.5	2.33x
Gmail homepage	2.5	6.5	2.60x
YouTube homepage	3.7	5.7	1.54x
Google video	5.7	6.3	1.11x
Firefox search box search	1.0	3.0	3.00x
Kame homepage	2.4	6.4	2.67x
KDDI homepage	1.4	3.4	2.43x
MIC homepage	2.7	3.1	1.15x

IE 9 on Windows 7



	IPv6 (seconds)	IPv6 to IPv4 fallback (seconds)	Slowdown
Google homepage	0.9	3.3	3.67x
Gmail homepage	1.5	5.4	3.60x
YouTube homepage	3.0	4.8	1.60x
Google video	2.3	5.0	2.17x
IE search box search	1.8	3.8	2.11x
Kame homepage	2.2	5.1	2.32x
KDDI homepage	1.6	3.8	2.38x
MIC homepage	3.7	4.6	1.24x

Possible solutions

Google

Non-solution: AAAA filtering



- AAAA filtering by the ISP is not a long-term solution
 - Eliminates fallback problem...
 - ... but disables IPv6 even for IPv6-capable users
- Common implementation filters AAAA records over IPv4
- Unfortunately, most DNS queries are over IPv4
 - NTT home gateway prefers DNS over IPv4
 - PPPoE clients seem to prefer PPPoE DNS over LAN
 - B-FLETS / NGN PPPoE DNS is IPv4-only
- Once enabled, it will be hard to disable

Short-term solutions



- Filter out AAAA records in the NTT home gateway
 - If the user does not have Option 2/4...
 - ...and the address is not from the FLETS prefixes
- Provision new DNS servers for IPv6 NGN users
 - Via RADIUS, or static configuration
 - Unclear whether this is possible
 - Enable AAAA filtering on existing DNS servers
- Wide-scale rollout of IPv6 using option 4
 - Currently not possible; users need to contact NTT
 - Move B-FLETS users to NGN or filter AAAA for them

Website AAAA filtering



- Impact if nothing is done:
 - 33% of Japan suffers ~870ms latency, ~1% failure
 - ... on all connections to World IPv6 \$NEXT websites
 - ... permanently
- This may force websites to disable IPv6 for Japan
 - Google, Facebook, Yahoo! (intl.), possibly Akamai
- Yahoo! JAPAN:
 - Not currently on World IPv6 \$NEXT participant list
- Not all websites have the ability to filter AAAAs

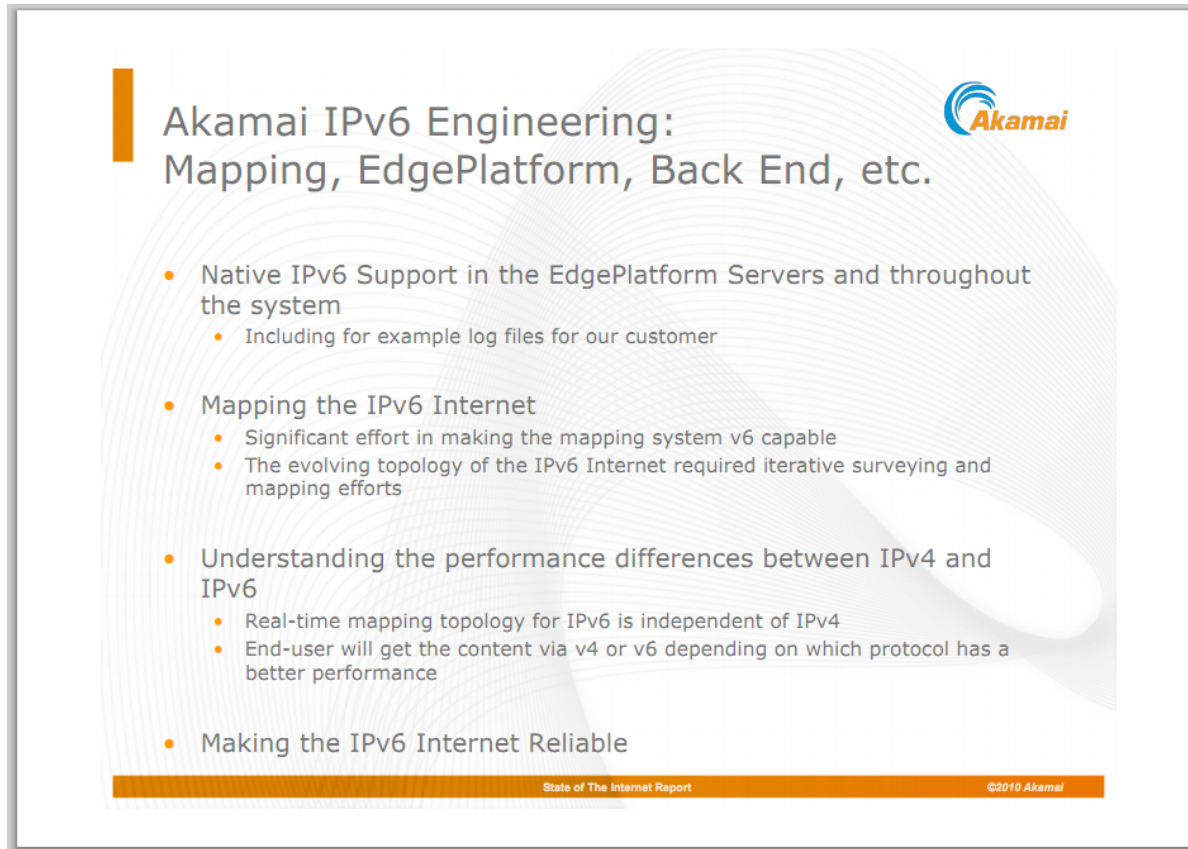
Walled Gardens

- In some locations, IPv6 has been deployed a long time.
- Television, voice, and internet combined
- Customers have IPv6 default route visible to their PC
- But limited to the ISP network – not the public network.

- Example: NTT Japan
 - › Will send TCP RST to fail IPv6 connections faster
 - › Windows will take 1+ second to fall back
 - › Concerned about World IPV6 \$NEXT_EVENT and user performance
 - › Some Japanese ISPs are considering filtering AAAA records ☹
 - › We are considering not serving AAAA records to Japan ☹☹☹



Source: TWNIC Global IPv6 Summit, 2011-11 [[slides](#)]



The slide features a background of a globe with latitude and longitude lines. The Akamai logo is in the top right corner. The title is 'Akamai IPv6 Engineering: Mapping, EdgePlatform, Back End, etc.' followed by a list of bullet points. At the bottom, there is an orange bar with the text 'State of The Internet Report' and '©2010 Akamai'.

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

State of The Internet Report ©2010 Akamai

Source: Netnod autumn meeting 2011 [[slides](#)]

1+ second latency could be seen as "worse performance"

Long term solutions



- Filtering mitigates but does not solve the problem
- Global IPv6 rollout
 - If everyone has working IPv6, there is no problem
- Redesign network to use Route Information Option instead of a default route
 - Might break devices in the field
- Better multihoming support on hosts
 - Currently unclear what this would look like
 - A default route is a default route
 - Even when it doesn't work for all destinations

Google™ Questions?
