

Measuring and combating IPv6 brokenness

Tore.Anderson@redpill-linpro.com
RIPE 61, Roma, November 2010
<http://fud.no/ipv6/>
@toreanderson

Introduction

- Content-heavy network operating in the Nordics
- Our customers are the actual content providers, e.g.:
 - **VG Multimedia**: Norway's largest web site
 - **A-pressen Digitale Medier**: ~70 regional news sites, in sum Norway's 4th largest web site
- IPv6 is the future so let's try to deploy it
 - But we were worried about Google's reports of brokenness (e.g. at RIPE 56)
- What's the cause of the problems? Google didn't say
- We decided to find out

Measurement setup



Tip: Want to try this on your own site? Check out Éric Vyncke's <http://www.vyncke.org/testv6/> !

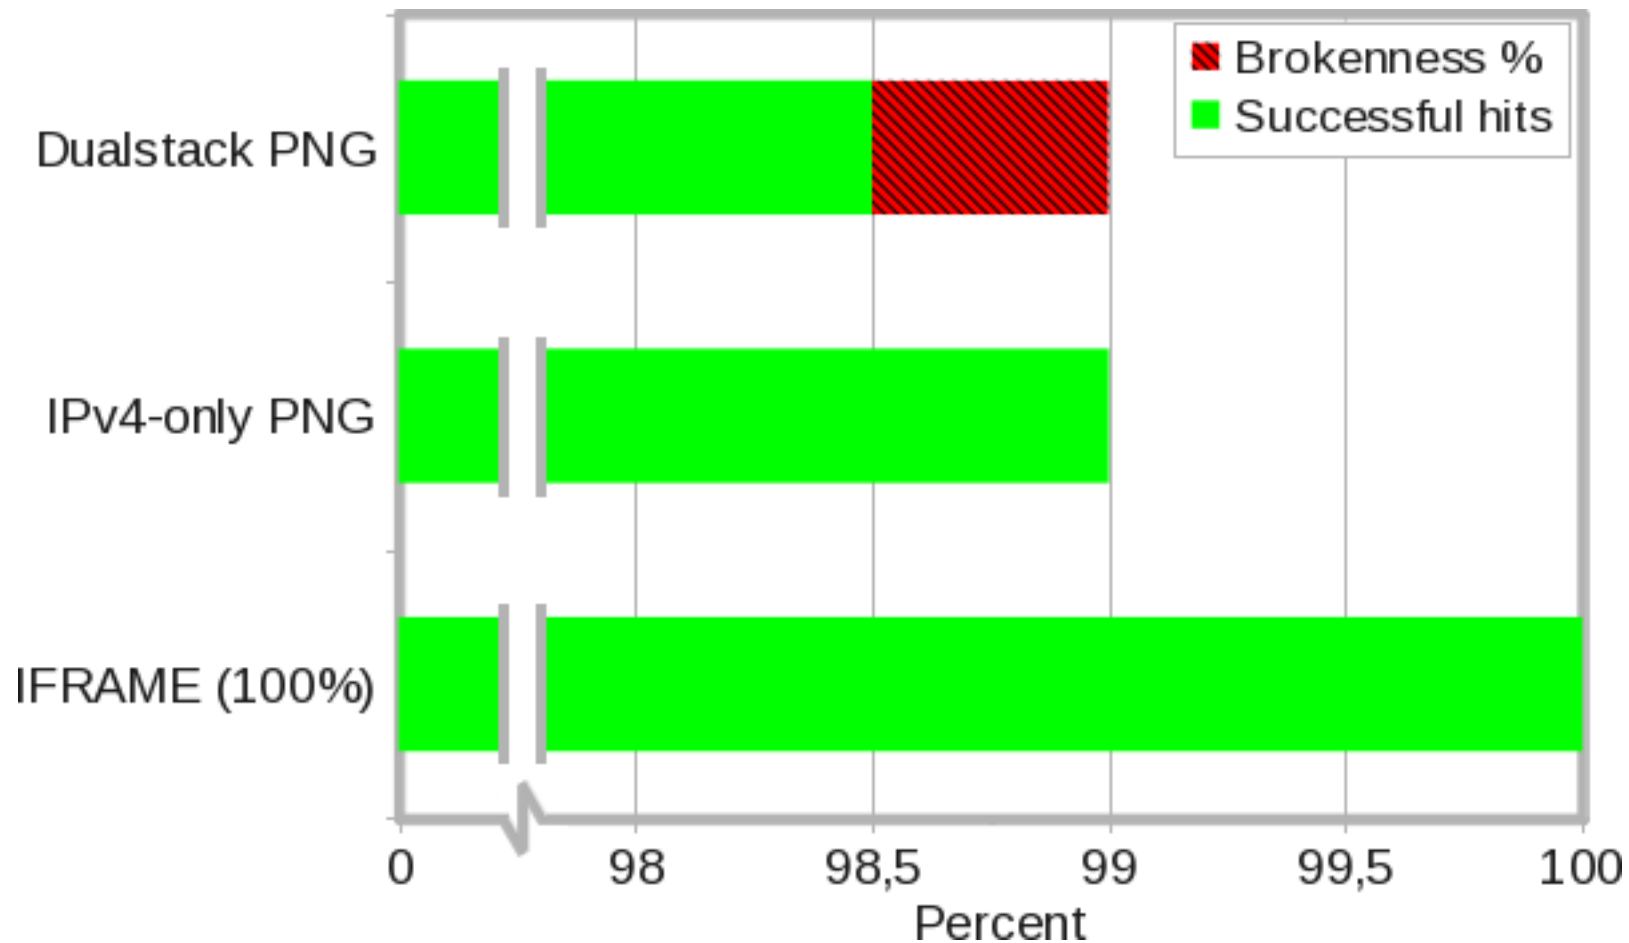
- Invisible IFRAME embedded in customer's HTML templates
- Single stack IPv4 only
- IMG links in random order

- 1x1.png
- IPv4-only

- 1x1.png
- Dual-stack

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

Definition of «brokenness»



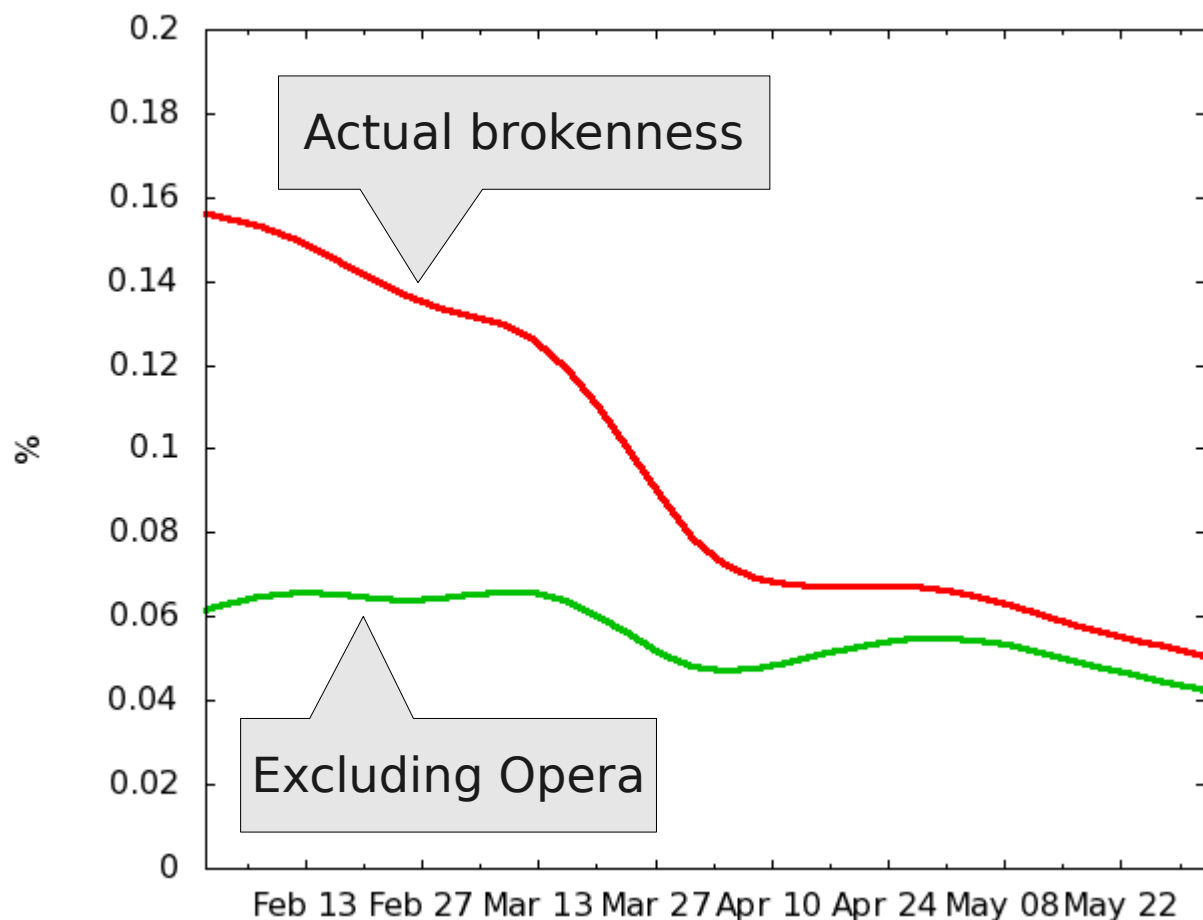
- The brokenness percentage is the spread, in percentage points, between the amount of successful hits to the IPv4-only PNG and to the dual-stacked PNG. In this example: **0.5%**.

Initial findings – Q4 2009

- 0,2-0,3% brokenness
 - ...a complete non-starter
- Certain sources of brokenness were standing out
 - Opera web browser on Windows
 - Mac OS X
 - Some networks (enterprises, universities), ISPs
- 70-80% of IPv6 traffic was 6to4 and Teredo
 - ..which runs on top on IPv4, so can't possibly be more reliable
 - There's no real reason to use either in preference to IPv4


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 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”

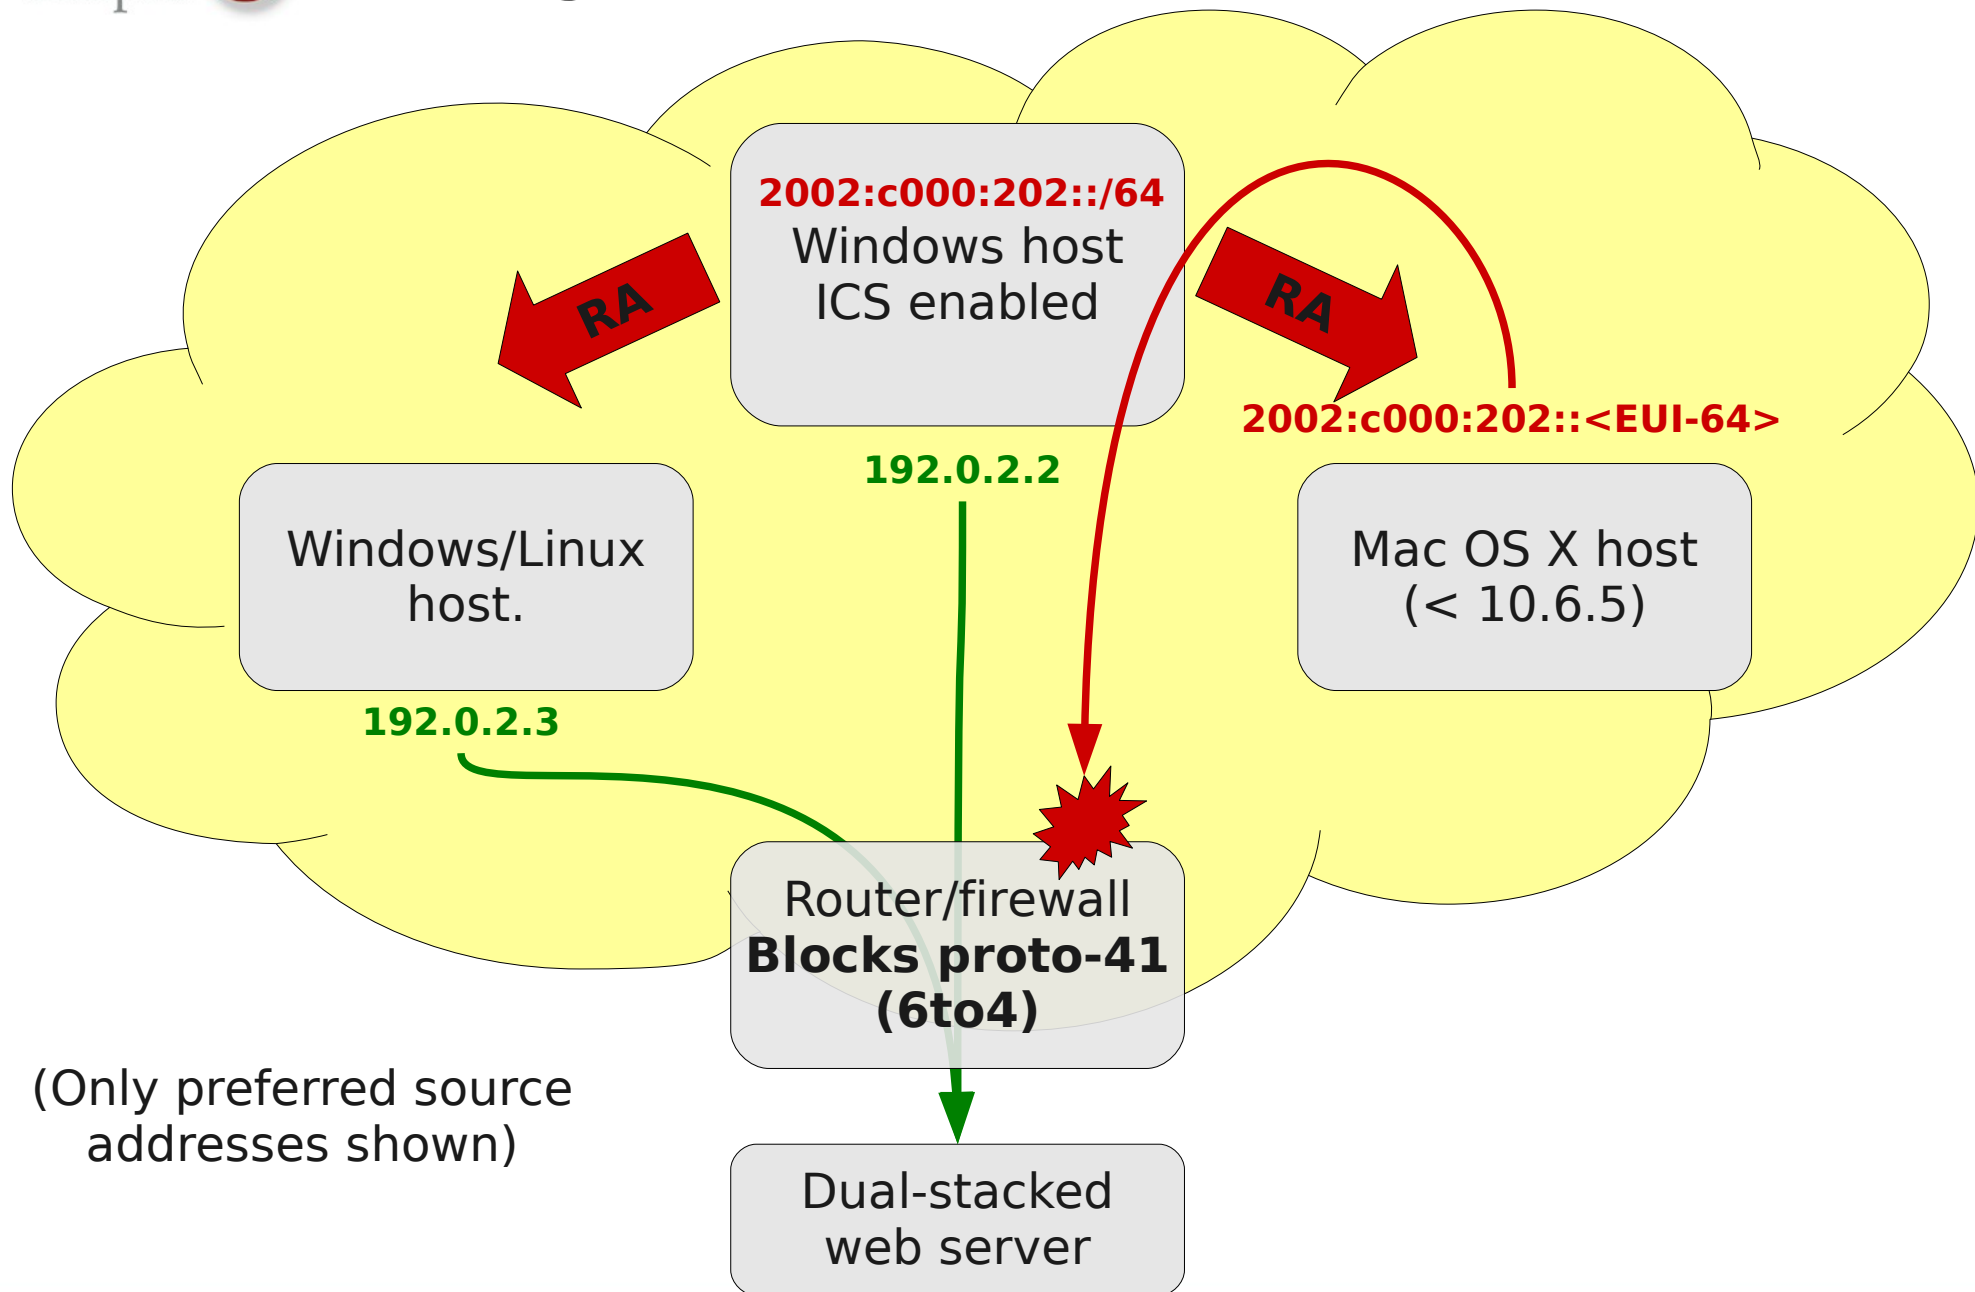


- Started nagging them about it
- Version 10.6.5, released 10th of November, de-prefers IPv6 completely if local 6to4 addresses are present
- No upgrade path for one-third of their users (running 10.4 and 10.5)

6to4-hostile networks and «Rogue RAs»

- Some networks block 6to4 traffic
 - ...but users will (unknowingly) try to use it anyway
- Unblocking it isn't always an option due to security concerns
- Rogue RA refers to a host announcing itself to the local network as an IPv6 router, typically using the 6to4 prefix
 - Most commonly Windows w/Internet Connection Sharing
 - Breaks dual-stack for all the Mac OS X hosts on the LAN
 - Brokenness of **10%** observed from certain networks
- Some workarounds exist
 - RA Guard, L2 ACLs, rfixd, more-specific route to 2002::/3, ...
 - Not applicable to every type of network deployment

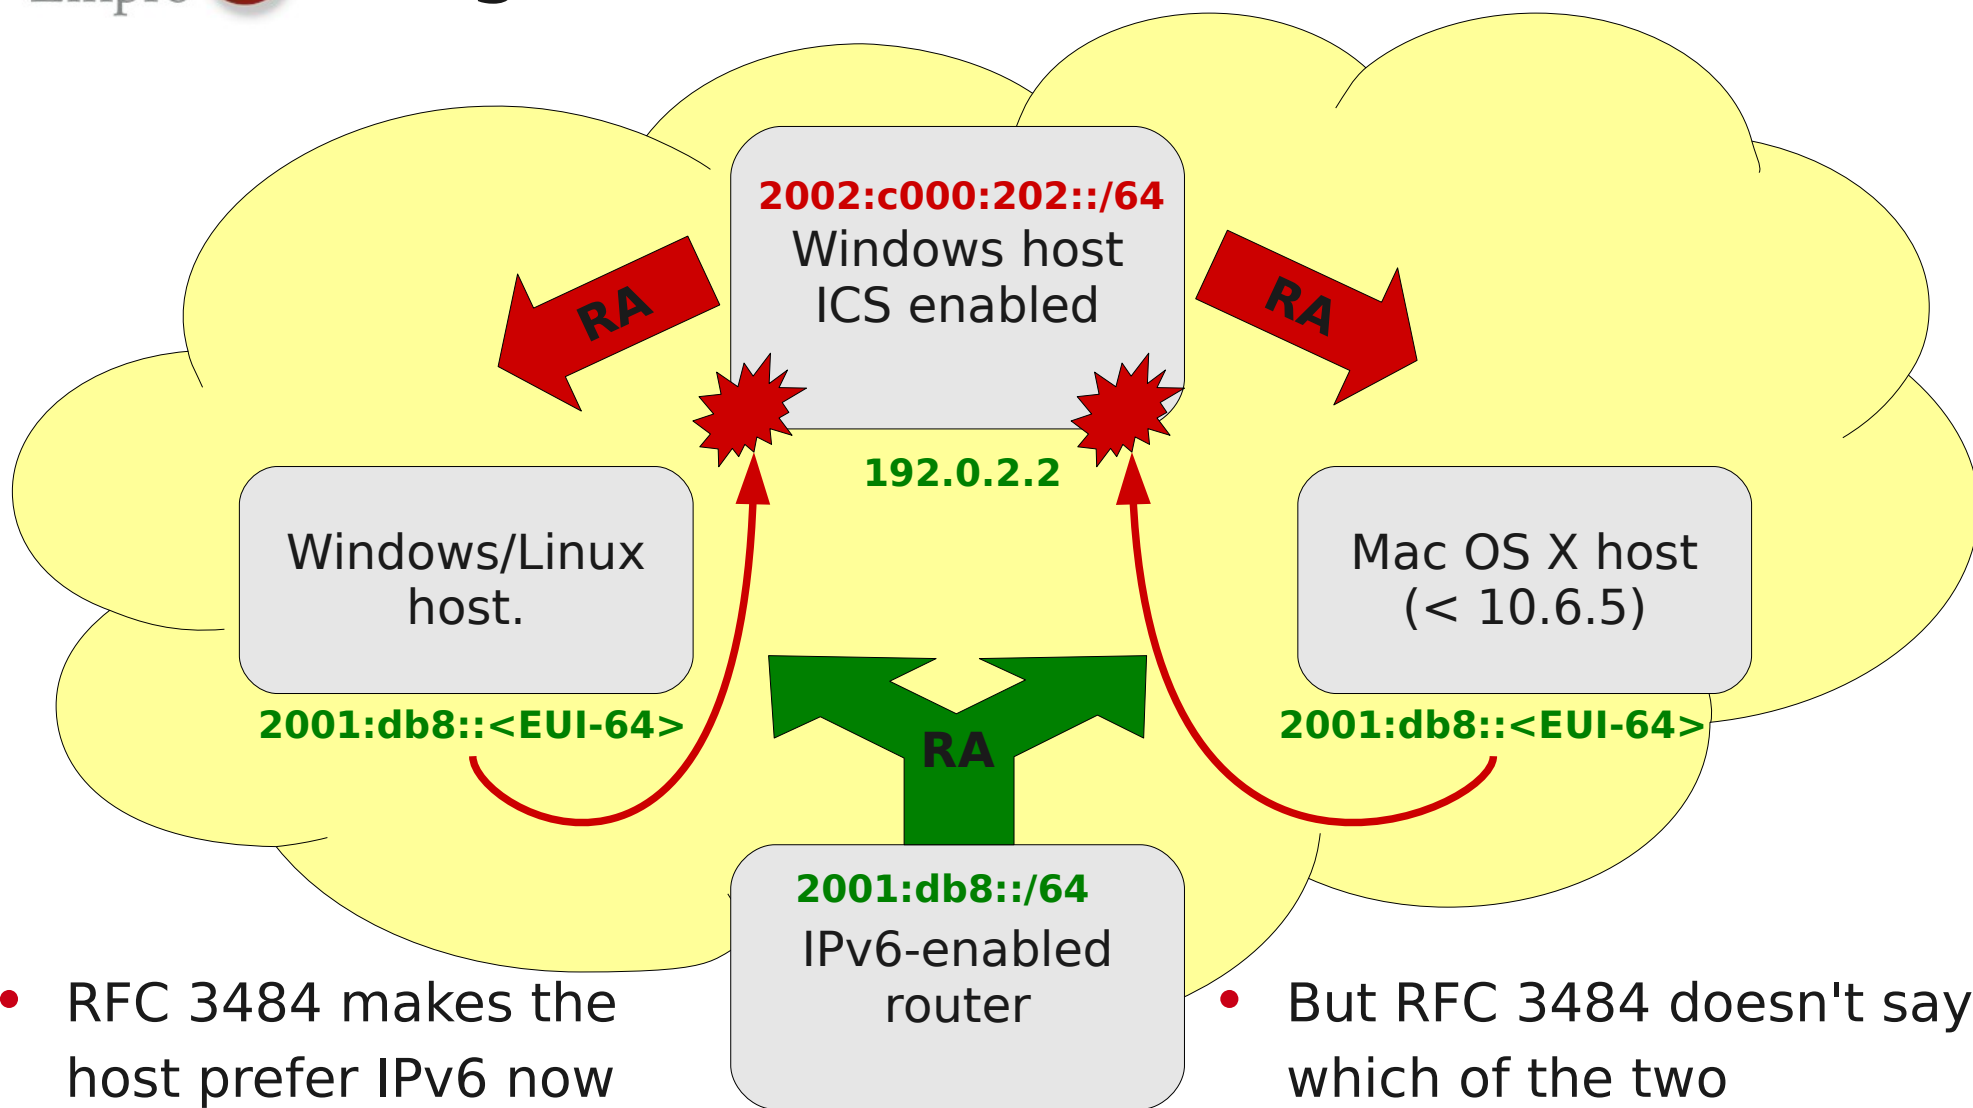
Rogue RA-infested network



Rogue RAs, continued

- Windows ICS will emit RAs for its local 6to4 prefix even if the shared interface (the “WAN” interface) isn't connected or active
 - Microsoft is looking into it, no word of any fix yet
 - Native IPv6 will stop Win 7 from going rogue, but not Vista
- 6to4+RA is also implemented in several CPE/HGWs
 - Microsoft has encouraged vendors to do so
(http://www.microsoft.com/whdc/device/network/ipv6_igd.msp)
- Deploying native IPv6 on such networks is **counter-productive**
 - Makes Windows and Linux randomly get in trouble too, probably due to lack of next-hop prioritisation
 - **20%** brokenness
 - Setting router priority to high has been ineffective

Rogue RA-infested network #2

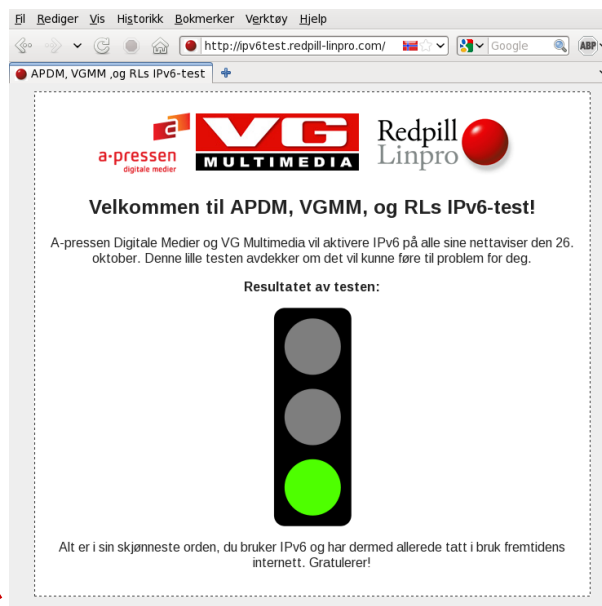


- RFC 3484 makes the host prefer IPv6 now that it has a native IPv6 source address

- But RFC 3484 doesn't say which of the two available default routes the host should use!

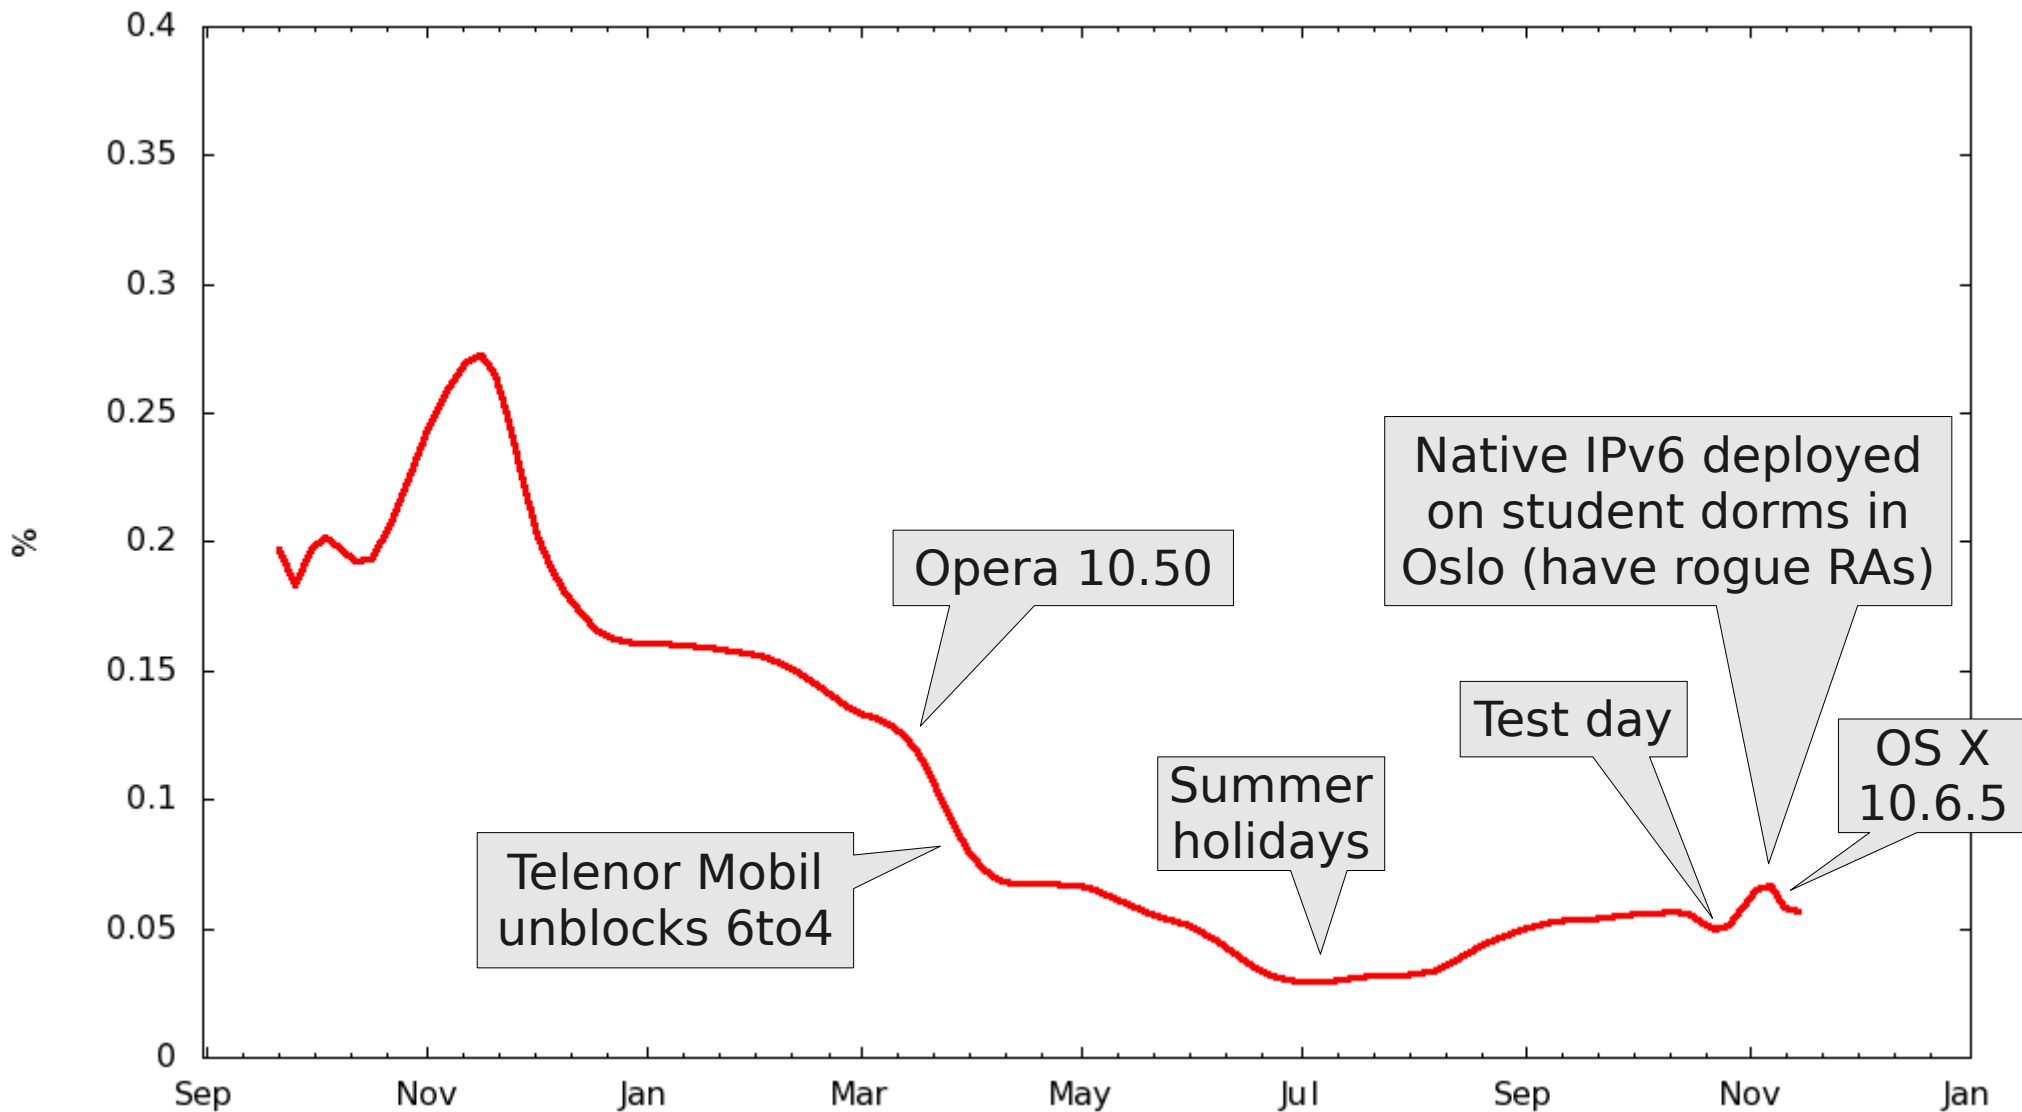
Production?

- We did a 24 hour production test, inspired by Heise.de
- Broken users are warned and redirected 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
- APDM and VGMM are no longer afraid of publishing quad As



- When we're in the long tail of upgrades to Mac OS X 10.6.5, **we will deploy**
- Which will end my measurements, sorry
- But they've served their purpose, I think

From the start and up until today



Brokenness over the last seven days: **0.058%**

Wrapping it up

- Many thanks to:
 - **Steinar H. Gunderson** from Google for helping out tremendously all along (his 20% project must have been me :-)
 - **Simen Graff Jensen** from **A-pressen Digitale Medier** and **Audun Ytterdal** from **VG Multimedia** for allowing me to use their readers as guinea pigs
 - The **ICSI Netalyzr** crew for providing an awesome tool for getting great debugging information out of non-technical users
 - And of course, everyone that fixed their software or networks: Apple, Debian, Fedora, Gentoo, Mandriva, openSUSE, Opera, Telenor Mobil, Ubuntu, UNINETT, University of Oslo, and many many more

Questions?