




IPv6 at Google

Lorenzo Colitti
lorenzo@google.com

Why?



IPv4 address space predictions (G. Huston)

Why IPv6? Cost

- Buying addresses will be expensive
- 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

Why IPv6? Opportunity

- We see a growing number of IPv6-only deployments
 - Comcast set-top boxes
 - free.fr set-top boxes
 - NTT's IPTV over IPv6
- 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

Why IPv6 at Google?

- When the day comes that users only have IPv6, Google needs to be there for them
- Serve current users better over IPv6
 - IPv6 can have lower latency and packet loss
 - We have user reports to prove it
 - AJAX applications break behind excessive NAT
 - Connections exhaust public IP port space
 - Growing number of IPv6-only client deployments
 - Set-top boxes, mobile, ...
- IPv6 is good for the Internet, and we want to help

The problem

IPv6 adoption



- Climbing, but still low
- We need faster growth than this

What?

Enable IPv6 for www.google.com?

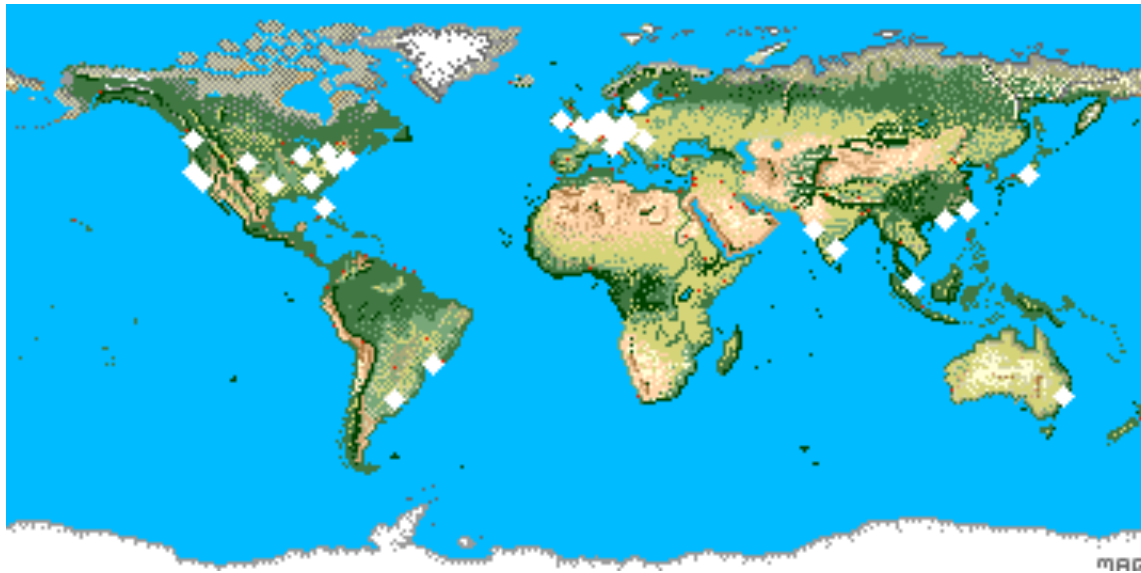
- We can't enable IPv6 for www.google.com today
 - ~0.1% users won't reach Google any more
 - Broken home gateways, DNS forwarders, ...
 - If you have a problem, can't reach Google to help fix it
 - 0.1% is a lot of users!
 - Many users would have higher latency
 - Long paths, suboptimal routing, tunnels...

So what have we done?

- IPv6 network rollout
- IPv6-only websites
 - ipv6.google.com (Mar 2008)
 - ipv6.google.cn, ipv6.google.co.jp
- IPv6 evangelism
 - Google IPv6 conference (Jan 2008)
 - IETF panels, blackout sessions, ...
 - Vendor outreach
- Google over IPv6 (Jan 2009)

Providing reliable services over IPv6

- Avoid bad routing by avoiding transit, prefix limits
- Peer with almost everybody
- Bring the network as close to the user as possible
- Serve IPv6 only to production-quality user networks



Google over IPv6

- Enables IPv6 access to Google for selected networks
- IPv6 access to most Google web properties
 - www, mail, calendar, docs, ... (no youtube yet)
 - Which ones do *you* and your users want?
- Requirements:
 - Good IPv6 connectivity to Google
 - Production-quality IPv6 network
 - Commitment to fix problems that break Google for users

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/>

Initial results

- Enthusiastic response:
 - Over 50 organizations participating
 - Universities, research institutions, an NREN
 - One large French access provider
 - > 250k unique IPv6 addresses per day
- Feedback so far has been positive
 - Some networks see better IPv6 routing than IPv4
 - Now enough IPv6 traffic that problems get reported
 - Allows participants to bypass IPv4 congestion
- Want to take part? Let us know!

google-ipv6@google.com

Lessons learned

Methodology

- Tap enthusiasm
 - IPv6 at Google started as a 20% project
 - Incredible 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

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
January 2008	Google IPv6 conference, Google available over IPv6 to attendees
March 2008	ipv6.google.com (IETF 72)
July 2008	ipv6.google.cn
October 2008	ipv6.google.co.jp
October 2008 - November 2008	First Google over IPv6 networks enabled. Google over IPv6 at RIPE / IETF / ...
January 2009	Google over IPv6 publicly available

And all this with a small core team

Design and operations

- Hardware support
 - Features mostly there, but not always well-tested
 - Might want to start with dedicated IPv6 devices
- Licensing
 - Some vendors charge separately for IPv6 support
 - Thanks to Cisco for doing this properly
- Design
 - As similar to IPv4 as possible, dual-stack if you can
 - Easy to support, easy to scale
 - IPv6 is not rocket science
 - Apart from ICMPv6, essentially == IPv4

Standards


- /127 is considered harmful by RFC 3627
 - Unfortunately, anything less than /127 is vulnerable to a packet amplification attack
- VRRP for IPv6 still in standards process

IPv6 licensing

- Some vendors (not Cisco) charge separately for IPv6
- Suppose it's \$10k per router:
 - Red tape blocks initial testing / deployment
 - Need to cut \$30k PO to try IPv6 on 3 routers
 - Bulk upgrade price blocks full rollouts
 - Have 100 routers? That will be \$1M, please...
- Charging separately for IPv6 *will* hinder adoption
 - The same goes for ISPs, exchanges, ...
- Thanks to Cisco for introducing IPv6 feature parity

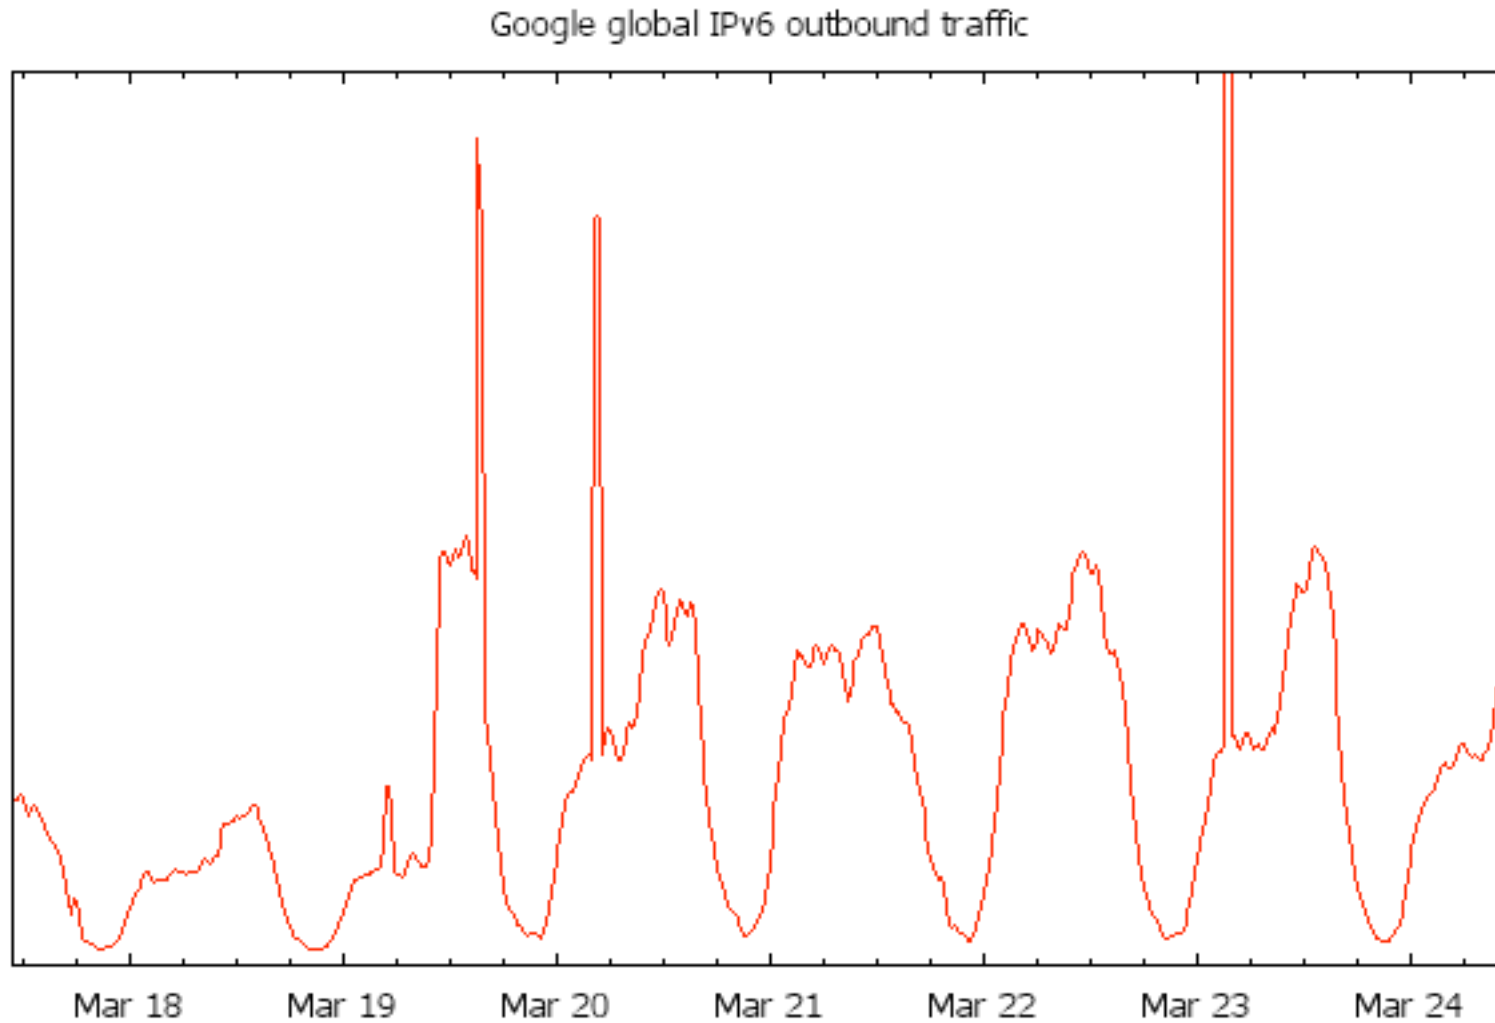
Statistics

IPv6 latency



6to4 and teredo latency penalty > 50ms

Traffic can appear overnight



(Maps enabled for Google over IPv6 users on 2009-03-19)



Questions?

Lorenzo Colitti
lorenzo@google.com