The path to IPv6

Marcus Hines
hines@google.com

Why?

No more IPv4... but no IPv6 yet either

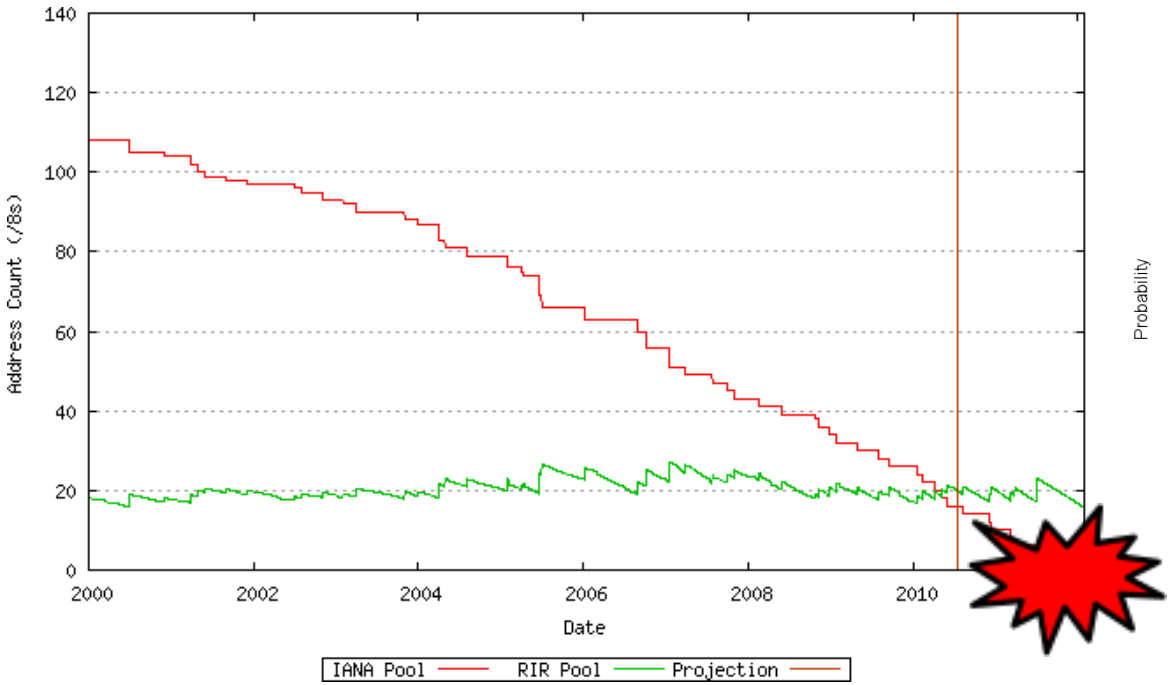

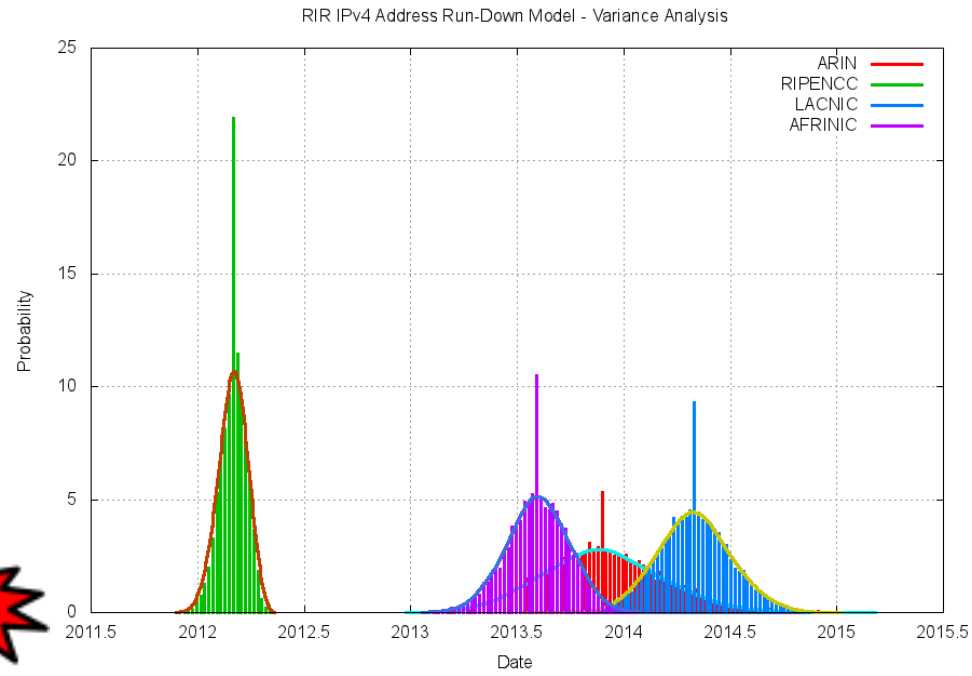

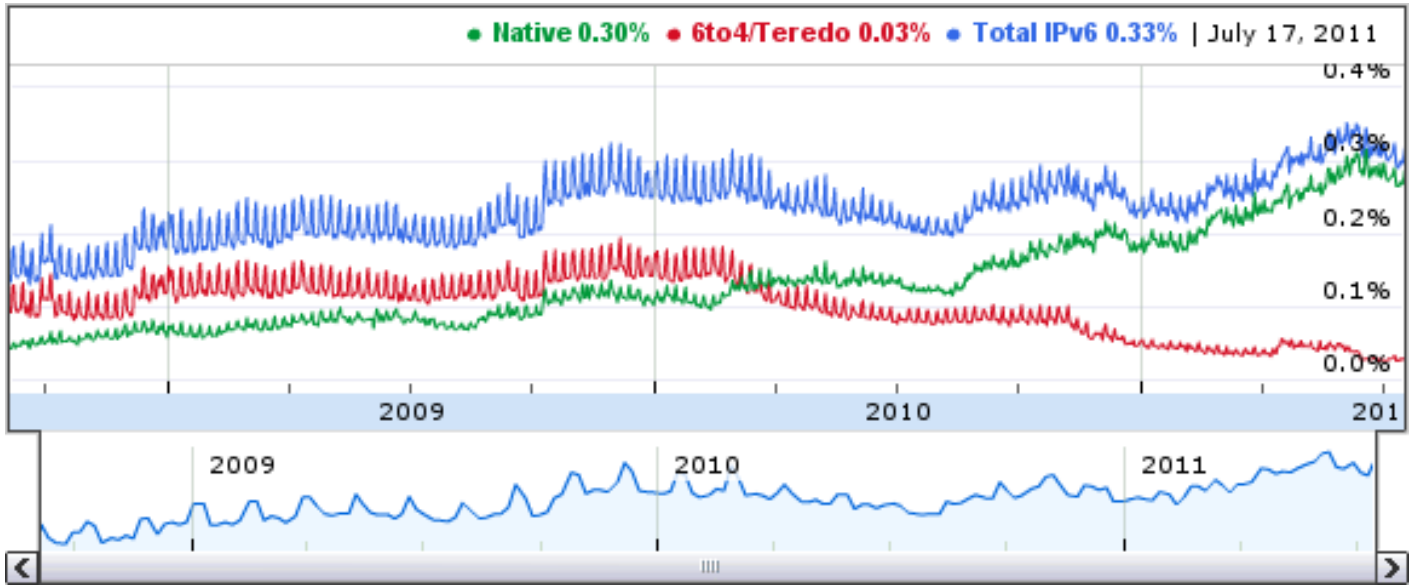

What's the alternative?

IPv4 today

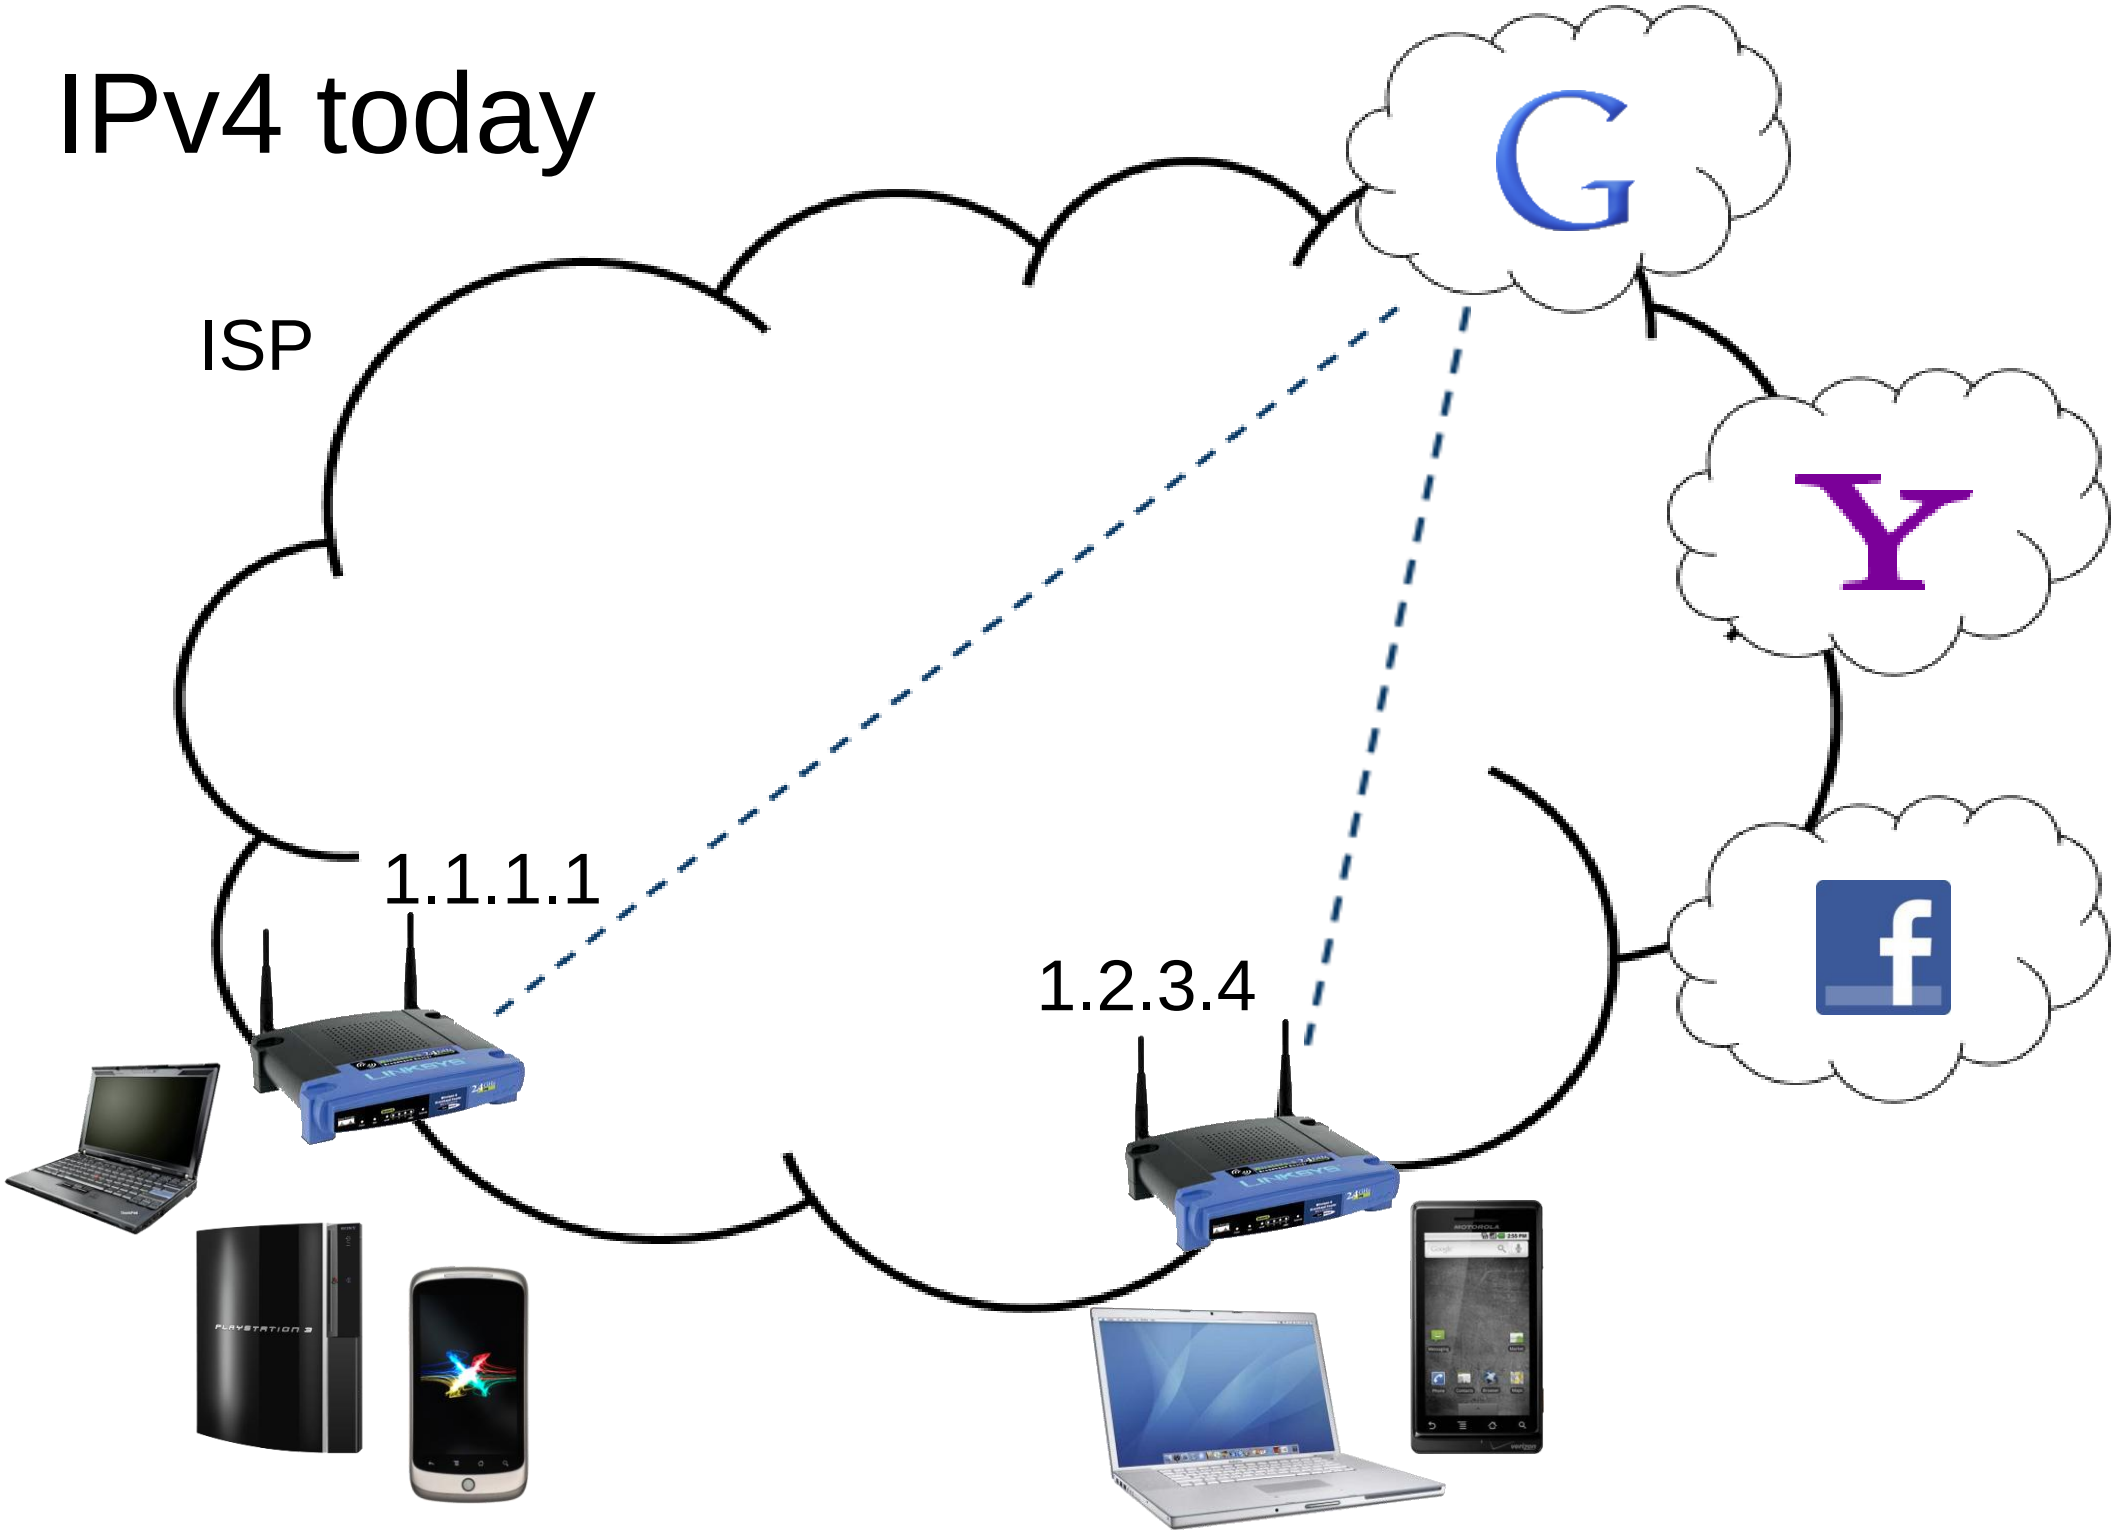

IPv4 with LSN

Impact on ISPs

- Expensive
  - Router / linecard resource use
    - Juniper: lose DPC ports, Cisco: CRS-1 blades
  - Logging TCP/UDP sessions for legal intercept
    - 5TB per month per 1M users (source: Yahoo! BB)
- Application breakage
  - VPN, VOIP, video streaming, gaming, P2P
  - UPnP doesn't work with CGN
  - Network complexity creates operation / support costs
- Will not get better over time
  - Will only get worse as more users are behind each IP

Impact on content

- 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?
- Port exhaustion and HTTP intercept
  - AJAX applications suffer

Impact on new applications

- The Internet was successful because of end-to-end
- Users *still want* end-to-end!
  - Skype, Bittorrent, cannot work 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 killer application of IPv6 is the survival of the open Internet as we know it

The Internet without IPv6

IP address sharing

TCP port overload

Content inspection/rewriting

Barriers to innovation

How?

Some misconceptions

- "We'll deploy IPv6 when users ask for it"
  - Users aren't asking for IPv4, only Internet access
    - Can we provide Internet access with no more IPv4?
- "We'll need IPv6 when IPv6-only <content|users> appear"
  - Nobody will go IPv6-only until it has >90% penetration
- "We will have to deploy IPv6 once IPv4 runs out"
  - No, carrier-grade NAT will work
  - It's just not very good, expensive, and no upgrade path
- "All our gear is IPv6 ready. We just need to turn it on"
  - "IPv6 ready" does not mean it will work
  - You'll only know it works when you turn it on

The way forward?

- IPv4 won't go away for at least a decade
- Large-scale NAT is inevitable
  - Long tail of content will not be IPv6-ready for years
- But it doesn't need to carry all traffic
  - A lot of of traffic is from a few content providers
- So, make big content available over IPv6
  - Long-tail content stays on IPv4
  - The LSNs can be smaller
    - Content providers don't suffer from LSN effects
    - ISPs save money
    - IPv6 matures

IPv6 brokenness

- To put content on IPv6, it needs to be as reliable as IPv4
- Currently, about 0.06% of Internet users experience connectivity problems when IPv6 is enabled on a web site
  - Mostly due to mis-configured / broken devices in home
  - If you have 1B users, 0.06% = 600k users
  - User doesn't know what's going on
    - "Everything else works"
    - This is unacceptable for a web site like Google
- How do we fix this?
  - Google over IPv6
  - Engage OS / browser vendors
  - World IPv6 Day

Google over IPv6

- We can't enable IPv6 for www.google.com
  - But we can enable IPv6 access for selected networks
- Most Google services are available
  - www, mail, news, docs, youtube, ...
- Requirements:
  - Good IPv6 connectivity to Google
  - Production-quality IPv6 network
  - Low brokenness

How it works

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

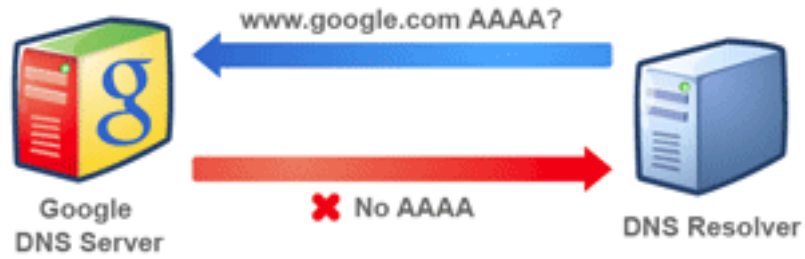

…but a DNS resolver with Google over IPv6 will receive an IPv6 address, and its users will be able to connect to Google web sites using IPv6.

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

Results so far

- Enthusiastic response:
  - O(100) 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
- Does not scale
  - We can't keep up with requests
  - Working on automating the process

World IPv6 Day

"World IPv6 day"

- Organized by the Internet society
- Google, Facebook, Yahoo!, Bing, Akamai, Limelight, ...
  - ... and about 500 others
- Turned on IPv6 for 24 hours on 2011-06-08

Preparation: ipv6test.google.com

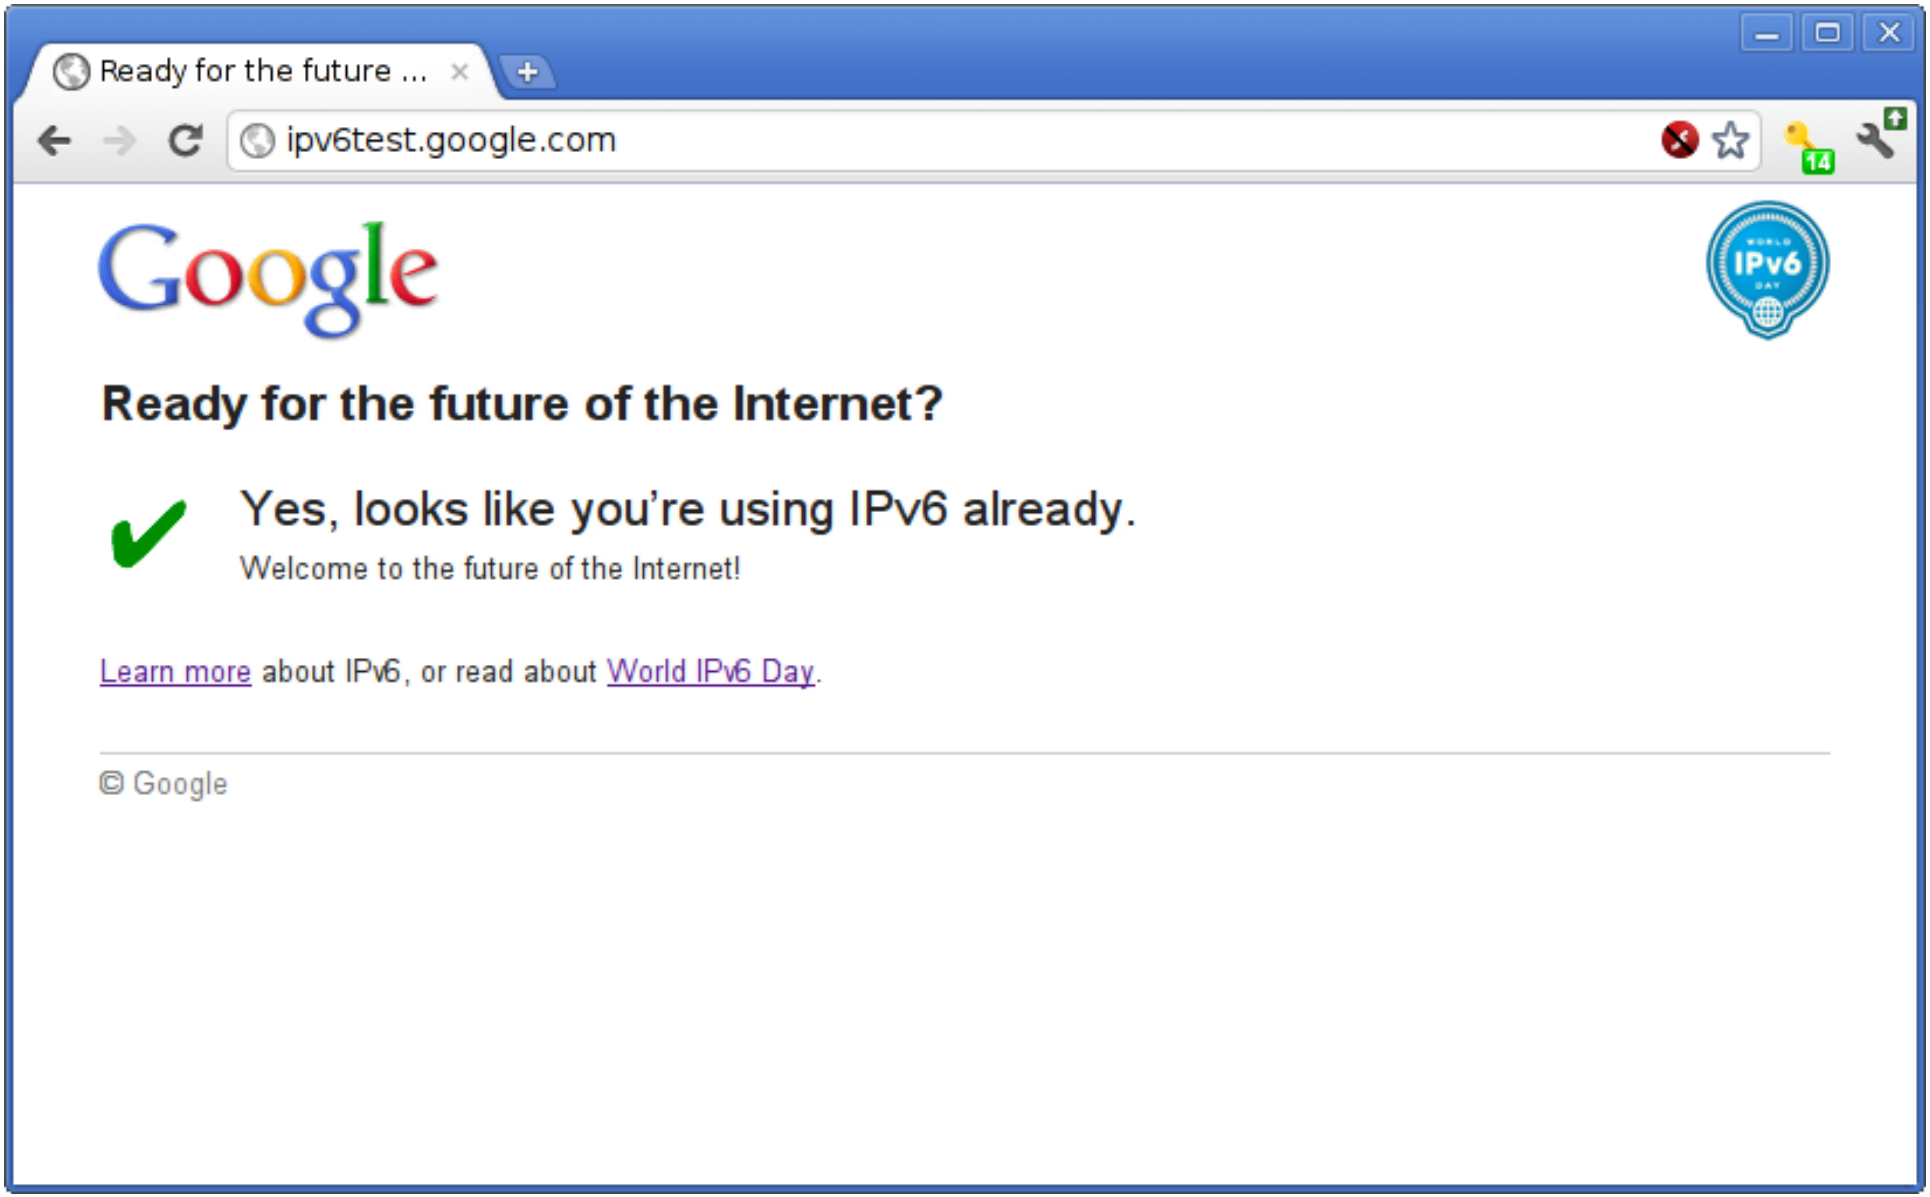

How it went

- Blissfully quiet, and publicly monitored
- Huge confidence-builder for the industry
  - CNET: "IPv6: Come on in, the water's fine"
- Caused lots of progress
  - 1000 websites participated
    - All the top 5, plus Bing, AOL, Microsoft, BBC, CNN...
  - Client rollouts (e.g., KDDI - IPv6 to Tokyo FTTH users)
- Some websites stayed on
  - developer.facebook.com, www.xbox.com
- Talk of "World IPv6 Week" in 2012

Technical progress

- A number of services became available over IPv6
  - Voice*, talk, mobile talk, analytics, tools, ...
  - Google Public DNS - 2001:4860:4860::{8888,8844}
  - Adwords and IMAP/SMTP (soon after the day)
  - Almost all Google services now available over IPv6
- Some services left IPv6 on for everyone
  - YouTube videos, mobile gtalk
- Implemented fast fallback from IPv6 to IPv4 in Chrome

*Site not actual SIP traffic

On the day: business as usual

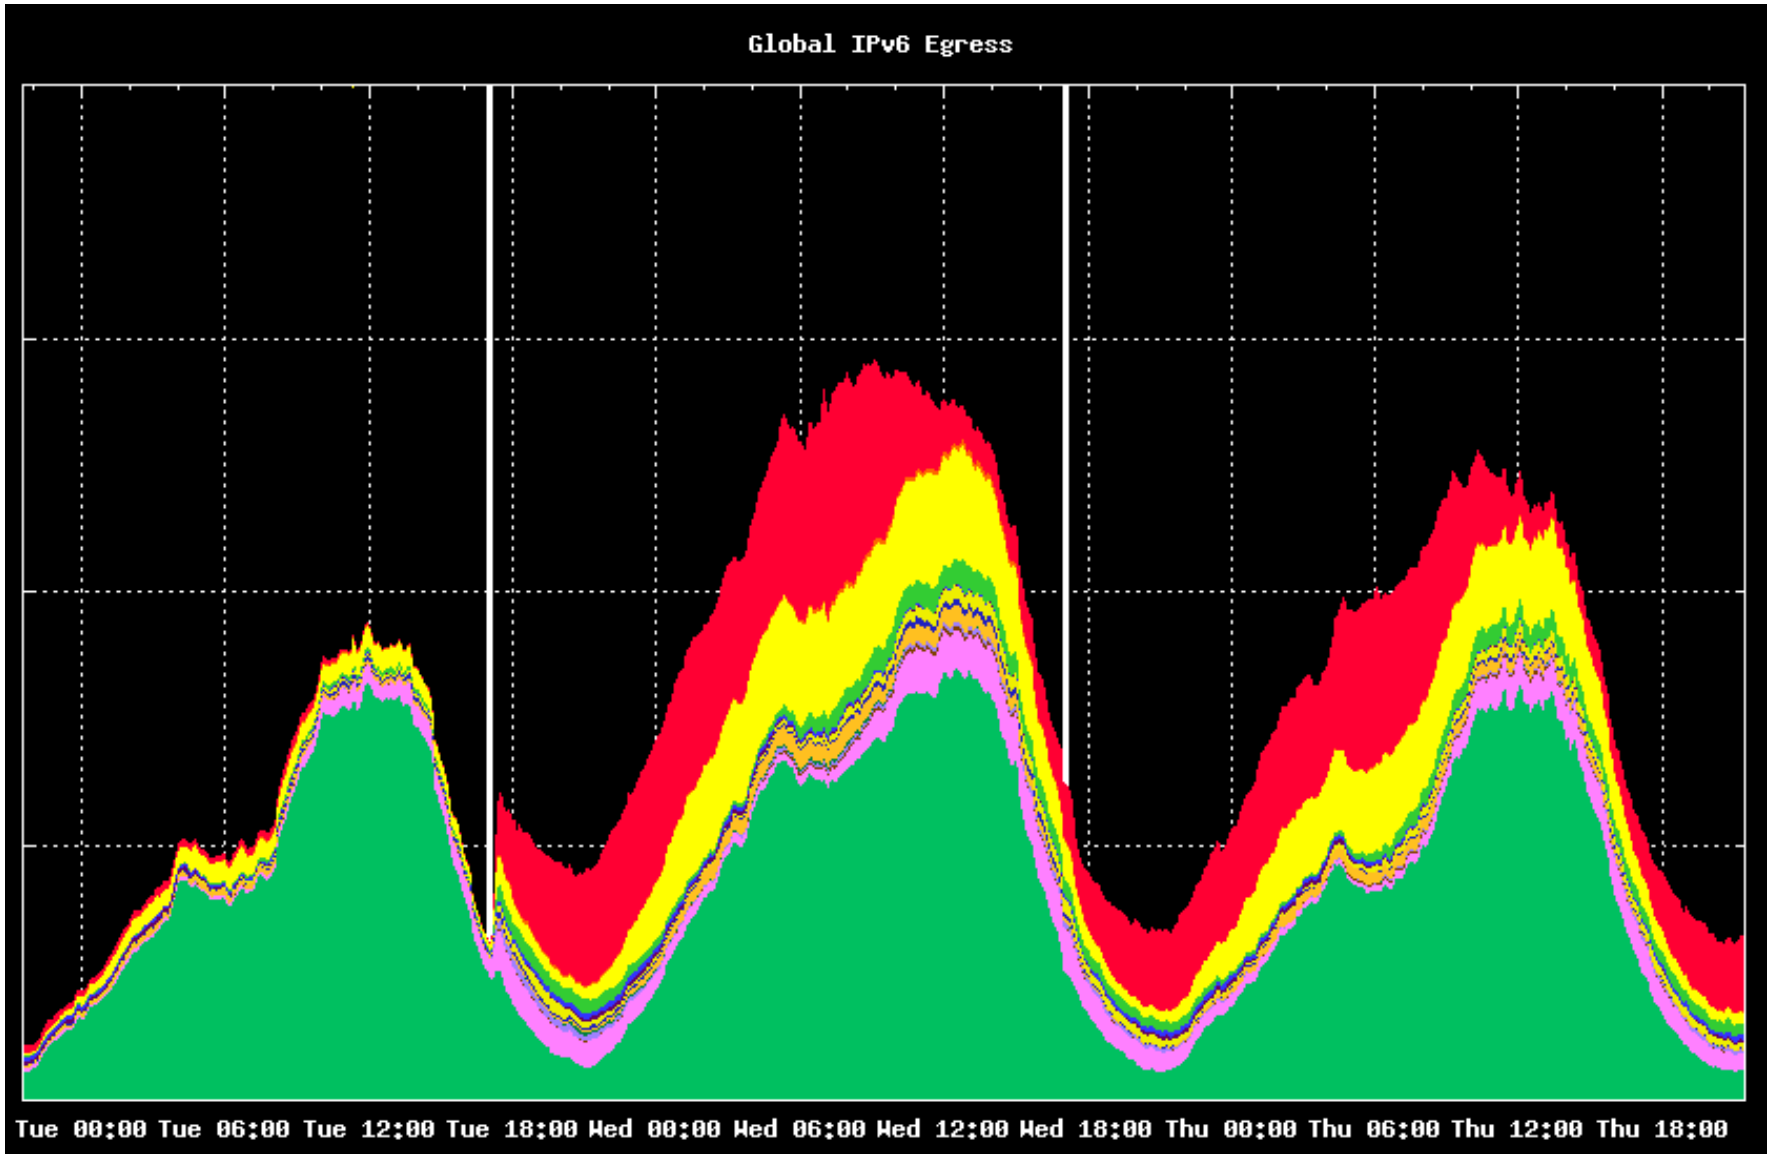

We were already serving 60% of the IPv6 Internet

Chrome IPv6 Brokenness

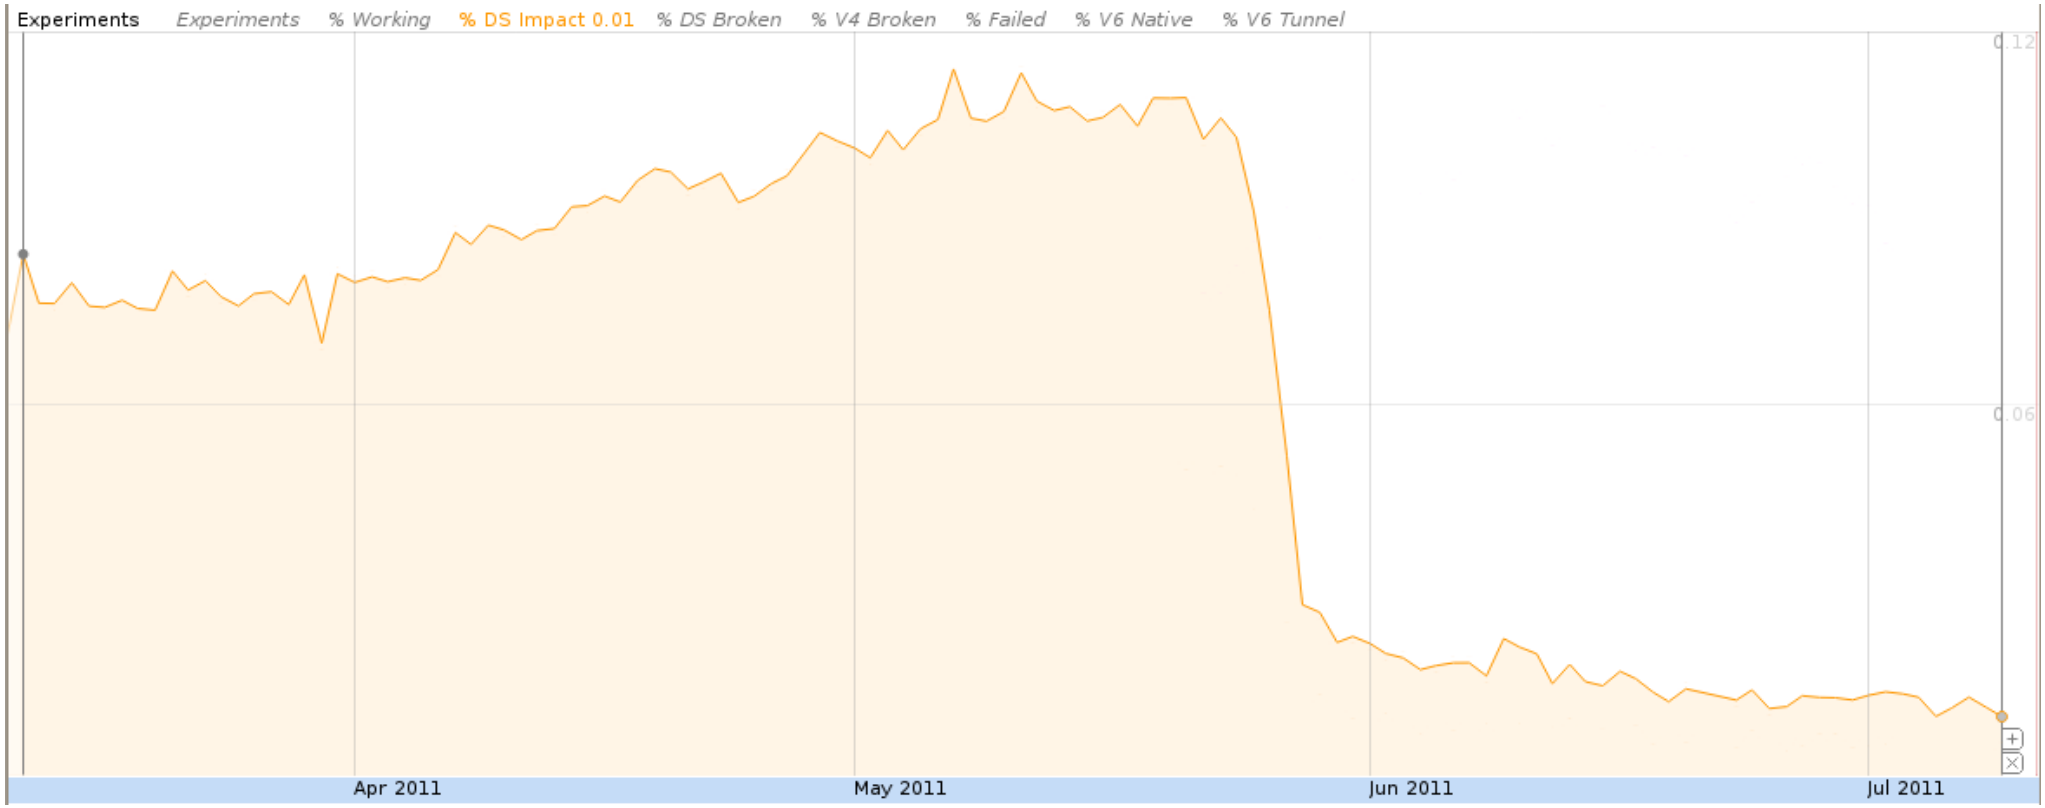

- 80-90% reduction in dual-stack brokenness
- Versions with fast-fallback 99.995% as reliable as IPv4
- Similar behavior will be in Firefox 7
- Apple also added robustness in OS X Lion
- If IE follows suit, users with recent browsers should be OK

Deploying IPv6

Do it for real

- Vendor claims and testbeds not enough
  - There **are** bugs lurking, and you need to find them
- The only way to find the gaps is to target real deployment
  - Same reliability as IPv4
  - Same performance as IPv4
  - Same operational complexity as IPv4
- Assume IPv6 needs to carry all traffic and start from there
  - Gather requirements and test it in the lab
  - Iterate with vendors until it works
    - Be prepared to hear "it's on the roadmap for 2013"
  - Deploy it!
    - If you don't deploy, you won't know it's really working

Guiding principles

- Make design as similar to IPv4 as possible
  - Principle of least surprise for NOC, other engineers, ...
- Dual stack everything
  - Scales better, no added maintenance / support load
  - Using IS-IS for IPv6? Might want to use it for IPv4
  - Using OSPFv3? Make sure implementation is proven
- Use IPv4 to carry IPv4 routes, IPv6 to carry IPv6 routes
  - Don't block convergence of one protocol on another
  - Avoid ::ffff:10.0.0.1 and ::10.0.0.1 as IPv6 next-hops

Before you start

- Check IPv6 feature parity for features you use
- Check you have enough resources to run IPv6
  - RAM, CPU: easy to check
  - CAM? Nexthops? ASIC cycles... ?
- Do you need to modify existing IPv4 configuration?
  - Do you need to change CAM profile?
  - Will you run out of CAM space? Have to rewrite ACLs?
- This is the easy part
- Produce a design and take it to the lab

Testing and iteration

- Implementations mostly work, but will have bugs
  - Nobody has really kicked the tires
- 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 a 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 in a fair number of datacenters

Operations

- Dispel notion that IPv6 is "experimental"
- IPv6 must be a production service
  - Monitored
  - Supported
  - Designed to the same quality standards as IPv4
- How to achieve this?
  - Make NOC aware of IPv6
  - Scale down, but don't skimp
  - Design as closely to IPv4 as possible
    - Make the principle of least surprise work for you

Google: case study

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

Development strategy

- Work from the outside, move in
- First the network, then the load-balancer, then the frontend, then the applications
  - But launch and iterate or nothing will ship
- "Address coercion" protects IPv4-only code from IPv6
  - Take IPv6 address
  - Remove user-modifiable bits
  - Hash into 224.0.0.0/3
- Sometimes not perfect
  - "Your last login was from 238.1.2.3"

Almost there

- 90% done, only 90% to go

Statistics

Adoption

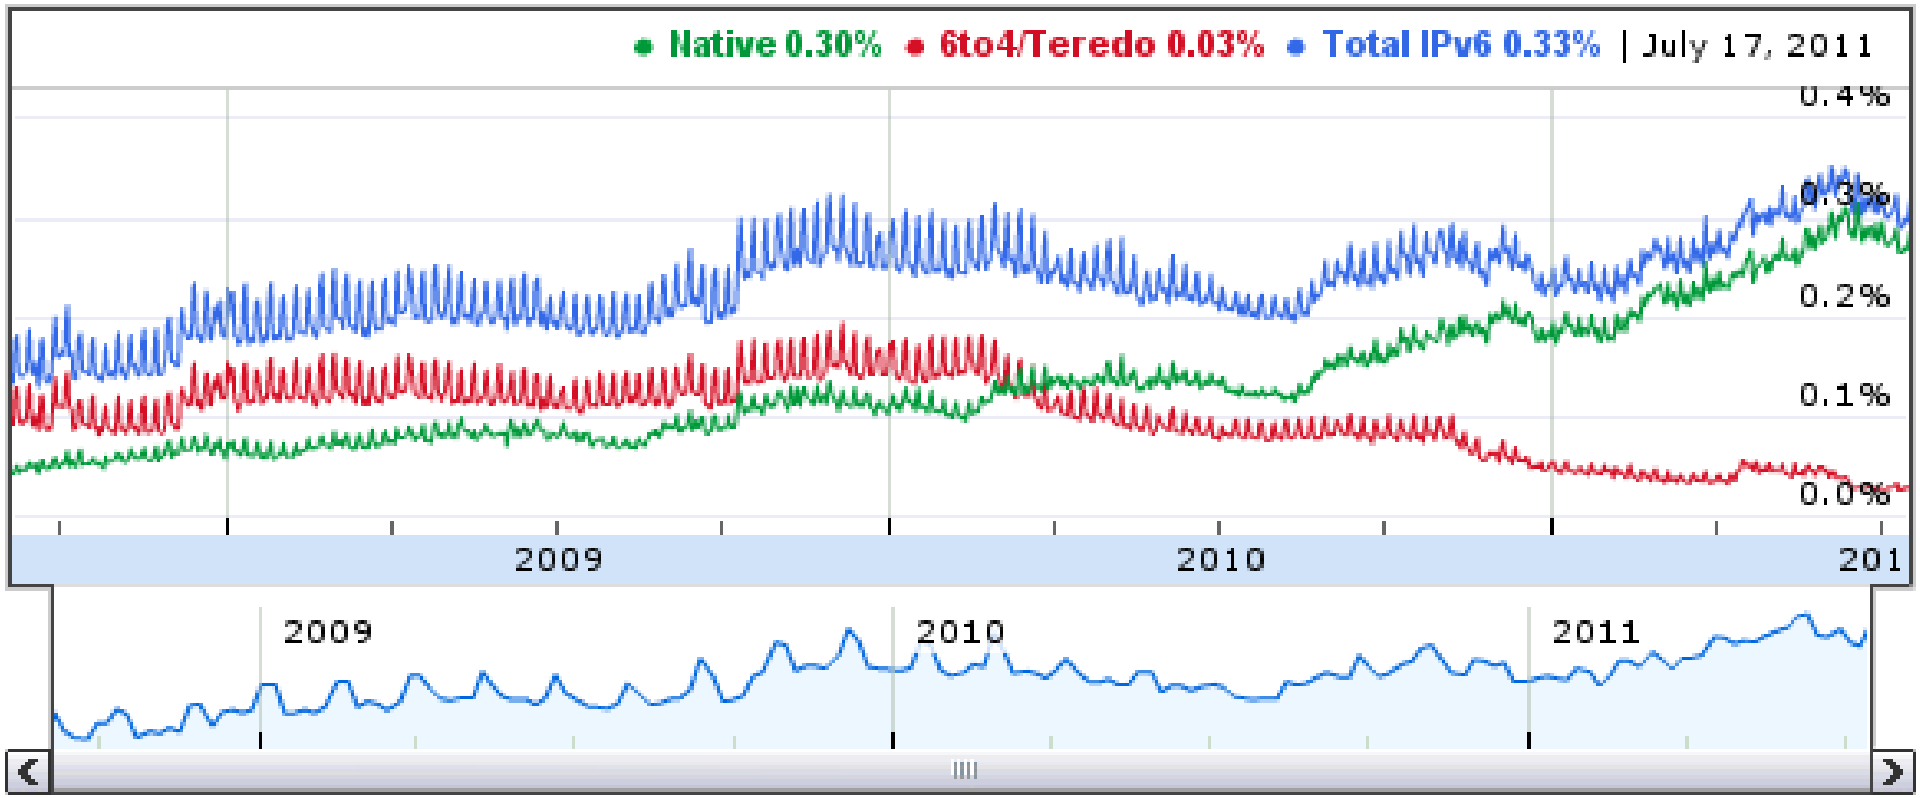

- Still 0.3%
  - France 3.4%, Japan 1.4% (since May)
  - Without France and Japan, 0.1%
  - Japan only joined the game recently

Traffic can appear overnight

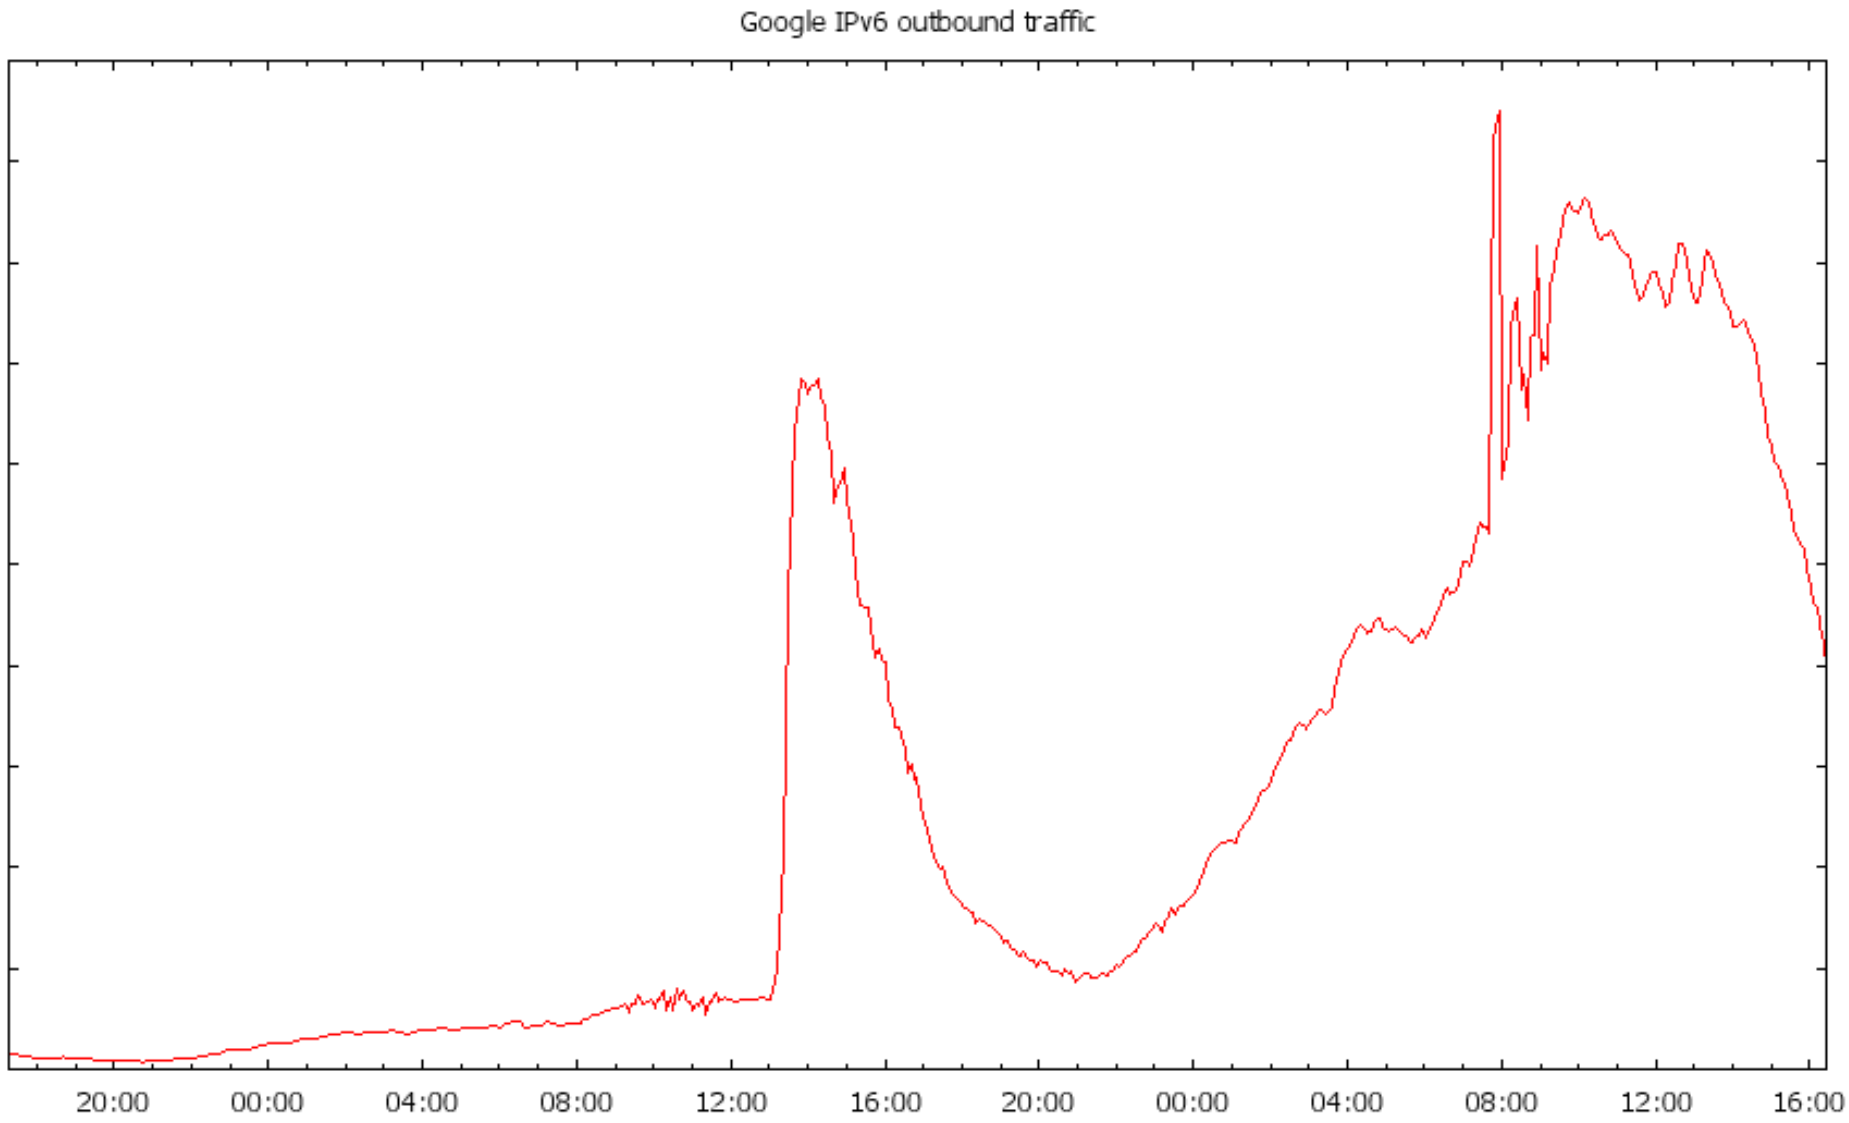

(IPv6 video launch 2010-01-28)

Questions?

Marcus Hines
hines@google.com