

IPv6 from the content perspective

Tore Anderson
CG Security and Networking
Redpill Linpro
IIS.se, Stockholm, June 2011

Introducing myself

- Working for Redpill Linpro in Oslo for the last 10 years
 - Before: UNIX sysadmin + jack of all trades
 - Now: Mainly IP/storage networking and data centres
- IPv6 became a professional hobby for me back in 2008
- These slides are available from: **<http://fud.no/talks>**
- Contact information:
 - tore.anderson@redpill-linpro.com
 - @toreanderson
 - +47 95 93 12 12


Introducing my employer

- Does pretty much anything that involves open-source software
- Offices all over the Nordic countries, customers world-wide
- **Managed Services** hosts and maintains customers' IT systems
 - Design and set up the customers' application stacks
 - Data centre hosting and internet connectivity in .SE and .NO
 - 24/7/365 server/OS/application maintenance and monitoring
 - Same SLA for IPv6 as for IPv4, of course
- **<http://www.redpill-linpro.se>** (or *.com .dk .no*)
- Owned by Sjöatta AP-fonden

Why bother with IPv6?

Good question.


Only approx. **1 out of 400** Norwegians have IPv6 connectivity
(And they all have IPv4 connectivity, too!)



(Connections using native IPv6 to A-pressen Digitale Medier and VG Nett, own data)

Very good question.

Only approx. **1 out of 400** Google users have IPv6 connectivity
(And they all have IPv4 connectivity, too!)




With thanks to Google




2 typical Norwegian internet users

Using mobile 3G broadband with GPS location tracking



With thanks to Digi.no

Happy IPv4 users




With thanks to Google Images





IPv4-only
web server
(owned by
a **happy**
content
provider)

Happy IPv4 users




With thanks to Google Images

Carrier Grade NAT



IPv4-only
web server
(owned by
a **unhappy**
content
provider)

Unhappy IPv4 users




With thanks to Google Images

Carrier Grade NAT

ISP's competing
web server

Happy IPv4 users



IPv4-only
web server
(owned by
a **very**
unhappy
content
provider)

With thanks to Google Images

- Performance impact, as CGNs are limited by session initiation rate
- Mobile carriers in Norway used to do CGN back in the WAP days
 - Then came the iPhone
 - The CGNs vanished overnight
- Geolocation will locate the CGN instead of the end user
- Abuse handling: which of the 100 people that are using 192.0.2.1 is the misbehaving one?
- And how do we blacklist him without causing collateral damage?
- **End-to-end IPv6 allows us to avoid all of the above problems**
- But we can't do it alone, the ISPs must do their part...



Photo by Fardin Waezi

*Extensive sharing? No problem, all payload will get there eventually!
(Picture from sales material of Honest Hank's Hardly Used MCs & CGNs.)*

Laying eggs

- For long, IPv6 suffered from a chicken & egg dilemma
- No longer, thanks to content providers big and small
- 400+ registered participants in World IPv6 Day
- Now we need IPv6-enabled end users to get the snowball rolling



venstre



a-pressen
digitale medier



XBOX 360




The Soap Trail

With thanks to all of the above

Dual-stack brokenness

IPv6 isn't entirely problem-free



- **Dual-stack brokenness** causes a bad user experience
- End-user's OS or web browser incorrectly thinks there's IPv6 connectivity
- IPv6 is tried first and fails; long timeouts before fail-over to IPv4
 - Live demo: <http://broken.redpill-linpro.se>
- Adding IPv6 results in a overall **less accessible** service than IPv4-only
 - Economic disincentive for content providers to deploy IPv6

Researching dual-stack brokenness

- In 2009 we enrolled two of our customers in an experiment
 - **VG Nett**
 - **A-pressen Digitale Medier**
- Purpose of the experiment was to:
 - Determine whether or not deploying IPv6 was (sufficiently) safe
 - Find out if there were any systemic failures
 - And if so, try to get them fixed
- Results openly shared with the ISP and content communities:
 - <http://fud.no/ipv6>

Measurement setup




- Invisible IFRAME embedded in customer's HTML templates
- Single stack IPv4 only
- IMG links in random order

- 1x1.png
- IPv4-only

- 1x1.png
- Dual-stack

Assumption: We should see the same amount of hits to the two 1x1 PNGs. If not, we're seeing brokenness.

Definition of «brokenness»




- The brokenness percentage is the spread, in percentage points, between the amount of successful hits to the IPv4-only PNG and to the dual-stacked PNG. In this example: **0.5%**.

Initial findings

- **0.2-0.3%** brokenness
- Certain sources of brokenness were standing out
 - Opera web browser, running on Windows
 - Mac OS X
 - Certain networks (enterprises, universities), some ISPs
- 70-80% of IPv6 traffic was «transitional IPv6»; 6to4 and Teredo
 - IPv6 tunnelled inside IPv4, so can't possibly be more reliable
 - Independent research from APNIC and RIPE NCC puts the lower bound on failures between **15 and 20%** (!!!!!)
 - There's no real reason to use either in preference to IPv4
- The results were considered by VG and APDM to be too broken, deployment therefore postponed until situation had improved
 - The complete lack of IPv6-enabled users didn't help either

Opera web browser on Windows


- Recent Windows will **automatically** enable 6to4 and/or Teredo
 - ..but de-prefers their use in the system resolver (RFC 3484)
- Opera, however, used its own built-in resolver



- Started nagging them about it
- Version 10.50, released the 22nd of March, fixed the problem
- Brokenness halved within a few weeks
- Also less 6to4/Teredo traffic

Mac OS X

- Mac OS X does not implement RFC 3484 and unconditionally preferred IPv6, including 6to4 and Teredo, above IPv4
- Does not automatically enable 6to4 but is duped by Rogue RAs




- Started nagging them about it
- Version 10.6.5, released 10th of November, de-prefers IPv6 completely if local 6to4 addresses are present
- No upgrade path for one-fourth of their users (running 10.4 and 10.5)

Rogue RAs


- Hosts that falsely announce themselves to the local network as IPv6 routers, often using the 6to4 prefix
 - Breaks dual-stack for all the old Mac OS X hosts on the LAN
 - Observed **10%** brokenness from certain campus networks
- Most common cause is Windows **Internet Connection Sharing**
 - Microsoft has not yet published a patch for this bug
- Routers that do 6to4 by default – championed by Microsoft
 - <http://msdn.microsoft.com/en-us/windows/hardware/gg463251.aspx#EZC>
- The IETF is about to deprecate 6to4 entirely – best to avoid it
 - <http://tools.ietf.org/html/draft-ietf-v6ops-6to4-to-historic>
 - <http://tools.ietf.org/html/draft-ietf-v6ops-6to4-advisory>

Production for VG and APDM

- In October we did a 24 hour production test, inspired by Heise.de
- Broken users were warned and sent to a test site which shows instructions on how to fix and/or get in touch with us for help
- The users didn't complain, but didn't really fix the problems either
- **AAAA records permanently deployed the 21th of December**




From the start until production



Brokenness over the last seven days before production: **0.024%**

Brokenness status right now



- 35% decrease in brokenness levels during the four months of production
- Known bugs in Opera and Firefox has been fixed
- We're expecting several more fixes (Windows, Mac OS X) in the coming months
- **World IPv6 Day** will hopefully also help out

(Dual-stack was turned off temporarily for a week in order to perform this new measurement)

World IPv6 Day

What is it?

- A copy of the IPv6 Day we had last October with VG and APDM
- Organised by non-profit ISOC
- Google, Yahoo, Bing, Facebook, Akamai, Limelight, as well as 400+ others participants, will deploy IPv6 access to their content simultaneously
- Takes place **right now!**
- Fate-sharing between competitors makes it easier to do
 - A broken user that experiences Google as being down will do so for Google's competitors Bing and Yahoo aswell
 - «The Internet doesn't work» instead of «Facebook doesn't work»
- <http://www.worldipv6day.org>
- Redpill Linpro participates as a facilitator for other participants
 - And been part of the «planning committee» from the start

So will the internet break?

- I bet today won't be much different from any other day
 - Every day so far in 2011 has been a «*Norwegian IPv6 Day*»
- My contacts tell me that so far it's been boring and anticlimactic
 - That's the best outcome possible!
 - Expect sites to stay IPv6-enabled in the future



[@auduny](#)
Audun Ytterdal

Vi har ikke fått inn en eneste klage på
trøbbel med [#ipv6](#) dualstack på [#vgnett](#)
ennå. Er det ingen der ute med problemer?
[#antiklimaks](#)

26 Oct via [TweetDeck](#) [☆ Favorite](#) [↻ Undo Retweet](#) [↩ Reply](#)

For the very few with problems

- <http://test-ipv6.com> , <http://ipv6test.google.com>
- Ensure browser and operating system is patched, especially:
 - Opera 11.10
 - Firefox 4.0
 - Mac OS X 10.6.5
 - Firmware of Apple AirPort/TimeCapsule to 7.5.2
 - Firmware of Cisco Linksys E-series to 1.0.04
- Windows users should disable tunneled IPv6
 - <http://go.microsoft.com/?linkid=9732984> (FixIt #50433)
- Windows users should also disable Internet Connection Sharing
 - Note: The ICS bug affects **other** hosts on the LAN, esp. Macs

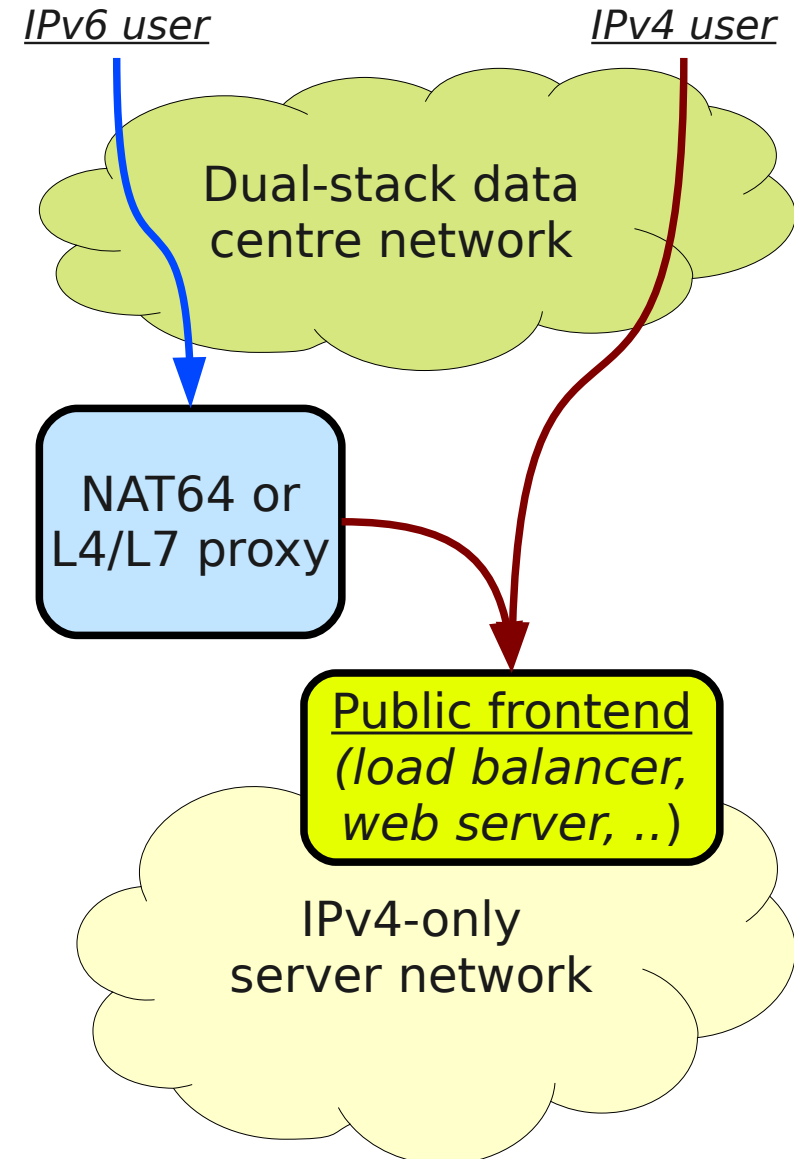
Still problems?

- The latest version of the Google Chrome browser glosses over brokenness by initiating IPv4 and IPv6 connections almost simultaneously, and picking the one that succeeds first
- There's a few other but less frequent problems too. I've documented all that I'm aware of in ARIN's IPv6 Wiki:
 - http://getipv6.info/index.php/Customer_problems_that_could_occur
- If all fails, simply disable IPv6 outright... :-(
 - Firefox: `about:config` -> `network.dns.disableIPv6` -> `True`
 - Windows: <http://go.microsoft.com/?linkid=9732985> (FixIt #50444)
 - Mac OS X: System Preferences -> Network -> Advanced -> TCP/IP -> Configure IPv6 -> Off
- Do let me know if you happen to come across a new and unknown bug!

Content deployment scenarios

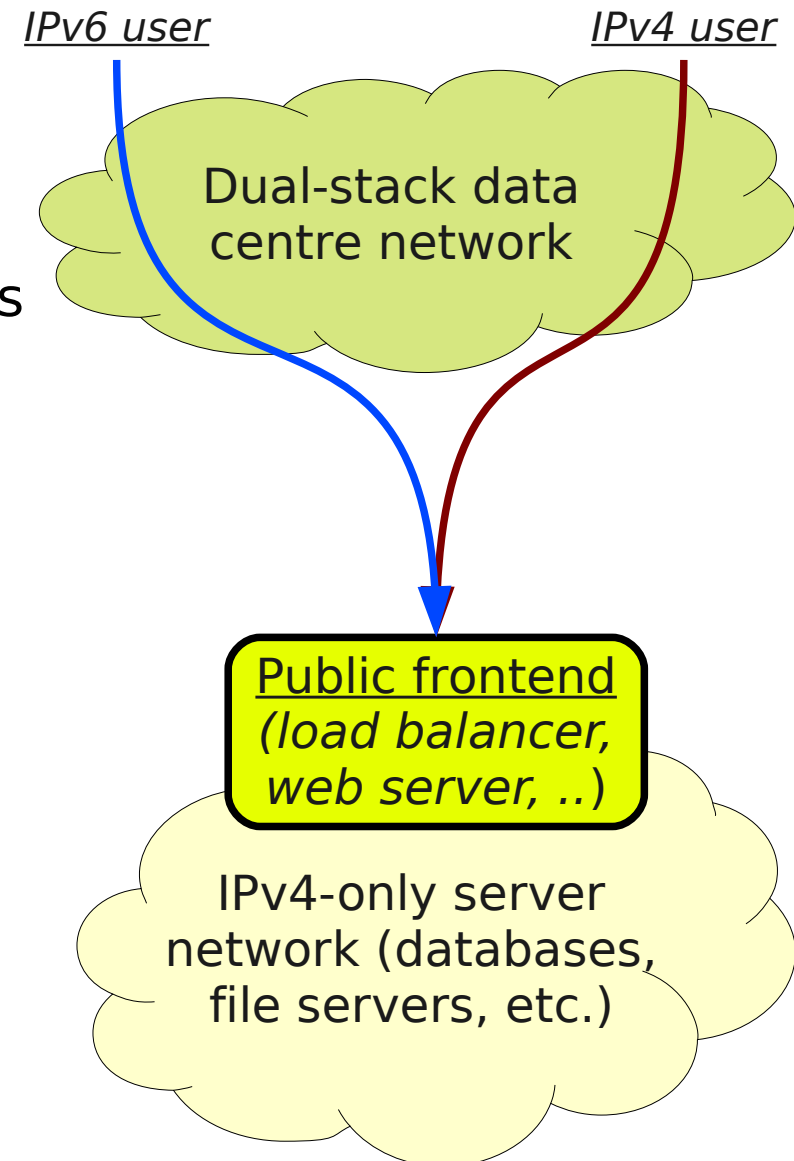
IPv4 + IPv6 through NAT64 or proxy

- IPv6 is routed to a system that translates incoming IPv6 traffic to IPv4
 - NAT64, Varnish, HAProxy, Nginx, ..
- Advantages
 - Simple to retrofit to existing installations
 - Translator function may be outsourced
- Disadvantages
 - Possible performance bottleneck
 - Loss of insight into client source address
 - L7 headers (e.g. **X-Forwarded-For**) might take care of that
 - No exit path from IPv4
- Høyre gets IPv6-as-a-service from us; their web servers are hosted by another provider



Dual-stacked frontend

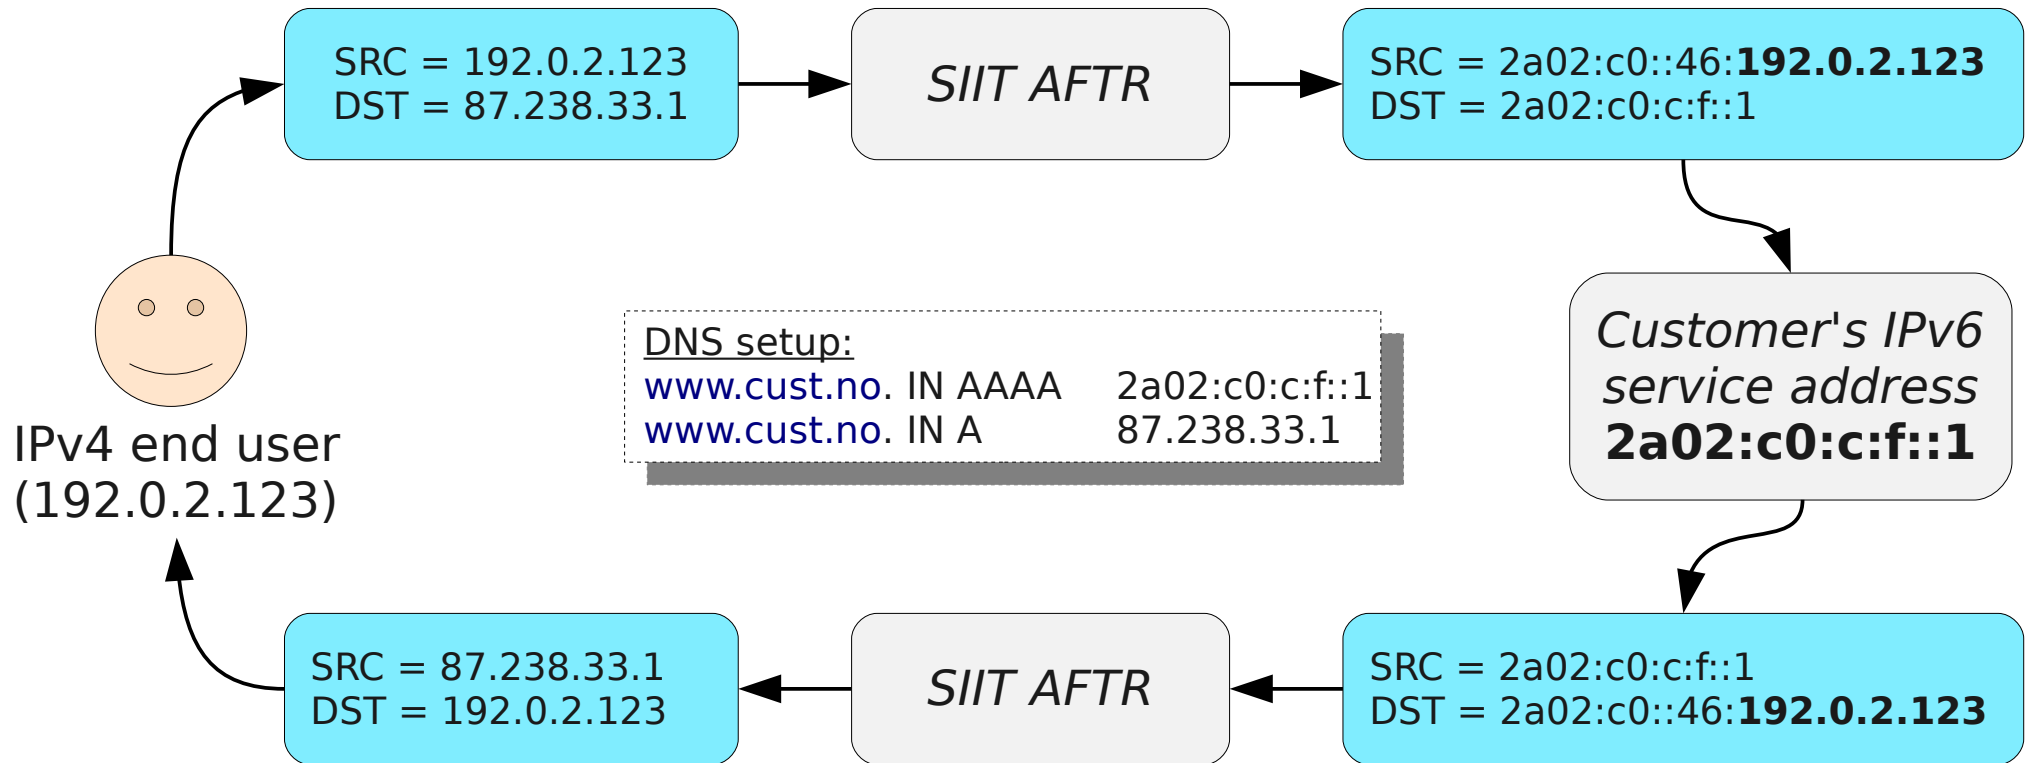
- Native dual-stack provisioned to the system's frontends
- Advantages
 - Simple to retrofit to existing installations
 - VG/APDM setup time: less than a day
- Disadvantages
 - IPv6 support required on the public frontend devices/software
 - Operational overhead due to dual-stack (monitoring, troubleshooting, firewall filters, etc.)
- Most sites we host use dual-stack for now
 - VG: F5 BIG-IP load balancer in front
 - APDM: Varnish & HAProxy in front



Or turn off IPv4 completely? (my favourite!)

- Dual-stack equals operational overhead
 - Twice the amount of ACLs to configure
 - Twice the amount of services to monitor
 - Twice the amount of OSPF adjacencies to maintain
 - RFC 5838 will solve this eventually though
 - Twice the amount of routes to carry in your IGP
- More things that can go wrong and disrupt service
- And I simply don't believe the «servers must remain dual-stacked for the next 10 or 20 years» mantra
 - IPv4-only end users can't be cut off quite yet, though

Stateless IP/ICMP translation (RFC 6145 ++)



- The Address Family Translation Router(s) does the following:
 - 1) it rewrites the IP destination field according to a static mapping (this is the ++ part)
 - 2) it rewrites the client's IP source field by prepending a predefined 96-bit IPv6 prefix
 - 3) the above two translations in reverse (in the server's outbound direction)
- The end user is not aware that the connection was translated
- The server can determine the IPv4 source address (if it wants to)

Advantages of IPv6-only + SIIT

- Minimal operational overhead compared to dual-stack operation
 - Translator(s) can be centrally located in the core data centre network
- Stateless per-packet operation, essentially no performance impact
 - Load balancing can be achieved with simple multipath routes
- The original IPv4 client address remains known to the application
 - Useful for Geo-location, ACLs, access logs, etc.
- Huge IPv4 address savings
 - One IPv4 address per **service** instead of one per server
 - Avoids unused addresses in a server LAN prefix – **100%** utilisation
- Forward-looking approach – why build services on a legacy foundation?
- I need a willing customer to partner with me in order to try
 - This has AFAIK never been done before, so we'd be pioneers :-)
- Working with Cisco to get the static mapping functionality into IOS-XE


Finishing words

Next steps for us

- Continue retro-fitting IPv6 capability to existing customers
- Hopefully do a pioneering IPv6-only deployment with a willing customer
 - And share the experiences with the community, of course!
- Identify remaining IPv4-only systems in our infrastructure and fix those
- Continue working to raise IPv6 awareness and accelerate deployment
- Wait for IPv6-enabled end users to appear...
 - There's still no compelling reason for content folks to enable IPv6
 - But there is no direct reason **not to**, either!
 - Hopefully the rest of the world will realise that today
 - I want this to be my last «IPv6 brokenness» presentation



This graph shows the percentage of networks (ASes) that announce an IPv6 prefix for a specified list of countries or groups of countries



- ISPs, large enterprises, and similar with IPv6 enabled in their core network
- .SE/.FI/.DK looks okay; better than most
- .NO/.NL is in a class of their own
- I challenge you to try to catch up with us :-)

With thanks to Emile Aben/RIPE NCC



Questions?