IPv6 Tunnels in Transition: Strategy and Resources

Tony Hill
Managing Director, IPv6Now

and

Kate Lance
Communications Manager, IPv6Now

IPv6 Tunnels…

Are they the bad guys?

What's a Tunnel?

Tunnelling occurs when a payload protocol is *encapsulated* by a delivery protocol - very common in communications. e.g.

- HTTP tunnelling: HTTP packets encapsulate TCP packets.
- IPv4 packets encapsulate TCP too - could be argued that IPv4 is a form of tunnelling!

What's an IPv6 Tunnel?

IPv6 data packets are:

- encapsulated inside IPv4 packets
- transferred over IPv4 networks
- decapsulated at IPv6 destination

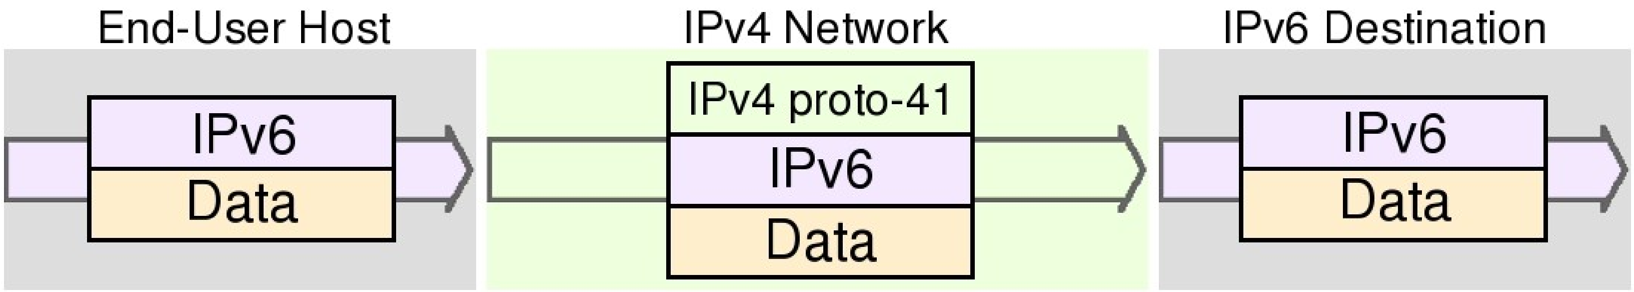

IPv6Now...

Declaration of interest...

We have been doing SLA-backed, carrier-grade tunnels since 2007

Personal Experience

- IPv6Now uses IPv6-only back office services
- I need access to IPv6 ***wherever I am***
- I use IPv6Now tunnels ***every day***
- I use IPv6Now tunnels ***all around the world***
- In four years, no major problems
- Any minor problems, I get fixed!

IPv6 Now

Transitioning to IPv6...

- Dual-stacked IPv4/IPv6 is the widely accepted step towards full native IPv6
- Must commercial-level IPv6 be delivered via dual-stack IPv6 for effective performance?
- Or can IPv6 tunnels be a useful part of a solution for business?

Dual-Stacked IPv6

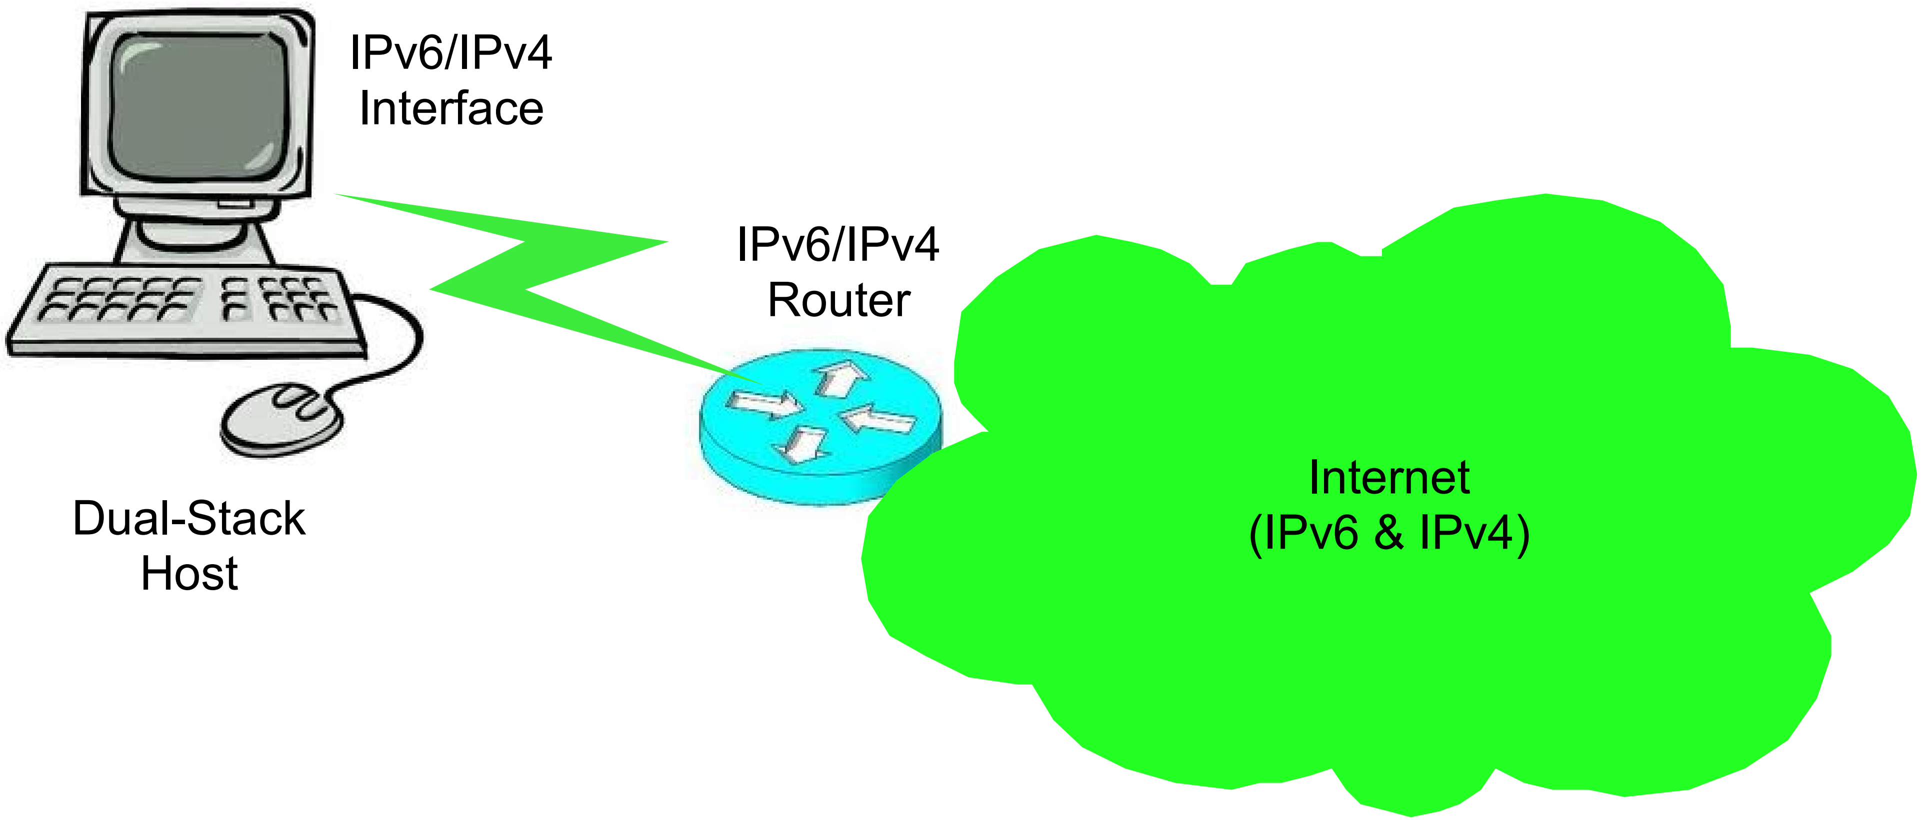

Tunnelled IPv6

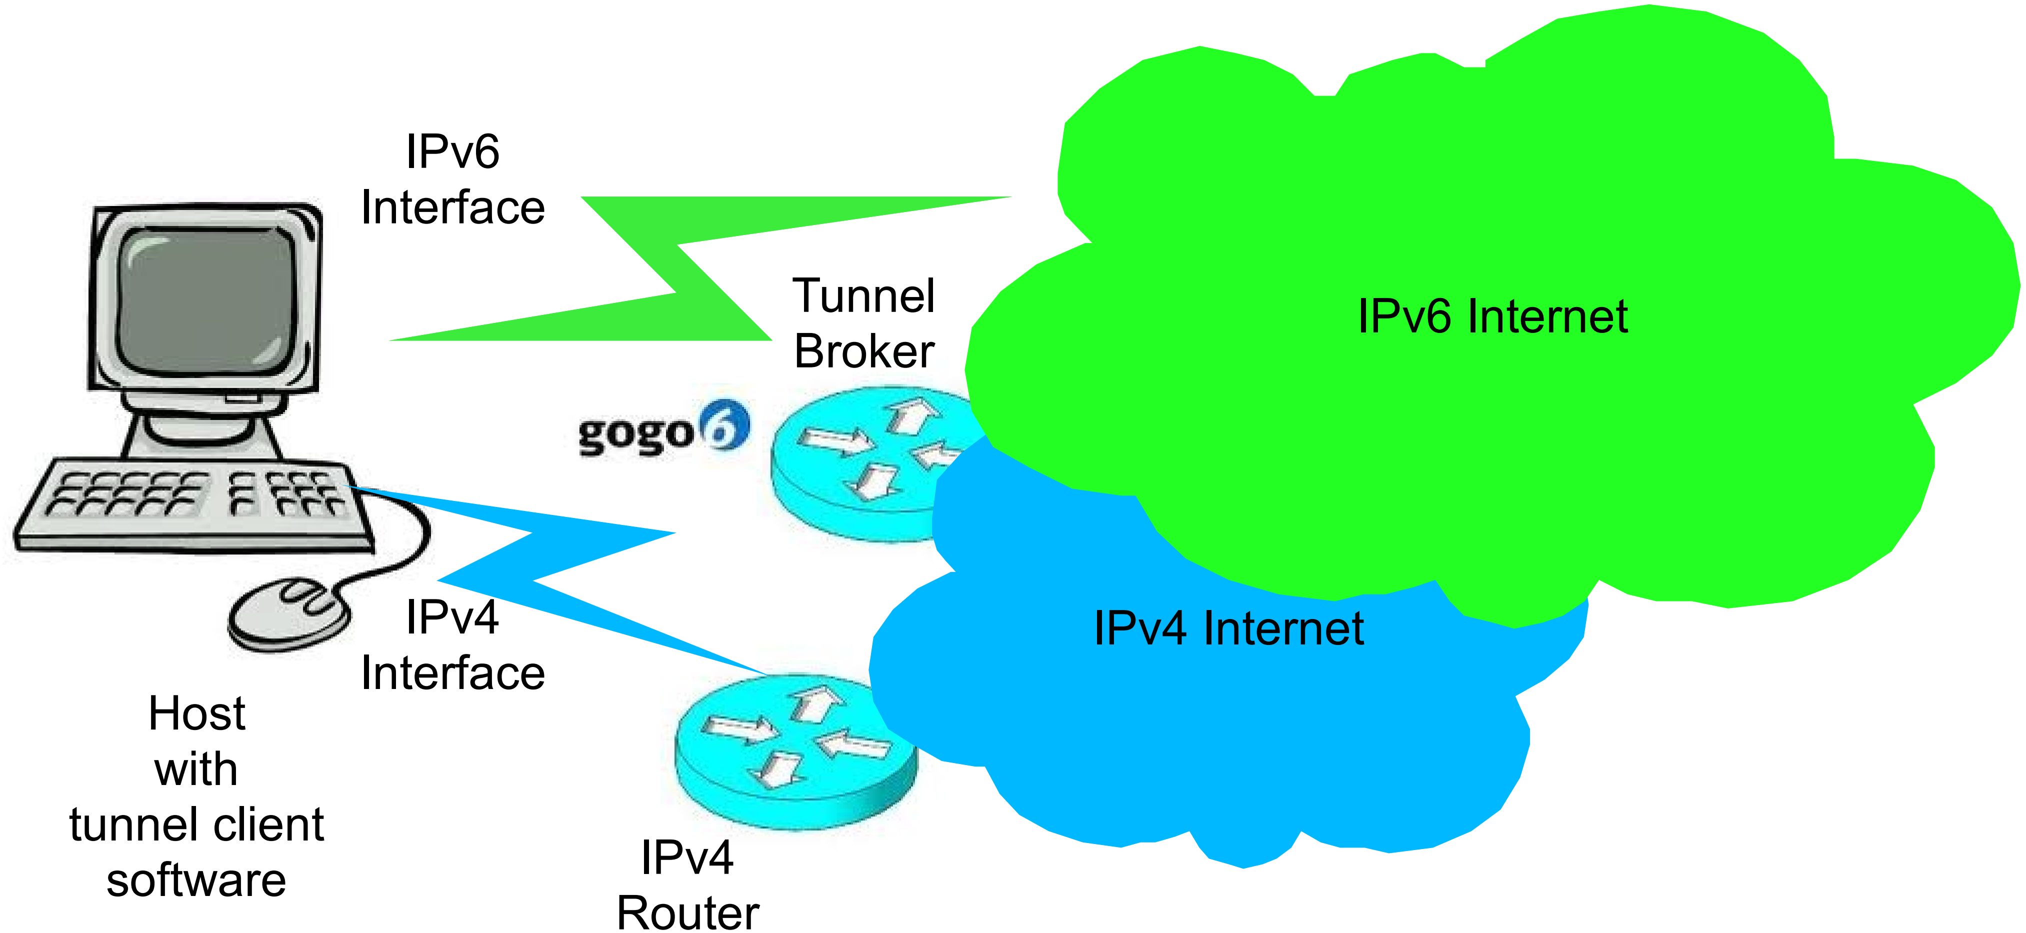

The Dual-Stack Hill of Cost = Pain

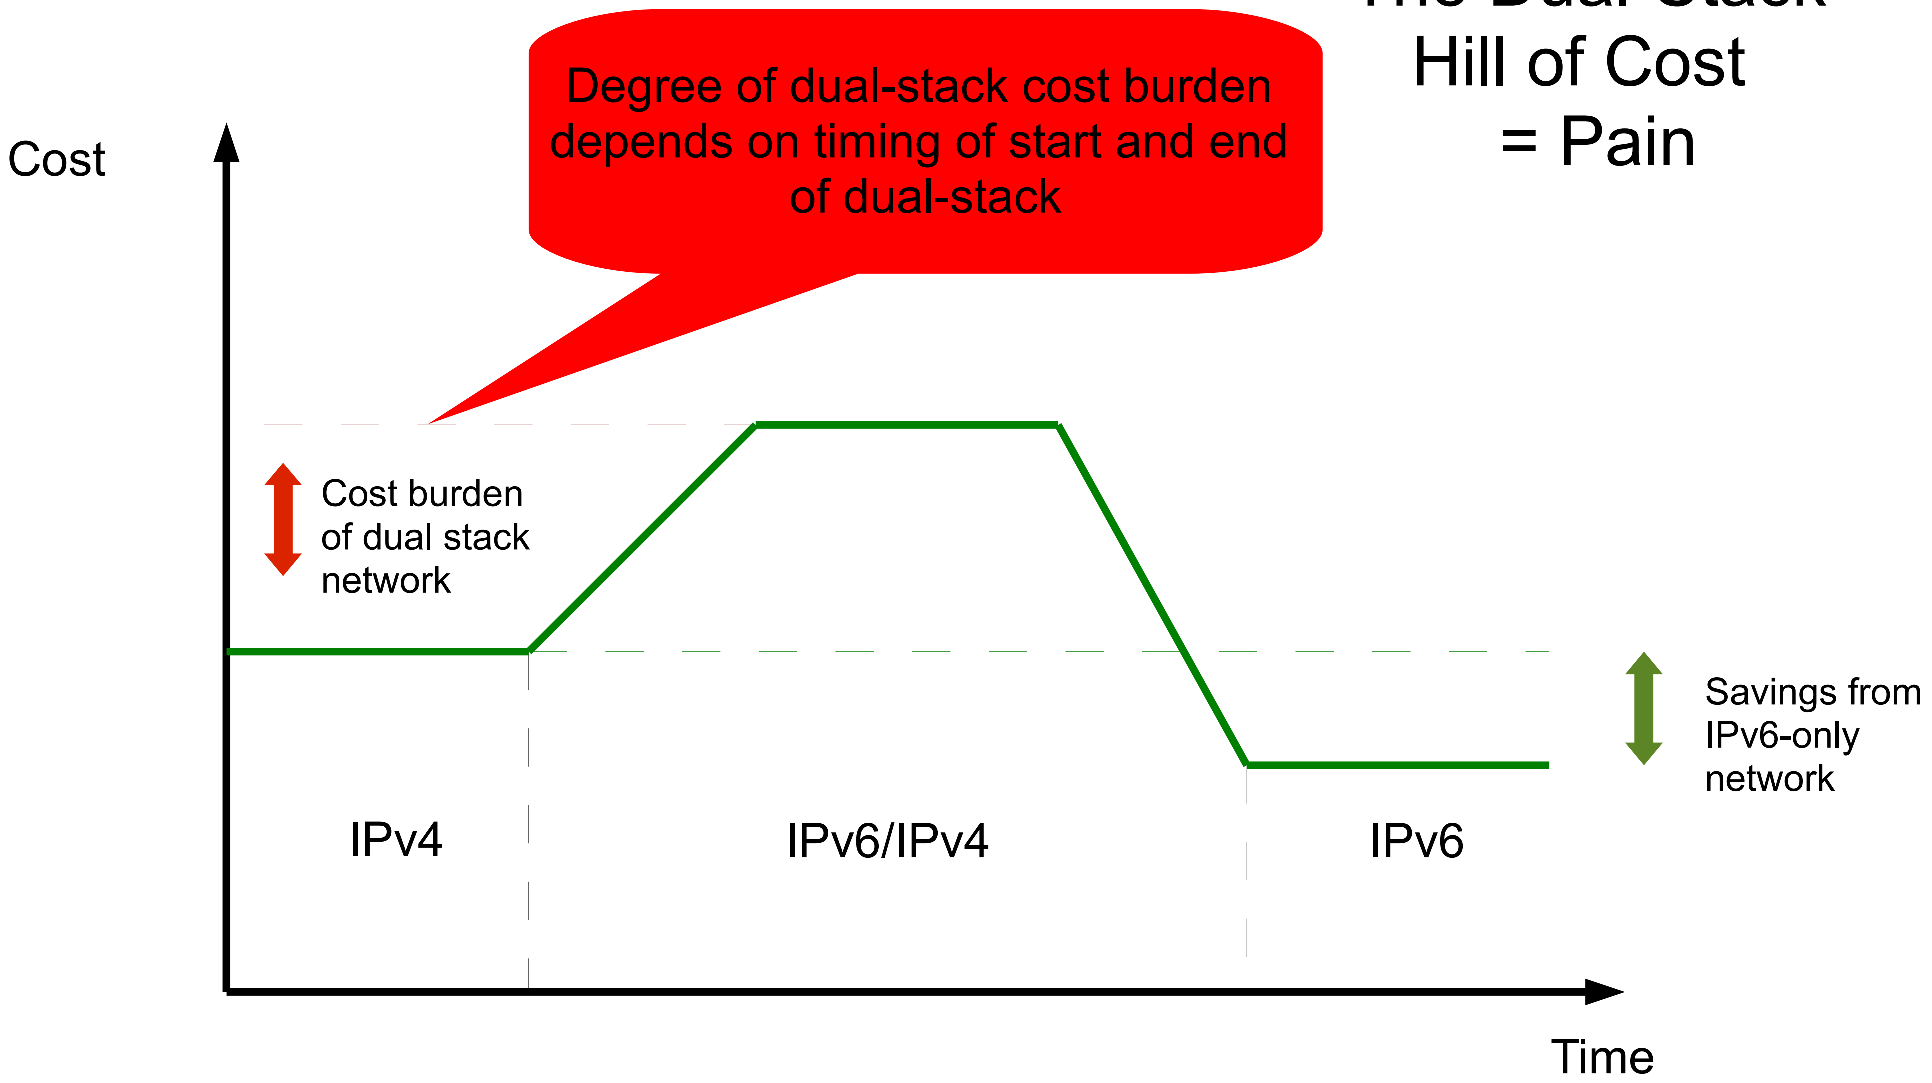

Can tunnels increase or decrease the dual-stack hill of pain?

Hint: *It depends on the type of tunnel*

Main Types of Tunnels

End-User

- 6in4
- TSP
- 6to4
- Teredo

Carrier-Grade

- 6RD
- MPLS-6PE
- TSP (Carrier Grade)
- DS-Lite

e.g. Teredo End-User Tunnel

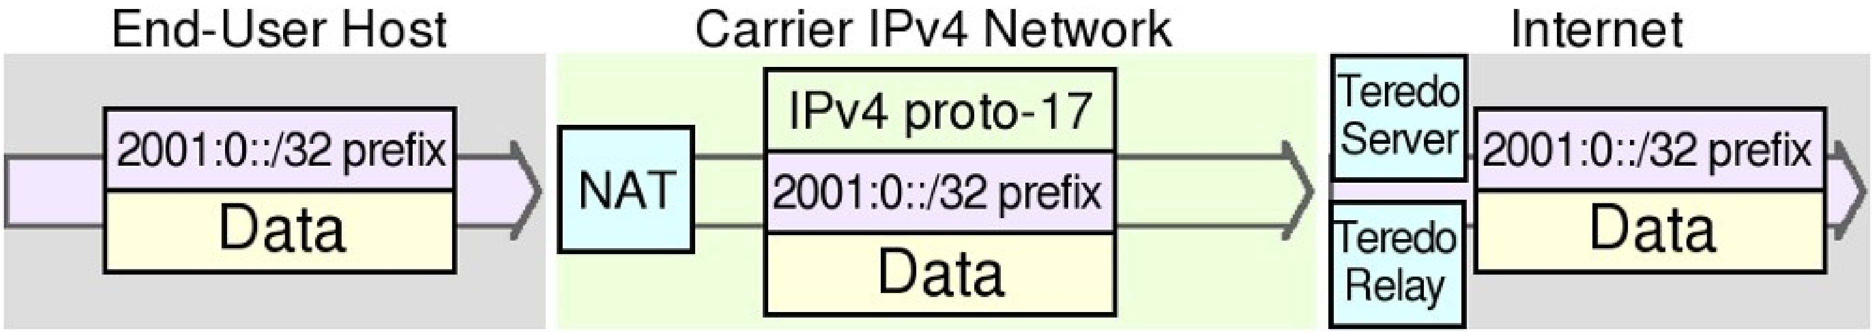

Teredo is an automatic tunnelling technique for hosts behind NATs:

- IPv6 address created from special prefix 2001:0::/32 plus values of IPv4 server & host addresses
- For hosts behind NATs allowing only IPv4 TCP/UDP
- IPv6 packets are encapsulated inside IPv4-UDP packets (protocol 17)
- Teredo servers and relays pass IPv6 traffic to destination
- Problems: relies on 3rd-party Teredo services, needs ICMP (sometimes-blocked) to negotiate NATs.

e.g. Carrier Grade TSP (Tunnel Setup Protocol)

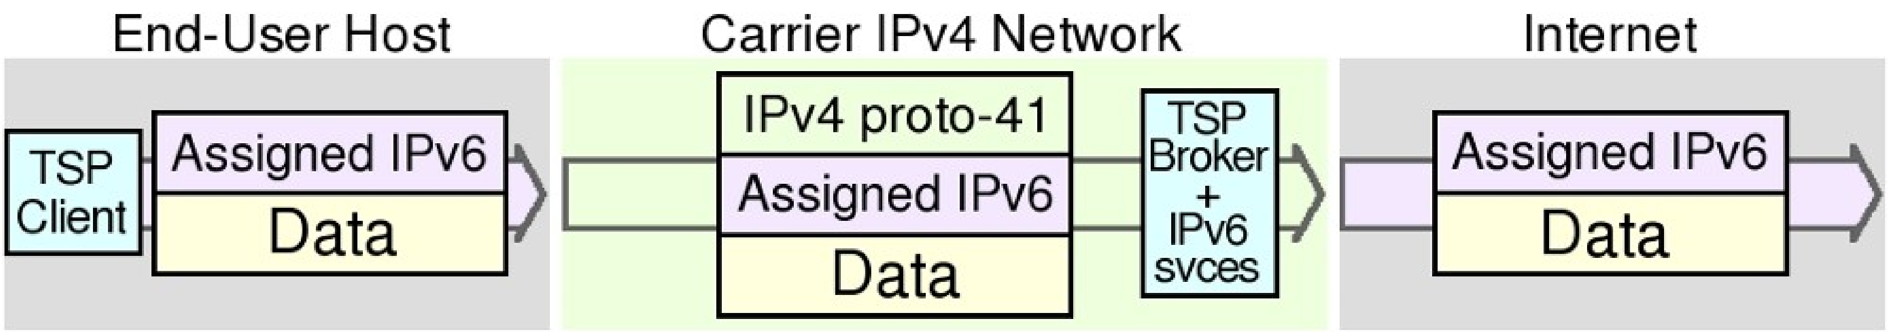

TSP runs on stand-alone Tunnel Broker devices using service provider's own IPv6 address space. SP's Broker assigns an IPv6 address to the authenticated client software endpoints.

Brokers usually have *integrated* security, monitoring, access, authentication, accounting, forward and reverse DNS, NAT functions, multiple encapsulation techniques, etc - avoid the drawbacks of unreliable 3rd-party services.

From *Easy IPv6: The Lookup Book*, 2nd Edition

End-User Tunnel Performance

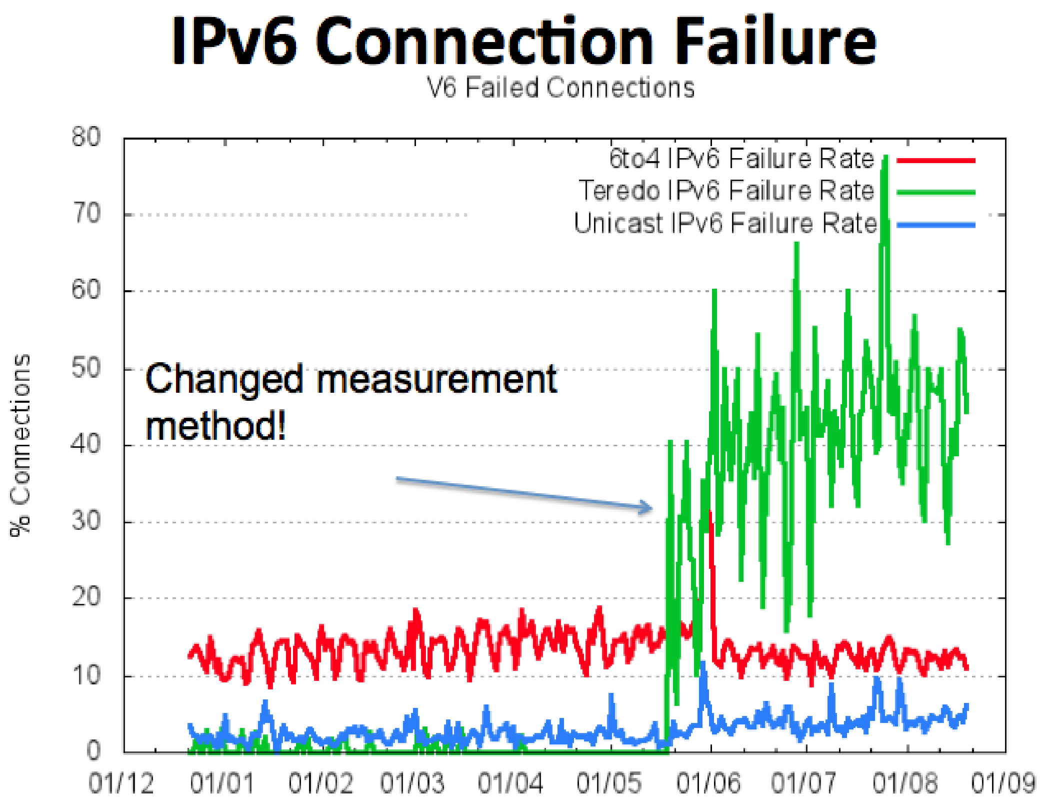

Forcing IPv6 connections to a test site when the end-user system is capable of using IPv6

IPv6 connection failure rates:

6to4 - 12-15%
Teredo - *about 45%!*

Source: Geoff Huston, APNIC32

Carrier-Grade Tunnel Performance

e.g. round trip times to *xkcd.com*

	16 Oct 2011		Mean
	traceroute to xkcd.com		**(ms)**
1	192.168.0.1	(192.168.0.1)	4
2	lns20.mel4.internode.on.net	(150.101.212.44)	22
3	te2-2.cor3.mel4.internode.on.net	(150.101.208.65)	26
4	gi0-0-5.bdr1.mel6.internode.on.net	(150.101.212.42)	200
5	te2-0-0.bdr1.cbr1.internode.on.net	(150.101.160.173)	226
6	te6-0-0.bdr1.syd4.internode.on.net	(150.101.160.170)	217
7	te0-0-0.bdr1.syd7.internode.on.net	(150.101.197.5)	216
8	pos5-0.bdr1.sjc2.internode.on.net	(203.16.213.162)	196
9	equinix-ix.sjc1.us.voxel.net	(206.223.116.4)	228
10	0.ge1-5.tsr1.ord1.us.voxel.net	(208.122.63.238)	279
11	910.te4-3.tsr1.lga3.us.voxel.net	(208.122.44.133)	292
12	0.ae59.tsr1.lga5.us.voxel.net	(208.122.44.202)	286
13	0.ae57.csr2.lga6.us.voxel.net	(208.122.44.210)	256
14	72.26.203.99	(72.26.203.99)	**260**
	traceroute6 to xkcd.com		
1	2406:a000::6:c	(2406:a000::6:c)	30
2	2406:a000::5:1	(2406:a000::5:1)	31
3	2406:a000:ffff:ffff::1	(2406:a000:ffff:ffff::1)	34
4	as4826.ipv6.sydney.pipenetworks.com	(2001:7fa:b::8)	34
5	ge-0-1-0-136.cor02.syd03.nsw.VOCUS.net.au	(2402:7800:0:1::35)	34
6	ten-1-3-0.cor01.syd03.nsw.VOCUS.net.au	(2402:7800:0:1::59)	35
7	pos-0-0-0.bdr01.sjc01.ca.vocusconnect.net	(2402:7800:0:1::82)	189
8	pos-0-1-1.bdr01.pao01.ca.VOCUS.net.au	(2402:7800:100:1::2a)	191
9	paix.ipv6.he.net	(2001:504:d::10)	192
10	paix-ix.pao1.us.voxel.net	(2001:504:d::3b)	194
11	0.ge1-5.tsr1.lga3.us.voxel.net	(2001:48c8::805)	262
12	910.te4-3.tsr1.lga3.us.voxel.net	(2001:48c8::829)	261
13	0.ae1.tsr1.lga5.us.voxel.net	(2001:48c8::822)	262
14	0.ae2.csr2.lga6.us.voxel.net	(2001:48c8::82e)	261
15	2001:48c8:1:d:0:23:5482:d026	(2001:48c8:1:d:0:23:5482:d026)	**259**

IPv4 = 260 ms

← Carrier-Grade Tunnel Broker

IPv6 = 259 ms

Carrier-Grade Tunnel Performance

Sample mean round trip times (ms) to *xkcd.com*

	IPv4	IPv6
16.10.11	260	259
16.10.11	255	270
15.10.11	298	267
30.05.11	296	267

Hardly statistically rigorous! But indicative of real Internet performance with a carrier-grade tunnel broker...

IPv6Now – behind the tunnels

- Simple client software implementation
- IPv6Now's own prefix space
- Multi-path upstream to the IPv6 Internet
- Independent IPv4 and IPv6 links
- Automatic transit of NAT-ed networks
- Integrated security, monitoring, access, authentication, accounting, multiple encapsulation techniques, etc

Commercial Examples

- Studentnet
- Capital Hill Consulting

Business example - Studentnet®

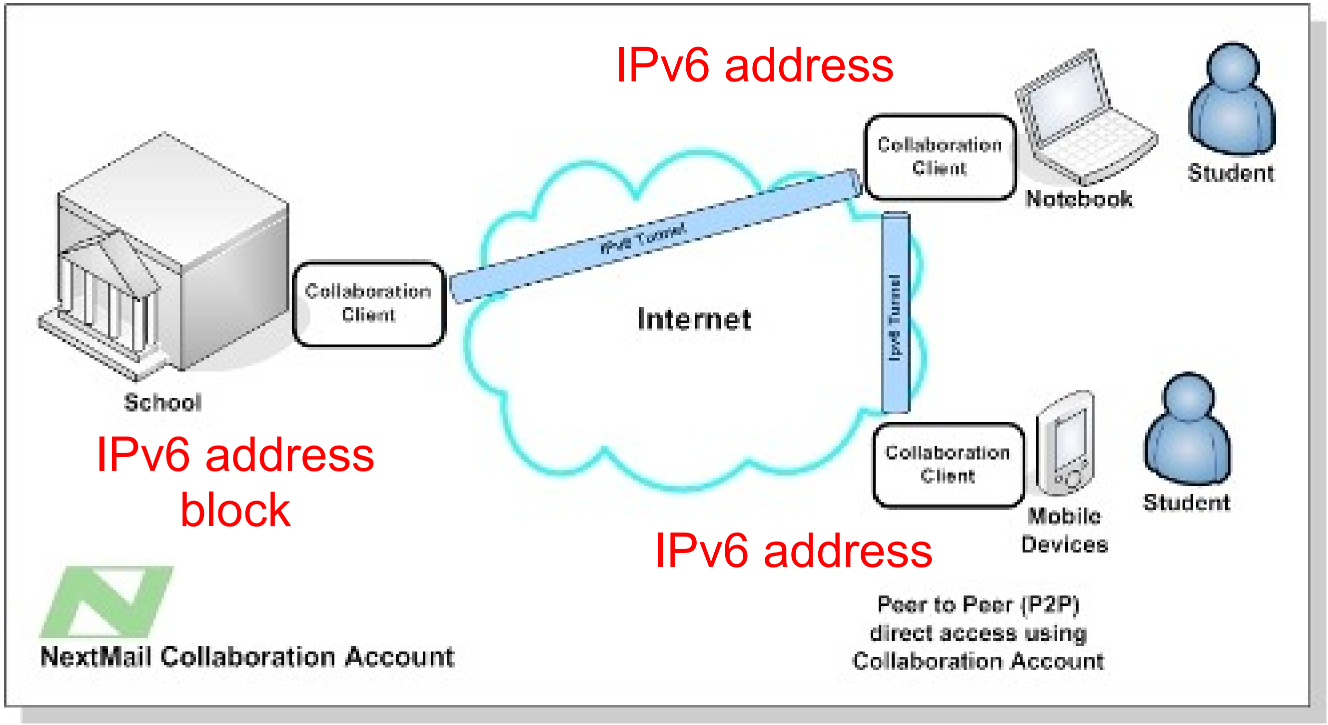

Business example – Capital Hill Consulting

IPv6 Tunnels in Transition

- Dual-stacking of networks is expensive in time, people, resources...
- *Then* you have to move to full IPv6 and discard that dual-stack investment!
- Carrier-grade tunnel brokers allow for immediate, scaleable, drop-in IPv6...
- Strategy: use tunnels during transition and direct scarce resources towards your full IPv6 implementation...

IPv6 Tunnels...

Are they the bad guys?

Maybe they're the good guys!

Thank you

Questions?

Tony Hill
ipv6now.com.au

tony@ipv6now.com.au
kate@ipv6now.com.au