



IPv6 at Google

a case study

Angus Lees
Site Reliability Engineer


Steinar H. Gunderson
Software Engineer

A Brief History



14 March 2005	Register with ARIN — 2001:4860::/32
August 2007	Network architecture and software engineering begins (all 20% time)
5 December 2007	Mark Townsley publicly suggests Google serve AAAAs by IETF 73 (which Google was sponsoring)
11 January 2008	First production IPv6 router
29 January 2008	First public demonstration of Google services transparently accessible over IPv6
12 March 2008	Launch ipv6.google.com for IPv4-blackout hour at IETF 71
23 October 2008	First external trusted tester of Google over IPv6
16 November 2008	IETF 73 conference network receives AAAAs for Google services
7 January 2009	Official public availability of Google over IPv6

IPv6 Maptiles Traffic



Global IPv6 Statistics


Measuring the current state of IPv6 for ordinary users

- There was **too little data** about IPv6 among clients
 - Existing measurements mostly on a small scale and/or only indirectly related to client IPv6 availability (e.g. IPv6 traffic percentage, IPv6-enabled ASNs)
 - Best existing number was probably 0.086% (Kevin Day, March 2008)
- General worry that turning on IPv6 would reveal all sorts of brokenness
 - Tunnels that people have forgotten
 - Suboptimal routing
 - Home routers doing evil things to AAAA queries
- We need to figure out **how common** IPv6 is among our users, how prevalent **brokenness** is, and how we can best serve our IPv6 users
 - Our question: What is the impact of adding an AAAA record to a website?

Methodology



- Enroll a small fraction of ordinary Google users into an “IPv6 experiment”, where their browser is asked to perform a background request
 - Involves users from all datacenters equally, but background request goes to one of two datacenters (one in the US, one in Europe)
 - Cryptographically signed to avoid easy injection of false data



- Recorded information:
 - **IPv4 and IPv6 addresses**, as applicable
 - Background request **latency**
 - Browser/OS details (**User-Agent** string)

Key Figures

Overview of connectivity and latency data




Connectivity Summary



- 0.26% of clients have IPv6 and will prefer it
 - Slowly increasing over time
- 0.08% of clients have broken IPv6 connectivity
 - That is, adding an AAAA record will make these clients unable to view your site
 - Due to statistical issues, this is a much less accurate figure (could easily be 0.06% or 0.1%), so take it with a grain of salt
- Probably at least a million distinct IPv6 hosts out there
 - Again, a number with statistical caveats


Connectivity Growth



IPv4 Latency




Latency distribution function,
clients visiting ipv4.ipv6-exp.l.google.com



IPv4 Latency (cont'd)



Latency distribution function,
IPv4-only and IPv4 to a dual-stack host



Not indicative of ordinary Google service latency

IPv6 vs IPv4 Latency




- We cannot directly graph IPv4 vs IPv6 latency
 - IPv6-enabled hosts are likely to have faster (or at least different) network connectivity overall (e.g. universities, power users, ...)
 - Need a way to remove inherent bias
- Solution: Find **pairs** of hits from the same /24 IPv4 network, discard all other data
 - Gives comparable (paired) data sets
- This means we are measuring relative latency for a different set of users, but the data is still indicative of what you can expect today

IPv6 vs IPv4 Latency (cont'd)



Latency distribution function,
IPv4 and IPv6 to a dual-stack host




Not indicative of ordinary Google service latency

Data Breakdowns

Drilling down for a more detailed look



IPv6 Connectivity By Weekday (UTC)



More working IPv6 at home?

IPv6 Connectivity By Country




- Geolocate by the client IPv4 address, then group by country
 - Some countries with relatively little Internet traffic removed

Country	IPv6 penetration
Russia	0.97%
Ukraine	0.76%
France	0.62%
Norway	0.50%
United States	0.46%
...	
China	0.31%
Japan	0.16%

Method of IPv6 Connectivity



- Based on the IPv6 address we can infer how the user gets IPv6 access
 - Unfortunately, no good way of distinguishing native from tunnels (based on the address alone)
 - Vista with Teredo prefers IPv4 by default, so probably undercounted



- Some countries stand out
 - United States, Canada: 95% 6to4
 - France: 95% native (almost all free.fr)
 - China: 70% native, 26% ISATAP

IPv6 Connectivity By OS



IPv6 penetration and connectivity type by Operating System

Ranked by overall IPv6 penetration

Operating System	IPv6 penetration	Native/...	6to4	Teredo	ISATAP
Mac OS	2.42%	9.5%	90.4%	0.04%	0%
Linux	0.96%	85.4%	13.4%	0.95%	0.25%
Windows Vista	0.37%	55.0%	43%	0.01%	2%
Windows Server 2k3	0.07%	75.0%	13.7%	5.8%	5.5%
Windows XP	0.04%	48.1%	29.5%	12.8%	9.5%
Windows 2000	<0.01%	68.1%	28.9%	3.0%	0%

50% of all IPv6 hits are from
Macs with 6to4

97% of all Teredo users are on Windows
(even undercounting Vista)

Overall Trends



- IPv6 prevalence is still low, but **growing** week by week
 - Large (and sometimes surprising) variations among individual countries
 - Still heavily influenced by single deployments (e.g. free.fr)
- It's **not that broken**
 - ~0.08% clients lost, ~150ms extra latency
- The default Operating System **policy matters** – a lot
 - Vista: 10x greater prevalence over XP (Vista defaults to enabling IPv6)
 - Mac OS: 8x greater prevalence than Vista (Airport Extreme, 6to4 enabled)
- 6to4 is by far the most common transition mechanism in use
 - Possibly due in part to the Airport Extreme
 - Consider running your own 6to4 relays for return packets (i.e. 2002::/16 routes to your own gear)

Observed Issues

A small sampling of IPv6 client issues



- The DNS resolver in some home gateways exhibits very broken behaviour
 - Some simply **drop AAAA** queries, leaving clients to timeout
 - Some actually **mangle AAAA** RRs into an A RR
 - Sufficiently problematic that the one DNS resolver library, c-ares, has a patch to treat a **malformed** response as **no** response, i.e. issue a subsequent A lookup^{1,2}
- Some DNS servers handle only A records and do bad things with AAAA requests
 - Some **don't reply**, some **return NXDOMAIN**, some **make a lame delegation**
 - For a while nytimes.com, among others, replied to all non-A requests with **referrals to the root** nameservers
 - Long timeouts as resolver libraries **resolve in circles** until they give up (actual error mode and timing was library-dependent)
 - For a complete discussion see RFC 4074

¹ <http://c-ares.haxx.se/mail/c-ares-archive-2009-01/0002.shtml>

² <http://c-ares.haxx.se/mail/c-ares-archive-2009-01/0006.shtml>

- Most browsers seem to either cleanly support or not support IPv6
 - Firefox has a [network.dns.disableIPv6](#) Boolean configuration option (defaulted to true for 2.x and many platforms)
- It's the "intermediate" support that's a problem
 - From the "IPv6 for Microsoft Windows" FAQ:

"URLs that use the format for literal IPv6 addresses ... are [not supported](#) by the Internet extensions DLL ... provided with Windows XP and Windows Server 2003."¹
 - Safari on 10.4 (Tiger) reportedly has the same problem
 - Google Chrome had a similar problem at launch time (subsequently fixed)
- Operating System default IPv6 policy is very important
 - Compare Vista Teredo traffic (~0%) with XP (~12.5%)

¹ <http://www.microsoft.com/technet/network/ipv6/ipv6faq.msp>

Thank You!

Q&A

Google