Skip the Transitions, Jump Straight into IPv6

Ivan Pepelnjak (@ioshints, ip@ioshints.info)
NIL Data Communications

Presentation @ 7. Slovenian IPv6 Summit organized by go6.si

Who is Ivan Pepelnjak (@ioshints)

- Networking engineer since 1985
- Consultant, blogger (blog.ioshints.info), book and webinar author
- Currently teaching “Scalable Web Application Design” at University of Ljubljana

Focus:

- Large-scale data centers and network virtualization
- Networking solutions for cloud computing
- Scalable application design
- Core IP routing/MPLS, IPv6, VPN

IPv6 Webinars on ipSpace.net

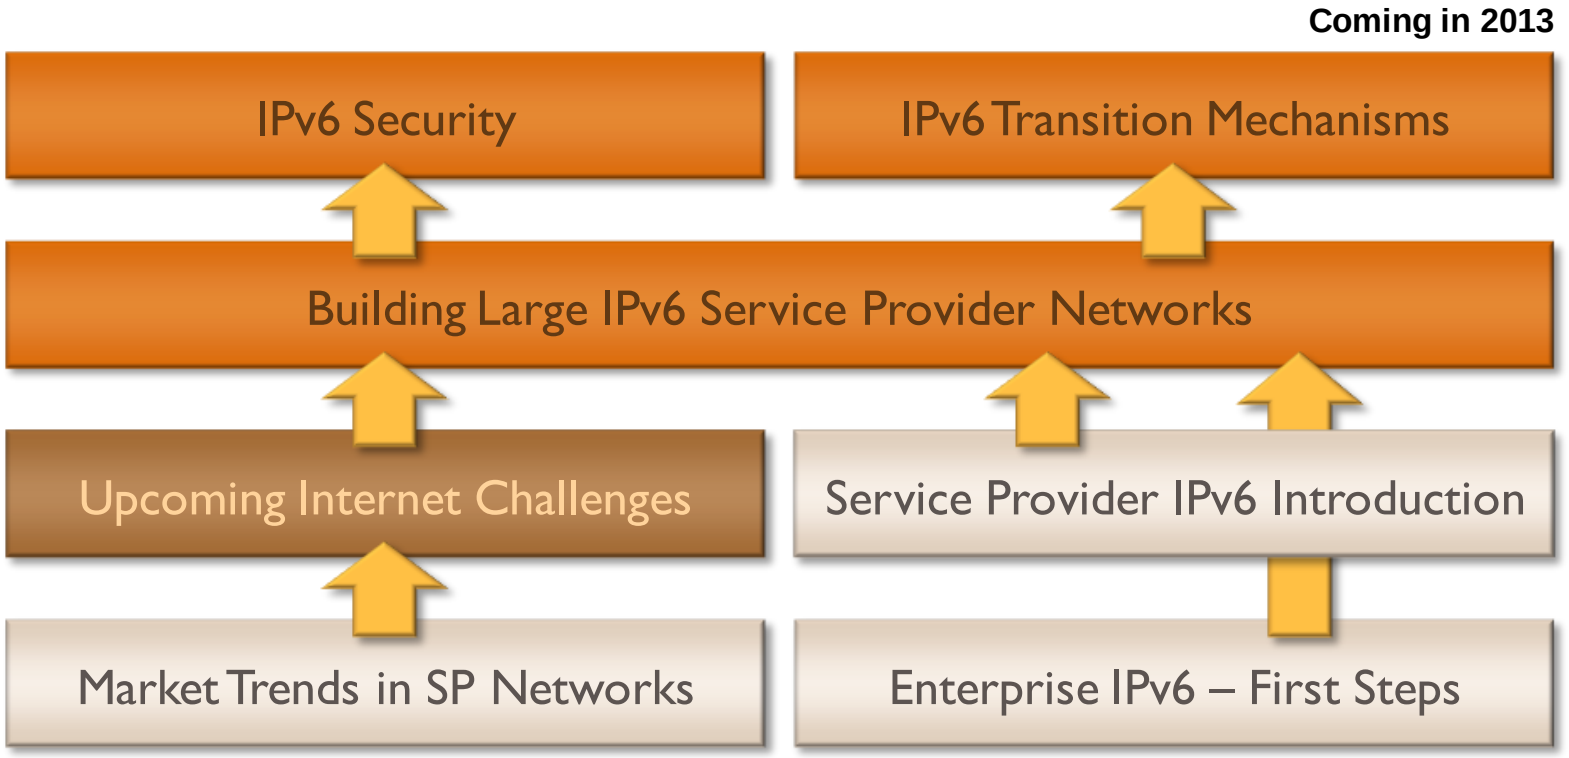

Availability

- Live sessions
- Recordings of individual webinars
- Yearly subscription

Other options

- Customized webinars
- ExpertExpress
- On-site workshops

More information @ http://www.ipSpace.net/Webinars

Past Predictions

- We'll run out of IPv4 addresses ✓
- IPv6-only mobile devices ✗
- Majority of the content will be on IPv6 ✓
- CGN will be expensive and thus avoided and/or neglected 

Special thanks to skype™

You must take control of your content

Content/Clients Dichotomy

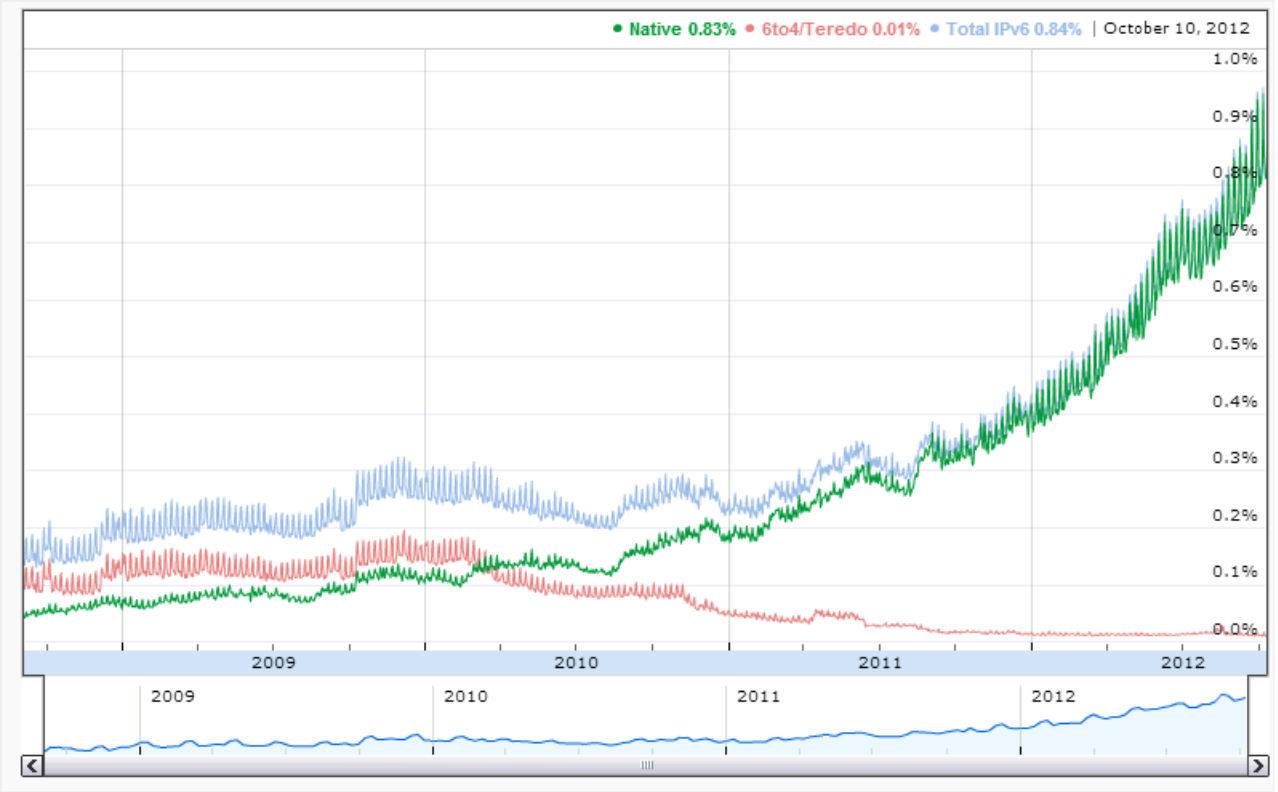

Source: google.com/ipv6/statistics

Content/Clients Dichotomy

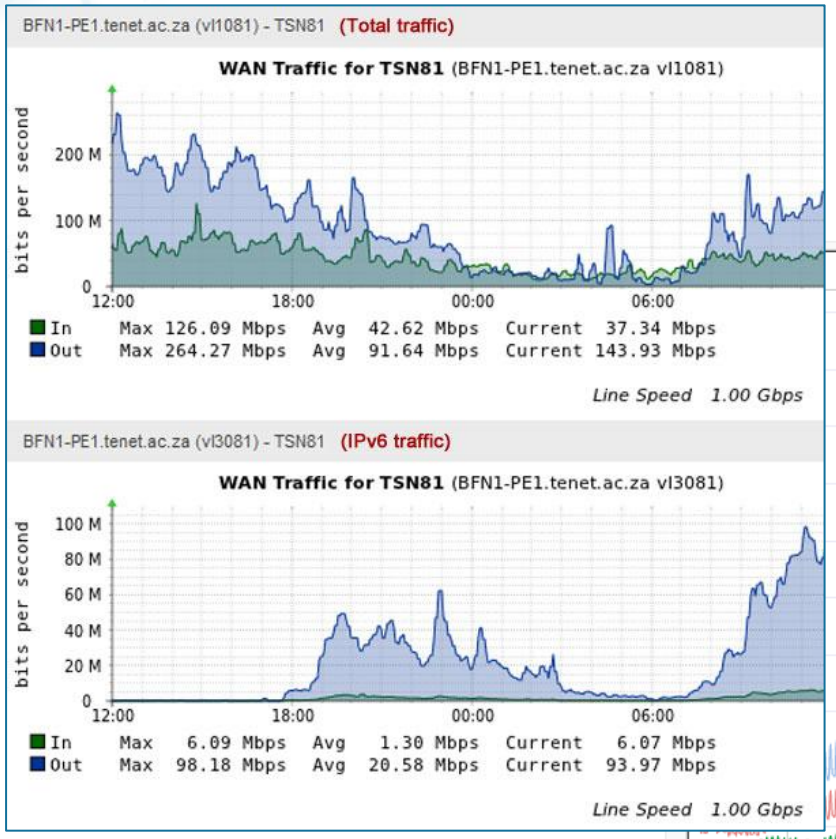

Source: http://mybroadband.co.za/news/internet/56241-shocking-ipv6-revelation-in-south-africa.html

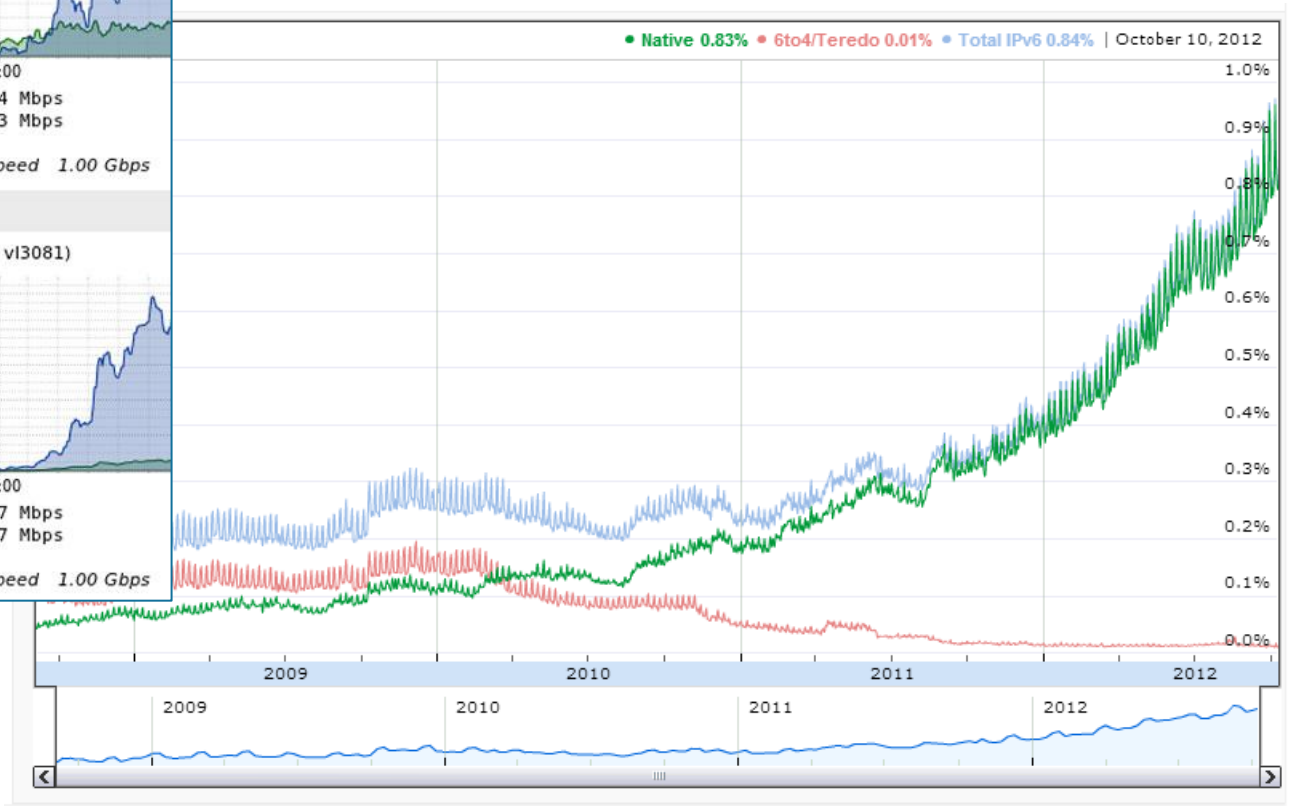

Source: google.com/ipv6/statistics

IPv6-Enabling a “Typical” Application Stack

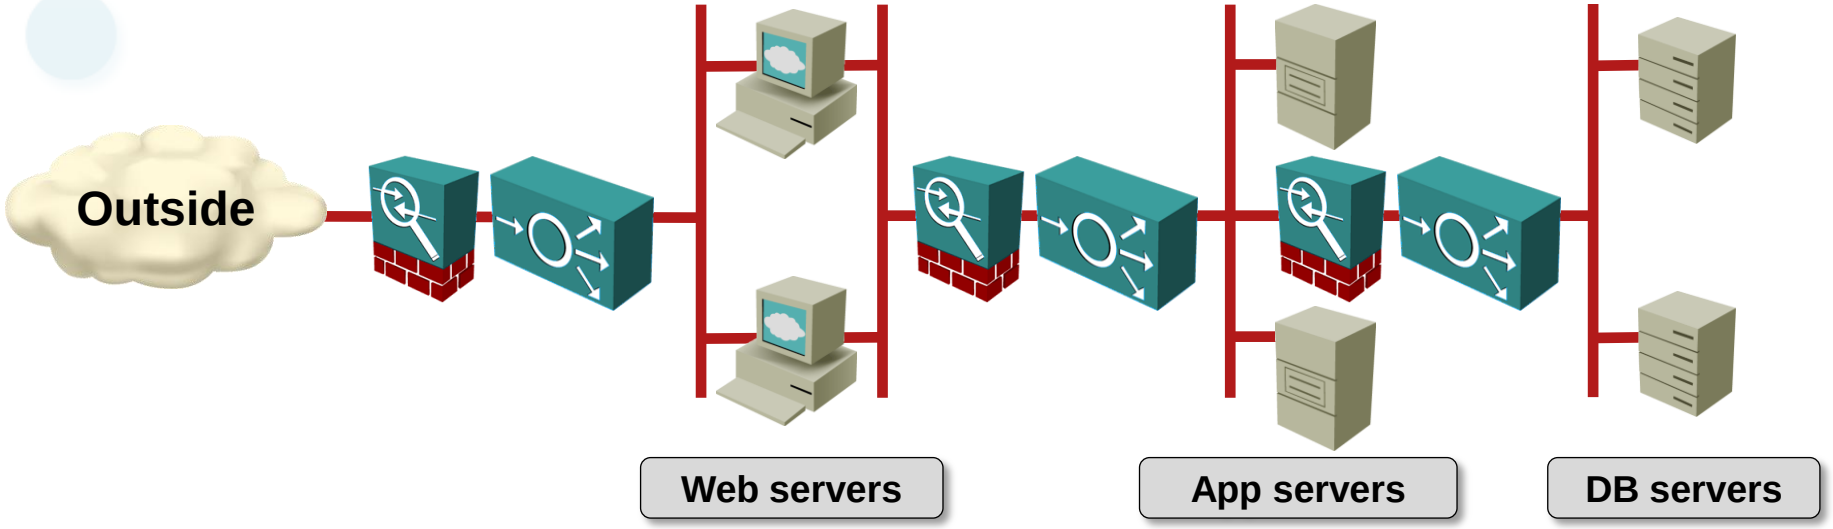

Typical “reasoning”

- Someone high enough asked us to make content available on IPv6
- We don’t know a thing about this new protocol
- Deploying IPv6 on load balancers or firewalls is too risky

IPv6-Enabling a "Typical" Application Stack

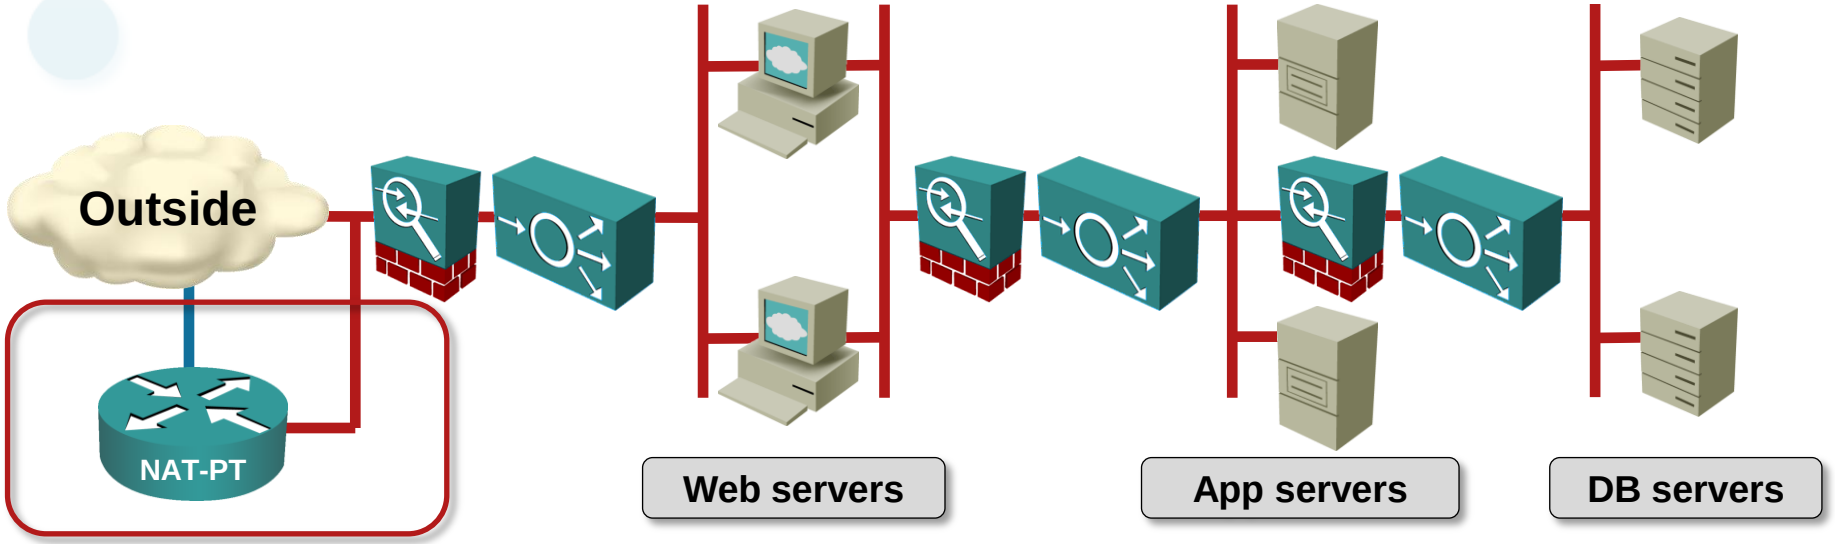

Typical "reasoning"

- Someone high enough asked us to make content available on IPv6
- We don't know a thing about this new protocol
- Deploying IPv6 on load balancers or firewalls is too risky
- Let's insert a NAT64 or NAT-PT box in the outside network

Don't even think about doing this!

Typical Steps

- IPv4 only

Losing control of user experience

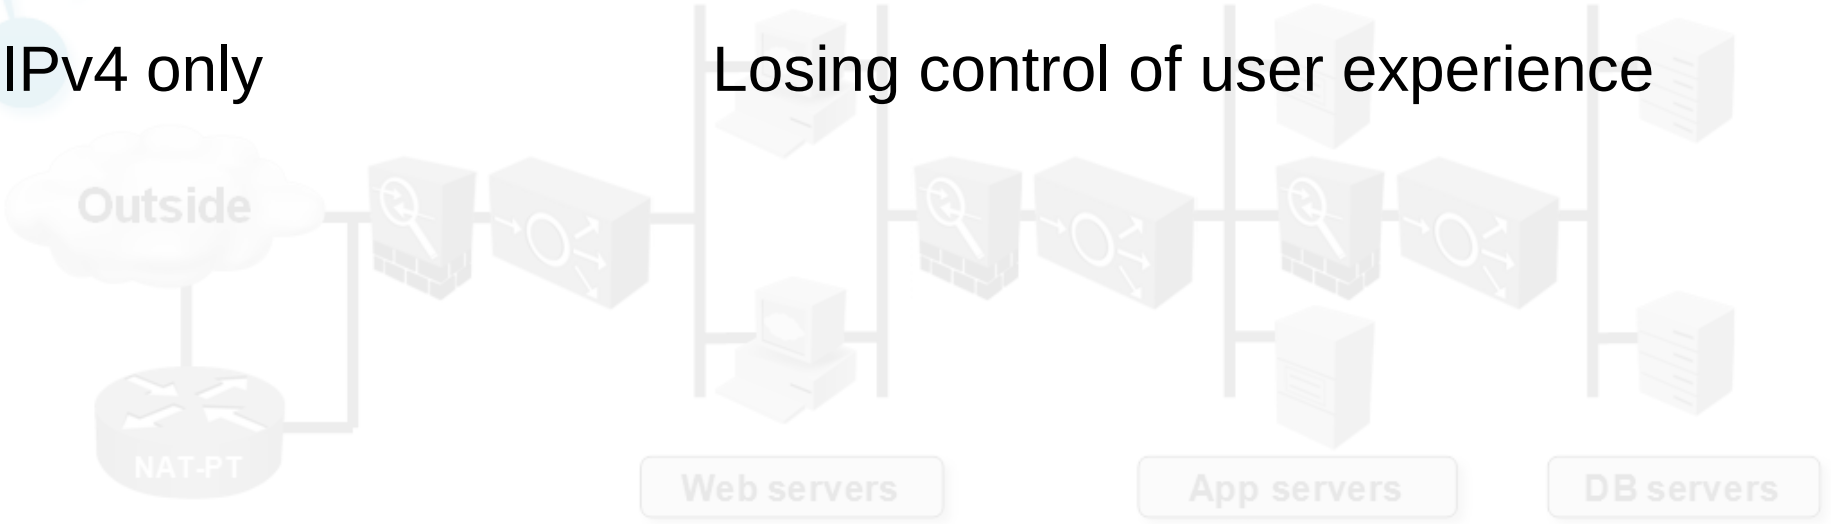

Typical Steps

- IPv4 only
- NAT64

Losing control of user experience

Why are we having performance issues?

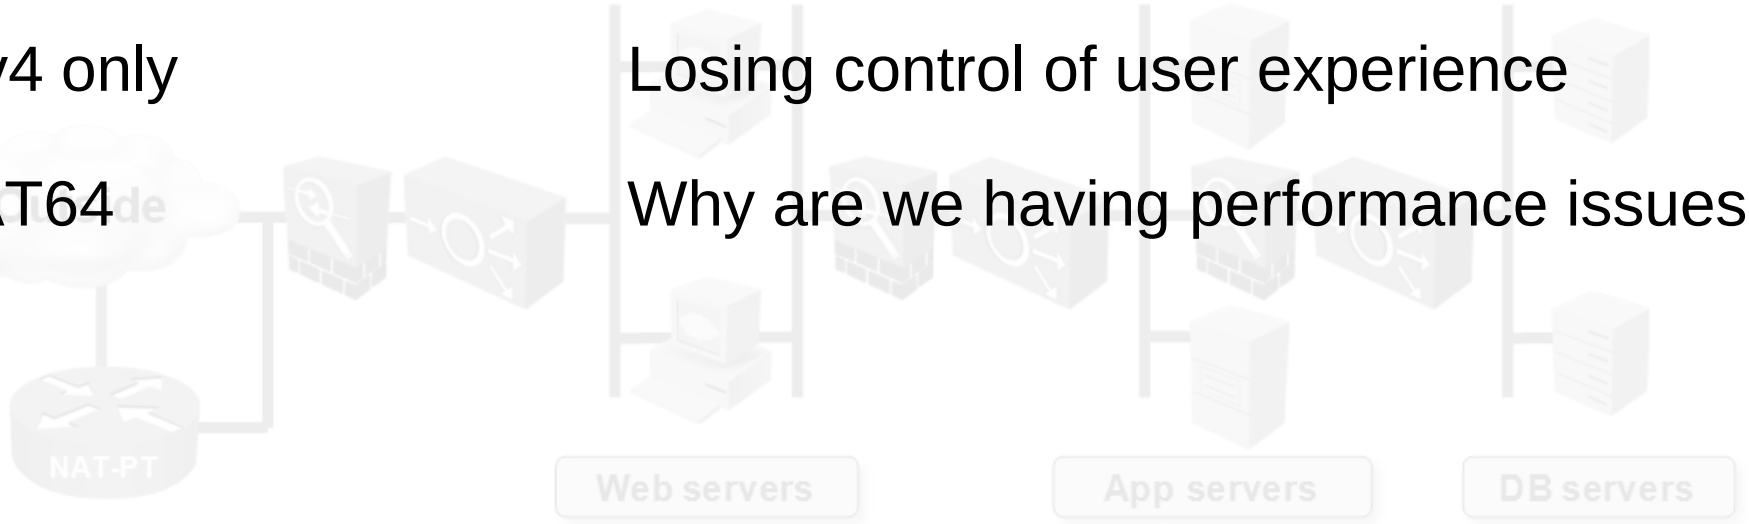

Typical Steps

- IPv4 only
- NAT64
- SLB64

Losing control of user experience

Why are we having performance issues?

Darn, we lost client IP addresses

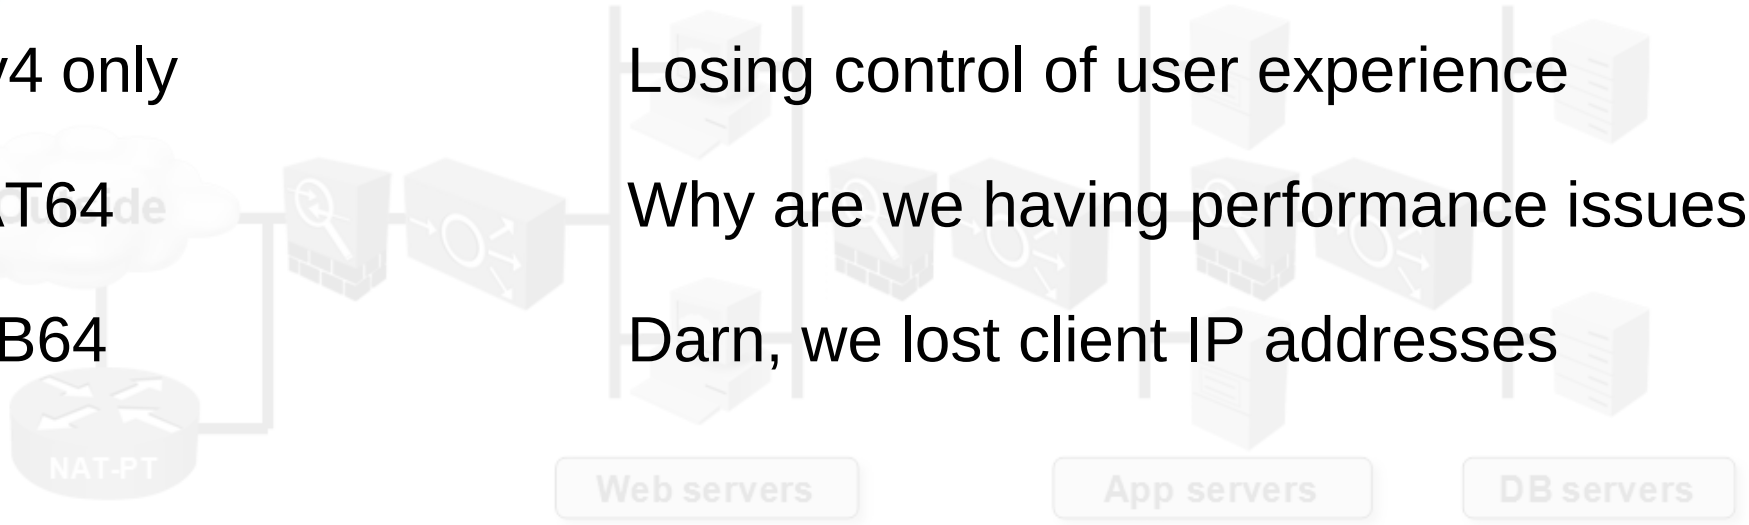

Typical Steps

- IPv4 only
- NAT64
- SLB64
- Dual-stack servers

Losing control of user experience

Why are we having performance issues?

Darn, we lost client IP addresses

Ouch, this is complex

Typical Steps

- IPv4 only — Losing control of user experience
- NAT64 — Why are we having performance issues?
- SLB64 — Darn, we lost client IP addresses
- Dual-stack servers — Ouch, this is complex
- IPv6-only servers with SLB46

Typical Steps

- IPv4 only — Losing control of user experience
- NAT64 — Why are we having performance issues?
- SLB64 — Darn, we lost client IP addresses
- Dual-stack servers — Ouch, this is complex
- IPv6-only servers with SLB46
- IPv6-only data center with NAT46

Typical Steps

- IPv4 only — Losing control of user experience
- NAT64 — Why are we having performance issues?
- SLB64 — Darn, we lost client IP addresses
- Dual-stack servers — Ouch, this is complex
- IPv6-only servers with SLB46
- IPv6-only data center with NAT46
- No IPv4 — ... in a universe far far away

Let Me Recap

- IPv4 only
- NAT64 in DMZ
- SLB64, IPv4-only servers
- SLB44, SLB66, dual-stack servers
- SLB46, IPv6-only servers
- NAT46, SLB66, IPv6-only servers
- IPv6 only

How many migrations do you want to do in the next 5 years?

Let Me Recap

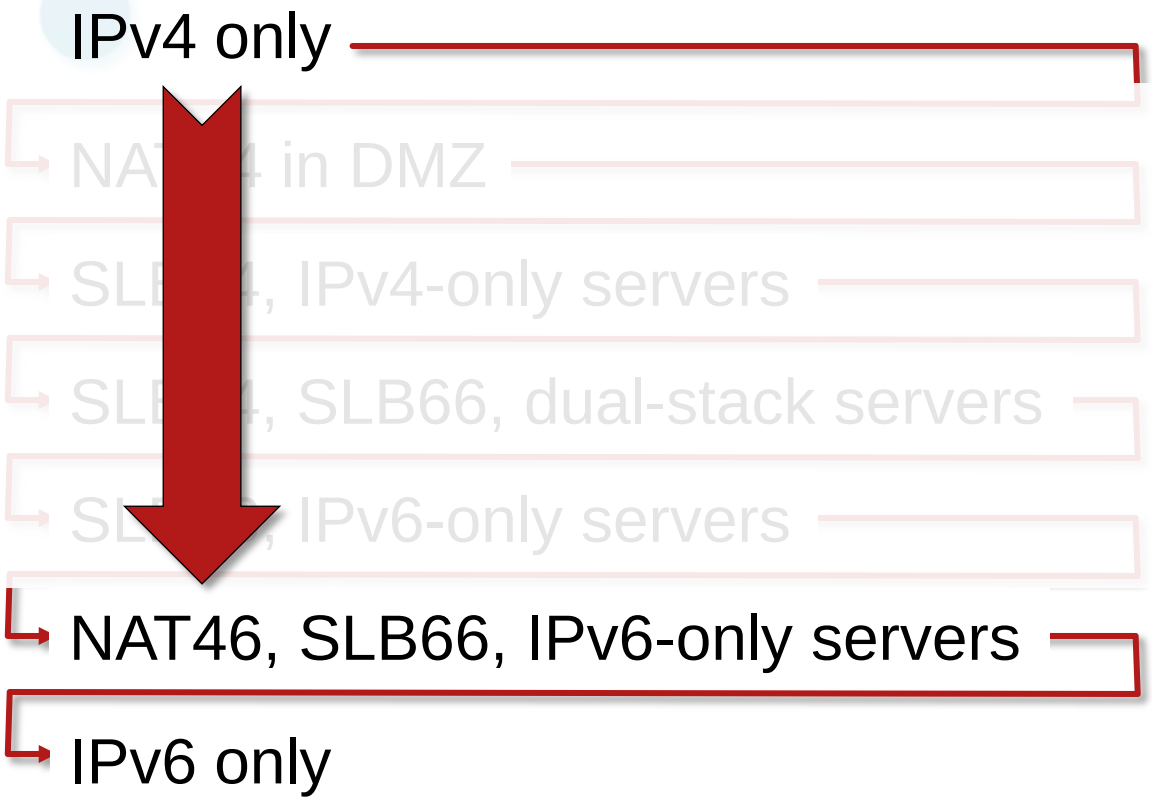

How many migrations do you want to do in the next 5 years?

Skip the Migrations: IPv6-Only Data Center

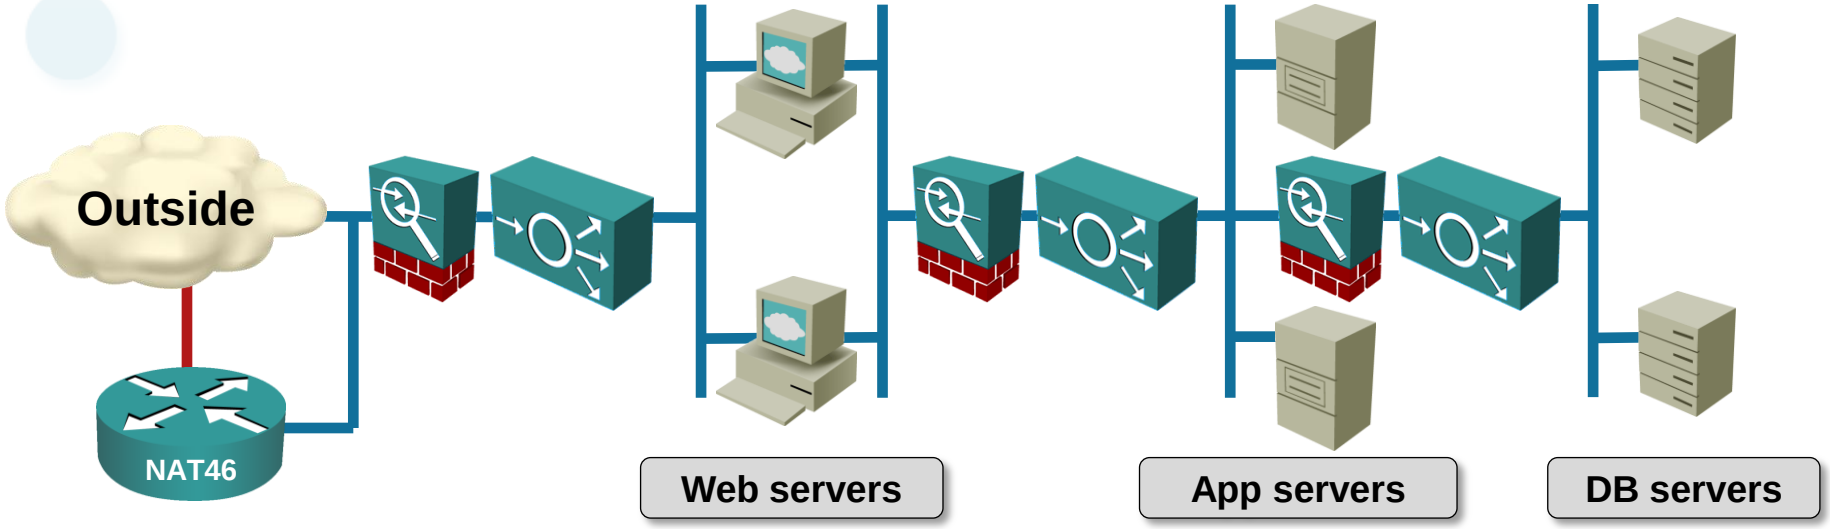

- IPv6-only data center, NAT46 on the edge
- Source IPv4 address mapped into source IPv6 address
- Stateless L3-only translation (easy scaling and redundancy)
- End-to-end visibility is retained, no problems with SSL termination

Can we do it?

IPv6-only Applications

- Some applications will never be IPv6-ready (ex: SNA applications in COBOL)
- Check back-end use of IP addresses
- Make sure you're using DNS names not IP addresses in your code
- Check IPv4 literals in your URLs

Component	IPv6-ready?
Operating system	✓
Web servers	✓
Programming languages	✓
Databases	✓
Clusters	✓
Proxy servers	✓
Caching servers	✓
Load balancers	✓

You're running out of excuses ;)

State of Data Center Infrastructure

Component	Cisco	Juniper	HP	Arista	Brocade	F5
Firewalls	✓	✓	✓			
Load balancers	✓				✓	✓
Core switches	✓	✓	✓	✓	Not on VDX	
ToR switches	✓	✓	✓	✓	✗	

Don’t trust me (or the vendors) – do your own performance tests

Big offenders: major virtualization vendors

- No IPv6-enabled virtual firewall (apart from ip6tables and VM appliances)
- No IPv6 support in VMware vShield or vCloud Director

Hint: Microsoft warns against disabling IPv6 on Windows 2008 servers

Do We Have the Magic NAT46 Box?

Short answer: not yet

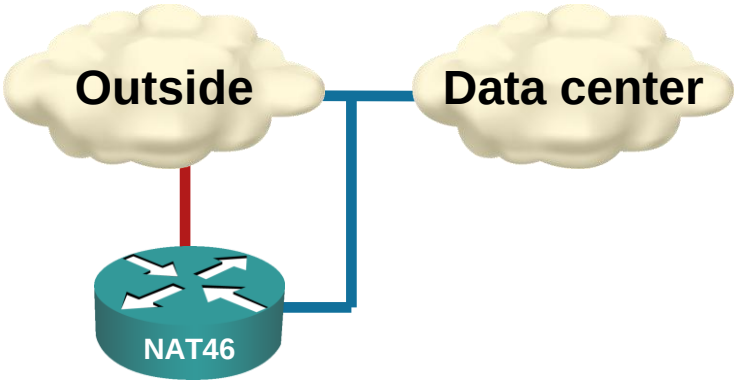

Implementation options:

- Stateless NAT64 ➔ routing challenges
- SLB46 with custom NAT rules ➔ per-session state
- TAYGA on Linux

Sample IPv6-Only Web Site

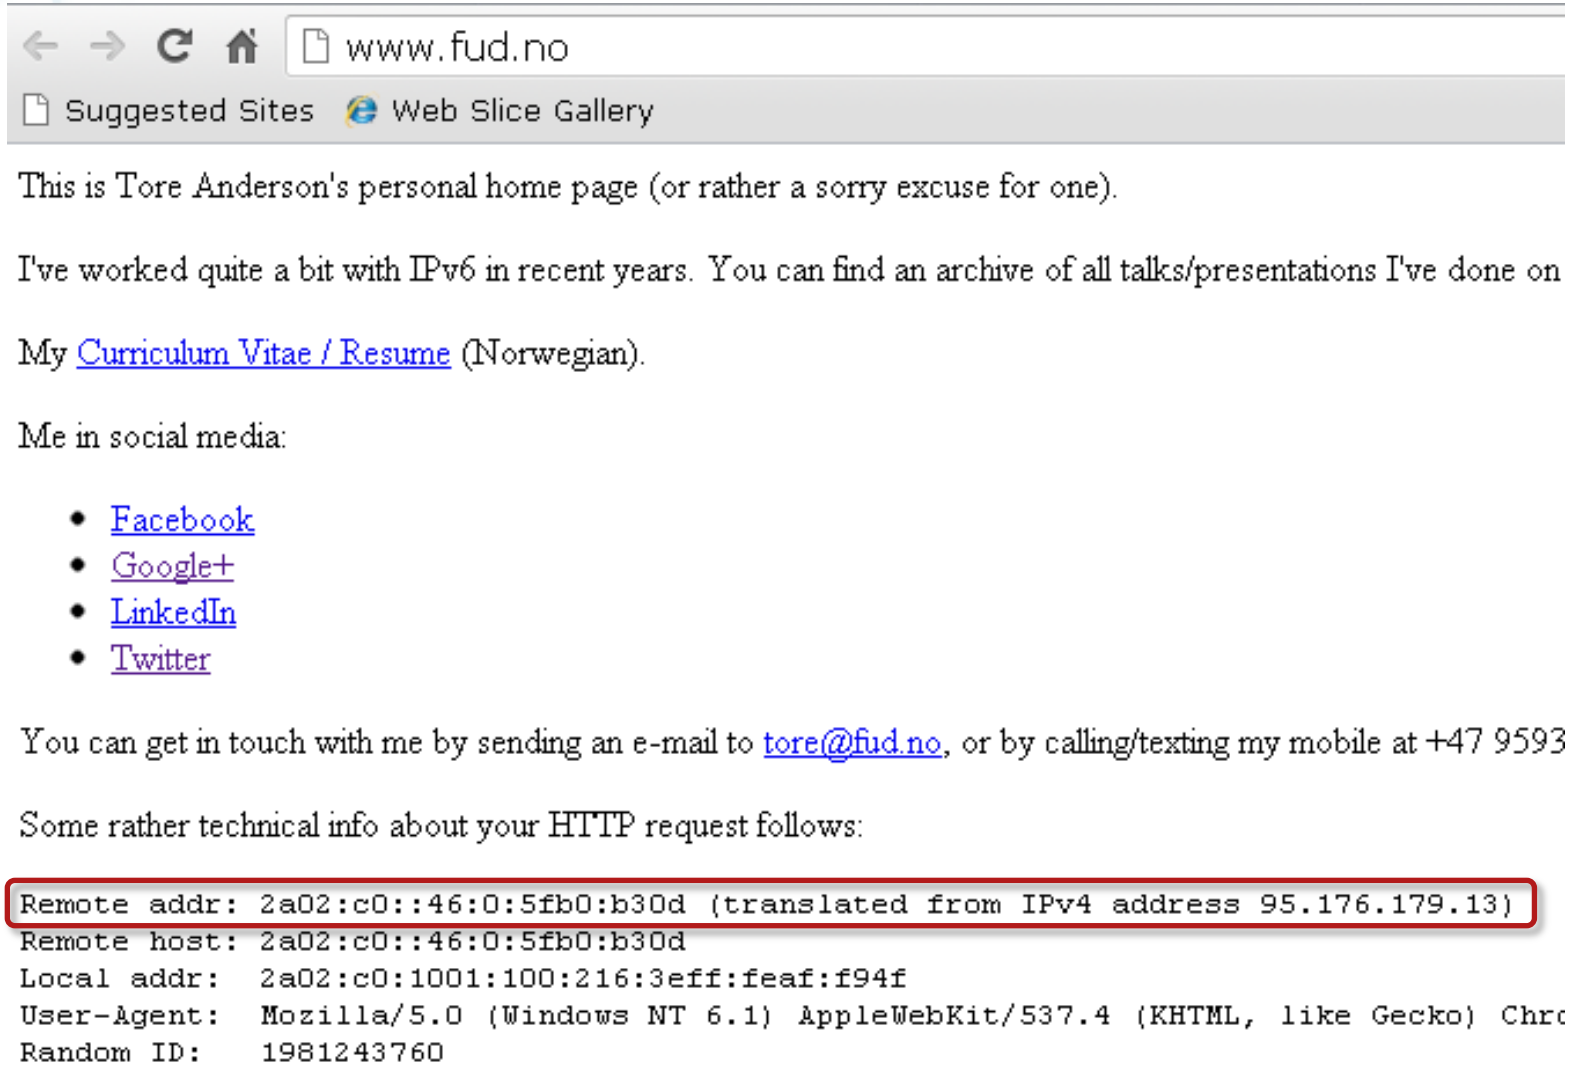

Conclusions

- The path to IPv6-only data center is usually long and winding
- Reduce the number of migration steps
- Aim for early deployment of IPv6-only data center
- We’re not there yet ... but you’ll never be unless you start moving

Special Thanks to

Tore Anderson, the original author of the IPv6-only DC concept

- toreanderson.no/talks/
- ripe64.ripe.net/archives/video/37/
- ripe64.ripe.net/presentations/67-20120417-RIPE64-The_Case_for_IPv6_Only_Data_Centres.pdf

Questions?

Send them to ip@ipspace.net or @ioshints