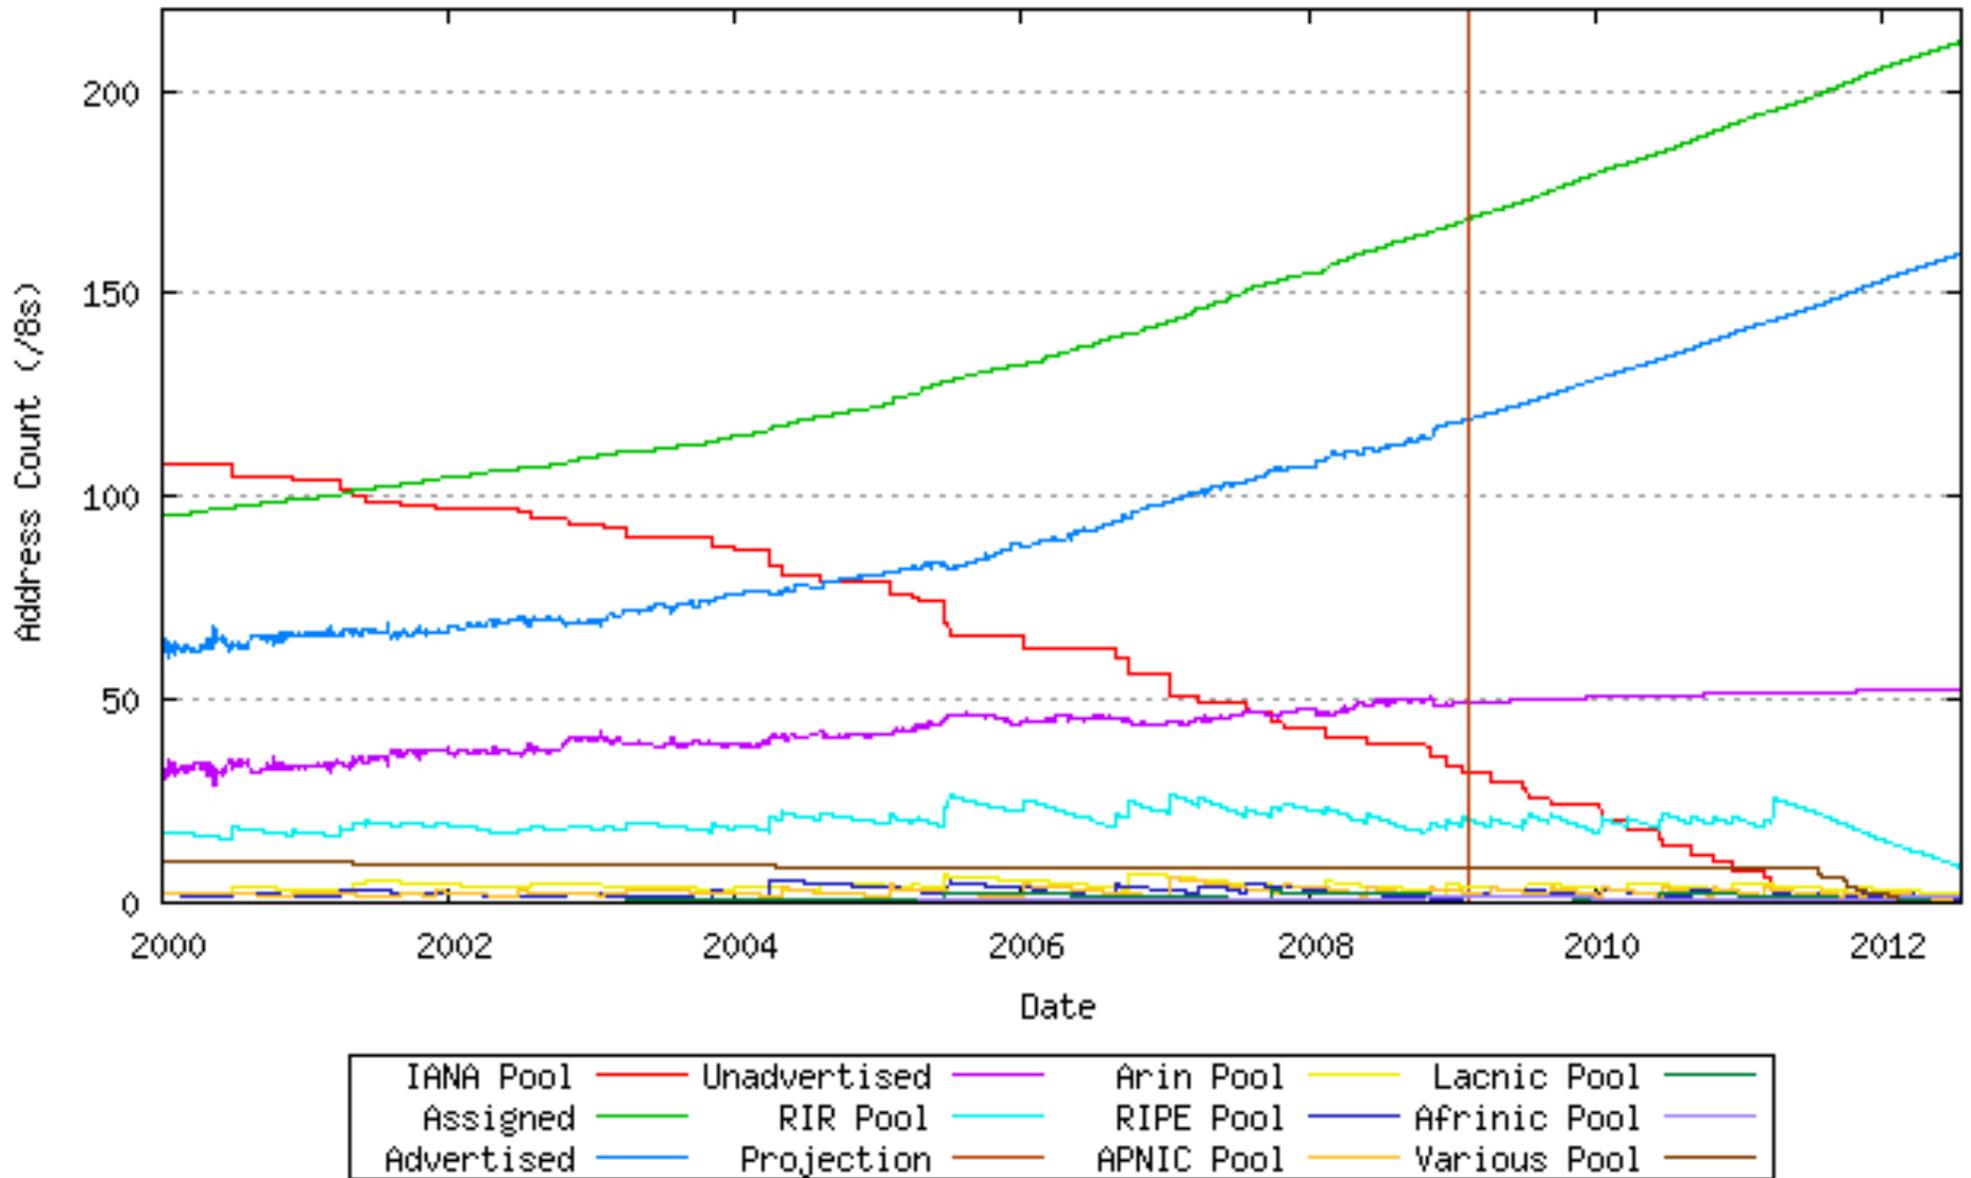




IPv6 at Google

Lorenzo Colitti
lorenzo@google.com

Why?

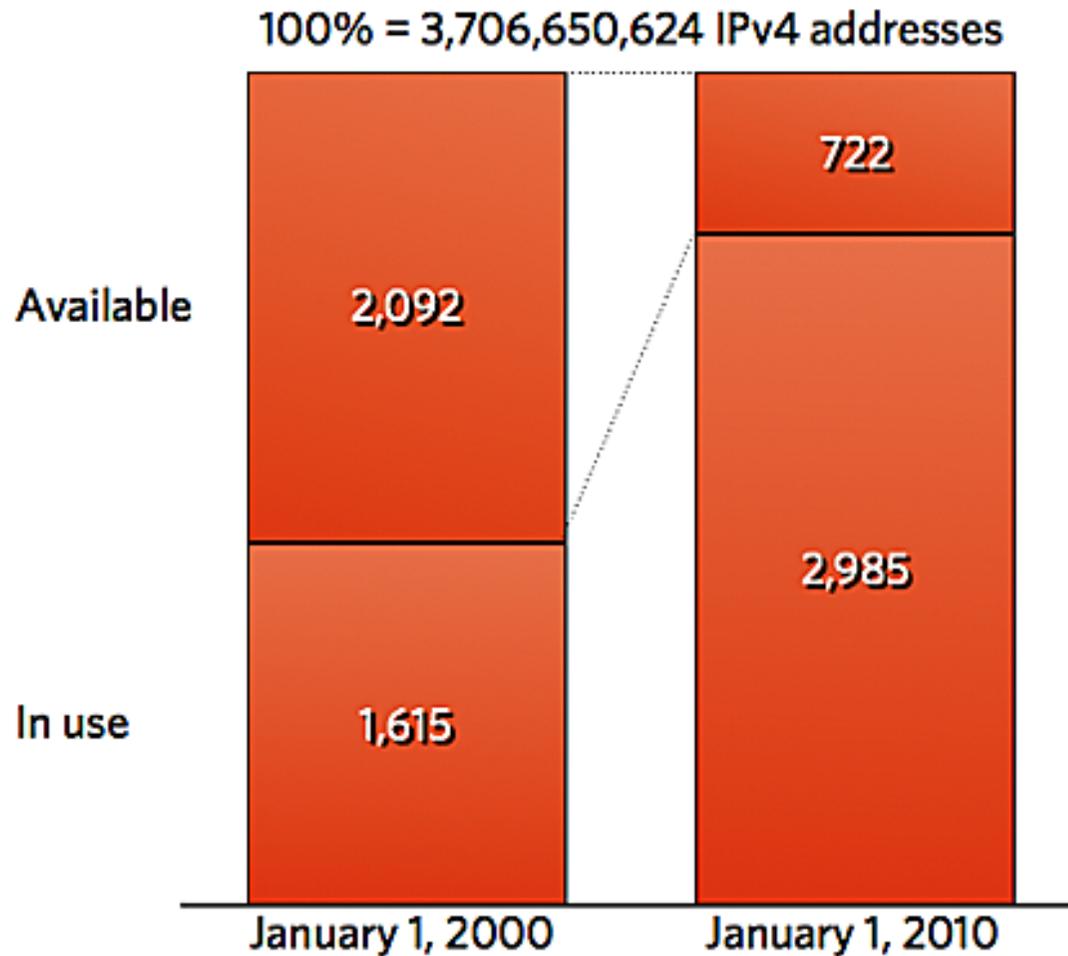


IPv4 address space predictions (G. Huston)

To put it into perspective...

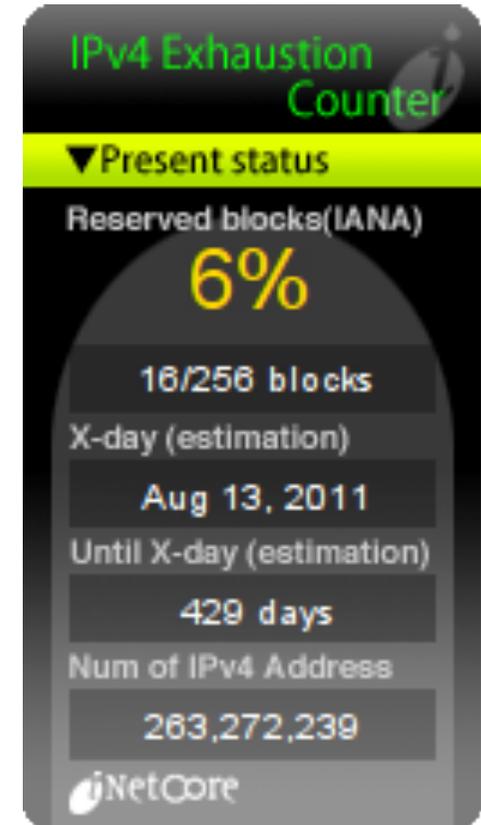
IPv4 address utilization: 2000 vs. 2010

Millions



Ijitsch van Beijnum, Ars Technica

Lorenzo Colitti



Takashi Arano, Inetcore



June 2010

Why IPv6? Cost

- Buying addresses will be expensive, more so over time
- Carrier-grade NAT may be expensive
 - Lots of session state memory
 - Session logging for legal reasons
 - Bandwidth
- Being behind a NAT is hard to manage
 - Can't fix problems without NAT operator's help
 - VPN, VOIP, video streaming, gaming, P2P
 - Expensive in operator time, support costs
- Network complexity creates operation / support costs

Why IPv6? Address semantics

- With carrier-grade NAT, users share IP addresses
- Less accurate geolocation
 - Content licensing for streaming, etc.
- Abuse identification / blocking
 - If an IPv4 address is spamming/hacking/...
 - If we block it, do we take out 100 users?
 - Different from blocking a proxy
- Port exhaustion and HTTP intercept
- IP-based authentication (nobody does this of course)

Why IPv6? New devices

- We see a growing number of IPv6-only deployments
 - Set-top boxes (free.fr, Comcast, ...)
 - NTT IPTV over IPv6
 - Mobile networks (LTE, NAT64, ...)
- There is simply not enough address space to assign IPv4 addresses to these devices
 - NAT is too expensive
 - CPU on home gateway
 - CPU on routers
- Want to talk to these devices? Need to use IPv6

Why IPv6? New applications

- The Internet was successful because of end-to-end
- Users *still want* end-to-end!
 - Skype, Bittorrent, ...
 - Neither of these could have been developed in the absence of public IP addresses
- What happens if this goes away?
 - Will the Internet become like TV?
 - Will the Internet become like the phone network?
 - Will any Internet communication require ISP support?

The search for the killer application

- Many are waiting for a "killer application" for IPv6
- This is a misconception
 - It's not "what can IPv6 can do better than IPv4?"
 - It's "can the Internet as we know it continue to operate using IPv4?"
- The killer application of IPv6 is the survival of the open Internet as we know it

So what have we done?

Enable IPv6 for www.google.com?

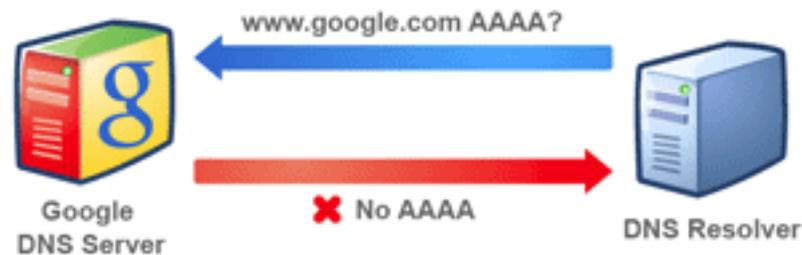
- We can't enable IPv6 for www.google.com today
 - ~0.05% users won't reach Google any more
 - Broken home gateways, DNS forwarders, ...
 - If you have a problem, can't reach Google to help fix it
 - 0.05% is a lot of users!
 - Many users would have higher latency
 - Long paths, suboptimal routing, tunnels...
 - Bad and incomplete routing is there, but you can't see or fix it if there's nowhere useful to go
 - This has gotten much better now

Google over IPv6

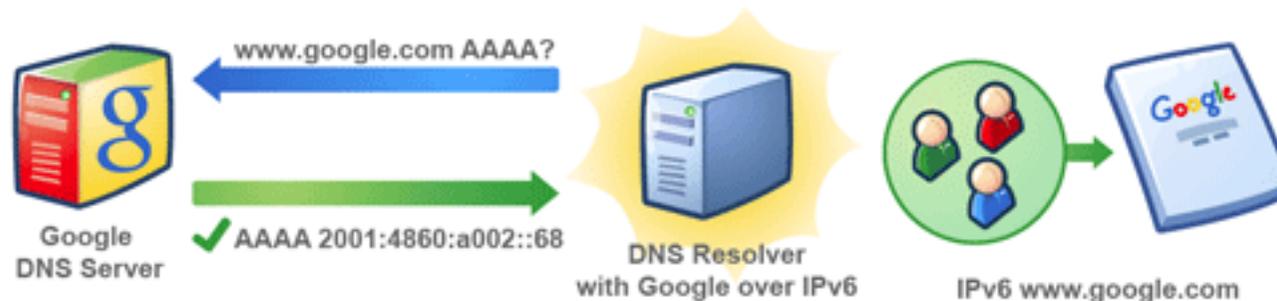
- Enables IPv6 access to Google for selected networks
- IPv6 access to most Google web properties
 - www, mail, news, docs, youtube, ...
- Requirements:
 - Good IPv6 connectivity to Google
 - Production-quality IPv6 network
 - Acceptable user breakage
 - Much harder than it sounds

How it works

Normally, if a DNS resolver requests an IPv6 address for a Google web site, it will not receive one...



...but a DNS resolver with Google over IPv6 will receive an IPv6 address, and its users will be able to connect to Google web sites using IPv6.



<http://www.google.com/ipv6/>

Results so far

- Enthusiastic response:
 - Over 80 organizations participating
 - ~ 75% of the native IPv6 Internet
- Feedback so far has been positive
 - Some networks see better IPv6 routing than IPv4
 - Now enough IPv6 traffic that problems get reported
- In unmanaged networks, support is a problem
 - Users with broken home networks can't reach Google
 - Won't know what happened: "everything else works"
 - May not call their ISP

Lessons learned

Testing and iteration

- Implementations mostly work, but will have bugs
 - Nobody has really kicked the tyres
- Don't expect something to work just because it's supported
- If you find a bug in the lab:
 - Report it
 - .. and keep testing!
 - There are many more bugs to find
- Work around it in the design
 - If you get to something that is supportable, trial it
 - That will help you find the hard bugs

For example...

- "If a firewall filter term has a 1-bit match in bits 32-64, and then term with a 2-bit match on bits 64-96, the second term will not match"
- "In particular circumstances, FIB and RIB may get out of sync due to race conditions in pushing updates"
- "If DAD triggers due to an interface loop, it requires removing config from the interface and putting it back"
- If a linux gets a packet too big on a receive-only interface with no route, it ignores it
- Are you going to find these in the lab?
 - We only saw the race condition after months in production on a largish number of datacenters

Methodology

- Tap enthusiasm
 - Started as 20% project, great influx of contributors
- Make it easy for contributors to get initial results
 - A pilot network is not expensive
 - Once network is up, internal applications follow
- Do it in stages
 - v6 needn't be as capable as v4 on day one
 - But it must be done properly
 - If it's not production-quality, it's no use to anyone
- Fold it into your normal upgrade cycles

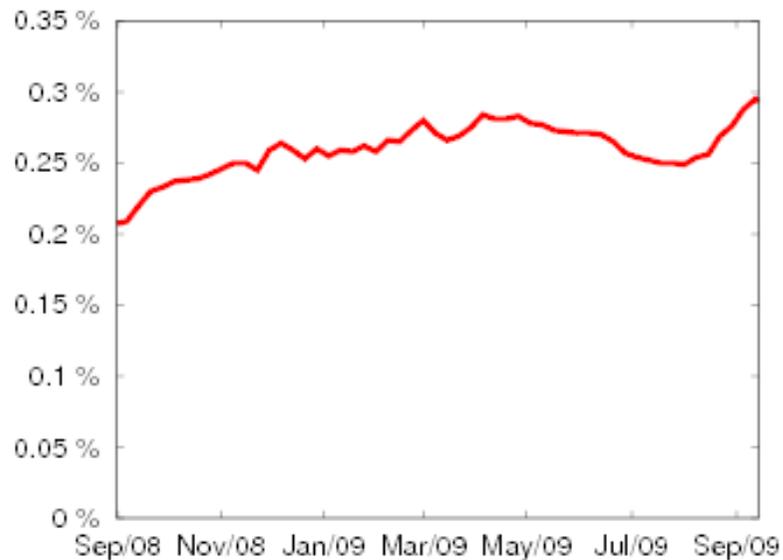
Timeline

April 2005	Obtain and announce address space
...	...
July 2007	Network architecture and software engineering begin (20%)
December 2007	Mark Townsley challenges Google to serve IPv6 by IETF 73
January 2008	First pilot router. Google IPv6 conference, Google over IPv6 for attendees
March 2008	ipv6.google.com (IETF 72)
Q3 2008	ipv6.google.cn, ipv6.google.co.jp
November 2008	First Google over IPv6 networks enabled. Google over IPv6 at RIPE / IETF / ...
January 2009	Google over IPv6 publicly available
March 2009	Google maps available over IPv6, 3x increase in traffic
August 2009	IPv6 enabled in Android (available on Droid and Nexus One)
February 2010	Youtube available over IPv6, 10x increase in traffic

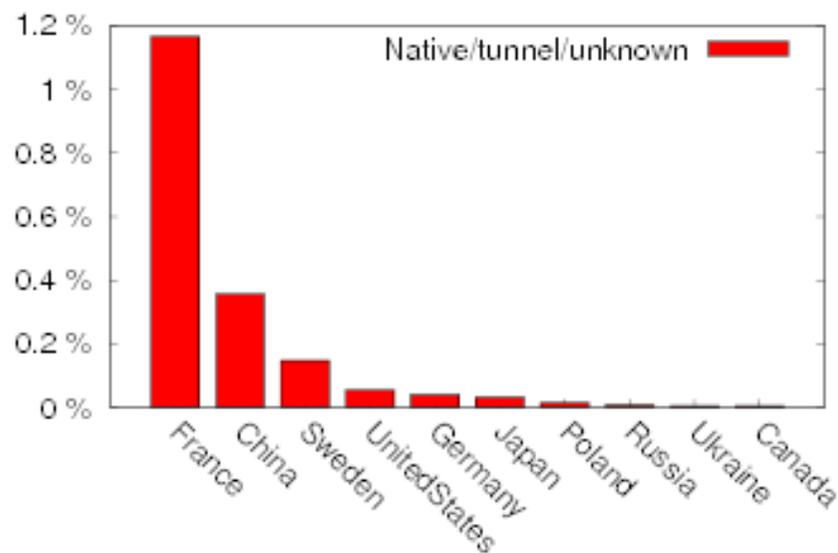
And all this with a small core team

Statistics

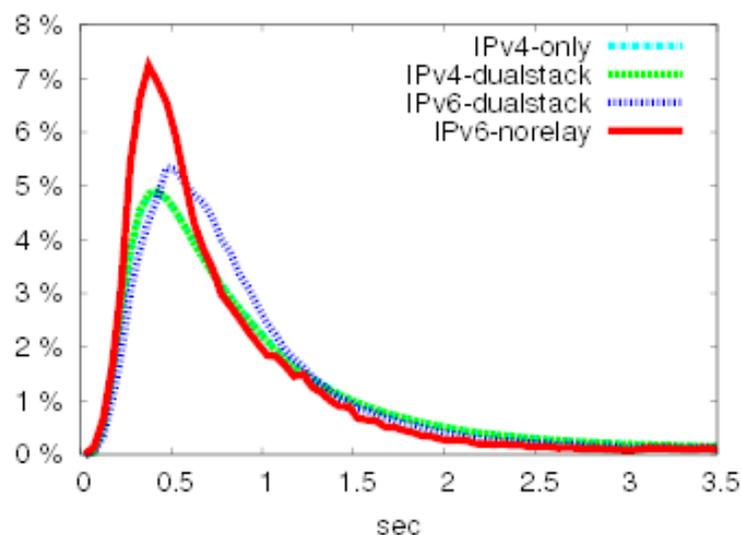
Some statistics



35% connectivity growth in one year



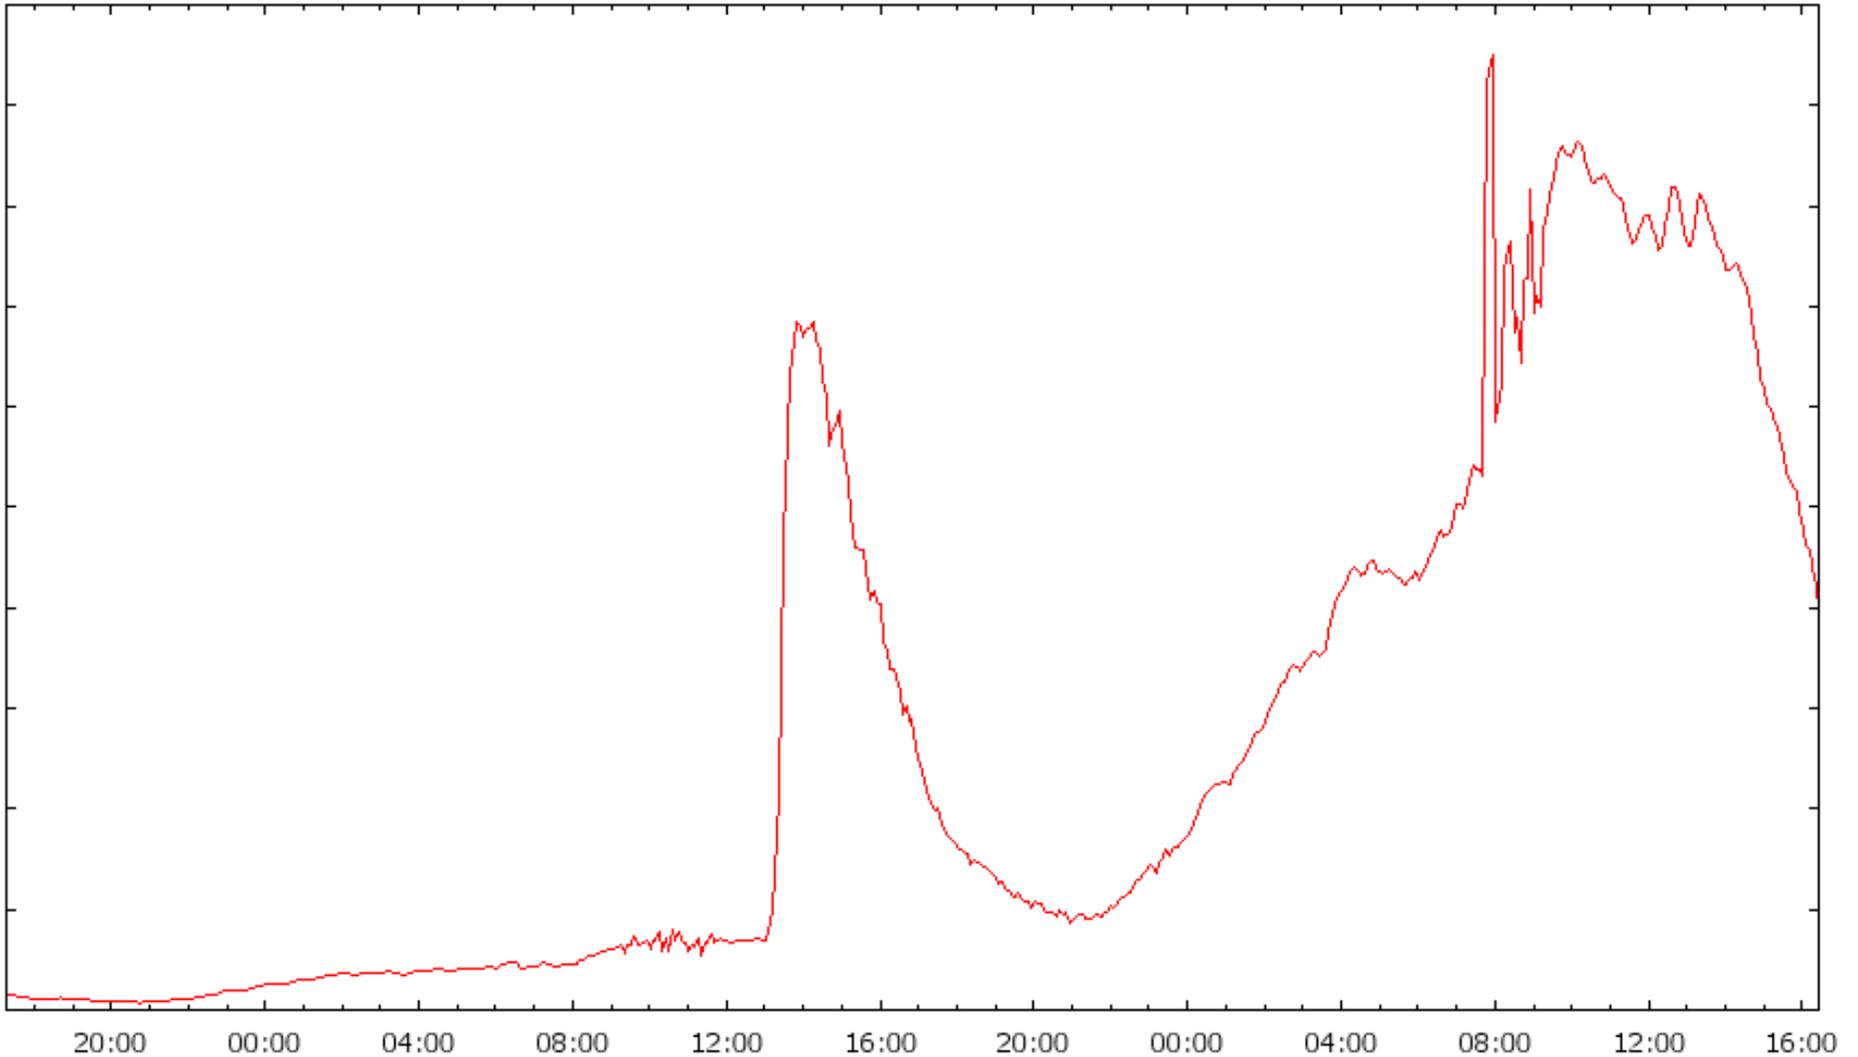
Most native IPv6 users are in France



6to4 / teredo latency penalty > 50ms

Traffic can appear overnight

Google IPv6 outbound traffic



(IPv6 video launch 2010-01-28)

Lorenzo Colitti



June 2010



Questions?

Lorenzo Colitti
lorenzo@google.com